
Machine Learning A Hands On Project Based Introdu

Machine Learning

Hands-On Machine Learning with Scikit-Learn and TensorFlow

Hands-On Data Science and Python Machine Learning

Hands-On Machine Learning for Cybersecurity

Hands-On Automated Machine Learning

The TensorFlow Workshop

Machine Learning

PyTorch Deep Learning Hands-On

Hands-On Mathematics for Deep Learning

Hands-On Unsupervised Learning Using Python

Hands-On Machine Learning with C++

Hands-On Machine Learning on Google Cloud Platform

Deep Reinforcement Learning Hands-On

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

Hands-On Machine Learning with R

Hands-On Machine Learning with ML.NET

Hands-On Unsupervised Learning with Python

Hands-On Reinforcement Learning with Python

Hands-On Machine Learning for Algorithmic Trading

Hands-On Machine Learning with Azure

AI Crash Course

Deep Learning By Example

Hands-On Deep Learning with Apache Spark

Deep Reinforcement Learning Hands-On

Hands-on Scikit-Learn for Machine Learning Applications

Hands-On Deep Learning Algorithms with Python

Machine Learning

Machine Learning

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

Hands-On Deep Learning for Games

Hands-on Machine Learning with JavaScript

Hands-On Python Deep Learning for the Web

Deep Learning for Coders with fastai and PyTorch

Pro Machine Learning Algorithms

Hands-On Artificial Intelligence with Java for Beginners

Hands-On Artificial Intelligence for IoT

The Art of Machine Learning

Programming Machine Learning

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

*Machine Learning A
Hands On Project Based
Introdu*

Downloaded from
intra.itu.edu by guest

CESAR CRAWFORD

Machine Learning Packt Publishing Ltd
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 Scikit-Learn and TensorFlow 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 Features
Become familiar with data processing, performance measuring, and model selection using various C++ libraries
Implement practical machine learning and deep learning techniques to build smart models
Deploy machine learning models to work on mobile and embedded devices
Book 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 learn
Explore how to load and preprocess various data types to suitable C++ data structures
Employ key machine learning algorithms with various C++ libraries
Understand the grid-search approach to find the best parameters for a machine learning model
Implement an algorithm for filtering anomalies in user data using Gaussian distribution
Improve collaborative filtering to deal with dynamic user preferences
Use C++ libraries and APIs to manage model structures and parameters
Implement a C++ program to solve image classification tasks with LeNet architecture
Who 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 Data Science and Python

Machine Learning Packt Publishing Ltd

Implement machine learning, cognitive services, and artificial intelligence solutions by leveraging Azure cloud technologies Key Features Learn advanced concepts in Azure ML and the Cortana Intelligence Suite architecture Explore ML Server using SQL Server and HDInsight

capabilities Implement various tools in Azure to build and deploy machine learning models Book 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 learn Discover the benefits of leveraging the cloud for ML and AI Use Cognitive Services APIs to build intelligent bots Build a model using canned algorithms from Microsoft and deploy it as a web service Deploy virtual machines in AI development scenarios Apply R, Python, SQL Server, and Spark in Azure Build and deploy deep learning solutions with CNTK, MMLSpark, and TensorFlow Implement model retraining in IoT, Streaming, and Blockchain solutions Explore best practices for integrating ML and AI functions with ADLA and logic apps Who 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 for Cybersecurity 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
Machine Learning

Understand basic to advanced deep learning algorithms, the mathematical principles behind them, and their practical applications. Key Features Get up-to-speed with building your own neural networks from scratch Gain insights into the mathematical principles behind deep learning algorithms Implement popular deep learning algorithms such as CNNs, RNNs, and more using TensorFlow Book Description Deep learning is one of the most popular domains in the AI space, allowing you to develop multi-layered models of varying complexities. This book introduces you to popular deep learning algorithms—from basic to advanced—and shows you how to implement them from scratch using TensorFlow. Throughout the book, you will gain insights into each algorithm, the mathematical principles behind it, and how to implement it in the best possible manner. The book starts by explaining how you can build your own neural networks, followed by introducing you to TensorFlow, the powerful Python-based library for machine learning and deep learning. Moving on, you will get up to speed with gradient descent variants,

such as NAG, AMSGrad, AdaDelta, Adam, and Nadam. The book will then provide you with insights into RNNs and LSTM and how to generate song lyrics with RNN. Next, you will master the math for convolutional and capsule networks, widely used for image recognition tasks. Then you learn how machines understand the semantics of words and documents using CBOW, skip-gram, and PV-DM. Afterward, you will explore various GANs, including InfoGAN and LSGAN, and autoencoders, such as contractive autoencoders and VAE. By the end of this book, you will be equipped with all the skills you need to implement deep learning in your own projects. What you will learn Implement basic-to-advanced deep learning algorithms Master the mathematics behind deep learning algorithms Become familiar with gradient descent and its variants, such as AMSGrad, AdaDelta, Adam, and Nadam Implement recurrent networks, such as RNN, LSTM, GRU, and seq2seq models Understand how machines interpret images using CNN and capsule networks Implement different types of generative adversarial network, such as CGAN, CycleGAN, and StackGAN Explore various types of autoencoder, such as Sparse autoencoders, DAE, CAE, and VAE Who this book is for If you are a machine learning engineer, data scientist, AI developer, or simply want to focus on neural networks and deep learning, this book is for you. Those who are completely new to deep learning, but have some experience in machine learning and Python programming, will also find the book very helpful. [The TensorFlow Workshop](#) Packt Publishing Ltd Discover the skill-sets required to implement various approaches to

Machine Learning with Python Key Features Explore unsupervised learning with clustering, autoencoders, restricted Boltzmann machines, and more Build your own neural network models using modern Python libraries Practical examples show you how to implement different machine learning and deep learning techniques Book Description Unsupervised learning is about making use of raw, untagged data and applying learning algorithms to it to help a machine predict its outcome. With this book, you will explore the concept of unsupervised learning to cluster large sets of data and analyze them repeatedly until the desired outcome is found using Python. This book starts with the key differences between supervised, unsupervised, and semi-supervised learning. You will be introduced to the best-used libraries and frameworks from the Python ecosystem and address unsupervised learning in both the machine learning and deep learning domains. You will explore various algorithms, techniques that are used to implement unsupervised learning in real-world use cases. You will learn a variety of unsupervised learning approaches, including randomized optimization, clustering, feature selection and transformation, and information theory. You will get hands-on experience with how neural networks can be employed in unsupervised scenarios. You will also explore the steps involved in building and training a GAN in order to process images. By the end of this book, you will have learned the art of unsupervised learning for different real-world challenges. What you will learn Use cluster algorithms to identify and optimize natural groups of data Explore advanced non-linear and hierarchical clustering in action Soft label

assignments for fuzzy c-means and Gaussian mixture models
 Detect anomalies through density estimation
 Perform principal component analysis using neural network models
 Create unsupervised models using GANs
 Who this book is for This book is intended for statisticians, data scientists, machine learning developers, and deep learning practitioners who want to build smart applications by implementing key building block unsupervised learning, and master all the new techniques and algorithms offered in machine learning and deep learning using real-world examples. Some prior knowledge of machine learning concepts and statistics is desirable.

Machine Learning Packt Publishing Ltd

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

PyTorch Deep Learning Hands-On

Packt Publishing Ltd

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

Hands-On Mathematics for Deep Learning Packt Publishing Ltd

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.

Hands-On Unsupervised Learning Using Python "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

C++ 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 Machine Learning on Google Cloud Platform](#) "O'Reilly Media, Inc."

You've decided to tackle machine learning - because you're job hunting, embarking on a new project, or just think self-driving cars are cool. But where to start? It's easy to be intimidated, even as a software developer. The good news is that it doesn't have to be that hard. Master machine learning by writing code one line at a time, from simple learning programs all the way to a true deep learning system. Tackle the hard topics by breaking them down so they're easier to understand, and build your confidence by getting your hands dirty. Peel away the obscurities of machine learning, starting from scratch and going all the way to deep learning. Machine learning can be intimidating, with its reliance on math and algorithms that most programmers don't encounter in their regular work. Take a hands-on approach,

writing the Python code yourself, without any libraries to obscure what's really going on. Iterate on your design, and add layers of complexity as you go. Build an image recognition application from scratch with supervised learning. Predict the future with linear regression. Dive into gradient descent, a fundamental algorithm that drives most of machine learning. Create perceptrons to classify data. Build neural networks to tackle more complex and sophisticated data sets. Train and refine those networks with backpropagation and batching. Layer the neural networks, eliminate overfitting, and add convolution to transform your neural network into a true deep learning system. Start from the beginning and code your way to machine learning mastery. What You Need: The examples in this book are written in Python, but don't worry if you don't know this language: you'll pick up all the Python you need very quickly. Apart from that, you'll only need your computer, and your code-adept brain.

Deep Reinforcement Learning Hands-On
Packt Publishing Ltd

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 Machine Learning with scikit-learn and Scientific Python Toolkits Createspace Independent Publishing Platform

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 R
John Wiley & Sons

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 Machine Learning with ML.NET](#)
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 Unsupervised Learning with Python Apress

Get into the world of smart data security using machine learning algorithms and Python libraries Key Features Learn machine learning algorithms and cybersecurity fundamentals Automate your daily workflow by applying use cases to many facets of security Implement smart machine learning solutions to detect various cybersecurity problems Book 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 learn Use machine learning algorithms with complex datasets to

implement cybersecurity concepts Implement machine learning algorithms such as clustering, k-means, and Naive Bayes to solve real-world problems Learn to speed up a system using Python libraries with NumPy, Scikit-learn, and CUDA Understand how to combat malware, detect spam, and fight financial fraud to mitigate cyber crimes Use TensorFlow in the cybersecurity domain and implement real-world examples Learn how machine learning and Python can be used in complex cyber issues Who 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 Reinforcement Learning with Python Packt Publishing Ltd
Machine Learning Createspace
Independent Publishing Platform

Hands-On Machine Learning for Algorithmic Trading Packt Publishing Ltd

Build smarter systems by combining artificial intelligence and the Internet of Things—two of the most talked about topics today Key Features Leverage the power of Python libraries such as TensorFlow and Keras to work with real-time IoT data Process IoT data and predict outcomes in real time to build smart IoT models Cover practical case studies on industrial IoT, smart cities, and home automation Book Description There are many applications that use data science and analytics to gain insights from terabytes of data. These apps, however, do not address the challenge of continually discovering patterns for IoT data. In Hands-On

Artificial Intelligence for IoT, we cover various aspects of artificial intelligence (AI) and its implementation to make your IoT solutions smarter. This book starts by covering the process of gathering and preprocessing IoT data gathered from distributed sources. You will learn different AI techniques such as machine learning, deep learning, reinforcement learning, and natural language processing to build smart IoT systems. You will also leverage the power of AI to handle real-time data coming from wearable devices. As you progress through the book, techniques for building models that work with different kinds of data generated and consumed by IoT devices such as time series, images, and audio will be covered. Useful case studies on four major application areas of IoT solutions are a key focal point of this book. In the concluding chapters, you will leverage the power of widely used Python libraries, TensorFlow and Keras, to build different kinds of smart AI models. By the end of this book, you will be able to build smart AI-powered IoT apps with confidence. What you will learn Apply different AI techniques including machine learning and deep learning using TensorFlow and Keras Access and process data from various distributed sources Perform supervised and unsupervised machine learning for IoT data Implement distributed processing of IoT data over Apache Spark using the MLLib and H2O.ai platforms Forecast time-series data using deep learning methods Implementing AI from case studies in Personal IoT, Industrial IoT, and Smart Cities Gain unique insights from data obtained from wearable devices and smart devices Who this book is for If you are a data science professional or a machine learning

developer looking to build smart systems for IoT, Hands-On Artificial Intelligence for IoT is for you. If you want to learn how popular artificial intelligence (AI) techniques can be used in the Internet of Things domain, this book will also be of benefit. A basic understanding of machine learning concepts will be required to get the best out of this book.

Hands-On Machine Learning with Azure Packt Publishing Ltd

Get started with TensorFlow fundamentals to build and train deep learning models with real-world data, practical exercises, and challenging activities

Key Features

- Understand the fundamentals of tensors, neural networks, and deep learning
- Discover how to implement and fine-tune deep learning models for real-world datasets
- Build your experience and confidence with hands-on exercises and activities

Book Description

Getting to grips with tensors, deep learning, and neural networks can be intimidating and confusing for anyone, no matter their experience level. The breadth of information out there, often written at a very high level and aimed at advanced practitioners, can make getting started even more challenging. If this sounds familiar to you, The TensorFlow Workshop is here to help. Combining clear explanations, realistic examples, and plenty of hands-on practice, it'll quickly get you up and running. You'll start off with the basics - learning how to load data into TensorFlow, perform tensor operations, and utilize common optimizers and activation functions. As

you progress, you'll experiment with different TensorFlow development tools, including TensorBoard, TensorFlow Hub, and Google Colab, before moving on to solve regression and classification problems with sequential models. Building on this solid foundation, you'll learn how to tune models and work with different types of neural network, getting hands-on with real-world deep learning applications such as text encoding, temperature forecasting, image augmentation, and audio processing. By the end of this deep learning book, you'll have the skills, knowledge, and confidence to tackle your own ambitious deep learning projects with TensorFlow.

What you will learn

- Get to grips with TensorFlow's mathematical operations
- Pre-process a wide variety of tabular, sequential, and image data
- Understand the purpose and usage of different deep learning layers
- Perform hyperparameter-tuning to prevent overfitting of training data
- Use pre-trained models to speed up the development of learning models
- Generate new data based on existing patterns using generative models

Who this book is for

This TensorFlow book is for anyone who wants to develop their understanding of deep learning and get started building neural networks with TensorFlow. Basic knowledge of Python programming and its libraries, as well as a general understanding of the fundamentals of data science and machine learning, will help you grasp the topics covered in this book more easily.

Best Sellers - Books :

- [How To Win Friends & Influence People \(dale Carnegie Books\) By Dale Carnegie](#)
- [The Democrat Party Hates America](#)
- [Haunting Adeline \(cat And Mouse Duet\)](#)
- [Remarkably Bright Creatures: A Read With Jenna Pick By Shelby Van Pelt](#)

- [The Summer I Turned Pretty \(summer I Turned Pretty, The\)](#)
- [Chicka Chicka Boom Boom \(board Book\) By Bill Martin Jr.](#)
- [The 5 Love Languages: The Secret To Love That Lasts By Gary Chapman](#)
- [The Alchemist, 25th Anniversary: A Fable About Following Your Dream By Paulo Coelho](#)
- [Outlive: The Science And Art Of Longevity By Peter Attia Md](#)
- [I Love You Like No Otter: A Funny And Sweet Board Book For Babies And Toddlers \(punderland\)](#)