



Learn deep learning for computer vision with Python Master linear regression and regularization techniques Classify objects such as flower species, handwritten digits, and pedestrians Explore the effective use of support vector machines, boosted decision trees, and random forests Get acquainted with neural networks and Deep Learning to address real-world problems Discover hidden structures in your data using k-means clustering Get to grips with data pre-processing and feature engineering In Detail Machine learning is no longer just a buzzword, it is all around us: from protecting your email, to automatically tagging friends in pictures, to predicting what movies you like. Computer vision is one of today's most exciting application fields of machine learning, with Deep Learning driving innovative systems such as self-driving cars and Google's DeepMind. OpenCV lies at the intersection of these topics, providing a comprehensive open-source library for classic as well as state-of-the-art computer vision and machine learning algorithms. In combination with Python Anaconda, you will have access to all the open-source computing libraries you could possibly ask for. Machine learning for OpenCV begins by introducing you to the essential concepts of statistical learning, such as classification and regression. Once all the basics are covered, you will start exploring various algorithms such as decision trees, support vector machines, and Bayesian networks, and learn how to combine them with other OpenCV functionality. As the book progresses, so will your machine learning skills, until you are ready to take on today's hottest topic in the field: Deep Learning. By the end of this book, you will be ready to take on your own machine learning problems, either by building on the existing source code or developing your own algorithm from scratch! Style and approach OpenCV machine learning connects the fundamental theoretical principles behind machine learning to their practical applications in a way that focuses on asking and answering the right questions. This book walks you through the key elements of OpenCV and its powerful machine learning classes, while demonstrating how to get to grips with a range of models.

If you want to learn how decision trees and random forests work, plus create your own, this visual book is for you. The fact is, decision tree and random forest algorithms are powerful and likely touch your life everyday. From online search to product development and credit scoring, both types of algorithms are at work behind the scenes in many modern applications and services. They are also used in countless industries such as medicine, manufacturing and finance to help companies make better decisions and reduce risk. Whether coded or scratched out by hand, both algorithms are powerful tools that can make a significant impact. This book is a visual introduction for beginners that unpacks the fundamentals of decision trees and random forests. If you want to dig into the basics with a visual twist plus create your own algorithms in Python, this book is for you. The four-volume set comprising LNCS volumes 5302/5303/5304/5305 constitutes the refereed proceedings of the 10th European Conference on Computer Vision, ECCV 2008, held in Marseille, France, in October 2008. The 243 revised papers presented were carefully reviewed and selected from a total of 871 papers submitted. The four books cover the entire range of current issues in computer vision. The papers are organized in topical sections on recognition, stereo, people and face recognition, object tracking, matching, learning and features, MRFs, segmentation, computational photography and active reconstruction.

Compressive Sampling and Learning Algorithms for Rotating Machines

The Elements of Statistical Learning

16th International Conference, ICFCA 2021, Strasbourg, France, June 29 – July 2, 2021, Proceedings

Learn Machine Learning With Data

Computer Vision - ECCV 2008

Made Easy To Understand Data Tree: How To Learn Machine Learning

Tree-based Machine Learning Algorithms

Learn to build expert NLP and machine learning projects using NLTK and other Python libraries About This Book Break text down into its component parts for spelling correction, feature extraction, and phrase transformation Work through NLP concepts with simple and easy-to-follow programming recipes Gain insights into the current and budding research topics of NLP Who This Book Is For If you are an NLP or machine learning enthusiast and an intermediate Python programmer who wants to quickly master NLTK for natural language processing, then this Learning Path will do you a lot of good. Students of linguistics and semantic/sentiment analysis professionals will find it invaluable. What You Will Learn The scope of natural language complexity and how they are processed by machines Clean and wrangle text using tokenization and chunking to help you process data better Tokenize text into sentences and sentences into words Classify text and perform sentiment analysis Implement string matching algorithms and normalization techniques Understand and implement the concepts of information retrieval and text summarization Find out how to implement various NLP tasks in Python In Detail Natural Language Processing is a field of computational linguistics and artificial intelligence that deals with human-computer interaction. It provides a seamless interaction between computers and human beings and gives computers the ability to understand human speech with the help of machine learning. The number of human-computer interaction instances are increasing so it's becoming imperative that computers comprehend all major natural languages. The first NLTK Essentials module is an introduction on how to build systems around NLP, with a focus on how to create a customized tokenizer and parser from scratch. You will learn essential concepts of NLP, be given practical insight into open source tool and libraries available in Python, shown how to analyze social media sites, and be given tools to deal with large scale text. This module also provides a workaround using some of the amazing capabilities of Python Libraries such as NLTK, scikit-learn, pandas, and NumPy. The second Python 3 Text Processing with NLTK 3 Cookbook module teaches you the essential techniques of text and language processing with simple, straightforward examples. This includes organizing text corpora, creating your own custom corpus, text classification with a focus on sentiment analysis, and distributed text processing methods. The third Mastering Natural Language Processing with Python module will help you become an expert and assist you in creating your own NLP projects using NLTK. You will be guided through model development with machine learning tools, shown how to create training data, and given insight into the best practices for designing and building NLP-based applications using Python. This Learning Path combines some of the best that Packt has to offer in one complete, curated package and is designed to help you quickly learn text processing with Python and NLTK. It includes content from the following Packt products: NLTK essentials by Nitin Hardeniya Python 3 Text Processing with NLTK 3 Cookbook by Jacob Perkins Mastering Natural Language Processing with Python by Deepti Chopra, Nisheeth Joshi, and Iti Mathur Style and approach This comprehensive course creates a smooth learning path that teaches you how to get started with Natural Language Processing using Python and NLTK. You'll learn to create effective NLP and machine learning projects using Python and NLTK. 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

Machine learning is the study of computer algorithms that improve automatically through experience and by the use of data. It is seen as a part of artificial intelligence. Machine learning algorithms build a model based on sample data, known as "training data", in order to make predictions or decisions without being explicitly programmed to do so. Machine learning algorithms are used in a wide variety of applications, such as in medicine, email filtering, speech recognition, and computer vision, where it is difficult or unfeasible to develop conventional algorithms to perform the needed tasks. If you are someone who learns by playing with the code and editing the data or equations to see what changes, then use those resources along with the book for a deeper understanding. The topics covered in this book are: -An overview of decision trees and random forests -A manual example of how a human would classify a dataset, compared to how a decision tree would work -How a decision tree works, and why it is prone to overfitting -How decision trees get combined to form a random forest -How to use that random forest to classify data and make predictions -How to determine how many trees to use in a random forest -Just where does the "randomness" come from -Out of Bag Errors & Cross-Validation - how good of a fit did the machine learning algorithm make? -Gini Criteria & Entropy Criteria - how to tell which split on a decision tree is best among many possible choices -And More

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

Computational Genomics with R

Machine Learning with Python Cookbook

A Unified Framework for Classification, Regression, Density Estimation, Manifold Learning and Semi-Supervised Learning

Data Mining with Decision Trees

Formal Concept Analysis

Decision Trees and Random Forests

The Essentials of Machine Learning in Finance and Accounting

Provides an extensive, up-to-date treatment of techniques used for machine condition monitoring Clear and concise throughout, this accessible book is the first to be wholly devoted to the field of condition monitoring for rotating machines using vibration signals. It covers various feature extraction, feature selection, and classification methods as well as their applications to machine vibration datasets. It also presents new methods including machine learning and compressive sampling, which help to improve safety, reliability, and performance. Condition Monitoring with Vibration Signals: Compressive Sampling and Learning Algorithms for Rotating Machines starts by introducing readers to Vibration Analysis Techniques and Machine Condition Monitoring (MCM). It then offers readers sections covering: Rotating Machine Condition Monitoring using Learning Algorithms; Classification Algorithms; and New Fault Diagnosis Frameworks designed for MCM. Readers will learn signal processing in the time-frequency domain, methods for linear subspace learning, and the basic principles of the learning method Artificial Neural Network (ANN). They will also discover recent trends of deep learning in the field of machine condition monitoring, new feature learning frameworks based on compressive sampling, subspace learning techniques for machine condition monitoring, and much more. Covers the fundamental as well as the state-of-the-art approaches to machine condition monitoringguiding readers from the basics of rotating machines to the generation of knowledge using vibration signals Provides new methods, including machine learning and compressive sampling, which offer significant improvements in accuracy with reduced computational costs Features learning algorithms that can be used for fault diagnosis and prognosis Includes previously and recently developed dimensionality reduction techniques and classification algorithms Condition Monitoring with Vibration Signals: Compressive Sampling and Learning Algorithms for Rotating Machines is an excellent book for research students, postgraduate students, industrial practitioners, and researchers. Get into the world of smart data security using machine learning algorithms and Python libraries Key FeaturesLearn machine learning algorithms and cybersecurity fundamentalsAutomate your daily workflow by applying use cases to many facets of securityImplement smart machine learning solutions to detect various cybersecurity problemsBook Description Cyber threats today are one of the costliest losses that an organization can face. In this book, we use the most efficient tool to solve the big problems that exist in the cybersecurity domain. The book begins by giving you the basics of ML in cybersecurity using Python and its libraries. You will explore various ML domains (such as time series analysis and ensemble modeling) to get your foundations right. You will implement various examples such as building system to identify malicious URLs, and building a program to detect fraudulent emails and spam. Later, you will learn how to make effective use of K-means algorithm to develop a solution to detect and alert you to any malicious activity in the network. Also learn how to implement biometrics and fingerprint to validate whether the user is a legitimate user or not. Finally, you will see how we change the game with TensorFlow and learn how deep learning is effective for creating models and training systems What you will learnUse machine learning algorithms with complex datasets to implement cybersecurity conceptsImplement machine learning algorithms such as clustering, k-means, and Naive Bayes to solve real-world problemsLearn to speed up a system using Python libraries with NumPy, Scikit-learn, and CUDAUnderstand how to combat malware, detect spam, and fight financial fraud to mitigate cyber crimesUse TensorFlow in the cybersecurity domain and implement real-world examplesLearn how machine learning and Python can be used in complex cyber issuesWho this book is for This book is for the data scientists, machine learning developers, security researchers, and anyone keen to apply machine learning to up-skill computer security. Having some working knowledge of Python and being familiar with the basics of machine learning and cybersecurity fundamentals will help to get the most out of the book

With the rise of different technologies in the Hi-tech world, Machine learning is becoming a hot topic, which everyone wants to learn about. Machine learning is a type of artificial intelligence program that you can use to give your computer the ability to learn without being completely programmed. Using algorithms that iteratively learn from data, machine learning allows computers to find hidden insights without being explicitly programmed where to look. Machine learning focuses deeply on developing computer programs that can change when exposed to new data. In addition to that, ML studies the construction of algorithms and how to make predictions on data. The topics covered in this book are: An introduction to Machine Learning Understanding Machine Learning How to use Machine Learning with you Data Analyzing needs Machine Learning examples Tips and tricks to teach your computer how to use Machine Learning Artificial Intelligence An overview of decision trees and random forests Supervised and Unsupervised learning Reinforcement Learning How to work with Clustering Algorithms Data Mining Machine learning techniques Data science How to use decision tree with Machine Learning And much, much more! Download this guide today and I guarantee you'll learn a lot from it.

In this practical book, four Cloudera data scientists present a set of self-contained patterns for performing large-scale data analysis with Spark. The authors bring Spark, statistical methods, and real-world data sets together to teach you how to approach analytics problems by example. You'll start with an introduction to Spark and its ecosystem, and then dive into patterns that apply common techniques—classification, collaborative filtering, and anomaly detection among others—to fields such as genomics, security, and finance. If you have an entry-level understanding of machine learning and statistics, and you program in Java, Python, or Scala, you'll find these patterns useful for working on your own data applications. Patterns include: Recommending music and the Audioscrobler data set Predicting forest cover with decision trees Anomaly detection in network traffic with K-means clustering Understanding Wikipedia with Latent Semantic Analysis Analyzing co-occurrence networks with GraphX Geospatial and temporal data analysis on the New York City Taxi Trips data Estimating financial risk through Monte Carlo simulation Analyzing genomics data and the BDG project Analyzing neuroimaging data with PySpark and Thunder

Machine Learning and Data Mining in Pattern Recognition

Decision Trees And Random Forests Work: Classification Machine Learning Algorithms

Patterns for Learning from Data at Scale

The Future Is Here!

Build 13 real-world projects with advanced numerical computations using the Python ecosystem

Decision Forests

Python Data Science Handbook

During the past decade there has been an explosion in computation and information technology. With it have come vast amounts of data in a variety of fields such as medicine, biology, finance, and marketing. The challenge of understanding these data has led to the development of new tools in the field of statistics, and spawned new areas such as data mining, machine learning, and bioinformatics. Many of these tools have common underpinnings but are often expressed with different terminology. This book describes the important ideas in these areas in a common conceptual framework. While the approach is statistical, the emphasis is on concepts rather than mathematics. Many examples are given, with a liberal use of color graphics. It should be a valuable resource for statisticians and anyone interested in data mining in science or industry. The book's coverage is broad, from supervised learning (prediction) to unsupervised learning. The many topics include neural networks, support vector machines, classification trees and boosting—the first comprehensive treatment of this topic in any book. This major new edition features many topics not covered in the original, including graphical models, random forests, ensemble methods, least angle regression & path algorithms for the lasso, non-negative matrix factorization, and spectral clustering. There is also a chapter on methods for “wide” data (p bigger than n), including multiple testing and false discovery rates. Trevor Hastie, Robert Tibshirani, and Jerome Friedman are professors of statistics at Stanford University. They are prominent researchers in this area: Hastie and Tibshirani developed generalized additive models and wrote a popular book of that title. Hastie co-developed much of the statistical modeling software and environment in R/S-PLUS and invented principal curves and surfaces. Tibshirani proposed the lasso and is co-author of the very successful An Introduction to the Bootstrap. Friedman is the co-inventor of many data-mining tools including CART, MARS, projection pursuit and gradient boosting.

"Learn intuitive machine learning techniques by exploring a classic problem. This Machine learning - decision trees and random forests online course will teach you cool machine learning techniques to predict survival probabilities aboard the Titanic - a Kaggle problem! In an age of decision fatigue and information overload, this course is a crisp yet thorough primer on two great machine learning techniques that help cut through the noise - decision trees and random forests."--Resource description page.

Theory and Applications

Machine Intelligence and Soft Computing

Statistics for Machine Learning

Hands-On Machine Learning for Cybersecurity

Condition Monitoring with Vibration Signals

Practical Solutions from Preprocessing to Deep Learning