

Python Machine Learning Cookbook Practical Soluti

If you ally compulsion such a referred **Python Machine Learning Cookbook Practical Soluti** books that will have the funds for you worth, acquire the categorically best seller from us currently from several preferred authors. If you desire to hilarious books, lots of novels, tale, jokes, and more fictions collections are next launched, from best seller to one of the most current released.

You may not be perplexed to enjoy every books collections Python Machine Learning Cookbook Practical Soluti that we will extremely offer. It is not as regards the costs. Its approximately what you craving currently. This Python Machine Learning Cookbook Practical Soluti, as one of the most functional sellers here will very be in the midst of the best options to review.



[Python Cookbook](#) Packt Publishing Ltd

Master the essential skills needed to recognize and solve complex problems with machine learning and deep learning. Using real-world examples that leverage the popular Python machine learning ecosystem, this book is your perfect companion for learning the art and science of machine learning to become a successful practitioner. The concepts, techniques, tools, frameworks, and methodologies used in this book will teach you how to think, design, build, and execute machine learning systems and projects successfully. Practical Machine Learning with Python follows a structured and comprehensive three-tiered approach packed with hands-on examples and code. Part 1 focuses on understanding machine learning concepts and tools. This includes machine learning basics with a broad overview of algorithms, techniques, concepts and applications, followed by a tour of the entire Python machine learning ecosystem. Brief guides for useful machine learning tools, libraries and frameworks are also covered. Part 2 details standard machine learning pipelines, with an emphasis on data processing analysis, feature engineering, and modeling. You will learn how to process, wrangle, summarize and visualize data in its various forms. Feature engineering and selection methodologies will be covered in detail with real-world datasets followed by model building, tuning, interpretation and deployment. Part 3 explores multiple real-world case studies spanning diverse domains and industries like retail, transportation, movies, music, marketing, computer vision and finance. For each case study, you will learn the application of various machine learning techniques and methods. The hands-on examples will help you become familiar with state-of-the-art machine learning tools and techniques and understand what algorithms are best suited for any problem. Practical Machine Learning with Python will empower you to start solving your own problems with machine learning today! What You'll Learn Execute end-to-end machine learning projects and systems Implement hands-on examples with industry standard, open source, robust machine learning tools and frameworks Review case studies depicting applications of machine learning and deep learning on diverse domains and industries Apply a wide range of machine learning models including regression, classification, and clustering. Understand and apply the latest models and methodologies from deep learning including CNNs, RNNs, LSTMs and transfer learning. Who This Book Is For IT professionals, analysts, developers, data scientists, engineers, graduate students

Deep Learning with Python Packt Publishing Ltd

From news and speeches to informal chatter on social media, natural language is one of the richest and most underutilized sources of data. Not only does it come in a constant stream, always changing and adapting in context; it also contains information that is not conveyed by traditional data sources. The key to unlocking natural language is through the creative application of text analytics. This practical book presents a data scientist's approach to building language-aware products with applied machine learning. You'll learn robust, repeatable, and scalable techniques for text analysis with Python, including contextual and linguistic feature engineering, vectorization, classification, topic modeling, entity resolution, graph analysis, and visual steering. By the end of the book, you'll be equipped with practical methods to solve any number of complex real-world problems. Preprocess and vectorize text into high-dimensional feature representations Perform document classification and topic modeling Steer the model selection process with visual diagnostics Extract key phrases, named entities, and graph structures to reason about data in text Build a dialog framework to enable chatbots and language-driven interaction Use Spark to scale processing power and neural networks to scale model complexity

Python Feature Engineering Cookbook Packt Publishing Ltd

Summary Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and

practical examples. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning has made remarkable progress in recent years. We went from near-unusable speech and image recognition, to near-human accuracy. We went from machines that couldn't beat a serious Go player, to defeating a world champion. Behind this progress is deep learning—a combination of engineering advances, best practices, and theory that enables a wealth of previously impossible smart applications. About the Book Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. You'll explore challenging concepts and practice with applications in computer vision, natural-language processing, and generative models. By the time you finish, you'll have the knowledge and hands-on skills to apply deep learning in your own projects. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image-classification models Deep learning for text and sequences Neural style transfer, text generation, and image generation About the Reader Readers need intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the Author François Chollet works on deep learning at Google in Mountain View, CA. He is the creator of the Keras deep-learning library, as well as a contributor to the TensorFlow machine-learning framework. He also does deep-learning research, with a focus on computer vision and the application of machine learning to formal reasoning. His papers have been published at major conferences in the field, including the Conference on Computer Vision and Pattern Recognition (CVPR), the Conference and Workshop on Neural Information Processing Systems (NIPS), the International Conference on Learning Representations (ICLR), and others. Table of Contents PART 1 -

FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of neural networks Getting started with neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE Deep learning for computer vision Deep learning for text and sequences Advanced deep-learning best practices Generative deep learning Conclusions appendix A - Installing Keras and its dependencies on Ubuntu appendix B - Running Jupyter notebooks on an EC2 GPU instance [Python Deep Learning Cookbook](#) Packt Publishing Ltd

Discover powerful ways to effectively solve real-world machine learning problems using key libraries including scikit-learn, TensorFlow, and PyTorch Key Features Learn and implement machine learning algorithms in a variety of real-life scenarios Cover a range of tasks catering to supervised, unsupervised and reinforcement learning techniques Find easy-to-follow code solutions for tackling common and not-so-common challenges Book Description This eagerly anticipated second edition of the popular Python Machine Learning Cookbook will enable you to adopt a fresh approach to dealing with real-world machine learning and deep learning tasks. With the help of over 100 recipes, you will learn to build powerful machine learning applications using modern libraries from the Python ecosystem. The book will also guide you on how to implement various machine learning algorithms for classification, clustering, and recommendation engines, using a recipe-based approach. With emphasis on practical solutions, dedicated sections in the book will help you to apply supervised and unsupervised learning techniques to real-world problems. Toward the concluding chapters, you will get to grips with recipes that teach you advanced techniques including reinforcement learning, deep neural networks, and automated machine learning. By the end of this book, you will be equipped with the skills you need to apply machine learning techniques and leverage the full capabilities of the Python ecosystem through real-world examples. What you will learn Use predictive modeling and apply it to real-world problems Explore data visualization techniques to interact with your data Learn how to build a recommendation engine Understand how to interact with text data and build models to analyze it Work with speech data and recognize spoken words using Hidden Markov Models Get well versed with reinforcement learning, automated ML, and transfer learning Work with image data and build systems for image recognition and biometric face recognition Use deep neural networks to build an optical character recognition system Who this book is for This book is for data scientists, machine learning developers, deep learning enthusiasts and Python programmers who want to solve real-world challenges using machine-learning techniques and algorithms. If you are facing challenges at work and want ready-to-use code solutions to cover key tasks in machine learning and the deep learning domain, then this book is what you need. Familiarity with Python programming and machine learning concepts will be useful.

[Python Machine Learning Cookbook](#) O'Reilly Media

Discover powerful ways to effectively solve real-world machine learning problems using

key libraries including scikit-learn, TensorFlow, and PyTorch Key Features Learn and implement machine learning algorithms in a variety of real-life scenarios Cover a range of tasks catering to supervised, unsupervised and reinforcement learning techniques Find easy-to-follow code solutions for tackling common and not-so-common challenges Book Description This eagerly anticipated second edition of the popular Python Machine Learning Cookbook will enable you to adopt a fresh approach to dealing with real-world machine learning and deep learning tasks. With the help of over 100 recipes, you will learn to build powerful machine learning applications using modern libraries from the Python ecosystem. The book will also guide you on how to implement various machine learning algorithms for classification, clustering, and recommendation engines, using a recipe-based approach. With emphasis on practical solutions, dedicated sections in the book will help you to apply supervised and unsupervised learning techniques to real-world problems. Toward the concluding chapters, you will get to grips with recipes that teach you advanced techniques including reinforcement learning, deep neural networks, and automated machine learning. By the end of this book, you will be equipped with the skills you need to apply machine learning techniques and leverage the full capabilities of the Python ecosystem through real-world examples. What you will learn Use predictive modeling and apply it to real-world problems Explore data visualization techniques to interact with your data Learn how to build a recommendation engine Understand how to interact with text data and build models to analyze it Work with speech data and recognize spoken words using Hidden Markov Models Get well versed with reinforcement learning, automated ML, and transfer learning Work with image data and build systems for image recognition and biometric face recognition Use deep neural networks to build an optical character recognition system Who this book is for This book is for data scientists, machine learning developers, deep learning enthusiasts and Python programmers who want to solve real-world challenges using machine-learning techniques and algorithms. If you are facing challenges at work and want ready-to-use code solutions to cover key tasks in machine learning and the deep learning domain, then this book is what you need. Familiarity with Python programming and machine learning concepts will be useful.

[Python Machine Learning](#) Packt Publishing Ltd

Over 85 recipes to help you complete real-world data science projects in R and Python About This Book Tackle every step in the data science pipeline and use it to acquire, clean, analyze, and visualize your data Get beyond the theory and implement real-world projects in data science using R and Python Easy-to-follow recipes will help you understand and implement the numerical computing concepts Who This Book Is For If you are an aspiring data scientist who wants to learn data science and numerical programming concepts through hands-on, real-world project examples, this is the book for you. Whether you are brand new to data science or you are a seasoned expert, you will benefit from learning about the structure of real-world data science projects and the programming examples in R and Python. What You Will Learn Learn and understand the installation procedure and environment required for R and Python on various platforms Prepare data for analysis by implement various data science concepts such as acquisition, cleaning and munging through R and Python Build a predictive model and an exploratory model Analyze the results of your model and create reports on the acquired data Build various tree-based methods and Build random forest In Detail As increasing amounts of data are generated each year, the need to analyze and create value out of it is more important than ever. Companies that know what to do with their data and how to do it well will have a competitive advantage over companies that don't. Because of this, there will be an increasing demand for people that possess both the analytical and technical abilities to extract valuable insights from data and create valuable solutions that put those insights to use. Starting with the basics, this book covers how to set up your numerical programming environment, introduces you to the data science pipeline, and guides you through several data projects in a step-by-step format. By sequentially working through the steps in each chapter, you will quickly familiarize yourself with the process and learn how to apply it to a variety of situations with examples using the two most popular programming languages for data analysis—R and Python. Style and approach This step-by-step guide to data science is full of hands-on examples of real-world data science tasks. Each recipe focuses on a particular task involved in the data

science pipeline, ranging from readying the dataset to analytics and visualization
[Practical Data Science Cookbook](#) Packt Publishing Ltd

Leverage the power of deep learning and Keras to develop smarter and more efficient data models
Key Features
Understand different neural networks and their implementation using Keras
Explore recipes for training and fine-tuning your neural network models
Put your deep learning knowledge to practice with real-world use-cases, tips, and tricks
Book Description
Keras has quickly emerged as a popular deep learning library. Written in Python, it allows you to train convolutional as well as recurrent neural networks with speed and accuracy. The Keras Deep Learning Cookbook shows you how to tackle different problems encountered while training efficient deep learning models, with the help of the popular Keras library. Starting with installing and setting up Keras, the book demonstrates how you can perform deep learning with Keras in the TensorFlow. From loading data to fitting and evaluating your model for optimal performance, you will work through a step-by-step process to tackle every possible problem faced while training deep models. You will implement convolutional and recurrent neural networks, adversarial networks, and more with the help of this handy guide. In addition to this, you will learn how to train these models for real-world image and language processing tasks. By the end of this book, you will have a practical, hands-on understanding of how you can leverage the power of Python and Keras to perform effective deep learning
What you will learn
Install and configure Keras in TensorFlow
Master neural network programming using the Keras library
Understand the different Keras layers
Use Keras to implement simple feed-forward neural networks, CNNs and RNNs
Work with various datasets and models used for image and text classification
Develop text summarization and reinforcement learning models using Keras
Who this book is for
Keras Deep Learning Cookbook is for you if you are a data scientist or machine learning expert who wants to find practical solutions to common problems encountered while training deep learning models. A basic understanding of Python and some experience in machine learning and neural networks is required for this book.

Machine Learning Using TensorFlow Cookbook Apress

Solve different problems in modelling deep neural networks using Python, Tensorflow, and Keras with this practical guide
About This Book
* Practical recipes on training different neural network models and tuning them for optimal performance
* Use Python frameworks like TensorFlow, Caffe, Keras, Theano for Natural Language Processing, Computer Vision, and more
* A hands-on guide covering the common as well as the not so common problems in deep learning using Python
Who This Book Is For
This book is intended for machine learning professionals who are looking to use deep learning algorithms to create real-world applications using Python. Thorough understanding of the machine learning concepts and Python libraries such as NumPy, SciPy and scikit-learn is expected. Additionally, basic knowledge in linear algebra and calculus is desired.
What You Will Learn
* Implement different neural network models in Python
* Select the best Python framework for deep learning such as PyTorch, Tensorflow, MXNet and Keras
* Apply tips and tricks related to neural networks internals, to boost learning performances
* Consolidate machine learning principles and apply them in the deep learning field
* Reuse and adapt Python code snippets to everyday problems
* Evaluate the cost/benefits and performance implication of each discussed solution
In Detail
Deep Learning is revolutionizing a wide range of industries. For many applications, deep learning has proven to outperform humans by making faster and more accurate predictions. This book provides a top-down and bottom-up approach to demonstrate deep learning solutions to real-world problems in different areas. These applications include Computer Vision, Natural Language Processing, Time Series, and Robotics.
The Python Deep Learning Cookbook presents technical solutions to the issues presented, along with a detailed explanation of the solutions. Furthermore, a discussion on corresponding pros and cons of implementing the proposed solution using one of the popular frameworks like TensorFlow, PyTorch, Keras and CNTK is provided. The book includes recipes that are related to the basic concepts of neural networks. All techniques s, as well as classical networks topologies. The main purpose of this book is to provide Python programmers a detailed list of recipes to apply deep learning to common and not-so-common scenarios.
Style and approach
Unique blend of independent recipes arranged in the most logical manner

[Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow](#) O'Reilly Media

Get hands-on experience in creating state-of-the-art reinforcement learning agents using TensorFlow and RLlib to solve complex real-world business and industry problems with the help of expert tips and best practices
Key Features
Understand how large-scale state-of-the-art RL algorithms and approaches work
Apply RL to solve complex problems in marketing, robotics, supply chain, finance, cybersecurity, and more
Explore tips and best practices from experts that will enable you to overcome real-world RL challenges
Book Description
Reinforcement learning (RL) is a field of artificial intelligence (AI) used for creating self-learning autonomous agents. Building on a strong theoretical foundation, this book takes a practical approach and uses examples inspired by real-world industry problems to teach you about state-of-the-art RL. Starting with bandit problems, Markov decision processes, and dynamic programming, the book provides an in-depth review of the classical RL techniques, such as Monte Carlo methods and temporal-difference

learning. After that, you will learn about deep Q-learning, policy gradient algorithms, actor-critic methods, model-based methods, and multi-agent reinforcement learning. Then, you'll be introduced to some of the key approaches behind the most successful RL implementations, such as domain randomization and curiosity-driven learning. As you advance, you ' ll explore many novel algorithms with advanced implementations using modern Python libraries such as TensorFlow and Ray ' s RLlib package. You ' ll also find out how to implement RL in areas such as robotics, supply chain management, marketing, finance, smart cities, and cybersecurity while assessing the trade-offs between different approaches and avoiding common pitfalls. By the end of this book, you ' ll have mastered how to train and deploy your own RL agents for solving RL problems. What you will learn
Model and solve complex sequential decision-making problems using RL
Develop a solid understanding of how state-of-the-art RL methods work
Use Python and TensorFlow to code RL algorithms from scratch
Parallelize and scale up your RL implementations using Ray's RLlib package
Get in-depth knowledge of a wide variety of RL topics
Understand the trade-offs between different RL approaches
Discover and address the challenges of implementing RL in the real world
Who this book is for
This book is for expert machine learning practitioners and researchers looking to focus on hands-on reinforcement learning with Python by implementing advanced deep reinforcement learning concepts in real-world projects. Reinforcement learning experts who want to advance their knowledge to tackle large-scale and complex sequential decision-making problems will also find this book useful. Working knowledge of Python programming and deep learning along with prior experience in reinforcement learning is required.

Machine Learning with Python Cookbook Packt Publishing Ltd

Have you come across the terms machine learning and neural networks in most articles you have recently read? Do you also want to learn how to build a machine learning model that will answer your questions within a blink of your eyes? If you responded yes to any of the above questions, you have come to the right place.

Artificial Intelligence with Python Cookbook Packt Publishing Ltd

Solve different problems in modelling deep neural networks using Python, Tensorflow, and Keras with this practical guide
About This Book
Practical recipes on training different neural network models and tuning them for optimal performance
Use Python frameworks like TensorFlow, Caffe, Keras, Theano for Natural Language Processing, Computer Vision, and more
A hands-on guide covering the common as well as the not so common problems in deep learning using Python
Who This Book Is For
This book is intended for machine learning professionals who are looking to use deep learning algorithms to create real-world applications using Python. Thorough understanding of the machine learning concepts and Python libraries such as NumPy, SciPy and scikit-learn is expected. Additionally, basic knowledge in linear algebra and calculus is desired.
What You Will Learn
Implement different neural network models in Python
Select the best Python framework for deep learning such as PyTorch, Tensorflow, MXNet and Keras
Apply tips and tricks related to neural networks internals, to boost learning performances
Consolidate machine learning principles and apply them in the deep learning field
Reuse and adapt Python code snippets to everyday problems
Evaluate the cost/benefits and performance implication of each discussed solution
In Detail
Deep Learning is revolutionizing a wide range of industries. For many applications, deep learning has proven to outperform humans by making faster and more accurate predictions. This book provides a top-down and bottom-up approach to demonstrate deep learning solutions to real-world problems in different areas. These applications include Computer Vision, Natural Language Processing, Time Series, and Robotics. The Python Deep Learning Cookbook presents technical solutions to the issues presented, along with a detailed explanation of the solutions. Furthermore, a discussion on corresponding pros and cons of implementing the proposed solution using one of the popular frameworks like TensorFlow, PyTorch, Keras and CNTK is provided. The book includes recipes that are related to the basic concepts of neural networks. All techniques s, as well as classical networks topologies. The main purpose of this book is to provide Python programmers a detailed list of recipes to apply deep learning to common and not-so-common scenarios.
Style and approach
Unique blend of independent recipes arranged in the most logical manner

Apache Spark Deep Learning Cookbook Packt Publishing Ltd

Whether you ' re a software engineer aspiring to enter the world of deep learning, a veteran data scientist, or a hobbyist with a simple dream of making the next viral

AI app, you might have wondered where to begin. This step-by-step guide teaches you how to build practical deep learning applications for the cloud, mobile, browsers, and edge devices using a hands-on approach. Relying on years of industry experience transforming deep learning research into award-winning applications, Anirudh Koul, Siddha Ganju, and Meher Kasam guide you through the process of converting an idea into something that people in the real world can use. Train, tune, and deploy computer vision models with Keras, TensorFlow, Core ML, and TensorFlow Lite
Develop AI for a range of devices including Raspberry Pi, Jetson Nano, and Google Coral
Explore fun projects, from Silicon Valley ' s Not Hotdog app to 40+ industry case studies
Simulate an autonomous car in a video game environment and build a miniature version with reinforcement learning
Use transfer learning to train models in minutes
Discover 50+ practical tips for maximizing model accuracy and speed, debugging, and scaling to millions of users
[Machine Learning with Amazon SageMaker Cookbook](#) Packt Publishing Ltd

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 Aur é lien G é ron 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

Mastering Reinforcement Learning with Python Packt Publishing Ltd

Comprehensive recipes to give you valuable insights on Transformers, Reinforcement Learning, and more
Key Features
Deep Learning solutions from Kaggle Masters and Google Developer Experts
Get to grips with the fundamentals including variables, matrices, and data sources
Learn advanced techniques to make your algorithms faster and more accurate
Book Description
The independent recipes in Machine Learning Using TensorFlow Cookbook will teach you how to perform complex data computations and gain valuable insights into your data. Dive into recipes on training models, model evaluation, sentiment analysis, regression analysis, artificial neural networks, and deep learning - each using Google ' s machine learning library, TensorFlow. This cookbook covers the fundamentals of the TensorFlow library, including variables, matrices, and various data sources. You ' ll discover real-world implementations of Keras and TensorFlow and learn how to use estimators to train linear models and boosted trees, both for classification and regression. Explore the practical applications of a variety of deep learning architectures, such as recurrent neural networks and Transformers, and see how they can be used to solve computer vision and natural language processing (NLP) problems. With the help of this book, you will be proficient in using TensorFlow, understand deep learning from the basics, and be able to implement machine learning algorithms in real-world scenarios. What you will learn
Take TensorFlow into production
Implement and fine-tune Transformer models for various NLP tasks
Apply reinforcement learning algorithms using the TF-Agents framework
Understand linear regression techniques and use Estimators to train linear models
Execute neural networks and improve predictions on tabular data
Master convolutional neural networks and recurrent neural networks through practical recipes
Who this book is for
If you are a data scientist or a machine learning engineer, and you want to skip detailed theoretical explanations in favor of building production-ready machine learning models using TensorFlow, this book is for you. Basic familiarity with Python, linear algebra, statistics, and machine learning is necessary to make the most out of this book.

Applied Text Analysis with Python Packt Publishing Ltd

A solution-based guide to put your deep learning models into production with the power of Apache Spark
Key Features
Discover practical recipes for distributed deep learning

with Apache Spark Learn to use libraries such as Keras and TensorFlow Solve problems in order to train your deep learning models on Apache Spark Book Description With deep learning gaining rapid mainstream adoption in modern-day industries, organizations are looking for ways to unite popular big data tools with highly efficient deep learning libraries. As a result, this will help deep learning models train with higher efficiency and speed. With the help of the Apache Spark Deep Learning Cookbook, you ' ll work through specific recipes to generate outcomes for deep learning algorithms, without getting bogged down in theory. From setting up Apache Spark for deep learning to implementing types of neural net, this book tackles both common and not so common problems to perform deep learning on a distributed environment. In addition to this, you ' ll get access to deep learning code within Spark that can be reused to answer similar problems or tweaked to answer slightly different problems. You will also learn how to stream and cluster your data with Spark. Once you have got to grips with the basics, you ' ll explore how to implement and deploy deep learning models, such as Convolutional Neural Networks (CNN) and Recurrent Neural Networks (RNN) in Spark, using popular libraries such as TensorFlow and Keras. By the end of the book, you'll have the expertise to train and deploy efficient deep learning models on Apache Spark. What you will learn Set up a fully functional Spark environment Understand practical machine learning and deep learning concepts Apply built-in machine learning libraries within Spark Explore libraries that are compatible with TensorFlow and Keras Explore NLP models such as Word2vec and TF-IDF on Spark Organize dataframes for deep learning evaluation Apply testing and training modeling to ensure accuracy Access readily available code that may be reusable Who this book is for If you ' re looking for a practical and highly useful resource for implementing efficiently distributed deep learning models with Apache Spark, then the Apache Spark Deep Learning Cookbook is for you. Knowledge of the core machine learning concepts and a basic understanding of the Apache Spark framework is required to get the best out of this book. Additionally, some programming knowledge in Python is a plus.

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

Deep learning doesn ' t have to be intimidating. Until recently, this machine-learning method required years of study, but with frameworks such as Keras and Tensorflow, software engineers without a background in machine learning can quickly enter the field. With the recipes in this cookbook, you ' ll learn how to solve deep-learning problems for classifying and generating text, images, and music. Each chapter consists of several recipes needed to complete a single project, such as training a music recommending system. Author Douwe Osinga also provides a chapter with half a dozen techniques to help you if you ' re stuck. Examples are written in Python with code available on GitHub as a set of Python notebooks. You ' ll learn how to: Create applications that will serve real users Use word embeddings to calculate text similarity Build a movie recommender system based on Wikipedia links Learn how AIs see the world by visualizing their internal state Build a model to suggest emojis for pieces of text Reuse pretrained networks to build an inverse image search service Compare how GANs, autoencoders and LSTMs generate icons Detect music styles and index song collections

[Hands-On Unsupervised Learning Using Python](#) BPB Publications

Deep learning doesn't have to be intimidating. Until recently, this machine-learning method required years of study, but with frameworks such as Keras and Tensorflow, software engineers without a background in machine learning can quickly enter the field. With the recipes in this cookbook, you'll learn how to solve deep-learning problems for classifying and generating text, images, and music. Each chapter consists of several recipes needed to complete a single project, such as training a music recommending system. Author Douwe Osinga also provides a chapter with half a dozen techniques to help you if you're stuck. Examples are written in Python with code available on GitHub as a set of Python notebooks. You'll learn how to: Create applications that will serve real users Use word embeddings to calculate text similarity Build a movie recommender system based on Wikipedia links Learn how AIs see the world by visualizing their internal state Build a model to suggest emojis for pieces of text Reuse pretrained networks to build an inverse image search service Compare how GANs, autoencoders and LSTMs generate icons Detect music styles and index song collections

Python Machine Learning Packt Publishing Ltd

Providing code examples in python, this book introduces the concepts of machine learning with mathematical explanations and programming fundamentals. --

Machine Learning for Time-Series with Python Packt Publishing

Explore machine learning concepts using the latest numerical computing library —

TensorFlow — with the help of this comprehensive cookbook About This Book

Your quick guide to implementing TensorFlow in your day-to-day machine

learning activities Learn advanced techniques that bring more accuracy and speed

to machine learning Upgrade your knowledge to the second generation of machine

learning with this guide on TensorFlow Who This Book Is For This book is ideal

for data scientists who are familiar with C++ or Python and perform machine

learning activities on a day-to-day basis. Intermediate and advanced machine

learning implementers who need a quick guide they can easily navigate will find it useful. What You Will Learn Become familiar with the basics of the TensorFlow machine learning library Get to know Linear Regression techniques with TensorFlow Learn SVMs with hands-on recipes Implement neural networks and improve predictions Apply NLP and sentiment analysis to your data Master CNN and RNN through practical recipes Take TensorFlow into production In Detail TensorFlow is an open source software library for Machine Intelligence. The independent recipes in this book will teach you how to use TensorFlow for complex data computations and will let you dig deeper and gain more insights into your data than ever before. You'll work through recipes on training models, model evaluation, sentiment analysis, regression analysis, clustering analysis, artificial neural networks, and deep learning – each using Google's machine learning library TensorFlow. This guide starts with the fundamentals of the TensorFlow library which includes variables, matrices, and various data sources. Moving ahead, you will get hands-on experience with Linear Regression techniques with TensorFlow. The next chapters cover important high-level concepts such as neural networks, CNN, RNN, and NLP. Once you are familiar and comfortable with the TensorFlow ecosystem, the last chapter will show you how to take it to production. Style and approach This book takes a recipe-based approach where every topic is explicated with the help of a real-world example.

[TensorFlow Machine Learning Cookbook](#) "O'Reilly Media, Inc."

Machine learning has become an integral part of many commercial applications and research projects, but this field is not exclusive to large companies with extensive research teams. If you use Python, even as a beginner, this book will teach you practical ways to build your own machine learning solutions. With all the data available today, machine learning applications are limited only by your imagination. You ' ll learn the steps necessary to create a successful machine-learning application with Python and the scikit-learn library. Authors Andreas Müller and Sarah Guido focus on the practical aspects of using machine learning algorithms, rather than the math behind them. Familiarity with the NumPy and matplotlib libraries will help you get even more from this book. With this book, you ' ll learn: Fundamental concepts and applications of machine learning Advantages and shortcomings of widely used machine learning algorithms How to represent data processed by machine learning, including which data aspects to focus on Advanced methods for model evaluation and parameter tuning The concept of pipelines for chaining models and encapsulating your workflow Methods for working with text data, including text-specific processing techniques Suggestions for improving your machine learning and data science skills