A Reinforcement Learning Model Of Selective Visual Attention

Build Machine Learning models with a sound statistical understanding. About This Book Learn about the statistics behind powerful predictive models with p-value, ANOVA, and F- statistics. Implement statistical computations programmatically for supervised and unsupervised learning through K-means clustering. Master the statistical aspect of Machine Learning with the help of this example-rich guide to R and Python. Who This Book Is For This book is intended for developers with little to no background in statistics, who want to implement Machine Learning in their systems. Some programming knowledge in R or Python will be useful. What You Will Learn Understand the Statistical and Machine Learning fundamentals necessary to build models Understand the major differences and parallels between the statistical way and the Machine Learning way to solve problems Learn how to prepare data and feed models by using the appropriate Machine Learning algorithms from the more-than-adequate R and Python packages Analyze the results and tune the model appropriately to your own predictive goals Understand the concepts of required statistics for Machine Learning Introduce yourself to necessary fundamentals required for building supervised & unsupervised deep learning models Learn reinforcement learning and its application in the field of artificial intelligence domain In Detail Complex statistics in Machine Learning worry a lot of developers. Knowing statistics helps you build strong Machine Learning models that are optimized for a given problem statement. This book will teach you all it takes to perform complex statistical computations required for Machine Learning. You will gain information on statistics behind supervised learning, unsupervised learning, reinforcement learning, and more. Understand the real-world examples that discuss the statistical side of Machine Learning and familiarize yourself with it. You will also design programs for performing tasks such as model, parameter fitting, regression, classification, density collection, and more. By the end of the book, you will have mastered the required statistics for Machine Learning and will be able to apply your new skills to any sort of industry problem. Style and approach This practical, step-by-step guide will give you an understanding of the Statistical and Machine Learning fundamentals you'll need to build models.

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.

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.

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

Reinforcement learning models learn the optimal policy by interacting with the environment and observing the states and rewards. If the rewards that the model observes are noisy then learning an optimal policy becomes a difficult task. We present an approach that can enhance the performance of reinforcement learning models in the presence of noisy rewards. Along with the standard reinforcement learning model, we propose to use a noise filter which estimates the true reward that the model should receive. The

noise filter is designed using a non-linear approximator. Through various experiments we demonstrate that this approach improves the performance of the model in the presence of noisy rewards

Reinforcement learning has developed as a successful learning approach for domains that are not fully understood and that are too complex to be described in closed form. However, reinforcement learning does not scale well to large and continuous problems. Furthermore, acquired knowledge specific to the learned task, and transfer of knowledge to new tasks is crucial. In this book the author investigates whether deficiencies of reinforcement learning can be overcome by suitable abstraction methods. He discusses various forms of spatial abstraction, in particular qualitative abstraction, a form of representing knowledge that has been thoroughly investigated and successfully applied in spatial cognition research. With his approach, he exploits spatial structures and structural similarity to support the learning process by abstracting from less important features and stressing the essential ones. The author demonstrates his learning approach and the transferability of knowledge by having his system learn in a virtual robot simulation system and consequently transfer the acquired knowledge to a physical robot. The approach is influenced by findings from cognitive science. The book is suitable for researchers working in artificial intelligence, in particular knowledge representation, learning, spatial cognition, and robotics.

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

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. 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 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. Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

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.

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.

Get to grips with the basics of Keras to implement fast and efficient deep-learning models About This Book Implement various deep-learning algorithms in Keras and see how deep-learning can be used in games See how various deep-learning models and practical use-cases can be implemented using Keras A practical, hands-on guide with real-world examples to give you a strong foundation in Keras Who This Book Is For If you are a data scientist with experience in machine learning or an AI programmer with some exposure to neural networks, you will find this book a useful entry point to deep-learning with Keras. A knowledge of Python is required for this book. What You Will Learn Optimize step-by-step functions on a large neural network using the Backpropagation Algorithm Fine-tune a neural network to improve the quality of results Use deep learning for image and audio processing Use Recursive Neural Tensor Networks (RNTNs) to outperform standard word embedding in special cases Identify problems for which Recurrent Neural Network (RNN) solutions are suitable Explore the process required to implement Autoencoders Evolve a deep neural network using reinforcement learning In Detail This book starts by introducing you to supervised learning algorithms such as simple linear regression, the classical multilayer perceptron and more sophisticated deep convolutional networks. You will also explore image processing with recognition of hand written digit images, classification of images into different categories, and advanced objects recognition with related image annotations. An example of identification of salient points for face detection is also provided. Next you will be introduced to Recurrent Networks, which are optimized for processing sequence data such as text, audio or time series. Following that, you will learn about unsupervised learning algorithms such as Autoencoders and the very popular Generative Adversarial Networks (GAN). You will also explore non-traditional uses of neural networks as Style Transfer. Finally, you will look at Reinforcement Learning and its application to AI game playing, another popular direction of research and application of neural networks. Style and approach This book is an easy-to-follow guide full of examples and real-world applications to help you gain an in-depth understanding of Keras. This book will showcase more than twenty working Deep Neural Networks coded in Python using Keras.

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.

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

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

Get to grips with the essentials of deep learning by leveraging the power of PythonKey Features Your one-stop solution to get started with the essentials of deep learning and neural network modeling Train different kinds of neural networks to tackle various problems in Natural Language Processing, computer vision, speech recognition, and more Covers popular Python libraries such as Tensorflow, Keras, and more, along with tips on training, deploying and optimizing your deep learning models in the best possible manner Deep Learning a trending topic in the field of Artificial Intelligence today and can be considered to be an advanced form of machine learning, which is quite tricky to master. This book will help you take your first steps in training efficient deep learning models and applying them in various practical scenarios. You will model, train, and deploy different kinds of neural networks such as Convolutional Neural Network, Recurrent Neural Network, and will see some of their applications in real-world domains including computer vision, natural language processing, speech recognition, and so on. You will build practical projects such as chatbots, implement reinforcement learning to build smart games, and develop expert systems for image captioning and processing. Popular Python library such as TensorFlow is used in this book to build the models. This book also covers solutions for different problems you might come across while training models, such as noisy datasets, small datasets, and more. This book does not assume any prior knowledge of deep learning. By the end of this book, you will have a firm understanding of the basics of deep learning and neural network modeling, along with their practical applications.What you will learn Get to grips with the core concepts of deep learning and neural networks Set up deep learning library such as TensorFlow Fine-tune your deep learning models for NLP and Computer Vision applications Unify different information sources, such as images, text, and speech through deep learning Optimize and fine-tune your deep learning models for better performance Train a deep reinforcement learning model that plays a game better than humans Learn how to make your models get the best out of your GPU or CPU Who This Book Is For Aspiring data scientists and machine learning experts who have limited or no exposure to deep learning will find this book to be very useful. If you are looking for a resource that gets you up and running with the fundamentals of deep learning and neural networks, this book is for you. As the models in the book are trained using the popular Python-based libraries such as Tensorflow and Keras, it would be useful to have sound programming knowledge of Python. Table of Contents 1. What is artificial intelligence 2. Why is the artificial intelligence important ? 3. Applications of Machine Learning 4. Semantics, Probability and IA 5. Numerical Computation 6. Sequence Modeling, Recurrent and Recursive Nets 7. Autoencoders 8. Markov Chains, Monte Carlo Methods, and Machine Learning

Motivated learning is an emerging research field in artificial intelligence and cognitive modelling. Computational models of motivation extend reinforcement learning to adaptive, multitask learning in complex, dynamic environments – the goal being to understand how machines can develop new skills and achieve goals that were not predefined by human engineers. In particular, this book describes how motivated reinforcement learning agents can be used in computer games for the design of non-player characters that can adapt their behaviour in response to unexpected changes in their environment. This book covers the design, application and evaluation of computational models of motivation in reinforcement learning. The authors start with overviews of motivation and reinforcement learning, then describe models for motivated reinforcement learning. The performance of these models is demonstrated by applications in simulated game scenarios and a live, open-ended virtual world. Researchers in artificial intelligence, machine learning and artificial life will benefit from this book, as will practitioners working on complex, dynamic systems – in particular multiuser, online games.

A reinforcement learning model of gaze following.

Data Analytics - 7 BOOK BUNDLE!! Book 1: Data Analytics For Beginners In this book you will learn: What is Data Analytics Types of Data Analytics Evolution of Data Analytics Big Data Defined Data Mining Data Visualization Cluster Analysis And of course much more! Book 2: Deep Learning With Keras In this book you will learn: Deep Neural Network Neural Network Elements Keras Models Sequential Model Functional API Model Keras Layers Core Keras Layers Convolutional Keras Layers Recurrent Keras Layers Deep Learning Algorithms Supervised Learning Algorithms Applications of Deep Learning Models Automatic Speech and Image Recognition Natural Language Processing And of course much more! Book 3: Analyzing Data With Power BI In this book you will learn: Basics of data analysis processes Fundamental data analysis algorithms Basic of data and text mining, data visualization, and business intelligence Techniques used for analysing quantitative data Basic data analysis tasks Conceptual, logical, and physical data models Power BI service and data modelling Creating reports and visualizations in Power BI And of course much more! Book 4: Reinforcement Learning With Python In this book you will learn: Types of fundamental machine learning algorithms in comparison to reinforcement learning Essentials of reinforcement learning process Marko decision processes and basic parameters How to integrate reinforcement learning algorithm using OpenAI Gym How to integrate Monte Carlo methods for prediction Monte Carlo tree search And much, much more... Book 5: Artificial Intelligence Python In this book you will learn: Different artificial intelligence approaches and goals How to define AI system Basic AI techniques Reinforcement learning And much, much more... Book 6: Text Analytics With Python In this book you will learn: Text analytics process How to build a corpus and analyze sentiment Named entity extraction with Groningen meaning bank corpus How to train your system Getting started with NLTK How to search syntax and tokenize sentences Automatic text summarization Stemming word and topic modeling with NLTK And much, much more... Book 7: Convolutional Neural Networks In Python In this book you will learn: Architecture of convolutional neural networks Solving computer vision tasks using convolutional neural networks Python and computer vision Automatic image and speech recognition Theano and TenroeFlow image recognition And of course much more! Download this book bundle NOW and SAVE money!!

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.

Human-in-the-Loop Machine Learning lays out methods for humans and machines to work together effectively. Summary Most machine learning systems that are deployed in the world today learn from human feedback. However, most machine learning courses focus almost exclusively on the algorithms, not the human-computer interaction part of the systems. This can leave a big knowledge gap for data scientists working in real-world machine learning, where data scientists spend more time on data management than on building algorithms. Human-in-the-Loop Machine Learning is a practical guide to optimizing the entire machine learning process, including techniques for annotation, active learning, transfer learning, and using machine learning to optimize every step of the process. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Machine learning applications perform better with human feedback. Keeping the right people in the loop improves the accuracy of models, reduces errors in data, lowers costs, and helps you ship models faster. About the book Human-in-the-Loop Machine Learning lays out methods for humans and machines to work together effectively. You'll find best practices on selecting sample data for human feedback, quality control for human annotations, and designing annotation interfaces. You'll learn to create training data for labeling, object detection, and semantic segmentation, sequence labeling, and more. The book starts with the basics and progresses to advanced techniques like transfer learning and self-supervision within annotation workflows. What's inside Identifying the right training and evaluation data Finding and managing people to annotate data Selecting annotation quality control strategies Designing interfaces to improve accuracy and efficiency About the author Robert (Munro) Monarch is a data scientist and engineer who has built machine learning data for companies such as Apple, Amazon, Google, and IBM. He holds a PhD from Stanford. Robert holds a PhD from Stanford focused on Human-in-the-Loop machine learning for healthcare and disaster response, and is a disaster response professional in addition to being a machine learning professional. A worked example throughout this text is classifying disaster-related messages from real disasters that Robert has helped respond to in the past. Table of Contents PART 1 - FIRST STEPS 1 Introduction to human-in-the-loop machine learning 2 Getting started with human-in-the-loop machine learning PART 2 - ACTIVE LEARNING 3 Uncertainty sampling 4 Diversity sampling 5 Advanced active learning 6 Applying active learning to different machine learning tasks PART 3 - ANNOTATION 7 Working with the people annotating your data 8 Quality control for data annotation 9 Advanced data annotation and augmentation 10 Annotation quality for different machine learning tasks PART 4 - HUMAN–COMPUTER INTERACTION FOR MACHINE LEARNING 11 Interfaces for data annotation 12 Human-in-the-loop machine learning products

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.

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

This book introduces machine learning methods in finance. It presents a unified treatment of machine learning and various statistical and computational disciplines in quantitative finance, such as financial econometrics and discrete time stochastic control, with an emphasis on how theory and hypothesis tests inform the choice of algorithm for financial data modeling and decision making. With the trend towards increasing computational resources and larger datasets, machine learning has grown into an important skillset for the finance industry. This book is written for advanced graduate students and academics in financial econometrics, mathematical finance and applied statistics, in addition to quants and data scientists in the field of quantitative finance. Machine Learning in Finance: From Theory to Practice is divided into

three parts, each part covering theory and applications. The first presents supervised learning for cross-sectional data from both a Bayesian and frequentist perspective. The more advanced material places a firm emphasis on neural networks, including deep learning, as well as Gaussian processes, with examples in investment management and derivative modeling. The second part presents supervised learning for time series data, arguably the most common data type used in finance with examples in trading, stochastic volatility and fixed income modeling. Finally, the third part presents reinforcement learning and its applications in trading, investment and wealth management. Python code examples are provided to support the readers' understanding of the methodologies and applications. The book also includes more than 80 mathematical and programming exercises, with worked solutions available to instructors. As a bridge to research in this emergent field, the final chapter presents the frontiers of machine learning in finance from a researcher's perspective, highlighting how many well-known concepts in statistical physics are likely to emerge as important methodologies for machine learning in finance.

A Reinforcement Learning Model of Gaze Following

Buy the Paperback Version of this Book and get the Kindle Book Version for FREEDo want to learn how machine learning and neural networks work quickly and simply? Do you want to know how to build a machine learning model and you have no programming skill? Do you want to get started with learning data science? This book is going to guide you to the basics and the principles behind machine learning. Machine learning is an active research domain and includes several different approaches. This book is going to help you understand the different approaches of machine learning and neural networks. It will guide you through the steps you need to build a machine learning model. Machine learning implies programming. This book will teach you Python programming. This book does not require any pre-programming skills. It will help to get you started in Python programming, as well as how to use Python libraries to analyze data and apply machine learning. Overall, this book is a go-to guide for getting started in machine learning modeling using Python programming. Once you get through the book, you will be able to develop your own machine learning models using Python. Through this book, you will learn: - Principles of machine learning - Types of machine learning: supervised, unsupervised, semi-supervised, and reinforcement learning - Advantages of each type of machine learning - Principle and types of neural networks - Steps to develop and fit artificial neural network model - Getting started and installing Python - Tools and platforms for Python programming - How to use pandas, NumPy and matplotlib Python libraries - How to develop a simple linear and logistic machine learning model - How to develop and train a multi-layer artificial neural network two ways: from scratch and using the Python libraries Even if you don't have any background in machine learning and Python programming, this book will give you the tools to develop machine learning models. Would you like to know more? Scroll to the top of the page and select the BUY NOW button.

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.

Reinforcement learning is a mathematical framework for developing computer agents that can learn an optimal behavior by relating generic reward signals with its past actions. With numerous successful applications in business intelligence, plant control, and gaming, the RL framework is ideal for decision making in unknown environments with large amounts of data. Supplying an up-to-date and accessible introduction to the field, Statistical Reinforcement Learning: Modern Machine Learning Approaches presents fundamental concepts and practical algorithms of statistical reinforcement learning from the modern machine learning viewpoint. It covers various types of RL approaches, including model-based and model-free approaches, policy iteration, and policy search methods. The book covers approaches recently introduced in the data mining and machine learning fields to provide a systematic bridge between RL and data mining/machine learning researchers. It presents state-of-the-art results, including dimensionality reduction in RL and risk-sensitive RL. Numerous illustrative examples are included to help readers understand the intuition and usefulness of reinforcement learning techniques. Book jacket.

Take a systematic approach to understanding the fundamentals of machine learning and deep learning from the ground up and how they are applied in practice. You will use this comprehensive guide for building and deploying learning models to address complex use cases while leveraging the computational resources of Google Cloud Platform. Author Ekaba Bisong shows you how machine learning tools and techniques are used to predict or classify events based on a set of interactions between variables known as features or attributes in a particular dataset. He teaches you how deep learning extends the machine learning algorithm of neural networks to learn complex tasks that are difficult for computers to perform, such as recognizing faces and understanding languages. And you will know how to leverage cloud computing to accelerate data science and machine learning deployments. Building Machine Learning and Deep Learning Models on Google Cloud Platform is divided into eight parts that cover the fundamentals of machine learning and deep learning, the concept of data science and cloud services, programming for data science using the Python stack, Google Cloud Platform (GCP) infrastructure and products, advanced analytics on GCP, and deploying end-to-end machine learning solution pipelines on GCP. What You’ll Learn Understand the principles and fundamentals of machine learning and deep learning, the algorithms, how to use them, when to use them, and how to interpret your results Know the programming concepts relevant to machine and deep learning design and development using the Python stack Build and interpret machine and deep learning models Use Google Cloud Platform tools and services to develop and deploy large-scale machine learning and deep learning products Be aware of the different facets and design choices to consider when modeling a learning problem Productionalize machine learning models into software products Who This Book Is For Beginners to the practice of data science and applied machine learning, data scientists at all levels, machine learning engineers, Google Cloud Platform data engineers/architects, and software developers

Do want to learn how machine learning and neural networks work quickly and simply? Do you want to know how to build a machine learning model and you have no programming skill? Do you know a bit of Python coding and want to learn more about how this deep learning works? This bundle is going to guide you to the basics and the principles behind machine learning. Machine learning is an active research domain and includes several different approaches. This bundle is going to help you understand the different approaches of machine learning and neural networks. It will guide you through the steps you need to build a machine learning model. This bundle is intended to address all these questions. In fact, the aim of this bundle is providing the absolute beginners or other programmers that has no experience with Python programming the basic and fundamental tools of the Python language. Through this bundle, you will learn: Principles of machine learning Types of machine learning: supervised, unsupervised, semi-supervised, and reinforcement learning Advantages of each type of machine learning Principle and types of

neural networks Steps to develop and fit artificial neural network model Getting started and installing Python Tools and platforms for Python programming How to use pandas, NumPy and matplotlib Python libraries How to develop a simple linear and logistic machine learning model How to develop and train a multi-layer artificial neural network two ways: from scratch and using the Python libraries When to use each type of machine learning The general concept of artificial neural networks Activation function in artificial neural network and to choose an activation function within an artificial neural network The 5 main types of artificial neural network The best function to be used to train artificial neural networks. the 2 main concepts to know in the training process of the artificial neural network the main variants and algorithms for the formation of an artificial neural network and a machine learning model in general. The basics of the three main Python languages that will help you get the work done--including TensorFlow, Keras, and PyTorch How to install the three Python libraries to help you get started How to install and use magic command in Ipython Functionalities of NumPy library for numerical programming Functionalities of Pandas library for data analysis Even if you don't have any background in machine learning and Python programming, this book will give you the tools to develop machine learning models. Would you like to know more? Scroll to the top of the page and select the BUY NOW button!

Machine learning is a field of Artificial intelligence that provides algorithms those can learn and improve from experiences. Machine learning algorithms are turned as integral parts of today’s digital life. Its applications include recommender systems, targeted campaigns, text categorization, computer vision and auto security systems etc. Machine learning also considered as essential part of data science due to its capability of providing predictive analytics, capability in handling variety of data and suitability for big data applications. Its capability for predictive analytics resulted of its general structure that is building statistical models out of training data. In other hand easy scalability advantage of machine learning algorithms is making them to be suitable for big data applications. The different types of learning algorithms includes supervised learning, unsupervised learning, reinforcement learning, feature learning, rule based learning, Robot or expert system learning, sparse dictionary and anomaly detection. These learning algorithms can be realized by computing models artificial neural networks, decision trees, support vector machines, regression analysis, Bayesian networks, Genetic algorithms and soft computing. The familiar tools to implement machine learning algorithms include Python, R, Matlab, Scala, Clojure and Ruby. Involving of such open source programming languages, tools and social network communities makes the machine learning most progressing filed of computer science. The machine learning life cycle includes defining project objectives, explore the types and format, modeling data to fit for machine learning algorithms, deciding suitable machine learning model and implement and decide best result from data for decision making. These days, machine learning is observing great interest by the society and it has turned as one of the significant responsibility of top level managers to transform their business in the profitable means by exploring its basic functionalities. The world is at the sheer of realizing a situation where machines will work in agreement with human being to work together, operation, and advertise their services in a novel way which is targeted, valuable, and well-versed. In order to achieve this, they can influence machine learning distinctiveness. Dr. Raghuram Bhukya

Recent Advances in Reinforcement Learning addresses current research in an exciting area that is gaining a great deal of popularity in the Artificial Intelligence and Neural Network communities. Reinforcement learning has become a primary paradigm of machine learning. It applies to problems in which an agent (such as a robot, a process controller, or an information-retrieval engine) has to learn how to behave given only information about the success of its current actions. This book is a collection of important papers that address topics including the theoretical foundations of dynamic programming approaches, the role of prior knowledge, and methods for improving performance of reinforcement-learning techniques. These papers build on previous work and will form an important resource for students and researchers in the area. Recent Advances in Reinforcement Learning is an edited volume of peer-reviewed original research comprising twelve invited contributions by leading researchers. This research work has also been published as a special issue of Machine Learning (Volume 22, Numbers 1, 2 and 3).

Do want to learn how machine learning and neural networks work quickly and simply? Do you want to know how to build a machine learning model and you have no programming skill? Do you know a bit of Python coding and want to learn more about how this deep learning works? This bundle is going to guide you to the basics and the principles behind machine learning. Machine learning is an active research domain and includes several different approaches. This bundle is going to help you understand the different approaches of machine learning and neural networks. It will guide you through the steps you need to build a machine learning model. This bundle is intended to address all these questions. In fact, the aim of this book is providing the absolute beginners or other programmers that has no experience with Python programming the basic and fundamental tools of the Python language. Through this bundle, you will learn: - Principles of machine learning - Types of machine learning: supervised, unsupervised, semi-supervised, and reinforcement learning - Advantages of each type of machine learning - Principle and types of neural networks - Steps to develop and fit artificial neural network model - Getting started and installing Python - Tools and platforms for Python programming - How to use pandas, NumPy and matplotlib Python libraries - How to develop a simple linear and logistic machine learning model - How to develop and train a multi-layer artificial neural network two ways: from scratch and using the Python libraries - When to use each type of machine learning - The general concept of artificial neural networks - Activation function in artificial neural network and to choose an activation function within an artificial neural network - The 5 main types of artificial neural network - The best function to be used to train artificial neural networks. - the 2 main concepts to know in the training process of the artificial neural network - the main variants and algorithms for the formation of an artificial neural network and a machine learning model in general. - The basics of the three main Python languages that will help you get the work done-including TensorFlow, Keras, and PyTorch; - How to install the three Python libraries to help you get started; How to install and use magic command in Ipython Functionalities of NumPy library for numerical programming Functionalities of Pandas library for data analysis Even if you don't have any background in machine learning and Python programming, this book will give you the tools to develop machine learning models.

A new edition of an introductory text in machine learning that gives a unified treatment of machine learning problems and solutions. The goal of machine learning is to program computers to use example data or past experience to solve a given problem. Many successful applications of machine learning exist already, including systems that analyze past sales data to predict customer behavior, optimize robot behavior so that a task can be completed using minimum resources, and extract knowledge from bioinformatics data. The second edition of Introduction to Machine Learning is a comprehensive textbook on the subject, covering a broad array of topics not usually included in introductory machine learning texts. In order to present a unified treatment of machine learning problems and solutions, it discusses many methods from different fields, including statistics, pattern recognition,

neural networks, artificial intelligence, signal processing, control, and data mining. All learning algorithms are explained so that the student can easily move from the equations in the book to a computer program. The text covers such topics as supervised learning, Bayesian decision theory, parametric methods, multivariate methods, multilayer perceptrons, local models, hidden Markov models, assessing and comparing classification algorithms, and reinforcement learning. New to the second edition are chapters on kernel machines, graphical models, and Bayesian estimation; expanded coverage of statistical tests in a chapter on design and analysis of machine learning experiments; case studies available on the Web (with downloadable results for instructors); and many additional exercises. All chapters have been revised and updated. Introduction to Machine Learning can be used by advanced undergraduates and graduate students who have completed courses in computer programming, probability, calculus, and linear algebra. It will also be of interest to engineers in the field who are concerned with the application of machine learning methods.

Do want to learn how machine learning and neural networks work quickly and simply? Do you want to know how to build a machine learning model and you have no programming skill? Do you want to get started with learning data science? This bundle is going to guide you to the basics and the principles behind machine learning. Machine learning is an active research domain and includes several different approaches. This bundle is going to help you understand the different approaches of machine learning and neural networks. It will guide you through the steps you need to build a machine learning model. Machine learning implies programming. This bundle will teach you Python programming. This bundle does not require any pre-programming skills. It will help to get you started in Python programming, as well as how to use Python libraries to analyze data and apply machine learning. Overall, this bundle is a go-to guide for getting started in machine learning modeling using Python programming. Once you get through the bundle, you will be able to develop your own machine learning models using Python. Through this bundle, you will learn: - Principles of machine learning - Types of machine learning: supervised, unsupervised, semi-supervised, and reinforcement learning - Advantages of each type of machine learning - Principle and types of neural networks - Steps to develop and fit artificial neural network model - Getting started and installing Python - Tools and platforms for Python programming - How to use pandas, NumPy and matplotlib Python libraries - How to develop a simple linear and logistic machine learning model - How to develop and train a multi-layer artificial neural network two ways: from scratch and using the Python libraries - When to use each type of machine learning - The general concept of artificial neural networks - Activation function in artificial neural network and to choose an activation function within an artificial neural network - The 5 main types of artificial neural network - The best function to be used to train artificial neural networks. - the 2 main concepts to know in the training process of the artificial neural network - the main variants and algorithms for the formation of an artificial neural network and a machine learning model in general. Even if you don't have any background in machine learning and Python programming, this book will give you the tools to develop machine learning models.

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.

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 modelfree 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.

Do want to learn how machine learning and neural networks work quickly and simply? Do you want to know how to build a machine learning model and you have no programming skill? Do you want to get started with learning data science? This bundle is going to guide you to the basics and the principles behind machine learning. Machine learning is an active research domain and includes several different approaches. This bundle is going to help you understand the different

approaches of machine learning and neural networks. It will guide you through the steps you need to build a machine learning model. Machine learning implies programming. This bundle will teach you Python programming. This bundle does not require any pre-programming skills. It will help to get you started in Python programming, as well as how to use Python libraries to analyze data and apply machine learning. Overall, this bundle is a go-to guide for getting started in machine learning modeling using Python programming. Once you get through the bundle, you will be able to develop your own machine learning models using Python. Through this bundle, you will learn: Principles of machine learning Types of machine learning: supervised, unsupervised, semi-supervised, and reinforcement learning Advantages of each type of machine learning Principle and types of neural networks Steps to develop and fit artificial neural network model Getting started and installing Python Tools and platforms for Python programming How to use pandas, NumPy and matplotlib Python libraries How to develop a simple linear and logistic machine learning model How to develop and train a multi-layer artificial neural network two ways: from scratch and using the Python libraries When to use each type of machine learning The general concept of artificial neural networks Activation function in artificial neural network and to choose an activation function within an artificial neural network The 5 main types of artificial neural network The best function to be used to train artificial neural networks. the 2 main concepts to know in the training process of the artificial neural network the main variants and algorithms for the formation of an artificial neural network and a machine learning model in general. Even if you don't have any background in machine learning and Python programming, this book will give you the tools to develop machine learning models. Would you like to know more? Scroll to the top of the page and select the BUY NOW button!!

A practical guide to mastering reinforcement learning algorithms using Keras Key Features Build projects across robotics, gaming, and finance fields, putting reinforcement learning (RL) into action Get to grips with Keras and practice on real-world unstructured datasets Uncover advanced deep learning algorithms such as Monte Carlo, Markov Decision, and Q-learning Book Description Reinforcement learning has evolved a lot in the last couple of years and proven to be a successful technique in building smart and intelligent AI networks. Keras Reinforcement Learning Projects installs human-level performance into your applications using algorithms and techniques of reinforcement learning, coupled with Keras, a faster experimental library. The book begins with getting you up and running with the concepts of reinforcement learning using Keras. You'll learn how to simulate a random walk using Markov chains and select the best portfolio using dynamic programming (DP) and Python. You'll also explore projects such as forecasting stock prices using Monte Carlo methods, delivering vehicle routing application using Temporal Distance (TD) learning algorithms, and balancing a Rotating Mechanical System using Markov decision processes. Once you've understood the basics, you'll move on to Modeling of a Segway, running a robot control system using deep reinforcement learning, and building a handwritten digit recognition model in Python using an image dataset. Finally, you'll excel in playing the board game Go with the help of Q-Learning and reinforcement learning algorithms. By the end of this book, you'll not only have developed hands-on training on concepts, algorithms, and techniques of reinforcement learning but also be all set to explore the world of AI. What you will learn Practice the Markov decision process in prediction and betting evaluations Implement Monte Carlo methods to forecast environment behaviors Explore TD learning algorithms to manage warehouse operations Construct a Deep Q-Network using Python and Keras to control robot movements Apply reinforcement concepts to build a handwritten digit recognition model using an image dataset Address a game theory problem using Q-Learning and OpenAI Gym Who this book is for Keras Reinforcement Learning Projects is for you if you are data scientist, machine learning developer, or AI engineer who wants to understand the fundamentals of reinforcement learning by developing practical projects. Sound knowledge of machine learning and basic familiarity with Keras is useful to get the most out of this book