Intelligenza Artificiale Con Python Interamente I

KEMP ATKINSON

Python Machine Learning By Example World Scientific
This book introduces the essentials of Python for the emerging fields of Machine Learning (ML) and Artificial Intelligence (AI). The authors explore the use of Python's advanced module features and apply them in probability, statistical testing, signal processing, financial forecasting, and various other applications. This includes mathematical operations with array data structures, Data Manipulation, Data Cleaning, machine learning, Data pipeline, probability density functions, interpolation, visualization, and other high-performance benefits using the core scientific packages NumPy, Pandas, SciPy, Sklearn/Scikit learn and Matplotlib. Readers will gain a deep understanding with problem-solving experience on these powerful platforms when dealing with engineering and scientific problems related to Machine Learning and Artificial Intelligence. Several examples of real problems using these techniques are provided along with examples. The authors also focus on the best practices in the industry on using Python for AI and ML. Deployment on a cloud infrastructure is described in detail (with code) to emphasize real scenarios.
Monetizing Machine Learning Youcanprint
Master the practical aspects of implementing deep learning solutions with PyTorch, using a hands-on approach to understanding both theory and practice. This updated edition will prepare you for applying deep learning to real world problems with a sound theoretical foundation and practical know-how with PyTorch, a platform developed by Facebook's Artificial Intelligence Research Group. You'll start with a perspective on how and why deep learning with PyTorch has emerged as an path-breaking framework with a set of tools and techniques to solve real-world problems. Next, the book will ground you with the mathematical fundamentals of linear algebra, vector calculus, probability and optimization. Having established this foundation, you'll move on to key components and functionality of PyTorch including layers, loss functions and optimization algorithms. You'll also gain an understanding of Graphical Processing Unit (GPU) based computation, which is essential for training deep learning models. All the key architectures in deep learning are covered, including feedforward networks, convolution neural networks, recurrent neural networks, long short-term memory networks, autoencoders and generative adversarial networks. Backed by a number of tricks of the trade for training and optimizing deep learning models, this edition of Deep Learning with Python explains the best practices in taking these models to production with PyTorch. What You'll Learn Review machine learning fundamentals such as overfitting, underfitting, and regularization. Understand deep learning fundamentals such as feed-forward networks, convolution neural networks, recurrent neural networks, automatic differentiation, and stochastic gradient descent. Apply in-depth linear algebra with PyTorch Explore PyTorch fundamentals and its building blocks Work with tuning and optimizing models Who This Book Is For Beginners with a working knowledge of Python who want to understand Deep Learning in a practical, hands-on manner.
Machine Learning For Dummies Springer Nature
New edition of the bestselling guide to artificial intelligence with Python, updated to Python 3.x, with seven new chapters that cover RNNs, AI and Big Data, fundamental use cases, chatbots, and

more. Key FeaturesCompletely updated and revised to Python 3.xNew chapters for AI on the cloud, recurrent neural networks, deep learning models, and feature selection and engineeringLearn more about deep learning algorithms, machine learning data pipelines, and chatbotsBook Description Artificial Intelligence with Python, Second Edition is an updated and expanded version of the bestselling guide to artificial intelligence using the latest version of Python 3.x. Not only does it provide you an introduction to artificial intelligence, this new edition goes further by giving you the tools you need to explore the amazing world of intelligent apps and create your own applications. This edition also includes seven new chapters on more advanced concepts of Artificial Intelligence, including fundamental use cases of AI; machine learning data pipelines; feature selection and feature engineering; AI on the cloud; the basics of chatbots; RNNs and DL models; and AI and Big Data. Finally, this new edition explores various real-world scenarios and teaches you how to apply relevant AI algorithms to a wide swath of problems, starting with the most basic AI concepts and progressively building from there to solve more difficult challenges so that by the end, you will have gained a solid understanding of, and when best to use, these many artificial intelligence techniques. What you will learnUnderstand what artificial intelligence, machine learning, and data science areExplore the most common artificial intelligence use casesLearn how to build a machine learning pipelineAssimilate the basics of feature selection and feature engineeringIdentify the differences between supervised and unsupervised learningDiscover the most recent advances and tools offered for AI development in the cloudDevelop automatic speech recognition systems and chatbotsApply AI algorithms to time series dataWho this book is for The intended audience for this book is Python developers who want to build real-world Artificial Intelligence applications. Basic Python programming experience and awareness of machine learning concepts and techniques is mandatory.

Deep Learning with Python Createspace Independent Publishing Platform

In this revolutionary book, a renowned computer scientist explains the importance of teaching children the basics of computing and how it can prepare them to succeed in the ever-evolving tech world. Computers have completely changed the way we teach children. We have Mindstorms to thank for that. In this book, pioneering computer scientist Seymour Papert uses the invention of LOGO, the first child-friendly programming language, to make the case for the value of teaching children with computers. Papert argues that children are more than capable of mastering computers, and that teaching computational processes like de-bugging in the classroom can change the way we learn everything else. He also shows that schools saturated with technology can actually improve socialization and interaction among students and between students and teachers. Technology changes every day, but the basic ways that computers can help us learn remain. For thousands of teachers and parents who have sought creative ways to help children learn with computers, Mindstorms is their bible.

Practical Machine Learning for Data Analysis Using Python BPB Publications

DescriptionThis book provides the concept of machine learning with mathematical explanation and programming examples. Every chapter starts with fundamentals of the technique and working example on the real-world dataset. Along with the advice on applying algorithms, each technique is provided with advantages and disadvantages on the data.In this book we provide code examples in python. Python is the most suitable and worldwide accepted language for this. First, it is free and open source. It contains very good support from open community. It contains a lot of library, so you don't need to code everything. Also, it is scalable for large amount of data and suitable for big data technologies.This book:Covers all major areas in Machine Learning.Topics are discussed with graphical explanations.Comparison of different Machine Learning methods to solve any problem.Methods to handle real-world noisy data before applying any Machine Learning algorithm.Python code example for each concept discussed.Jupyter notebook scripts are provided with dataset used to test and try the algorithms ContentsIntroduction to Machine Learning Understanding Python Feature Engineering Data VisualisationBasic and Advanced Regression techniquesClassification Un Supervised LearningText AnalysisNeural Network and Deep Learning Recommendation System Time Series Analysis

The Object-Oriented Approach to Problem Solving and Machine Learning with Python Apogeo Editore

The book teaches students to model a scientific problem and write a computer program in C language to solve that problem. It introduces the basics of C language, and then describes and discusses algorithms commonly used in scientific applications (e.g. searching, graphs, statistics, equation solving, Monte Carlo methods etc.).

Machine Learning with Python Roland Bind

★☆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. Machine learning is an incredibly dense topic. It's hard to imagine condensing it into an easily readable and digestible format. However, this book aims to do exactly that. Machine learning and artificial intelligence have been used in different machines and applications to improve the user's experience. One can also use machine learning to make data analysis and predicting the output for some data sets easy. All you need to do is choose the right algorithm, train the model and test the model before you apply it on any real-world tool. It is that simple isn't it? ★★Apart from this, you will also learn more about★★ ♦ The Different Types Of Learning Algorithm That You Can Expect To Encounter ♦ The Numerous Applications Of Machine Learning And Deep Learning ♦ The Best Practices For Picking Up Neural Networks ♦ What Are The Best Languages And Libraries To Work With ♦ The Various Problems That You Can Solve With Machine Learning Algorithms ♦ And much more... Well, you can do it faster if you use Python. This language has made it easy for any user, even an amateur, to build a strong machine learning model since it has numerous directories and libraries that make it easy for one to build a model. Do you want to know how to build a machine learning model and a neural network? So, what are you waiting for? Grab a copy of this book now!

Machine Learning with Python for PC, Raspberry Pi, and Maixduino "O'Reilly Media, Inc."

This is a translation of one of Guenon's most significant works. It contains chapters with titles such as: The Science of Letters; Symbolic Weapons; and The Symbolism of the Zodiac among the Pythagoreans. The work aims to give new meaning to so many of the objects involved in daily life.

Machine Learning and Deep Learning Using Python and TensorFlow Packt Publishing Ltd

This beginner-oriented book will help you understand and perform anomaly detection by learning cutting-edge machine learning and deep learning techniques. This updated second edition focuses on supervised, semi-supervised, and unsupervised approaches to anomaly detection. Over the course of the book, you will learn how to use Keras and PyTorch in practical applications. It also introduces new chapters on GANs and transformers to reflect the latest trends in deep learning. Beginning Anomaly Detection Using Python-Based Deep Learning begins with an introduction to anomaly detection, its importance, and its applications. It then covers core data science and machine learning modeling concepts before delving into traditional machine learning algorithms such as OC-SVM and Isolation Forest for anomaly detection using scikit-learn. Following this, the authors explain the essentials of machine learning and deep learning, and how to implement multilayer perceptrons for supervised anomaly detection in both Keras and PyTorch. From here, the focus shifts to the applications of deep learning models for anomaly detection, including various types of autoencoders, recurrent neural networks (via LSTM), temporal convolutional networks, and transformers, with the latter three architectures applied to time-series anomaly detection. This edition has a new chapter on GANs (Generative Adversarial Networks), as well as new material covering transformer architecture in the context of time-series anomaly detection. After completing this book, you will have a thorough understanding of anomaly detection as well as an assortment of methods to approach it in various contexts, including time-series data. Additionally, you will have gained an introduction to scikit-learn, GANs, transformers, Keras, and PyTorch, empowering you to create your own machine learning- or deep learning-based anomaly detectors. What You Will Learn Understand what anomaly detection is, why it it is important, and how it is applied Grasp the core concepts of machine learning. Master traditional machine learning approaches to anomaly detection using scikit-kearn. Understand deep learning in Python using Keras and PyTorch Process data through pandas and evaluate your model's performance using metrics like F1-score, precision, and recall Apply deep learning to supervised, semi-supervised, and unsupervised anomaly detection tasks for tabular datasets and time series applications Who This Book Is For Data scientists and machine learning engineers of all levels of experience interested in learning the basics of deep learning applications in anomaly detection.

Deep Learning con Python Apress

Implement machine learning and deep learning methodologies to build smart, cognitive AI projects using Python Key FeaturesA go-to guide to help you master AI algorithms and concepts8 real-world projects tackling different challenges in healthcare, e-commerce, and surveillanceUse TensorFlow, Keras, and other Python libraries to implement smart AI applicationsBook Description This book will be a perfect companion if you want to build insightful projects from leading AI domains using Python. The book covers detailed implementation of projects from all the core disciplines of AI. We start by covering the basics of how to create smart systems using machine learning and deep learning techniques. You will assimilate various neural network architectures such as CNN, RNN, LSTM, to solve critical new world challenges. You will learn to train a model to detect diabetic retinopathy conditions in the human eye and create an intelligent system for performing a video-to-text translation. You will use the transfer learning technique in the healthcare domain and implement style transfer using GANs. Later you will learn to build AI-based recommendation systems, a mobile app for sentiment analysis and a powerful chatbot for carrying customer services. You will implement AI techniques in the cybersecurity domain to generate Captchas. Later you will train and build autonomous vehicles to self-drive using reinforcement learning. You will be using libraries from the Python ecosystem such as TensorFlow, Keras and more to bring the core aspects of machine learning, deep learning, and AI. By the end of this book, you will be skilled to build your own smart models for tackling any kind of AI problems without any hassle. What you will learnBuild an intelligent machine translation system using seq-2-seq neural translation machinesCreate AI applications using GAN and deploy smart mobile apps using TensorFlowTranslate videos into text using CNN and RNNImplement smart AI Chatbots, and integrate and extend them in several domainsCreate smart reinforcement, learning-based applications using Q-LearningBreak and generate CAPTCHA using Deep Learning and Adversarial Learning Who this book is for This book is intended for data scientists, machine learning professionals, and deep learning practitioners who are ready to extend their knowledge and potential in AI. If you want to build real-life smart systems to play a crucial role in every complex domain, then this book is what you need. Knowledge of Python programming and a familiarity with basic machine learning and deep learning concepts are expected to help you get the most out of the book

Building Machine Learning Systems Using Python Apress

Artificial intelligence takes many shapes and forms. At this point in its evolution, machine learning and deep learning are two of the most common shapes it takes. This is primarily because we are at a point where we have discovered how to create networks of information that can actually be filtered and processed just as a normal human cognitive process would be. Beyond all those shapes and forms of AI, though, this entire concept is built based on a few basic ideas: Information is power Neural networks can imitate the human brain Programmers can create machine programs that enable them to filter information in a specific way, that allows them to draw conclusions and grow their learning based on that Hopefully, the book at hand will help you gain a better understanding of the grand scheme of deep learning - and even more specifically, how deep learning connects to Python, one of the single most popular programming languages of the moment. I have structured the book in a way that allows you to make sense of everything in the most logical way possible. What are you waiting for?.. Grab a copy of this book and start to explore the world of Deep Learning with Python!

Python for Civil and Structural Engineers Academic Press

Understand the principles and practices of machine learning and deep learning This hands-on guide lays out machine learning and deep learning techniques and technologies in a style that is approachable, using just the basic math required. Written by a pair of experts in the field, Machine Learning and Deep Learning Using Python and TensorFlow contains case studies in several industries, including banking, insurance, e-commerce, retail, and healthcare. The book shows how to utilize machine learning and deep learning functions in today’s smart devices and apps. You will get download links for datasets, code, and sample projects referred to in the text. Coverage includes: Machine learning and deep learning concepts Python programming and statistics fundamentals Regression and logistic regression Decision trees Model selection and cross-validation Cluster analysis Random forests and boosting Artificial neural networks TensorFlow and Keras Deep learning hyperparameters Convolutional neural networks Recurrent neural networks and long short-term memory

Python Machine Learning Packt Publishing Ltd

This book explains how to use the programming language Python to develop machine learning and deep learning tasks.

Deep Learning with Python John Wiley & Sons

This book provides a full presentation of the current concepts and available techniques to make “machine learning” systems more explainable. The approaches presented can be applied to almost all the current “machine learning” models: linear and logistic regression, deep learning neural

networks, natural language processing and image recognition, among the others. Progress in Machine Learning is increasing the use of artificial agents to perform critical tasks previously handled by humans (healthcare, legal and finance, among others). While the principles that guide the design of these agents are understood, most of the current deep-learning models are "opaque" to human understanding. Explainable AI with Python fills the current gap in literature on this emerging topic by taking both a theoretical and a practical perspective, making the reader quickly capable of working with tools and code for Explainable AI. Beginning with examples of what Explainable AI (XAI) is and why it is needed in the field, the book details different approaches to XAI depending on specific context and need. Hands-on work on interpretable models with specific examples leveraging Python are then presented, showing how intrinsic interpretable models can be interpreted and how to produce “human understandable” explanations. Model-agnostic methods for XAI are shown to produce explanations without relying on ML models internals that are “opaque.” Using examples from Computer Vision, the authors then look at explainable models for Deep Learning and prospective methods for the future. Taking a practical perspective, the authors demonstrate how to effectively use ML and XAI in science. The final chapter explains Adversarial Machine Learning and how to do XAI with adversarial examples.

Beginning Anomaly Detection Using Python-Based Deep Learning Packt Publishing Ltd

Demystify the complexity of machine learning techniques and create evolving, clever solutions to solve your problems Key FeaturesMaster supervised, unsupervised, and semi-supervised ML algorithms and their implementation Build deep learning models for object detection, image classification, similarity learning, and moreBuild, deploy, and scale end-to-end deep neural network models in a production environmentBook Description This Learning Path is your complete guide to quickly getting to grips with popular machine learning algorithms. You'll be introduced to the most widely used algorithms in supervised, unsupervised, and semi-supervised machine learning, and learn how to use them in the best possible manner. Ranging from Bayesian models to the MCMC algorithm to Hidden Markov models, this Learning Path will teach you how to extract features from your dataset and perform dimensionality reduction by making use of Python-based libraries. You'll bring the use of TensorFlow and Keras to build deep learning models, using concepts such as transfer learning, generative adversarial networks, and deep reinforcement learning. Next, you'll learn the advanced features of TensorFlow1.x, such as distributed TensorFlow with TF clusters, deploy production models with TensorFlow Serving. You'll implement different techniques related to object classification, object detection, image segmentation, and more. By the end of this Learning Path, you'll have obtained in-depth knowledge of TensorFlow, making you the go-to person for solving artificial intelligence problems This Learning Path includes content from the following Packt products: Mastering Machine Learning Algorithms by Giuseppe BonaccorsoMastering TensorFlow 1.x by Armando FandangoDeep Learning for Computer Vision by Rajalingappaa ShanmugamaniWhat you will learnExplore how an ML model can be trained, optimized, and evaluatedWork with Autoencoders and Generative Adversarial NetworksExplore the most important Reinforcement Learning techniquesBuild end-to-end deep learning (CNN, RNN, and Autoencoders) modelsWho this book is for This Learning Path is for data scientists, machine learning engineers, artificial intelligence engineers who want to delve into complex machine learning algorithms, calibrate models, and improve the predictions of the trained model. You will encounter the advanced intricacies and complex use cases of deep learning and AI. A basic knowledge of programming in Python and some understanding of machine learning concepts are required to get the best out of this Learning Path.

Machine Learning with Python BPB Publications

Want to build apps for Android devices? This book is the perfect way to master the fundamentals. Written by an expert who's taught this mobile platform to hundreds of developers in large organizations, this gentle introduction shows experienced object-oriented programmers how to use Android’s basic building blocks to create user interfaces, store data, connect to the network, and more. You'll build a Twitter-like application throughout the course of this book, adding new features with each chapter. Along the way, you'll also create your own toolbox of code patterns to help you program any type of Android application with ease. Get an overview of the Android platform and discover how it fits into the mobile ecosystem Learn about the Android stack, including its application framework, and the structure and distribution of application packages (APK) Set up your Android development environment and get started with simple programs Use Android’s building blocks—Activities, Intents, Services, Content Providers, and Broadcast Receivers Learn how to build basic Android user interfaces and organize UI elements in Views and Layouts Build a service that uses a background process to update data in your application Get an introduction to Android Interface Definition Language (AIDL) and the Native Development Kit (NDK)

Practical Machine Learning with Python McGraw Hill Professional

One of Mark Cuban’s top reads for better understanding A.I. (inc.com, 2021) Your comprehensive entry-level guide to machine learning While machine learning expertise doesn’t quite mean you can create your own Turing Test-proof android—as in the movie Ex Machina—it is a form of artificial intelligence and one of the most exciting technological means of identifying opportunities and solving problems fast and on a large scale. Anyone who masters the principles of machine learning is mastering a big part of our tech future and opening up incredible new directions in careers that include fraud detection, optimizing search results, serving real-time ads, credit-scoring, building accurate and sophisticated pricing models—and way, way more. Unlike most machine learning books, the fully updated 2nd Edition of Machine Learning For Dummies doesn't assume you have years of experience using programming languages such as Python (R source is also included in a downloadable form with comments and explanations), but lets you in on the ground floor, covering the entry-level materials that will get you up and running building models you need to perform practical tasks. It takes a look at the underlying—and fascinating—math principles that power machine learning but also shows that you don't need to be a math whiz to build fun new tools and apply them to your work and study. Understand the history of AI and machine learning Work with Python 3.8 and TensorFlow 2.x (and R as a download) Build and test your own models Use the latest datasets, rather than the worn out data found in other books Apply machine learning to real problems Whether you want to learn for college or to enhance your business or career performance, this friendly beginner's guide is your best introduction to machine learning, allowing you to become quickly confident using this amazing and fast-developing technology that's impacting lives for the better all over the world.

Essentials of Python for Artificial Intelligence and Machine Learning Springer

Take your machine learning skills to the next level by mastering Deep Learning concepts and algorithms using Python.About This Book* Explore and create intelligent systems using cutting-edge deep learning techniques* Implement deep learning algorithms and work with revolutionary libraries in Python* Get real-world examples and easy-to-follow tutorials on Theano, TensorFlow, H2O and moreWho This Book Is ForThis book is for Data Science practitioners as well as aspirants who have a basic foundational understanding of Machine Learning concepts and some programming experience with Python. A mathematical background with a conceptual understanding of calculus and statistics is also desired.What You Will Learn* Get a practical deep dive into deep learning algorithms* Explore deep learning further with Theano, Caffe, Keras, and TensorFlow* Learn about two of the most powerful techniques at the core of many practical deep learning implementations: Auto-Encoders and Restricted Boltzmann Machines* Dive into Deep Belief Nets and Deep Neural Networks* Discover more deep learning algorithms with Dropout and Convolutional Neural Networks* Get to know device strategies so you can use deep learning algorithms and libraries in the real worldIn DetailWith an increasing interest in AI around the world, deep learning has attracted a great deal of public attention. Every day, deep learning algorithms are used broadly across different industries.The book will give you all the practical information available on the subject, including the best practices, using real-world use cases. You will learn to recognize and extract information to increase predictive accuracy and optimize results.Starting with a quick recap of important machine learning concepts, the book will delve straight into deep learning principles using Sci-kit learn. Moving ahead, you will learn to use the latest open source libraries such as Theano, Keras, Google's TensorFlow, and H20. Use this guide to uncover the difficulties of pattern recognition, scaling data with greater accuracy and discussing deep learning algorithms and techniques.Whether you want to dive deeper into Deep Learning, or want to investigate how to get more out of this powerful technology, you'll find everything inside.Style and approachPython Machine Learning by example follows practical hands on approach. It walks you through the key elements of Python and its powerful machine learning libraries with the help of real world projects.

Explainable AI with Python Apress

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

Python Deep Learning Data Science

In un mondo in cui compiti complessi e ripetitivi possono essere svolti da automi con estrema precisione ed efficienza, la programmazione di robot è un tema più che mai attuale. Questo libro mostra come l'utilizzo combinato di Raspberry Pi e Python possa essere un ottimo punto di partenza per avventurarsi in questo mondo. Si comincia introducendo le basi della robotica e da qui si passa velocemente alla progettazione e realizzazione di un primo robot controllato da remoto. Quindi si procede aggiungendo funzionalità e controlli, sensori e sistemi per rilevare dati, motori, servomotori e fotocamere, per passare infine alla scrittura del codice che permette al robot di svolgere alcuni compiti e agire in autonomia grazie a funzioni di intelligenza artificiale di base. Una guida passo-passo corredata da immagini ed esempi, adatta non solo a chi desidera applicare le proprie competenze software a un progetto hardware, ma anche agli appassionati con conoscenze base di programmazione che vogliono imparare a progettare, costruire e programmare robot.