
Machine Learning With Python The Ultimate Guide F

If you ally infatuation such a referred **Machine Learning With Python The Ultimate Guide F** book that will find the money for you worth, get the extremely best seller from us currently from several preferred authors. If you want to witty books, lots of novels, tale, jokes, and more fictions collections are moreover launched, from best seller to one of the most current released.

You may not be perplexed to enjoy every ebook collections Machine Learning With Python The Ultimate Guide F that we will enormously offer. It is not in this area the costs. Its not quite what you obsession currently. This Machine Learning With Python The Ultimate Guide F, as one of the most practicing sellers here will enormously be in the middle of the best options to review.

*Machine Learning With Python The
Ultimate Guide F*

2019-07-19

REEVES BRADFORD

Python Machine Learning Packt Publishing Ltd

This is a tutorial-driven and practical, but well-grounded book showcasing good Machine Learning practices. There will be an emphasis on using existing technologies instead of showing how to write your own implementations of algorithms. This book is a scenario-based, example-driven tutorial. By the end of the book you will have learnt critical aspects of Machine Learning Python projects and experienced the power of ML-based systems by actually working on them. This book primarily targets Python developers who want to learn about and build Machine Learning into their projects, or who want to pro.

Artificial Intelligence with Python Packt Publishing

Take tiny steps to enter the big world of data science through this interesting guide About This Book Learn the fundamentals of

machine learning and build your own intelligent applications Master the art of building your own machine learning systems with this example-based practical guide Work with important classification and regression algorithms and other machine learning techniques Who This Book Is For This book is for anyone interested in entering the data science stream with machine learning. Basic familiarity with Python is assumed. What You Will Learn Exploit the power of Python to handle data extraction, manipulation, and exploration techniques Use Python to visualize data spread across multiple dimensions and extract useful features Dive deep into the world of analytics to predict situations correctly Implement machine learning classification and regression algorithms from scratch in Python Be amazed to see the algorithms in action Evaluate the performance of a machine learning model and optimize it Solve interesting real-world problems using machine learning and Python as the journey unfolds In Detail Data science and machine learning are some of the top buzzwords in the technical world today. A resurging

interest in machine learning is due to the same factors that have made data mining and Bayesian analysis more popular than ever. This book is your entry point to machine learning. This book starts with an introduction to machine learning and the Python language and shows you how to complete the setup. Moving ahead, you will learn all the important concepts such as, exploratory data analysis, data preprocessing, feature extraction, data visualization and clustering, classification, regression and model performance evaluation. With the help of various projects included, you will find it intriguing to acquire the mechanics of several important machine learning algorithms – they are no more obscure as they thought. Also, you will be guided step by step to build your own models from scratch. Toward the end, you will gather a broad picture of the machine learning ecosystem and best practices of applying machine learning techniques. Through this book, you will learn to tackle data-driven problems and implement your solutions with the powerful yet simple language, Python. Interesting and easy-to-follow examples, to name some, news topic classification, spam email detection, online ad click-through prediction, stock prices forecast, will keep you glued till you reach your goal. Style and approach This book is an enticing journey that starts from the very basics and gradually picks up pace as the story unfolds. Each concept is first succinctly defined in the larger context of things, followed by a detailed explanation of their application. Every concept is explained with the help of a project that solves a real-world problem, and involves hands-on work—giving you a deep insight into the world of machine learning. With simple yet rich language—Python—you will understand and be able to

implement the examples with ease.

Applied Reinforcement Learning with Python Packt Publishing Ltd

This book primarily targets Python developers who want to learn and use Python's machine learning capabilities and gain valuable insights from data to develop effective solutions for business problems.

Python 3 for Machine Learning Nova Science Publishers

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

Machine Learning with Python for Everyone Packt Publishing Ltd

Gain the confidence you need to apply machine learning in your daily work. With this practical guide, author Matthew Kirk shows you how to integrate and test machine learning algorithms in your code, without the academic subtext. Featuring graphs and highlighted code examples throughout, the book features tests with Python's Numpy, Pandas, Scikit-Learn, and SciPy data science libraries. If you're a software engineer or business analyst interested in data science, this book will help you: Reference real-world examples to test each algorithm through engaging, hands-on exercises Apply test-driven development (TDD) to write and run tests before you start coding Explore techniques for improving your machine-learning models with data extraction and feature development Watch out for the risks of machine learning, such as underfitting or overfitting data Work with K-Nearest Neighbors, neural networks, clustering, and other algorithms

Introduction to Machine Learning with Python "O'Reilly Media, Inc."

Explore Machine Learning Techniques, Different Predictive Models, and its Applications KEY FEATURES ● Extensive coverage of real examples on implementation and working of ML models. ● Includes different strategies used in Machine Learning by leading data scientists. ● Focuses on Machine Learning concepts and their evolution to algorithms. DESCRIPTION This book covers basic concepts of Machine Learning, various learning paradigms, different architectures and algorithms used in these paradigms.

You will learn the power of ML models by exploring different predictive modeling techniques such as Regression, Clustering, and Classification. You will also get hands-on experience on methods and techniques such as Overfitting, Underfitting, Random Forest, Decision Trees, PCA, and Support Vector Machines. In this book real life examples with fully working of Python implementations are discussed in detail. At the end of the book you will learn about the unsupervised learning covering Hierarchical Clustering, K-means Clustering, Dimensionality Reduction, Anomaly detection, Principal Component Analysis.

WHAT YOU WILL LEARN ● Learn to perform data engineering and analysis. ● Build prototype ML models and production ML models from scratch. ● Develop strong proficiency in using scikit-learn and Python. ● Get hands-on experience with Random Forest, Logistic Regression, SVM, PCA, and Neural Networks. WHO THIS BOOK IS FOR This book is meant for beginners who want to gain knowledge about Machine Learning in detail. This book can also be used by Machine Learning users for a quick reference for fundamentals in Machine Learning. Readers should have basic knowledge of Python and Scikit-Learn before reading the book.

TABLE OF CONTENTS 1. Introduction to Machine Learning 2. Linear Regression 3. Classification Using Logistic Regression 4. Overfitting and Regularization 5. Feasibility of Learning 6. Support Vector Machine 7. Neural Network 8. Decision Trees 9. Unsupervised Learning 10. Theory of Generalization 11. Bias and Fairness in ML

Deep Learning with Python John Wiley & Sons

Unlock deeper insights into Machine Learning with this vital guide to cutting-edge predictive analytics About This Book Leverage

Python's most powerful open-source libraries for deep learning, data wrangling, and data visualization Learn effective strategies and best practices to improve and optimize machine learning systems and algorithms Ask – and answer – tough questions of your data with robust statistical models, built for a range of datasets Who This Book Is For If you want to find out how to use Python to start answering critical questions of your data, pick up Python Machine Learning – whether you want to get started from scratch or want to extend your data science knowledge, this is an essential and unmissable resource. What You Will Learn Explore how to use different machine learning models to ask different questions of your data Learn how to build neural networks using Keras and Theano Find out how to write clean and elegant Python code that will optimize the strength of your algorithms Discover how to embed your machine learning model in a web application for increased accessibility Predict continuous target outcomes using regression analysis Uncover hidden patterns and structures in data with clustering Organize data using effective pre-processing techniques Get to grips with sentiment analysis to delve deeper into textual and social media data In Detail Machine learning and predictive analytics are transforming the way businesses and other organizations operate. Being able to understand trends and patterns in complex data is critical to success, becoming one of the key strategies for unlocking growth in a challenging contemporary marketplace. Python can help you deliver key insights into your data – its unique capabilities as a language let you build sophisticated algorithms and statistical models that can reveal new perspectives and answer key questions that are vital for success. Python Machine Learning

gives you access to the world of predictive analytics and demonstrates why Python is one of the world's leading data science languages. If you want to ask better questions of data, or need to improve and extend the capabilities of your machine learning systems, this practical data science book is invaluable. Covering a wide range of powerful Python libraries, including scikit-learn, Theano, and Keras, and featuring guidance and tips on everything from sentiment analysis to neural networks, you'll soon be able to answer some of the most important questions facing you and your organization. Style and approach Python Machine Learning connects the fundamental theoretical principles behind machine learning to their practical application in a way that focuses you on asking and answering the right questions. It walks you through the key elements of Python and its powerful machine learning libraries, while demonstrating how to get to grips with a range of statistical models.

Building Machine Learning Systems with Python Addison-Wesley Professional

Create, execute, modify, and share machine learning applications with Python and TensorFlow 2.0 in the Jupyter Notebook environment. This book breaks down any barriers to programming machine learning applications through the use of Jupyter Notebook instead of a text editor or a regular IDE. You'll start by learning how to use Jupyter Notebooks to improve the way you program with Python. After getting a good grounding in working with Python in Jupyter Notebooks, you'll dive into what TensorFlow is, how it helps machine learning enthusiasts, and how to tackle the challenges it presents. Along the way, sample programs created using Jupyter Notebooks allow you to apply

concepts from earlier in the book. Those who are new to machine learning can dive in with these easy programs and develop basic skills. A glossary at the end of the book provides common machine learning and Python keywords and definitions to make learning even easier. What You Will Learn Program in Python and TensorFlow Tackle basic machine learning obstacles Develop in the Jupyter Notebooks environment Who This Book Is For Ideal for Machine Learning and Deep Learning enthusiasts who are interested in programming with Python using Tensorflow 2.0 in the Jupyter Notebook Application. Some basic knowledge of Machine Learning concepts and Python Programming (using Python version 3) is helpful.

Machine Learning Engineering with Python "O'Reilly Media, Inc."

For many researchers, Python is a first-class tool mainly because of its libraries for storing, manipulating, and gaining insight from data. Several resources exist for individual pieces of this data science stack, but only with the Python Data Science Handbook do you get them all—IPython, NumPy, Pandas, Matplotlib, Scikit-Learn, and other related tools. Working scientists and data crunchers familiar with reading and writing Python code will find this comprehensive desk reference ideal for tackling day-to-day issues: manipulating, transforming, and cleaning data; visualizing different types of data; and using data to build statistical or machine learning models. Quite simply, this is the must-have reference for scientific computing in Python. With this handbook, you'll learn how to use: IPython and Jupyter: provide computational environments for data scientists using Python NumPy: includes the ndarray for efficient storage and manipulation of dense data arrays in Python Pandas: features the

DataFrame for efficient storage and manipulation of labeled/columnar data in Python Matplotlib: includes capabilities for a flexible range of data visualizations in Python Scikit-Learn: for efficient and clean Python implementations of the most important and established machine learning algorithms Machine Learning Concepts with Python and the Jupyter Notebook Environment Apress

Build strong foundation for entering the world of Machine Learning and data science with the help of this comprehensive guide About This Book Get started in the field of Machine Learning with the help of this solid, concept-rich, yet highly practical guide. Your one-stop solution for everything that matters in mastering the whats and whys of Machine Learning algorithms and their implementation. Get a solid foundation for your entry into Machine Learning by strengthening your roots (algorithms) with this comprehensive guide. Who This Book Is For This book is for IT professionals who want to enter the field of data science and are very new to Machine Learning. Familiarity with languages such as R and Python will be invaluable here. What You Will Learn Acquaint yourself with important elements of Machine Learning Understand the feature selection and feature engineering process Assess performance and error trade-offs for Linear Regression Build a data model and understand how it works by using different types of algorithm Learn to tune the parameters of Support Vector machines Implement clusters to a dataset Explore the concept of Natural Processing Language and Recommendation Systems Create a ML architecture from scratch. In Detail As the amount of data continues to grow at an almost incomprehensible rate, being able to understand and process

data is becoming a key differentiator for competitive organizations. Machine learning applications are everywhere, from self-driving cars, spam detection, document search, and trading strategies, to speech recognition. This makes machine learning well-suited to the present-day era of Big Data and Data Science. The main challenge is how to transform data into actionable knowledge. In this book you will learn all the important Machine Learning algorithms that are commonly used in the field of data science. These algorithms can be used for supervised as well as unsupervised learning, reinforcement learning, and semi-supervised learning. A few famous algorithms that are covered in this book are Linear regression, Logistic Regression, SVM, Naive Bayes, K-Means, Random Forest, TensorFlow, and Feature engineering. In this book you will also learn how these algorithms work and their practical implementation to resolve your problems. This book will also introduce you to the Natural Processing Language and Recommendation systems, which help you run multiple algorithms simultaneously. On completion of the book you will have mastered selecting Machine Learning algorithms for clustering, classification, or regression based on for your problem. Style and approach An easy-to-follow, step-by-step guide that will help you get to grips with real -world applications of Algorithms for Machine Learning.

Machine Learning with Python "O'Reilly Media, Inc."

Applied machine learning with a solid foundation in theory. Revised and expanded for TensorFlow 2, GANs, and reinforcement learning. Purchase of the print or Kindle book includes a free eBook in the PDF format. Key Features Third edition of the bestselling, widely acclaimed Python machine

learning book Clear and intuitive explanations take you deep into the theory and practice of Python machine learning Fully updated and expanded to cover TensorFlow 2, Generative Adversarial Network models, reinforcement learning, and best practices Book Description Python Machine Learning, Third Edition is a comprehensive guide to machine learning and deep learning with Python. It acts as both a step-by-step tutorial, and a reference you'll keep coming back to as you build your machine learning systems. Packed with clear explanations, visualizations, and working examples, the book covers all the essential machine learning techniques in depth. While some books teach you only to follow instructions, with this machine learning book, Raschka and Mirjalili teach the principles behind machine learning, allowing you to build models and applications for yourself. Updated for TensorFlow 2.0, this new third edition introduces readers to its new Keras API features, as well as the latest additions to scikit-learn. It's also expanded to cover cutting-edge reinforcement learning techniques based on deep learning, as well as an introduction to GANs. Finally, this book also explores a subfield of natural language processing (NLP) called sentiment analysis, helping you learn how to use machine learning algorithms to classify documents. This book is your companion to machine learning with Python, whether you're a Python developer new to machine learning or want to deepen your knowledge of the latest developments. What you will learn Master the frameworks, models, and techniques that enable machines to 'learn' from data Use scikit-learn for machine learning and TensorFlow for deep learning Apply machine learning to image classification, sentiment analysis, intelligent web applications, and more Build

and train neural networks, GANs, and other models Discover best practices for evaluating and tuning models Predict continuous target outcomes using regression analysis Dig deeper into textual and social media data using sentiment analysis Who this book is for If you know some Python and you want to use machine learning and deep learning, pick up this book. Whether you want to start from scratch or extend your machine learning knowledge, this is an essential resource. Written for developers and data scientists who want to create practical machine learning and deep learning code, this book is ideal for anyone who wants to teach computers how to learn from data.

Mastering Machine Learning with Python in Six Steps

Packt Publishing Ltd

This practical guide provides nearly 200 self-contained recipes to help you solve machine learning challenges you may encounter in your daily work. If you're comfortable with Python and its libraries, including pandas and scikit-learn, you'll be able to address specific problems such as loading data, handling text or numerical data, model selection, and dimensionality reduction and many other topics. Each recipe includes code that you can copy and paste into a toy dataset to ensure that it actually works. From there, you can insert, combine, or adapt the code to help construct your application. Recipes also include a discussion that explains the solution and provides meaningful context. This cookbook takes you beyond theory and concepts by providing the nuts and bolts you need to construct working machine learning applications. You'll find recipes for: Vectors, matrices, and arrays Handling numerical and categorical data, text, images, and dates and times Dimensionality reduction using feature extraction or

feature selection Model evaluation and selection Linear and logical regression, trees and forests, and k-nearest neighbors Support vector machines (SVM), naïve Bayes, clustering, and neural networks Saving and loading trained models

Building Machine Learning Systems with Python - Second Edition
John Wiley & Sons

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 *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow* Manning Publications Understand the key aspects and challenges of machine learning interpretability, learn how to overcome them with interpretation

methods, and leverage them to build fairer, safer, and more reliable models Key Features Learn how to extract easy-to-understand insights from any machine learning model Become well-versed with interpretability techniques to build fairer, safer, and more reliable models Mitigate risks in AI systems before they have broader implications by learning how to debug black-box models Book Description Do you want to understand your models and mitigate risks associated with poor predictions using machine learning (ML) interpretation? Interpretable Machine Learning with Python can help you work effectively with ML models. The first section of the book is a beginner's guide to interpretability, covering its relevance in business and exploring its key aspects and challenges. You'll focus on how white-box models work, compare them to black-box and glass-box models, and examine their trade-off. The second section will get you up to speed with a vast array of interpretation methods, also known as Explainable AI (XAI) methods, and how to apply them to different use cases, be it for classification or regression, for tabular, time-series, image or text. In addition to the step-by-step code, the book also helps the reader to interpret model outcomes using examples. In the third section, you'll get hands-on with tuning models and training data for interpretability by reducing complexity, mitigating bias, placing guardrails, and enhancing reliability. The methods you'll explore here range from state-of-the-art feature selection and dataset debiasing methods to monotonic constraints and adversarial retraining. By the end of this book, you'll be able to understand ML models better and enhance them through interpretability tuning. What you will learn Recognize the importance of interpretability in business Study models that are

intrinsically interpretable such as linear models, decision trees, and Naïve Bayes Become well-versed in interpreting models with model-agnostic methods Visualize how an image classifier works and what it learns Understand how to mitigate the influence of bias in datasets Discover how to make models more reliable with adversarial robustness Use monotonic constraints to make fairer and safer models Who this book is for This book is for data scientists, machine learning developers, and data stewards who have an increasingly critical responsibility to explain how the AI systems they develop work, their impact on decision making, and how they identify and manage bias. Working knowledge of machine learning and the Python programming language is expected.

Thoughtful Machine Learning with Python Packt Publishing Ltd Equipped with the latest updates, this third edition of Python Machine Learning By Example provides a comprehensive course for ML enthusiasts to strengthen their command of ML concepts, techniques, and algorithms.

Deep Learning with Python Apress

Build real-world Artificial Intelligence applications with Python to intelligently interact with the world around you About This Book Step into the amazing world of intelligent apps using this comprehensive guide Enter the world of Artificial Intelligence, explore it, and create your own applications Work through simple yet insightful examples that will get you up and running with Artificial Intelligence in no time Who This Book Is For This book is for Python developers who want to build real-world Artificial Intelligence applications. This book is friendly to Python beginners, but being familiar with Python would be useful to play

around with the code. It will also be useful for experienced Python programmers who are looking to use Artificial Intelligence techniques in their existing technology stacks. What You Will Learn Realize different classification and regression techniques Understand the concept of clustering and how to use it to automatically segment data See how to build an intelligent recommender system Understand logic programming and how to use it Build automatic speech recognition systems Understand the basics of heuristic search and genetic programming Develop games using Artificial Intelligence Learn how reinforcement learning works Discover how to build intelligent applications centered on images, text, and time series data See how to use deep learning algorithms and build applications based on it In Detail Artificial Intelligence is becoming increasingly relevant in the modern world where everything is driven by technology and data. It is used extensively across many fields such as search engines, image recognition, robotics, finance, and so on. We will explore various real-world scenarios in this book and you'll learn about various algorithms that can be used to build Artificial Intelligence applications. During the course of this book, you will find out how to make informed decisions about what algorithms to use in a given context. Starting from the basics of Artificial Intelligence, you will learn how to develop various building blocks using different data mining techniques. You will see how to implement different algorithms to get the best possible results, and will understand how to apply them to real-world scenarios. If you want to add an intelligence layer to any application that's based on images, text, stock market, or some other form of data, this exciting book on Artificial Intelligence will definitely be your

guide! Style and approach This highly practical book will show you how to implement Artificial Intelligence. The book provides multiple examples enabling you to create smart applications to meet the needs of your organization. In every chapter, we explain an algorithm, implement it, and then build a smart application.

Advanced Machine Learning with Python Packt Publishing Ltd

Deep learning is the most interesting and powerful machine learning technique right now. Top deep learning libraries are available on the Python ecosystem like Theano and TensorFlow. Tap into their power in a few lines of code using Keras, the best-of-breed applied deep learning library. In this Ebook, learn exactly how to get started and apply deep learning to your own machine learning projects.

Machine Learning with Pytorch and Scikit-Learn "O'Reilly Media, Inc."

This book is designed to provide the reader with basic Python 3 programming concepts related to machine learning. The first four chapters provide a fast-paced introduction to Python 3, NumPy, and Pandas. The fifth chapter introduces the fundamental concepts of machine learning. The sixth chapter is devoted to machine learning classifiers, such as logistic regression, k-NN, decision trees, random forests, and SVMs. The final chapter includes material on NLP and RL. Keras-based code samples are included to supplement the theoretical discussion. The book also contains separate appendices for regular expressions, Keras, and TensorFlow 2. Features: Provides the reader with basic Python 3 programming concepts related to machine learning Includes separate appendices for regular expressions, Keras, and

TensorFlow 2

Python Machine Learning BPB Publications

Discover the practical aspects of implementing deep-learning solutions using the rich Python ecosystem. This book bridges the gap between the academic state-of-the-art and the industry state-of-the-practice by introducing you to deep learning frameworks such as Keras, Theano, and Caffe. The practicalities of these frameworks is often acquired by practitioners by reading source code, manuals, and posting questions on community forums, which tends to be a slow and a painful process. Deep Learning with Python allows you to ramp up to such practical know-how in a short period of time and focus more on the domain, models, and algorithms. This book briefly covers the mathematical prerequisites and fundamentals of deep learning, making this book a good starting point for software developers who want to get started in deep learning. A brief survey of deep learning architectures is also included. Deep Learning with Python also introduces you to key concepts of automatic differentiation and GPU computation which, while not central to deep learning, are critical when it comes to conducting large scale experiments. What You Will Learn Leverage deep learning frameworks in Python namely, Keras, Theano, and Caffe Gain the fundamentals of deep learning with mathematical prerequisites Discover the practical considerations of large scale experiments Take deep learning models to production Who This Book Is For Software developers who want to try out deep learning as a practical solution to a particular problem. Software developers in a data science team who want to take deep learning models developed by data scientists to production.

Deep Learning for Coders with fastai and PyTorch Mercury Learning and Information

Practical Machine Learning for Data Analysis Using Python is a problem solver's guide for creating real-world intelligent systems. It provides a comprehensive approach with concepts, practices, hands-on examples, and sample code. The book teaches readers the vital skills required to understand and solve different problems with machine learning. It teaches machine learning techniques necessary to become a successful practitioner, through the presentation of real-world case studies in Python machine learning ecosystems. The book also focuses on building a foundation of machine learning knowledge to solve different real-world case studies across various fields, including biomedical signal analysis, healthcare, security, economics, and finance.

Moreover, it covers a wide range of machine learning models, including regression, classification, and forecasting. The goal of the book is to help a broad range of readers, including IT professionals, analysts, developers, data scientists, engineers, and graduate students, to solve their own real-world problems. Offers a comprehensive overview of the application of machine learning tools in data analysis across a wide range of subject areas Teaches readers how to apply machine learning techniques to biomedical signals, financial data, and healthcare data Explores important classification and regression algorithms as well as other machine learning techniques Explains how to use Python to handle data extraction, manipulation, and exploration techniques, as well as how to visualize data spread across multiple dimensions and extract useful features