

Hands On Machine Learning With Scikit Learn Keras And Tensorflow 2nd Edition

Hands On Machine Learning With Scikit Learn Keras And Tensorflow 2nd Edition

Downloaded from staging.editor.seenic.io by Kade & Koch

HANDS ON MACHINE LEARNING WITH SCIKIT LEARN KERAS AND TENSORFLOW 2ND EDITION DOWNLOAD AND INSTALL PDF

Invite to our library, where you can easily download and install Hands On Machine Learning With Scikit Learn Keras And Tensorflow 2nd Edition to enhance your discovering and research study experience. Our large collection of PDF documents can give valuable instructional sources that cater to different topics and passions. We recognize the significance of accessing details swiftly and quickly, so we aim to make the process of **downloading and install Hands On Machine Learning With Scikit Learn Keras And Tensorflow 2nd Edition PDF** from our system simple and hassle-free. With simply a couple of clicks, you can open a globe of knowledge from our library with no barriers. Join us in discovering our comprehensive collection and start your PDF downloads today!

EXPLORING OUR CONSIDERABLE COLLECTION INCLUDING HANDS ON MACHINE LEARNING WITH SCIKIT LEARN KERAS AND TENSORFLOW 2ND EDITION

At our platform, we take satisfaction in our comprehensive collection of PDF data including Hands On Machine Learning With Scikit Learn Keras And Tensorflow 2nd Edition that cater to numerous passions and areas of research study. Whether you are seeking to increase your knowledge or performing research, we have a large range of PDFs that make certain to satisfy your needs.

Our PDF files Hands On Machine Learning With Scikit Learn Keras And Tensorflow 2nd Edition are carefully curated and selected to provide useful insights and details to our individuals. We have actually worked together with professionals in various areas to make certain that our collection remains up-to-date and pertinent.

From scientific study documents to academic resources, our PDF data cover a variety of subjects and topics. With simple accessibility to our collection, you can swiftly browse through and discover the PDF Hands On Machine Learning With Scikit Learn Keras And Tensorflow 2nd Edition that passion you the most.

Our system is committed to providing you with a smooth and reliable method to enhance your understanding and research study experience. We understand the relevance of having dependable and important resources at your disposal, which's why our PDF collection is continuously growing and increasing.

So whether you're a pupil, specialist or just curious, exploring our comprehensive collection of PDF documents Hands On Machine Learning With Scikit Learn Keras And Tensorflow 2nd Edition makes certain to provide you with useful insights and knowledge. Beginning surfing today to discover interesting new study possibilities!

EASY STEPS TO DOWNLOADING HANDS ON MACHINE LEARNING WITH SCIKIT LEARN KERAS AND TENSORFLOW 2ND EDITION PDF

At our system, our company believe in making the procedure of downloading and install PDF data Hands On Machine Learning With Scikit Learn Keras And Tensorflow 2nd Edition fast and easy. Below's exactly how you can access and download and install PDFs free of charge:

Step 1: Check out our substantial collection of PDF files to locate the one you need.

Action 2: Click on the download switch alongside the PDF Hands On Machine Learning With Scikit Learn Keras And Tensorflow 2nd Edition you intend to save.

Step 3: Wait for the PDF data Hands On Machine Learning With Scikit Learn Keras And Tensorflow 2nd Edition to download and install to your gadget. This need to just take a couple of secs.

And that's it! You can now access Hands On Machine Learning With Scikit Learn Keras And Tensorflow 2nd Edition PDF data offline at any moment and share it with others if you want.

Our company believe that knowing and investigating ought to be an easy and obtainable experience for all. That's why we offer our solution free of charge, ensuring that you can access the details you need with no barriers.

BOOST YOUR LEARNING AND STUDY

At our platform, our company believe that education and learning must come to all. That's why we offer a substantial collection of PDF downloads including **Hands On Machine Learning With Scikit Learn Keras And Tensorflow 2nd Edition** that accommodate a wide range of passions and subjects. Our educational sources are perfect for pupils, specialists, and any person looking to increase their expertise.

With our PDF downloads, you can access important info on numerous topics, including history, science, technology, and off training course Hands On Machine Learning With Scikit Learn Keras And Tensorflow 2nd Edition. Our resources are ideal for research objectives and can aid you strengthen your understanding of intricate subjects.

Our library is regularly growing, and we aim to include brand-new and appropriate material frequently. With our easy to use interface, you can quickly navigate our platform and discover the most recent academic resources.

By downloading Hands On Machine Learning With Scikit Learn Keras And Tensorflow 2nd Edition, you can raise your discovering and research study undertakings and gain useful understandings that can profit you in your individual and expert life.

So, what are you waiting on? Begin discovering our collection today and unlock a globe of expertise at your fingertips.

VERDICT

At our platform, we strive to supply a problem-free and free service that enables you to download Hands On Machine Learning With Scikit Learn Keras And Tensorflow 2nd Edition from our large library easily. Our straightforward interface makes certain that you can access the info you require with no complications or challenges.

Whether you're a trainee, professional, or merely interested, our PDF downloads offer valuable educational sources that can enhance your expertise and understanding of various subjects. By exploring our considerable collection, you can increase your discovering and research study endeavors and boost your understanding of the world around you.

So why wait? Start downloading **Hands On Machine Learning With Scikit Learn Keras And Tensorflow 2nd Edition** and begin exploring our library today and unlock a world of understanding at your fingertips. Whether you're seeking to expand your perspectives or carry out study, our straightforward and cost-free service is here to sustain you every step of the way.

Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow "O'Reilly Media, Inc."

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

Hands-On Machine Learning with R CRC Press

Hands-on Machine Learning with R provides a practical and applied approach to learning and developing intuition into today's most popular machine learning methods. This book serves as a practitioner's guide to the machine learning process and is meant to help the reader learn to apply the machine learning stack within R, which includes using various R packages such as glmnet, h2o, ranger, xgboost, keras, and others to effectively model and gain insight from their data. The book favors a hands-on approach, providing an intuitive understanding of machine learning concepts through concrete examples and just a little bit of theory. Throughout this book, the reader will be exposed to the entire machine learning process including feature engineering, resampling, hyperparameter tuning, model evaluation, and interpretation. The reader will be exposed to powerful algorithms such as regularized regression, random forests, gradient boosting machines, deep learning, generalized low rank models, and more! By favoring a hands-on approach and using real world data, the reader will gain an intuitive understanding of the architectures and engines that drive these algorithms and packages, understand when and how to tune the various hyperparameters, and be able to interpret model results. By the end of this book, the reader should have a firm grasp of R's machine learning stack and be able to implement a systematic approach for producing high quality modeling results. Features: · Offers a practical and applied introduction to the most popular machine learning methods. · Topics covered include feature engineering, resampling, deep learning and more. · Uses a hands-on approach and real world data.

Hands-on Machine Learning with JavaScript Packt Publishing Ltd

A definitive guide to creating an intelligent web application with the best of machine learning and JavaScript Key Features Solve complex computational problems in browser with JavaScript Teach your browser how to learn from rules using the power of machine learning Understand discoveries on web interface and API in machine learning Book Description In over 20 years of existence, JavaScript has been pushing beyond the boundaries of web evolution with proven existence on servers, embedded devices, Smart TVs, IoT, Smart Cars, and more. Today, with the added advantage of machine learning research and support for JS libraries, JavaScript makes your browsers smarter than ever with the ability to learn patterns and reproduce them to become a part of innovative products and applications. Hands-on Machine Learning with JavaScript presents various avenues of machine learning in a practical and objective way, and helps implement them using the JavaScript language. Predicting behaviors, analyzing feelings, grouping data, and building neural models are some of the skills you will build from this book. You will learn how to train your machine learning models and work with different kinds of data. During this journey, you will come across use cases such as face detection, spam filtering, recommendation systems, character recognition, and more. Moreover, you will learn how to work with deep neural networks and guide your applications to gain insights from data. By the end of this book, you'll have gained hands-on knowledge on evaluating and implementing the right model, along with choosing from different JS libraries, such as NaturalNode, brain, harthur, classifier, and many more to design smarter applications. What you will learn Get an overview of state-of-the-art machine learning Understand the pre-processing of data handling, cleaning, and preparation Learn Mining and Pattern Extraction with JavaScript Build your own model for classification, clustering, and prediction Identify the most appropriate model for each type of problem Apply machine learning techniques to real-world applications Learn how JavaScript can be a powerful language for machine learning Who this book is for This book is for you if you are a JavaScript developer who wants to implement machine learning to make applications smarter, gain insightful information from the data, and enter the field of machine learning without switching to another language. Working knowledge of JavaScript language is expected to get the most out of the book.

Hands-On Machine Learning with scikit-learn and Scientific Python Toolkits Packt Publishing Ltd

Integrate scikit-learn with various tools such as NumPy, pandas, imbalanced-learn, and scikit-surprise and use it to solve real-world machine learning problems Key Features Delve into machine learning with this comprehensive guide to scikit-learn and scientific Python Master the art of data-driven problem-solving with hands-on examples Foster your theoretical and practical knowledge of supervised and unsupervised machine learning algorithms Book Description Machine learning is applied everywhere, from business to research and academia, while scikit-learn is a versatile library that is popular among machine learning practitioners. This book serves as a practical guide for anyone looking to provide hands-on machine learning solutions with scikit-learn and Python toolkits. The book begins with an explanation of machine learning concepts and fundamentals, and strikes a balance between theoretical concepts and their applications. Each chapter covers a different set of algorithms, and shows you how to use them to solve real-life problems. You'll also learn about various key supervised and unsupervised machine learning algorithms using practical examples. Whether it is an instance-based learning algorithm, Bayesian estimation, a deep neural network, a

tree-based ensemble, or a recommendation system, you'll gain a thorough understanding of its theory and learn when to apply it. As you advance, you'll learn how to deal with unlabeled data and when to use different clustering and anomaly detection algorithms. By the end of this machine learning book, you'll have learned how to take a data-driven approach to provide end-to-end machine learning solutions. You'll also have discovered how to formulate the problem at hand, prepare required data, and evaluate and deploy models in production. What you will learn Understand when to use supervised, unsupervised, or reinforcement learning algorithms Find out how to collect and prepare your data for machine learning tasks Tackle imbalanced data and optimize your algorithm for a bias or variance tradeoff Apply supervised and unsupervised algorithms to overcome various machine learning challenges Employ best practices for tuning your algorithm's hyper parameters Discover how to use neural networks for classification and regression Build, evaluate, and deploy your machine learning solutions to production Who this book is for This book is for data scientists, machine learning practitioners, and anyone who wants to learn how machine learning algorithms work and to build different machine learning models using the Python ecosystem. The book will help you take your knowledge of machine learning to the next level by grasping its ins and outs and tailoring it to your needs. Working knowledge of Python and a basic understanding of underlying mathematical and statistical concepts is required.

Machine Learning *CRC Press*

Traditional books on machine learning can be divided into two groups- those aimed at advanced undergraduates or early postgraduates with reasonable mathematical knowledge and those that are primers on how to code algorithms. The field is ready for a text that not only demonstrates how to use the algorithms that make up machine learning methods, but

Hands-On Machine Learning with Azure *Packt Publishing Ltd*

Implement machine learning, cognitive services, and artificial intelligence solutions by leveraging Azure cloud technologies Key FeaturesLearn advanced concepts in Azure ML and the Cortana Intelligence Suite architectureExplore ML Server using SQL Server and HDInsight capabilitiesImplement various tools in Azure to build and deploy machine learning modelsBook Description Implementing Machine learning (ML) and Artificial Intelligence (AI) in the cloud had not been possible earlier due to the lack of processing power and storage. However, Azure has created ML and AI services that are easy to implement in the cloud. Hands-On Machine Learning with Azure teaches you how to perform advanced ML projects in the cloud in a cost-effective way. The book begins by covering the benefits of ML and AI in the cloud. You will then explore Microsoft's Team Data Science Process to establish a repeatable process for successful AI development and implementation. You will also gain an understanding of AI technologies available in Azure and the Cognitive Services APIs to integrate them into bot applications. This book lets you explore prebuilt templates with Azure Machine Learning Studio and build a model using canned algorithms that can be deployed as web services. The book then takes you through a preconfigured series of virtual machines in Azure targeted at AI development scenarios. You will get to grips with the ML Server and its capabilities in SQL and HDInsight. In the concluding chapters, you'll integrate patterns with other non-AI services in Azure. By the end of this book, you will be fully equipped to implement smart cognitive actions in your models. What you will learnDiscover the benefits of leveraging the cloud for ML and AIUse Cognitive Services APIs to build intelligent botsBuild a model using canned algorithms from Microsoft and deploy it as a web serviceDeploy virtual machines in AI development scenariosApply R, Python, SQL Server, and Spark in AzureBuild and deploy deep learning solutions with CNTK, MMLSpark, and TensorFlowImplement model retraining in IoT, Streaming, and Blockchain solutionsExplore best practices for integrating ML and AI functions with ADLA and logic appsWho this book is for If you are a data scientist or developer familiar with Azure ML and cognitive services and want to create smart models and make sense of data in the cloud, this book is for you. You'll also find this book useful if you want to bring powerful machine learning services into your cloud applications. Some experience with data manipulation and processing, using languages like SQL, Python, and R, will aid in understanding the concepts covered in this book

Hands-On Machine Learning with C++ *Packt Publishing Ltd*

Implement supervised and unsupervised machine learning algorithms using C++ libraries such as PyTorch C++ API, Caffe2, Shogun, Shark-ML, mlpack, and dlib with the help of real-world examples and datasets Key FeaturesBecome familiar with data processing, performance measuring, and model selection using various C++ librariesImplement practical machine learning and deep learning techniques to build smart modelsDeploy machine learning models to work on mobile and embedded devicesBook Description C++ can make your machine learning models run faster and more efficiently. This handy guide will help you learn the fundamentals of machine learning (ML), showing you how to use C++ libraries to get the most out of your data. This book makes machine learning with C++ for beginners easy with its example-based approach, demonstrating how to implement supervised and unsupervised ML algorithms through real-world examples. This book will get you hands-on with tuning and optimizing a model for different use cases, assisting you with model selection and the measurement of performance. You'll cover techniques such as product recommendations, ensemble learning, and anomaly detection using modern C++ libraries such as PyTorch C++ API, Caffe2, Shogun, Shark-ML, mlpack, and dlib. Next, you'll explore neural networks and deep learning using examples such as image classification and sentiment analysis, which will help you solve various problems. Later, you'll learn how to handle production and deployment challenges on mobile and cloud platforms, before discovering how to export and import models using the ONNX format. By the end of this C++ book, you will have real-world machine learning and C++ knowledge, as well as the skills to use C++ to build powerful ML systems. What you will learnExplore how to load and preprocess various data types to suitable C++ data structuresEmploy key machine learning algorithms with various C++ librariesUnderstand the grid-search approach to find the best parameters for a machine learning modelImplement an algorithm for filtering anomalies in user data using Gaussian distributionImprove collaborative filtering to deal with dynamic user preferencesUse C++ libraries and APIs to manage model structures and parametersImplement a C++ program to solve image classification tasks with LeNet architectureWho this book is for You will find this C++ machine learning book useful if you want to get started with machine learning algorithms and techniques using the popular C++ language. As well as being a useful first course in machine learning with C++, this book will also appeal to data analysts, data scientists, and machine learning developers who are looking to implement different machine learning models in production using varied datasets and examples. Working knowledge of the C++ programming language is mandatory to get started with this book.

Hands-On Machine Learning for Cybersecurity *Packt Publishing Ltd*

Get into the world of smart data security using machine learning algorithms and Python libraries Key FeaturesLearn machine learning algorithms and cybersecurity fundamentalsAutomate your daily workflow by applying use cases to many facets of securityImplement smart machine learning solutions to detect various cybersecurity problemsBook Description Cyber threats today are one of

the costliest losses that an organization can face. In this book, we use the most efficient tool to solve the big problems that exist in the cybersecurity domain. The book begins by giving you the basics of ML in cybersecurity using Python and its libraries. You will explore various ML domains (such as time series analysis and ensemble modeling) to get your foundations right. You will implement various examples such as building system to identify malicious URLs, and building a program to detect fraudulent emails and spam. Later, you will learn how to make effective use of K-means algorithm to develop a solution to detect and alert you to any malicious activity in the network. Also learn how to implement biometrics and fingerprint to validate whether the user is a legitimate user or not. Finally, you will see how we change the game with TensorFlow and learn how deep learning is effective for creating models and training systems What you will learnUse machine learning algorithms with complex datasets to implement cybersecurity conceptsImplement machine learning algorithms such as clustering, k-means, and Naive Bayes to solve real-world problemsLearn to speed up a system using Python libraries with NumPy, Scikit-learn, and CUDAUnderstand how to combat malware, detect spam, and fight financial fraud to mitigate cyber crimesUse TensorFlow in the cybersecurity domain and implement real-world examplesLearn how machine learning and Python can be used in complex cyber issuesWho this book is for This book is for the data scientists, machine learning developers, security researchers, and anyone keen to apply machine learning to up-skill computer security. Having some working knowledge of Python and being familiar with the basics of machine learning and cybersecurity fundamentals will help to get the most out of the book

Hands-On Machine Learning with TensorFlow.js *Packt Publishing Ltd*

Get hands-on with the browser-based JavaScript library for training and deploying machine learning models effectively Key Features Build, train and run machine learning models in the browser using TensorFlow.js Create smart web applications from scratch with the help of useful examples Use flexible and intuitive APIs from TensorFlow.js to understand how machine learning algorithms function Book Description TensorFlow.js is a framework that enables you to create performant machine learning (ML) applications that run smoothly in a web browser. With this book, you will learn how to use TensorFlow.js to implement various ML models through an example-based approach. Starting with the basics, you'll understand how ML models can be built on the web. Moving on, you will get to grips with the TensorFlow.js ecosystem to develop applications more efficiently. The book will then guide you through implementing ML techniques and algorithms such as regression, clustering, fast Fourier transform (FFT), and dimensionality reduction. You will later cover the Bellman equation to solve Markov decision process (MDP) problems and understand how it is related to reinforcement learning. Finally, you will explore techniques for deploying ML-based web applications and training models with TensorFlow Core. Throughout this ML book, you'll discover useful tips and tricks that will build on your knowledge. By the end of this book, you will be equipped with the skills you need to create your own web-based ML applications and fine-tune models to achieve high performance. What you will learn Use the t-SNE algorithm in TensorFlow.js to reduce dimensions in an input dataset Deploy tfjs-converter to convert Keras models and load them into TensorFlow.js Apply the Bellman equation to solve MDP problems Use the k-means algorithm in TensorFlow.js to visualize prediction results Create tf.js packages with Parcel, Webpack, and Rollup to deploy web apps Implement tf.js backend frameworks to tune and accelerate app performance Who this book is for This book is for web developers who want to learn how to integrate machine learning techniques with web-based applications from scratch. This book will also appeal to data scientists, machine learning practitioners, and deep learning enthusiasts who are looking to perform accelerated, browser-based machine learning on Web using TensorFlow.js. Working knowledge of JavaScript programming language is all you need to get started.

Hands-On Deep Learning Algorithms with Python *Packt Publishing Ltd*

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.

Hands-On Machine Learning with Scikit-Learn and TensorFlow *"O'Reilly Media, Inc."*

Graphics in this book are printed in black and white. 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 Apply practical code examples without acquiring excessive machine learning theory or algorithm details

Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, 2nd Edition

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.

Hands-On Machine Learning for Algorithmic Trading *Packt Publishing Ltd*

Explore effective trading strategies in real-world markets using NumPy, spaCy, pandas, scikit-learn, and Keras Key FeaturesImplement machine learning algorithms to build, train, and validate algorithmic modelsCreate your own algorithmic design process to apply probabilistic machine learning approaches to trading decisionsDevelop neural networks for algorithmic trading to perform time series forecasting and smart analyticsBook Description The explosive growth of digital data has

boosted the demand for expertise in trading strategies that use machine learning (ML). This book enables you to use a broad range of supervised and unsupervised algorithms to extract signals from a wide variety of data sources and create powerful investment strategies. This book shows how to access market, fundamental, and alternative data via API or web scraping and offers a framework to evaluate alternative data. You'll practice the ML workflow from model design, loss metric definition, and parameter tuning to performance evaluation in a time series context. You will understand ML algorithms such as Bayesian and ensemble methods and manifold learning, and will know how to train and tune these models using pandas, statsmodels, sklearn, PyMC3, xgboost, lightgbm, and catboost. This book also teaches you how to extract features from text data using spaCy, classify news and assign sentiment scores, and to use gensim to model topics and learn word embeddings from financial reports. You will also build and evaluate neural networks, including RNNs and CNNs, using Keras and PyTorch to exploit unstructured data for sophisticated strategies. Finally, you will apply transfer learning to satellite images to predict economic activity and use reinforcement learning to build agents that learn to trade in the OpenAI Gym. What you will learnImplement machine learning techniques to solve investment and trading problemsLeverage market, fundamental, and alternative data to research alpha factorsDesign and fine-tune supervised, unsupervised, and reinforcement learning modelsOptimize portfolio risk and performance using pandas, NumPy, and scikit-learnIntegrate machine learning models into a live trading strategy on QuantopianEvaluate strategies using reliable backtesting methodologies for time seriesDesign and evaluate deep neural networks using Keras, PyTorch, and TensorFlowWork with reinforcement learning for trading strategies in the OpenAI GymWho this book is for Hands-On Machine Learning for Algorithmic Trading is for data analysts, data scientists, and Python developers, as well as investment analysts and portfolio managers working within the finance and investment industry. If you want to perform efficient algorithmic trading by developing smart investigating strategies using machine learning algorithms, this is the book for you. Some understanding of Python and machine learning techniques is mandatory.

Hands-On Deep Learning with R *Packt Publishing Ltd*

Explore and implement deep learning to solve various real-world problems using modern R libraries such as TensorFlow, MXNet, H2O, and Deepnet Key FeaturesUnderstand deep learning algorithms and architectures using R and determine which algorithm is best suited for a specific problemImprove models using parameter tuning, feature engineering, and ensemblingApply advanced neural network models such as deep autoencoders and generative adversarial networks (GANs) across different domainsBook Description Deep learning enables efficient and accurate learning from a massive amount of data. This book will help you overcome a number of challenges using various deep learning algorithms and architectures with R programming. This book starts with a brief overview of machine learning and deep learning and how to build your first neural network. You'll understand the architecture of various deep learning algorithms and their applicable fields, learn how to build deep learning models, optimize hyperparameters, and evaluate model performance. Various deep learning applications in image processing, natural language processing (NLP), recommendation systems, and predictive analytics will also be covered. Later chapters will show you how to tackle recognition problems such as image recognition and signal detection, programmatically summarize documents, conduct topic modeling, and forecast stock market prices. Toward the end of the book, you will learn the common applications of GANs and how to build a face generation model using them. Finally, you'll get to grips with using reinforcement learning and deep reinforcement learning to solve various real-world problems. By the end of this deep learning book, you will be able to build and deploy your own deep learning applications using appropriate frameworks and algorithms. What you will learnDesign a feedforward neural network to see how the activation function computes an outputCreate an image recognition model using convolutional neural networks (CNNs)Prepare data, decide hidden layers and neurons and train your model with the backpropagation algorithmApply text cleaning techniques to remove uninformative text using NLPBuild, train, and evaluate a GAN model for face generationUnderstand the concept and implementation of reinforcement learning in RWho this book is for This book is for data scientists, machine learning engineers, and deep learning developers who are familiar with machine learning and are looking to enhance their knowledge of deep learning using practical examples. Anyone interested in increasing the efficiency of their machine learning applications and exploring various options in R will also find this book useful. Basic knowledge of machine learning techniques and working knowledge of the R programming language is expected.

Hands-On Machine Learning with C++

This book will help you explore how to implement different well-known machine learning algorithms with various C++ frameworks and libraries. You will cover basic to advanced machine learning concepts with practical and easy to follow examples. By the end of the book, you will be able to build various machine learning models with ease.

Hands-On Machine Learning with ML.NET *Packt Publishing Ltd*

Create, train, and evaluate various machine learning models such as regression, classification, and clustering using ML.NET, Entity Framework, and ASP.NET Core Key FeaturesGet well-versed with the ML.NET framework and its components and APIs using practical examplesLearn how to build, train, and evaluate popular machine learning algorithms with ML.NET offeringsExtend your existing machine learning models by integrating with TensorFlow and other librariesBook Description Machine learning (ML) is widely used in many industries such as science, healthcare, and research and its popularity is only growing. In March 2018, Microsoft introduced ML.NET to help .NET enthusiasts in working with ML. With this book, you'll explore how to build ML.NET applications with the various ML models available using C# code. The book starts by giving you an overview of ML and the types of ML algorithms used, along with covering what ML.NET is and why you need it to build ML apps. You'll then explore the ML.NET framework, its components, and APIs. The book will serve as a practical guide to helping you build smart apps using the ML.NET library. You'll gradually become well versed in how to implement ML algorithms such as regression, classification, and clustering with real-world examples and datasets. Each chapter will cover the practical implementation, showing you how to implement ML within .NET applications. You'll also learn to integrate TensorFlow in ML.NET applications. Later you'll discover how to store the regression model housing price prediction result to the database and display the real-time predicted results from the database on your web application using ASP.NET Core Blazor and SignalR. By the end of this book, you'll have learned how to confidently perform basic to advanced-level machine learning tasks in ML.NET. What you will learnUnderstand the framework, components, and APIs of ML.NET using C#Develop regression models using ML.NET for employee attrition and file classificationEvaluate classification models for sentiment prediction of restaurant reviewsWork with clustering models for file type classificationsUse anomaly detection to find anomalies in both network traffic and login historyWork with ASP.NET Core Blazor to create an ML.NET enabled web applicationIntegrate pre-trained TensorFlow and ONNX models in a WPF ML.NET application for image classification and object detectionWho this book is for If you are a .NET developer who wants to implement machine

learning models using ML.NET, then this book is for you. This book will also be beneficial for data scientists and machine learning developers who are looking for effective tools to implement various machine learning algorithms. A basic understanding of C# or .NET is mandatory to grasp the concepts covered in this book effectively.

Machine Learning *John Wiley & Sons*

Dig deep into the data with a hands-on guide to machine learning with updated examples and more! Machine Learning: Hands-On for Developers and Technical Professionals provides hands-on instruction and fully-coded working examples for the most common machine learning techniques used by developers and technical professionals. The book contains a breakdown of each ML variant, explaining how it works and how it is used within certain industries, allowing readers to incorporate the presented techniques into their own work as they follow along. A core tenant of machine learning is a strong focus on data preparation, and a full exploration of the various types of learning algorithms illustrates how the proper tools can help any developer extract information and insights from existing data. The book includes a full complement of Instructor's Materials to facilitate use in the classroom, making this resource useful for students and as a professional reference. At its core, machine learning is a mathematical, algorithm-based technology that forms the basis of historical data mining and modern big data science. Scientific analysis of big data requires a working knowledge of machine learning, which forms predictions based on known properties learned from training data. Machine Learning is an accessible, comprehensive guide for the non-mathematician, providing clear guidance that allows readers to: Learn the languages of machine learning including Hadoop, Mahout, and Weka Understand decision trees, Bayesian networks, and artificial neural networks Implement Association Rule, Real Time, and Batch learning Develop a strategic plan for safe, effective, and efficient machine learning By learning to construct a system that can learn from data, readers can increase their utility across industries. Machine learning sits at the core of deep dive data analysis and visualization, which is increasingly in demand as companies discover the goldmine hiding in their existing data. For the tech professional involved in data science, Machine Learning: Hands-On for Developers and Technical Professionals provides the skills and techniques required to dig deeper.

Hands-On Reinforcement Learning with Python *Packt Publishing Ltd*

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.

Hands-On Reinforcement Learning for Games *Packt Publishing Ltd*

Explore reinforcement learning (RL) techniques to build cutting-edge games using Python libraries such as PyTorch, OpenAI Gym, and TensorFlow Key FeaturesGet to grips with the different reinforcement and DRL algorithms for game developmentLearn how to implement components such as artificial agents, map and level generation, and audio generationGain insights into cutting-edge RL research and understand how it is similar to artificial general researchBook 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 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 learnUnderstand how deep learning can be integrated into an RL agentExplore basic to advanced algorithms commonly used in game developmentBuild agents that can learn and solve problems in all types of environmentsTrain a Deep Q-Network (DQN) agent to solve the CartPole balancing problemDevelop game AI agents by understanding the mechanism behind complex AIIntegrate all the concepts learned into new projects or gaming agentsWho 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.

Hands-On Deep Learning with TensorFlow *Packt Publishing Ltd*

This book is your guide to exploring the possibilities in the field of deep learning, making use of Google's TensorFlow. You will learn about convolutional neural networks, and logistic regression while training models for deep learning to gain key insights into your data. About This Book Explore various possibilities with deep learning and gain amazing insights from data using Google's brainchild-- TensorFlow Want to learn what more can be done with deep learning? Explore various neural networks with the help of this comprehensive guide Rich in concepts, advanced guide on deep learning that will give you background to innovate in your environment Who This Book Is For If you are a data scientist who performs machine learning on a regular basis, are familiar with deep neural networks, and now want to gain expertise in working with convoluted neural networks, then this book is for you. Some familiarity with C++ or Python is assumed. What You Will Learn Set up your computing environment and install TensorFlow Build simple TensorFlow graphs for everyday computations Apply logistic regression for classification with TensorFlow Design and train a multilayer neural network with TensorFlow Intuitively understand convolutional neural networks for image recognition Bootstrap a neural network from simple to more accurate models See how to use TensorFlow with other types of networks Program networks with SciKit-Flow, a high-level interface to TensorFlow In Detail Dan Van Boxel's Deep Learning with TensorFlow is based on Dan's best-selling TensorFlow video course. With deep learning going mainstream, making sense of data and getting accurate results using deep networks is possible. Dan Van Boxel will be your guide to exploring the possibilities with deep learning; he will enable you to understand data like never before. With the efficiency and simplicity of TensorFlow, you will be able to process your data and gain insights that will change how you look at data. With Dan's guidance, you will dig deeper into the hidden layers of abstraction using raw data. Dan then shows you various complex algorithms for deep learning and various examples that use these deep neural networks. You will also learn how to train your machine to craft new features to make sense of deeper layers of data. In this book, Dan shares his knowledge across topics such as logistic regression, convolutional neural networks, recurrent neural networks, training deep networks, and high level interfaces. With the help of novel practical examples, you will become an ace at advanced multilayer networks, image recognition, and beyond. Style and

Approach This book is your go-to guide to becoming a deep learning expert in your organization. Dan helps you evaluate common and not-so-common deep neural networks with the help of insightful examples that you can relate to, and show how they can be exploited in the real world with complex raw data.

Hands-On Machine Learning on Google Cloud Platform *Packt Publishing Ltd*

Unleash Google's Cloud Platform to build, train and optimize machine learning models Key Features Get well versed in GCP pre-existing services to build your own smart models A comprehensive guide covering aspects from data processing, analyzing to building and training ML models A practical approach to produce your trained ML models and port them to your mobile for easy access Book Description Google Cloud Machine Learning Engine combines the services of Google Cloud Platform with the power and flexibility of TensorFlow. With this book, you will not only learn to build and train different complexities of machine learning models at scale but also host them in the cloud to make predictions. This book is focused on making the most of the Google Machine Learning Platform for large datasets and complex problems. You will learn from scratch how to create powerful machine learning based applications for a wide variety of problems by leveraging different data services from the Google Cloud Platform. Applications include NLP, Speech to text, Reinforcement learning, Time series, recommender systems, image classification, video content inference and many other. We will implement a wide variety of deep learning use cases and also make extensive use of data related services comprising the Google Cloud Platform ecosystem such as Firebase, Storage APIs, Datalab and so forth. This will enable you to integrate Machine Learning and data processing features into your web and mobile applications. By the end of this book, you will know the main difficulties that you may encounter and get appropriate strategies to overcome these difficulties and build efficient systems. What you will learn Use Google Cloud Platform to build data-based applications for dashboards, web, and mobile Create, train and optimize deep learning models for various data science problems on big data Learn how to leverage BigQuery to explore big datasets Use Google's pre-trained TensorFlow models for NLP, image, video and much more Create models and architectures for Time series, Reinforcement Learning, and generative models Create, evaluate, and optimize TensorFlow and Keras models for a wide range of applications Who this book is for This book is for data scientists, machine learning developers and AI developers who want to learn Google Cloud Platform services to build machine learning applications. Since the interaction with the Google ML platform is mostly done via the command line, the reader is supposed to have some familiarity with the bash shell and Python scripting. Some understanding of machine learning and data science concepts will be handy

Hands-On Deep Learning with Apache Spark *Packt Publishing Ltd*

Speed up the design and implementation of deep learning solutions using Apache Spark Key Features Explore the world of distributed deep learning with Apache Spark Train neural networks with deep learning libraries such as BigDL and TensorFlow Develop Spark deep learning applications to intelligently handle large and complex datasets Book Description Deep learning is a subset of machine learning where datasets with several layers of complexity can be processed. Hands-On Deep Learning with Apache Spark addresses the sheer complexity of technical and analytical parts and the speed at which deep learning solutions can be implemented on Apache Spark. The book starts with the fundamentals of Apache Spark and deep learning. You will set up Spark for deep learning, learn principles of distributed modeling, and understand different types of neural nets. You will then implement deep learning models, such as convolutional neural networks (CNNs), recurrent neural networks (RNNs), and long short-term memory (LSTM) on Spark. As you progress through the book, you will gain hands-on experience of what it takes to understand the complex datasets you are dealing with. During the course of this book, you will use popular deep learning frameworks, such as TensorFlow, Deeplearning4j, and Keras to train your distributed models. By the end of this book, you'll have gained experience with the implementation of your models on a variety of use cases. What you will learn Understand the basics of deep learning Set up Apache Spark for deep learning Understand the principles of distribution modeling and different types of neural networks Obtain an understanding of deep learning algorithms Discover textual analysis and deep learning with Spark Use popular deep learning frameworks, such as Deeplearning4j, TensorFlow, and Keras Explore popular deep learning algorithms Who this book is for If you are a Scala developer, data scientist, or data analyst who wants to learn how to use Spark for implementing efficient deep learning models, Hands-On Deep Learning with Apache Spark is for you. Knowledge of the core machine learning concepts and some exposure to Spark will be helpful.

Hands-On Deep Learning with Go *Packt Publishing Ltd*

Apply modern deep learning techniques to build and train deep neural networks using Gorgonia Key Features Gain a practical understanding of deep learning using Golang Build complex neural network models using Go libraries and Gorgonia Take your deep learning model from design to deployment with this handy guide Book Description Go is an open source programming language designed by Google for handling large-scale projects efficiently. The Go ecosystem comprises some really powerful deep learning tools such as DQN and CUDA. With this book, you'll be able to use these tools to train and deploy scalable deep learning models from scratch. This deep learning book begins by introducing you to a variety of tools and libraries available in Go. It then takes you through building neural networks, including activation functions and the learning algorithms that make neural networks tick. In addition to this, you'll learn how to build advanced architectures such as autoencoders, restricted Boltzmann machines (RBMs), convolutional neural networks (CNNs), recurrent neural networks (RNNs), and more. You'll also understand how you can scale model deployments on the AWS cloud infrastructure for training and inference. By the end of this book, you'll have mastered the art of building, training, and deploying deep learning models in Go to solve real-world problems. What you will learn Explore the Go ecosystem of libraries and communities for deep learning Get to grips with Neural Networks, their history, and how they work Design and implement Deep Neural Networks in Go Get a strong foundation of concepts such as Backpropagation and Momentum Build Variational Autoencoders and Restricted Boltzmann Machines using Go Build models with CUDA and benchmark CPU and GPU models Who this book is for This book is for data scientists, machine learning engineers, and AI developers who want to build state-of-the-art deep learning models using Go. Familiarity with basic machine learning concepts and Go programming is required to get the best out of this book.

Hands-On Deep Learning Architectures with Python *Packt Publishing Ltd*

Concepts, tools, and techniques to explore deep learning architectures and methodologies Key Features Explore advanced deep learning architectures using various datasets and frameworks Implement deep architectures for neural network models such as CNN, RNN, GAN, and many more Discover design patterns and different challenges for various deep learning architectures Book Description Deep learning architectures are composed of multilevel nonlinear operations that represent high-level abstractions; this allows you to learn useful feature representations from the data. This book will help you learn and implement deep learning architectures to resolve various

deep learning research problems. Hands-On Deep Learning Architectures with Python explains the essential learning algorithms used for deep and shallow architectures. Packed with practical implementations and ideas to help you build efficient artificial intelligence systems (AI), this book will help you learn how neural networks play a major role in building deep architectures. You will understand various deep learning architectures (such as AlexNet, VGG Net, GoogLeNet) with easy-to-follow code and diagrams. In addition to this, the book will also guide you in building and training various deep architectures such as the Boltzmann mechanism, autoencoders, convolutional neural networks (CNNs), recurrent neural networks (RNNs), natural language processing (NLP), GAN, and more—all with practical implementations. By the end of this book, you will be able to construct deep models using popular frameworks and datasets with the required design patterns for each architecture. You will be ready to explore the potential of deep architectures in today's world. What you will learn Implement CNNs, RNNs, and other commonly used architectures with Python Explore architectures such as VGGNet, AlexNet, and GoogLeNet Build deep learning architectures for AI applications such as face and image recognition, fraud detection, and many more Understand the architectures and applications of Boltzmann machines and autoencoders with concrete examples Master artificial intelligence and neural network concepts and apply them to your architecture Understand deep learning architectures for mobile and embedded systems Who this book is for If you're a data scientist, machine learning developer/engineer, or deep learning practitioner, or are curious about AI and want to upgrade your knowledge of various deep learning architectures, this book will appeal to you. You are expected to have some knowledge of statistics and machine learning algorithms to get the best out of this book

Analytical Skills for AI and Data Science *O'Reilly Media*

While several market-leading companies have successfully transformed their business models by following data- and AI-driven paths, the vast majority have yet to reap the benefits. How can your business and analytics units gain a competitive advantage by capturing the full potential of this predictive revolution? This practical guide presents a battle-tested end-to-end method to help you translate business decisions into tractable prescriptive solutions using data and AI as fundamental inputs. Author Daniel Vaughan shows data scientists, analytics practitioners, and others interested in using AI to transform their businesses not only how to ask the right questions but also how to generate value using modern AI technologies and decision-making principles. You'll explore several use cases common to many enterprises, complete with examples you can apply when working to solve your own issues. Break business decisions into stages that can be tackled using different skills from the analytical toolbox Identify and embrace uncertainty in decision making and protect against common human biases Customize optimal decisions to different customers using predictive and prescriptive methods and technologies Ask business questions that create high value through AI- and data-driven technologies

Deep Reinforcement Learning Hands-On *Packt Publishing Ltd*

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.

Hands-On Unsupervised Learning Using Python *"O'Reilly Media, Inc."*

Many industry experts consider unsupervised learning the next frontier in artificial intelligence, one that may hold the key to general artificial intelligence. Since the majority of the world's data is unlabeled, conventional supervised learning cannot be applied. Unsupervised learning, on the other hand, can be applied to unlabeled datasets to discover meaningful patterns buried deep in the data, patterns that may be near impossible for humans to uncover. Author Ankur Patel shows you how to apply unsupervised learning using two simple, production-ready Python frameworks: Scikit-learn and TensorFlow using Keras. With code and hands-on examples, data scientists will identify difficult-to-find patterns in data and gain deeper business insight, detect anomalies, perform automatic feature engineering and selection, and generate synthetic datasets. All you need is programming and some machine learning experience to get started. Compare the strengths and weaknesses of the different machine learning approaches: supervised, unsupervised, and reinforcement learning Set up and manage machine learning projects end-to-end Build an anomaly detection system to catch credit card fraud Clusters users into distinct and homogeneous groups Perform semisupervised learning Develop movie recommender systems using restricted Boltzmann machines Generate synthetic images using generative adversarial networks

Hands-On Deep Learning for Images with TensorFlow *Packt Publishing Ltd*

Explore TensorFlow's capabilities to perform efficient deep learning on images Key Features Discover image processing for machine vision Build an effective image classification system using the power of CNNs Leverage TensorFlow's capabilities to perform efficient deep learning Book Description TensorFlow is Google's popular offering for machine learning and deep learning, quickly becoming a favorite tool for performing fast, efficient, and accurate deep learning tasks. Hands-On Deep Learning for Images with TensorFlow shows you the practical implementations of real-world projects, teaching you how to leverage TensorFlow's capabilities to perform efficient image processing using the power of deep learning. With the help of this book, you will get to grips with the different paradigms of performing deep learning such as deep neural nets and convolutional neural

networks, followed by understanding how they can be implemented using TensorFlow. By the end of this book, you will have mastered all the concepts of deep learning and their implementation with TensorFlow and Keras. What you will learn

- Build machine learning models particularly focused on the MNIST digits
- Work with Docker and Keras to build an image classifier
- Understand natural language models to process text and images
- Prepare your dataset for machine learning
- Create classical, convolutional, and deep neural networks
- Create a RESTful image classification server

Who this book is for

Hands-On Deep Learning for Images with TensorFlow is for you if you are an application developer, data scientist, or machine learning practitioner looking to integrate machine learning into application software and master deep learning by implementing practical projects in TensorFlow. Knowledge of Python programming and basics of deep learning are required to get the best out of this book.

Hands-On Data Science and Python Machine Learning *Packt Publishing Ltd*

This book covers the fundamentals of machine learning with Python in a concise and dynamic manner. It covers data mining and large-scale machine learning using Apache Spark. About This Book

- Take your first steps in the world of data science by understanding the tools and techniques of data analysis
- Train efficient Machine Learning models in Python using the supervised and unsupervised learning methods
- Learn how to use Apache Spark for processing Big Data efficiently

Who This Book Is For

If you are a budding data scientist or a data analyst who wants to analyze and gain actionable insights from data using Python, this book is for you. Programmers with some experience in Python who want to enter the lucrative world of Data Science will also find this book to be very useful, but you don't need to be an expert Python coder or mathematician to get the most from this book.

What You Will Learn

- Learn how to clean your data and ready it for analysis
- Implement the popular clustering and regression methods in Python
- Train efficient machine learning models using decision trees and random forests
- Visualize the results of your analysis using Python's Matplotlib library
- Use Apache Spark's MLlib package to perform machine learning on large datasets

In Detail

Join Frank Kane, who worked on Amazon and IMDb's machine learning algorithms, as he guides you on your first steps into the world of data science. Hands-On Data Science and Python Machine Learning gives you the tools that you need to understand and explore the core topics in the field, and the confidence and practice to build and analyze your own machine learning models. With the help of interesting and easy-to-follow practical examples, Frank Kane explains potentially complex topics such as Bayesian methods and K-means clustering in a way that anybody can understand them. Based on Frank's successful data science course, Hands-On Data Science and Python Machine Learning empowers you to conduct data analysis and perform efficient machine learning using Python. Let Frank help you unearth the value in your data using the various data mining and data analysis techniques available in Python, and to develop efficient predictive models to predict future results. You will also learn how to perform large-scale machine learning on Big Data using Apache Spark. The book covers preparing your data for analysis, training machine learning models, and visualizing the final data analysis. Style and approach

This comprehensive book is a perfect blend of theory and hands-on code examples in Python which can be used for your reference at any time.

Hands-On Automated Machine Learning *Packt Publishing Ltd*

Automate data and model pipelines for faster machine learning applications

Key Features

- Build automated modules for different machine learning components
- Understand each component of a machine learning pipeline in depth
- Learn to use different open source AutoML and feature engineering platforms

Book Description

AutoML is designed to automate parts of Machine Learning. Readily available AutoML tools are making data science practitioners' work easy and are received well in the advanced analytics community. Automated Machine Learning covers the necessary foundation needed to create automated machine learning modules and helps you get up to speed with them in the most practical way possible. In this book, you'll learn how to automate different tasks in the machine learning pipeline such as data preprocessing, feature selection, model training, model optimization, and much more. In addition to this, it demonstrates how you can use the available automation libraries, such as auto-sklearn and MLBox, and create and extend your own custom AutoML components for Machine Learning. By the end of this book, you will have a clearer understanding of the different aspects of automated Machine Learning, and you'll be able to incorporate automation tasks using practical datasets. You can leverage your learning from this book to implement Machine Learning in your projects and get a step closer to winning various machine learning competitions. What you will learn

- Understand the fundamentals of Automated Machine Learning systems
- Explore auto-sklearn and MLBox for AutoML tasks
- Automate your preprocessing methods along with feature transformation
- Enhance feature selection and generation using the Python stack
- Assemble individual components of ML into a complete AutoML framework
- Demystify hyperparameter tuning to optimize your ML models
- Dive into Machine Learning concepts such as neural networks and autoencoders
- Understand the information costs and trade-offs associated with AutoML

Who this book is for

If you're a budding data scientist, data analyst, or Machine Learning enthusiast and are new to the concept of automated machine learning, this book is ideal for you. You'll also find this book useful if you're an ML engineer or data professional interested in developing quick machine learning pipelines for your projects. Prior exposure to Python programming will help you get the best out of this book.

Hands-On Machine Learning with Microsoft Excel 2019 *Packt Publishing Ltd*

A practical guide to getting the most out of Excel, using it for data preparation, applying machine learning models (including cloud services) and understanding the outcome of the data analysis. Key Features

- Use Microsoft's product Excel to build advanced forecasting models using varied examples
- Cover range of machine learning tasks such as data mining, data analytics, smart visualization, and more
- Derive data-driven techniques using Excel plugins and APIs without much code required

Book Description

We have made huge progress in teaching computers to perform difficult tasks, especially those that are repetitive and time-consuming for humans. Excel users, of all levels, can feel left behind by this innovation wave. The truth is that a large amount of the work needed to develop and use a machine learning model can be done in Excel. The book starts by giving a general introduction to machine learning, making every concept clear and understandable. Then, it shows every step of a machine learning project, from data collection, reading from different data sources, developing models, and visualizing the results using Excel features and offerings. In every chapter, there are several examples and hands-on exercises that will show the reader how to combine Excel functions, add-ins, and connections to databases and to cloud services to reach the desired goal: building a full data analysis flow. Different machine learning models are shown, tailored to the type of data to be analyzed. At the end of the book, the reader is presented with some advanced use cases using Automated Machine Learning, and artificial neural network, which simplifies the analysis task and represents the future of machine learning. What you will learn

- Use Excel to preview and cleanse datasets
- Understand correlations between variables and optimize the input to machine learning models
- Use and evaluate different machine learning models from Excel
- Understand the use of different visualizations
- Learn the basic concepts and calculations to understand how artificial

neural networks work

Learn how to connect Excel to the Microsoft Azure cloud

Get beyond proof of concepts and build fully functional data analysis flows

Who this book is for

This book is for data analysis, machine learning enthusiasts, project managers, and someone who doesn't want to code much for performing core tasks of machine learning. Each example will help you perform end-to-end smart analytics. Working knowledge of Excel is required.

The Self-Assembling Brain *Princeton University Press*

"In this book, Peter Robin Hiesinger explores historical and contemporary attempts to understand the information needed to make biological and artificial neural networks. Developmental neurobiologists and computer scientists with an interest in artificial intelligence - driven by the promise and resources of biomedical research on the one hand, and by the promise and advances of computer technology on the other - are trying to understand the fundamental principles that guide the generation of an intelligent system. Yet, though researchers in these disciplines share a common interest, their perspectives and approaches are often quite different. The book makes the case that "the information problem" underlies both fields, driving the questions that are driving forward the frontiers, and aims to encourage cross-disciplinary communication and understanding, to help both fields make progress. The questions that challenge researchers in these fields include the following. How does genetic information unfold during the years-long process of human brain development, and can this be a short-cut to create human-level artificial intelligence? Is the biological brain just messy hardware that can be improved upon by running learning algorithms in computers? Can artificial intelligence bypass evolutionary programming of "grown" networks? These questions are tightly linked, and answering them requires an understanding of how information unfolds algorithmically to generate functional neural networks. Via a series of closely linked "discussions" (fictional dialogues between researchers in different disciplines) and pedagogical "seminars," the author explores the different challenges facing researchers working on neural networks, their different perspectives and approaches, as well as the common ground and understanding to be found amongst those sharing an interest in the development of biological brains and artificial intelligent systems"--

Machine Learning in Python: Hands on Machine Learning with Python Tools, Concepts and Techniques *Bob Mather*

The ability to crunch data effectively can propel your career or business to great heights. Machine Learning is the most effective data analysis tool. While it is a complex topic, it can be broken down into simpler steps, as show in this book. We are using Python, which is a great programming language for beginners.

Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow *O'Reilly Media, Inc."*

Through a recent series of 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 bestselling book uses concrete examples, minimal theory, and production-ready Python frameworks (Scikit-Learn, Keras, and TensorFlow) to help you gain an intuitive understanding of the concepts and tools for building intelligent systems. With this updated third edition, author Aurélien Géron explores a range of techniques, starting with simple linear regression and progressing to deep neural networks. Numerous code examples and exercises throughout the book help you apply what you've learned. Programming experience is all you need to get started. Use Scikit-learn to track an example ML project end to end

- Explore several models, including support vector machines, decision trees, random forests, and ensemble methods
- Exploit unsupervised learning techniques such as dimensionality reduction, clustering, and anomaly detection
- Dive into neural net architectures, including convolutional nets, recurrent nets, generative adversarial networks, autoencoders, diffusion models, and transformers
- Use TensorFlow and Keras to build and train neural nets for computer vision, natural language processing, generative models, and deep reinforcement learning

Hands-On Deep Learning for Games *Packt Publishing Ltd*

Understand the core concepts of deep learning and deep reinforcement learning by applying them to develop games

Key Features

- Apply the power of deep learning to complex reasoning tasks by building a Game AI
- Exploit the most recent developments in machine learning and AI for building smart games
- Implement deep learning models and neural networks with Python

Book Description

The number of applications of deep learning and neural networks has multiplied in the last couple of years. Neural nets has enabled significant breakthroughs in everything from computer vision, voice generation, voice recognition and self-driving cars. Game development is also a key area where these techniques are being applied. This book will give an in depth view of the potential of deep learning and neural networks in game development. We will take a look at the foundations of multi-layer perceptron's to using convolutional and recurrent networks. In applications from GANs that create music or textures to self-driving cars and chatbots. Then we introduce deep reinforcement learning through the multi-armed bandit problem and other OpenAI Gym environments. As we progress through the book we will gain insights about DRL techniques such as Motivated Reinforcement Learning with Curiosity and Curriculum Learning. We also take a closer look at deep reinforcement learning and in particular the Unity ML-Agents toolkit. By the end of the book, we will look at how to apply DRL and the ML-Agents toolkit to enhance, test and automate your games or simulations. Finally, we will cover your possible next steps and possible areas for future learning. What you will learn

- Learn the foundations of neural networks and deep learning
- Use advanced neural network architectures in applications to create music, textures, self driving cars and chatbots
- Understand the basics of reinforcement and DRL and how to apply it to solve a variety of problems
- Working with Unity ML-Agents toolkit and how to install, setup and run the kit
- Understand core concepts of DRL and the differences between discrete and continuous action environments
- Use several advanced forms of learning in various scenarios from developing agents to testing games

Who this book is for

This books is for game developers who wish to create highly interactive games by leveraging the power of machine and deep learning. No prior knowledge of machine learning, deep learning or neural networks is required this book will teach those concepts from scratch. A good understanding of Python is required.

Grokking Deep Learning *Simon and Schuster*

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

Hands-On Transfer Learning with Python *Packt Publishing Ltd*

Deep learning simplified by taking supervised, unsupervised, and reinforcement learning to the next level using the Python ecosystem Key Features Build deep learning models with transfer learning principles in Python implement transfer learning to solve real-world research problems Perform complex operations such as image captioning neural style transfer Book Description Transfer learning is a machine learning (ML) technique where knowledge gained during training a set of problems can be used to solve other similar problems. The purpose of this book is two-fold; firstly, we focus on detailed coverage of deep learning (DL) and transfer learning, comparing and contrasting the two with easy-to-follow concepts and examples. The second area of focus is real-world examples and research problems using TensorFlow, Keras, and the Python ecosystem with hands-on examples. The book starts with the key essential concepts of ML and DL, followed by depiction and coverage of important DL architectures such as convolutional neural networks (CNNs), deep neural networks (DNNs), recurrent neural networks (RNNs), long short-term memory (LSTM), and capsule networks. Our focus then shifts to transfer learning concepts, such as model freezing, fine-tuning, pre-trained models including VGG, inception, ResNet, and how these systems perform better than DL models with practical examples. In the concluding chapters, we will focus on a multitude of real-world case studies and problems associated with areas such as computer vision, audio analysis and natural language processing (NLP). By the end of this book, you will be able to implement both DL and transfer learning principles in your own systems. What you will learn Set up your own DL environment with graphics processing unit (GPU) and Cloud support Delve into transfer learning principles with ML and DL models Explore various DL architectures, including CNN, LSTM, and capsule networks Learn about data and network representation and loss functions Get to grips with models and strategies in transfer learning Walk through potential challenges in building complex transfer learning models from scratch Explore real-world research problems related to computer vision and audio analysis Understand how transfer learning can be leveraged in NLP Who this book is for Hands-On Transfer Learning with Python is for data scientists, machine learning engineers, analysts and developers with an interest in data and applying state-of-the-art transfer learning methodologies to solve tough real-world problems. Basic proficiency in machine learning and Python is required.

Hands-on Scikit-Learn for Machine Learning Applications *Apress*

Aspiring data science professionals can learn the Scikit-Learn library along with the fundamentals of machine learning with this book. The book combines the Anaconda Python distribution with the popular Scikit-Learn library to demonstrate a wide range of supervised and unsupervised machine learning algorithms. Care is taken to walk you through the principles of machine learning through clear examples written in Python that you can try out and experiment with at home on your own

machine. All applied math and programming skills required to master the content are covered in this book. In-depth knowledge of object-oriented programming is not required as working and complete examples are provided and explained. Coding examples are in-depth and complex when necessary. They are also concise, accurate, and complete, and complement the machine learning concepts introduced. Working the examples helps to build the skills necessary to understand and apply complex machine learning algorithms. Hands-on Scikit-Learn for Machine Learning Applications is an excellent starting point for those pursuing a career in machine learning. Students of this book will learn the fundamentals that are a prerequisite to competency. Readers will be exposed to the Anaconda distribution of Python that is designed specifically for data science professionals, and will build skills in the popular Scikit-Learn library that underlies many machine learning applications in the world of Python. What You'll Learn Work with simple and complex datasets common to Scikit-Learn Manipulate data into vectors and matrices for algorithmic processing Become familiar with the Anaconda distribution used in data science Apply machine learning with Classifiers, Regressors, and Dimensionality Reduction Tune algorithms and find the best algorithms for each dataset Load data from and save to CSV, JSON, Numpy, and Pandas formats Who This Book Is For The aspiring data scientist yearning to break into machine learning through mastering the underlying fundamentals that are sometimes skipped over in the rush to be productive. Some knowledge of object-oriented programming and very basic applied linear algebra will make learning easier, although anyone can benefit from this book.

PyTorch Deep Learning Hands-On *Packt Publishing Ltd*

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.

Deep Learning for Coders with fastai and PyTorch *"O'Reilly Media, Inc."*

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