Real World Machine Learning

Recognizing the pretentiousness ways to acquire this books **Real World Machine Learning** is additionally useful. You have remained in right site to start getting this info. acquire the Real World Machine Learning join that we offer here and check out the link.

You could purchase guide Real World Machine Learning or acquire it as soon as feasible. You could speedily download this Real World Machine Learning after getting deal. So, considering you require the ebook swiftly, you can straight get it. Its therefore very simple and appropriately fats, isnt it? You have to favor to in this appearance

Real-World Natural Language Processing - Masato Hagiwara 2021-12-21

Real-world Natural Language Processing shows you how to build the practical NLP applications that are transforming the way humans and computers work together. In Real-world Natural Language Processing you will learn how to: Design, develop, and deploy useful NLP applications Create named entity taggers Build machine translation systems Construct language generation systems and chatbots Use advanced NLP concepts such as attention and transfer learning Real-world Natural Language Processing teaches you how to create practical NLP applications without getting bogged down in complex language theory and the mathematics of deep learning. In this engaging book, you'll explore the core tools and techniques required to build a huge range of powerful NLP apps, including chatbots, language detectors, and text classifiers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Training computers to interpret and generate speech and text is a monumental challenge, and the payoff for reducing labor and improving human/computer interaction is huge! Th e field of Natural Language Processing (NLP) is advancing rapidly, with countless new tools and practices. This unique book offers an innovative collection of NLP techniques with applications in machine translation, voice assistants, text generation, and more. About the book Real-world Natural Language Processing shows you how to build the practical NLP applications that are transforming the way humans and computers work together. Guided by clear explanations of each core NLP topic, you'll create many interesting applications including a sentiment analyzer and a chatbot. Along the way, you'll use Python and open source libraries like AllenNLP and HuggingFace Transformers to speed up your development process. What's inside Design, develop, and deploy useful NLP applications Create named entity taggers Build machine translation systems Construct language generation systems and chatbots About the reader For Python programmers. No prior machine learning knowledge assumed. About the author Masato Hagiwara received his computer science PhD from Nagoya University in 2009. He has interned at Google and Microsoft Research, and worked at Duolingo as a Senior Machine Learning Engineer. He now runs his own research and consulting company. Table of Contents PART 1 BASICS 1 Introduction to natural language processing 2 Your first NLP application 3 Word and document embeddings 4 Sentence classification 5 Sequential labeling and language modeling PART 2 ADVANCED MODELS 6 Sequence-to-sequence models 7 Convolutional neural networks 8 Attention and Transformer 9 Transfer learning with pretrained language models PART 3 PUTTING INTO PRODUCTION 10 Best practices in developing NLP applications 11 Deploying and serving NLP applications

Lifelong Machine Learning - Zhiyuan Chen 2018-08-14

Lifelong Machine Learning, Second Edition is an introduction to an advanced machine learning paradigm that continuously learns by accumulating past knowledge that it then uses in

future learning and problem solving. In contrast, the current dominant machine learning paradigm learns in isolation: given a training dataset, it runs a machine learning algorithm on the dataset to produce a model that is then used in its intended application. It makes no attempt to retain the learned knowledge and use it in subsequent learning. Unlike this isolated system, humans learn effectively with only a few examples precisely because our learning is very knowledge-driven: the knowledge learned in the past helps us learn new things with little data or effort. Lifelong learning aims to emulate this capability, because without it, an AI system cannot be considered truly intelligent. Research in lifelong learning has developed significantly in the relatively short time since the first edition of this book was published. The purpose of this second edition is to expand the definition of lifelong learning, update the content of several chapters, and add a new chapter about continual learning in deep neural networks—which has been actively researched over the past two or three years. A few chapters have also been reorganized to make each of them more coherent for the reader. Moreover, the authors want to propose a unified framework for the research area. Currently, there are several research topics in machine learning that are closely related to lifelong learning—most notably, multi-task learning, transfer learning, and meta-learning—because they also employ the idea of knowledge sharing and transfer. This book brings all these topics under one roof and discusses their similarities and differences. Its goal is to introduce this emerging machine learning paradigm and present a comprehensive survey and review of the important research results and latest ideas in the area. This book is thus suitable for students, researchers, and practitioners who are interested in machine learning, data mining, natural language processing, or pattern recognition. Lecturers can readily use the book for courses in any of these related fields.

Hands-On Machine Learning with C++ - Kirill Kolodiazhnyi 2020-05-15

Implement supervised and unsupervised machine learning algorithms using C++ libraries such as PyTorch C++ API, Caffe2, Shogun, Shark-ML, mlpack, and dlib with the help of real-world examples and datasets Key FeaturesBecome familiar with data processing, performance measuring, and model selection using various C++ librariesImplement practical machine learning and deep learning techniques to build smart modelsDeploy machine learning models to work on mobile and embedded devicesBook Description C++ can make your machine learning models run faster and more efficiently. This handy guide will help you learn the fundamentals of machine learning (ML), showing you how to use C++ libraries to get the most out of your data. This book makes machine learning with C++ for beginners easy with its example-based approach, demonstrating how to implement supervised and unsupervised ML algorithms through real-world examples. This book will get you hands-on with tuning and optimizing a model for different use cases, assisting you with model selection and the measurement of performance. You'll cover techniques such as product recommendations, ensemble learning, and anomaly detection using modern C++ libraries such as PyTorch C++ API, Caffe2, Shogun, Shark-ML, mlpack, and dlib. Next, you'll explore neural networks and deep learning using examples such as image classification and sentiment analysis, which will help you solve various problems. Later, you'll learn how to handle production and deployment challenges on mobile and cloud platforms, before discovering how to export and import models using the ONNX format. By the end of this C++ book, you will have real-world machine learning and C++ knowledge, as well as the skills to use C++ to build powerful ML systems. What you will learnExplore how to load and preprocess various data types to suitable C++ data structuresEmploy key machine learning algorithms with various C++ librariesUnderstand the grid-search approach to find the best parameters for a machine learning modelImplement an algorithm for filtering anomalies in user data using Gaussian distributionImprove collaborative filtering to deal with dynamic user preferencesUse C++ libraries and APIs to manage model structures and parametersImplement a C++ program to solve image classification tasks with LeNet architectureWho this book is for You will find this C++ machine learning book useful if you

want to get started with machine learning algorithms and techniques using the popular C++ language. As well as being a useful first course in machine learning with C++, this book will also appeal to data analysts, data scientists, and machine learning developers who are looking to implement different machine learning models in production using varied datasets and examples. Working knowledge of the C++ programming language is mandatory to get started with this book.

Real-World Machine Learning - Henrik Brink 2016-03-02

In a world where big data is the norm and near-real-time decisions are crucial, machine learning (ML) is a critical component of the data workflow. Machine learning systems can quickly crunch massive amounts of information to offer insights and make decisions in a way that matches or even surpasses human cognitive abilities. These systems use sophisticated computational and statistical tools to build models that can recognize and visualize patterns, predict outcomes, forecast values, and make recommendations. Real-World Machine Learning is a practical guide designed to teach developers the art of ML project execution. The book introduces the day-to-day practice of machine learning and prepares readers to successfully build and deploy powerful ML systems. Using the Python language and the R statistical package, it starts with core concepts like data acquisition and modeling, classification, and regression. Then it moves through the most important ML tasks, like model validation, optimization and feature engineering. It uses real-world examples that help readers anticipate and overcome common pitfalls. Along the way, they will discover scalable and online algorithms for large and streaming data sets. Advanced readers will appreciate the in-depth discussion of enhanced ML systems through advanced data exploration and pre-processing methods. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

Building Machine Learning Powered Applications - Emmanuel Ameisen 2020-01-21

Learn the skills necessary to design, build, and deploy applications powered by machine learning (ML). Through the course of this hands-on book, you'll build an example ML-driven application from initial idea to deployed product. Data scientists, software engineers, and product managers—including experienced practitioners and novices alike—will learn the tools, best practices, and challenges involved in building a real-world ML application step by step. Author Emmanuel Ameisen, an experienced data scientist who led an AI education program, demonstrates practical ML concepts using code snippets, illustrations, screenshots, and interviews with industry leaders. Part I teaches you how to plan an ML application and measure success. Part II explains how to build a working ML model. Part III demonstrates ways to improve the model until it fulfills your original vision. Part IV covers deployment and monitoring strategies. This book will help you: Define your product goal and set up a machine learning problem Build your first end-to-end pipeline quickly and acquire an initial dataset Train and evaluate your ML models and address performance bottlenecks Deploy and monitor your models in a production environment

MATLAB for Machine Learning - Giuseppe Ciaburro 2017-08-28

Extract patterns and knowledge from your data in easy way using MATLAB About This Book Get your first steps into machine learning with the help of this easy-to-follow guide Learn regression, clustering, classification, predictive analytics, artificial neural networks and more with MATLAB Understand how your data works and identify hidden layers in the data with the power of machine learning. Who This Book Is For This book is for data analysts, data scientists, students, or anyone who is looking to get started with machine learning and want to build efficient data processing and predicting applications. A mathematical and statistical background will really help in following this book well. What You Will Learn Learn the introductory concepts of machine learning. Discover different ways to transform data using SAS XPORT, import and export tools, Explore the different types of regression techniques such as simple & multiple linear regression, ordinary least squares estimation, correlations and how to apply them to your data. Discover the basics of classification methods and how to implement Naive Bayes algorithm and Decision Trees in the

Matlab environment. Uncover how to use clustering methods like hierarchical clustering to grouping data using the similarity measures. Know how to perform data fitting, pattern recognition, and clustering analysis with the help of MATLAB Neural Network Toolbox. Learn feature selection and extraction for dimensionality reduction leading to improved performance. In Detail MATLAB is the language of choice for many researchers and mathematics experts for machine learning. This book will help you build a foundation in machine learning using MATLAB for beginners. You'll start by getting your system ready with t he MATLAB environment for machine learning and you'll see how to easily interact with the Matlab workspace. We'll then move on to data cleansing, mining and analyzing various data types in machine learning and you'll see how to display data values on a plot. Next, you'll get to know about the different types of regression techniques and how to apply them to your data using the MATLAB functions. You'll understand the basic concepts of neural networks and perform data fitting, pattern recognition, and clustering analysis. Finally, you'll explore feature selection and extraction techniques for dimensionality reduction for performance improvement. At the end of the book, you will learn to put it all together into real-world cases covering major machine learning algorithms and be comfortable in performing machine learning with MATLAB. Style and approach The book takes a very comprehensive approach to enhance your understanding of machine learning using MATLAB. Sufficient real-world examples and use cases are included in the book to help you grasp the concepts quickly and apply them easily in your day-to-day work.

Applications of Machine Learning - Prashant Johri 2020-05-04

This book covers applications of machine learning in artificial intelligence. The specific topics covered include human language, heterogeneous and streaming data, unmanned systems, neural information processing, marketing and the social sciences, bioinformatics and robotics, etc. It also provides a broad range of techniques that can be successfully applied and adopted in different areas. Accordingly, the book offers an interesting and insightful read for scholars in the areas of computer vision, speech recognition, healthcare, business, marketing, and bioinformatics.

Intelligent Projects Using Python - Santanu Pattanayak 2019-01-31

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

Machine Learning Algorithms - Giuseppe Bonaccorso 2017-07-24

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.

Hands-On Genetic Algorithms with Python - Eyal Wirsansky 2020-01-31

Explore the ever-growing world of genetic algorithms to solve search, optimization, and AI-related tasks, and improve machine learning models using Python libraries such as DEAP, scikit-learn, and NumPy Key Features Explore the ins and outs of genetic algorithms with this fast-paced guide Implement tasks such as feature selection, search optimization, and cluster analysis using Python Solve combinatorial problems, optimize functions, and enhance the performance of artificial intelligence applications Book Description Genetic algorithms are a family of search, optimization, and learning algorithms inspired by the principles of natural evolution. By imitating the evolutionary process, genetic algorithms can overcome hurdles encountered in traditional search algorithms and provide high-quality solutions for a variety of problems. This book will help you get to grips with a powerful yet simple approach to applying genetic

algorithms to a wide range of tasks using Python, covering the latest developments in artificial intelligence. After introducing you to genetic algorithms and their principles of operation, you'll understand how they differ from traditional algorithms and what types of problems they can solve. You'll then discover how they can be applied to search and optimization problems, such as planning, scheduling, gaming, and analytics. As you advance, you'll also learn how to use genetic algorithms to improve your machine learning and deep learning models, solve reinforcement learning tasks, and perform image reconstruction. Finally, you'll cover several related technologies that can open up new possibilities for future applications. By the end of this book, you'll have hands-on experience of applying genetic algorithms in artificial intelligence as well as in numerous other domains. What you will learn Understand how to use state-of-the-art Python tools to create genetic algorithm-based applications Use genetic algorithms to optimize functions and solve planning and scheduling problems Enhance the performance of machine learning models and optimize deep learning network architecture Apply genetic algorithms to reinforcement learning tasks using OpenAI Gym Explore how images can be reconstructed using a set of semi-transparent shapes Discover other bio-inspired techniques, such as genetic programming and particle swarm optimization Who this book is for This book is for software developers, data scientists, and AI enthusiasts who want to use genetic algorithms to carry out intelligent tasks in their applications. Working knowledge of Python and basic knowledge of mathematics and computer science will help you get the most out of this book.

Python - Prateek Joshi 2017-06-16

TensorFlow Machine Learning Projects - Ankit Jain 2018-11-30

Implement TensorFlow's offerings such as TensorBoard, TensorFlow.js, TensorFlow Probability, and TensorFlow Lite to build smart automation projects Key FeaturesUse machine learning and deep learning principles to build real-world projectsGet to grips with TensorFlow's impressive range of module offeringsImplement projects on GANs, reinforcement learning, and capsule networkBook Description TensorFlow has transformed the way machine learning is perceived. TensorFlow Machine Learning Projects teaches you how to exploit the benefits—simplicity, efficiency, and flexibility—of using TensorFlow in various real-world projects. With the help of this book, you'll not only learn how to build advanced projects using different datasets but also be able to tackle common challenges using a range of libraries from the TensorFlow ecosystem. To start with, you'll get to grips with using TensorFlow for machine learning projects; you'll explore a wide range of projects using TensorForest and TensorBoard for detecting exoplanets, TensorFlow.js for sentiment analysis, and TensorFlow Lite for digit classification. As you make your way through the book, you'll build projects in various real-world domains, incorporating natural language processing (NLP), the Gaussian process, autoencoders, recommender systems, and Bayesian neural networks, along with trending areas such as Generative Adversarial Networks (GANs), capsule networks, and reinforcement learning. You'll learn how to use the TensorFlow on Spark API and GPU-accelerated computing with TensorFlow to detect objects, followed by how to train and develop a recurrent neural network (RNN) model to generate book scripts. By the end of this book, you'll have gained the required expertise to build full-fledged machine learning projects at work. What you will learnUnderstand the TensorFlow ecosystem using various datasets and techniquesCreate recommendation systems for quality product recommendationsBuild projects using CNNs, NLP, and Bayesian neural networksPlay Pac-Man using deep reinforcement learningDeploy scalable TensorFlow-based machine learning systemsGenerate your own book script using RNNsWho this book is for TensorFlow Machine Learning Projects is for you if you are a data analyst, data scientist, machine learning professional, or deep learning enthusiast with basic knowledge of TensorFlow. This book is also for you if you want to build end-to-end projects in the machine learning domain using supervised, unsupervised, and reinforcement learning techniques

Introducing Machine Learning - Dino Esposito 2020-02-05

Master machine learning concepts and develop real-world solutions Machine learning offers immense opportunities, and Introducing Machine Learning delivers practical knowledge to make the most of them. Dino and Francesco Esposito start with a quick overview of the foundations of artificial intelligence and the basic steps of any machine learning project. Next, they introduce Microsoft's powerful ML.NET library, including capabilities for data processing, training, and evaluation. They present families of algorithms that can be trained to solve real-life problems, as well as deep learning techniques utilizing neural networks. The authors conclude by introducing valuable runtime services available through the Azure cloud platform and consider the long-term business vision for machine learning. · 14-time Microsoft MVP Dino Esposito and Francesco Esposito help you · Explore what's known about how humans learn and how intelligent software is built · Discover which problems machine learning can address · Understand the machine learning pipeline: the steps leading to a deliverable model · Use AutoML to automatically select the best pipeline for any problem and dataset · Master ML.NET, implement its pipeline, and apply its tasks and algorithms · Explore the mathematical foundations of machine learning · Make predictions, improve decision-making, and apply probabilistic methods · Group data via classification and clustering · Learn the fundamentals of deep learning, including neural network design · Leverage AI cloud services to build better real-world solutions faster About This Book · For professionals who want to build machine learning applications: both developers who need data science skills and data scientists who need relevant programming skills · Includes examples of machine learning coding scenarios built using the ML.NET library

Machine Learning Bookcamp - Alexey Grigorev 2021-11-23

Time to flex your machine learning muscles! Take on the carefully designed challenges of the Machine Learning Bookcamp and master essential ML techniques through practical application. Summary In Machine Learning Bookcamp you will: Collect and clean data for training models Use popular Python tools, including NumPy, Scikit-Learn, and TensorFlow Apply ML to complex datasets with images Deploy ML models to a production-ready environment The only way to learn is to practice! In Machine Learning Bookcamp, you'll create and deploy Python-based machine learning models for a variety of increasingly challenging projects. Taking you from the basics of machine learning to complex applications such as image analysis, each new project builds on what you've learned in previous chapters. You'll build a portfolio of business-relevant machine learning projects that hiring managers will be excited to see. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Master key machine learning concepts as you build actual projects! Machine learning is what you need for analyzing customer behavior, predicting price trends, evaluating risk, and much more. To master ML, you need great examples, clear explanations, and lots of practice. This book delivers all three! About the book Machine Learning Bookcamp presents realistic, practical machine learning scenarios, along with crystal-clear coverage of key concepts. In it, you'll complete engaging projects, such as creating a car price predictor using linear regression and deploying a churn prediction service. You'll go beyond the algorithms and explore important techniques like deploying ML applications on serverless systems and serving models with Kubernetes and Kubeflow. Dig in, get your hands dirty, and have fun building your ML skills! What's inside Collect and clean data for training models Use popular Python tools, including NumPy, Scikit-Learn, and TensorFlow Deploy ML models to a production-ready environment About the reader Python programming skills assumed. No previous machine learning knowledge is required. About the author Alexey Grigorev is a principal data scientist at OLX Group. He runs DataTalks.Club, a community of people who love data. Table of Contents 1 Introduction to machine learning 2 Machine learning for regression 3 Machine learning for classification 4 Evaluation metrics for classification 5 Deploying machine learning models 6 Decision

trees and ensemble learning 7 Neural networks and deep learning 8 Serverless deep learning 9 Serving models with Kubernetes and Kubeflow

Artificial Intelligence and Machine Learning Fundamentals - Zsolt Nagy 2018-12-12

Create AI applications in Python and lay the foundations for your career in data science Key FeaturesPractical examples that explain key machine learning algorithmsExplore neural networks in detail with interesting examplesMaster core AI concepts with engaging activitiesBook Description Machine learning and neural networks are pillars on which you can build intelligent applications. Artificial Intelligence and Machine Learning Fundamentals begins by introducing you to Python and discussing AI search algorithms. You will cover in-depth mathematical topics, such as regression and classification, illustrated by Python examples. As you make your way through the book, you will progress to advanced AI techniques and concepts, and work on real-life datasets to form decision trees and clusters. You will be introduced to neural networks, a powerful tool based on Moore's law. By the end of this book, you will be confident when it comes to building your own AI applications with your newly acquired skills! What you will learnUnderstand the importance, principles, and fields of AIImplement basic artificial intelligence concepts with PythonApply regression and classification concepts to real-world problemsPerform predictive analysis using decision trees and random forestsCarry out clustering using the k-means and mean shift algorithmsUnderstand the fundamentals of deep learning via practical examplesWho this book is for Artificial Intelligence and Machine Learning Fundamentals is for software developers and data scientists who want to enrich their projects with machine learning. You do not need any prior experience in AI. However, it’s recommended that you have knowledge of high school-level mathematics and at least one programming language (preferably Python).

Interpretable Machine Learning with Python - Serg Masís 2021-03-26

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 FeaturesLearn how to extract easy-to-understand insights from any machine learning modelBecome well-versed with interpretability techniques to build fairer, safer, and more reliable modelsMitigate risks in AI systems before they have broader implications by learning how to debug black-box modelsBook 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 learnRecognize the importance of interpretability in businessStudy models that are intrinsically interpretable such as linear models, decision trees, and Naïve BayesBecome well-versed in interpreting models with model-agnostic methodsVisualize how an image classifier works and what it learnsUnderstand how to mitigate the influence of bias in datasetsDiscover how to make models more reliable with adversarial robustnessUse monotonic constraints to make fairer and safer modelsWho 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.

Python Machine Learning Cookbook - Prateek Joshi 2016-06-23

100 recipes that teach you how to perform various machine learning tasks in the real world About This Book Understand which algorithms to use in a given context with the help of this exciting recipe-based guide Learn about perceptrons and see how they are used to build neural networks Stuck while making sense of images, text, speech, and real estate? This guide will come to your rescue, showing you how to perform machine learning for each one of these using various techniques Who This Book Is For This book is for Python programmers who are looking to use machine-learning algorithms to create real-world applications. This book is friendly to Python beginners, but familiarity with Python programming would certainly be useful to play around with the code. What You Will Learn Explore classification algorithms and apply them to the income bracket estimation problem Use predictive modeling and apply it to real-world problems Understand how to perform market segmentation using unsupervised learning Explore data visualization techniques to interact with your data in diverse ways Find out 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 Analyze stock market data using Conditional Random Fields Work with image data and build systems for image recognition and biometric face recognition Grasp how to use deep neural networks to build an optical character recognition system In Detail Machine learning is becoming increasingly pervasive in the modern data-driven world. It is used extensively across many fields such as search engines, robotics, self-driving cars, and more. With this book, you will learn how to perform various machine learning tasks in different environments. We'll start by exploring a range of real-life scenarios where machine learning can be used, and look at various building blocks. Throughout the book, you'll use a wide variety of machine learning algorithms to solve real-world problems and use Python to implement these algorithms. You'll discover how to deal with various types of data and explore the differences between machine learning paradigms such as supervised and unsupervised learning. We also cover a range of regression techniques, classification algorithms, predictive modeling, data visualization techniques, recommendation engines, and more with the help of real-world examples. Style and approach You will explore various real-life scenarios in this book where machine learning can be used, and learn about different building blocks of machine learning using independent recipes in the book.

Practical Machine Learning for Data Analysis Using Python - Abdulhamit Subasi 2020-06-05

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

Recent Advances in Internet of Things and Machine Learning - Valentina E. Balas 2022-02-15

This book covers a domain that is significantly impacted by the growth of soft computing. Internet of Things (IoT)-related applications are gaining much attention with more and more devices which are getting connected, and they become the potential components of some smart applications. Thus, a global enthusiasm has sparked over various domains such as health, agriculture, energy, security, and retail. So, in this book, the main objective is to capture this multifaceted nature of IoT and machine learning in one single place. According to the contribution of each chapter, the book also provides a future direction for IoT and machine learning research. The objectives of this book are to identify different issues, suggest feasible solutions to those identified issues, and enable researchers and practitioners from both academia and industry to interact with each other regarding emerging technologies related to IoT and machine learning. In this book, we look for novel chapters that recommend new methodologies, recent advancement, system architectures, and other solutions to prevail over the limitations of IoT and machine learning.

Mastering .NET Machine Learning - Jamie Dixon 2016-03-29

Master the art of machine learning with .NET and gain insight into real-world applications About This Book Based on .NET framework 4.6.1, includes examples on ASP.NET Core 1.0 Set up your business application to start using machine learning techniques Familiarize the user with some of the more common .NET libraries for machine learning Implement several common machine learning techniques Evaluate, optimize and adjust machine learning models Who This Book Is For This book is targeted at .Net developers who want to build complex machine learning systems. Some basic understanding of data science is required. What You Will Learn Write your own machine learning applications and experiments using the latest .NET framework, including .NET Core 1.0 Set up your business application to start using machine learning. Accurately predict the future using regressions. Discover hidden patterns using decision trees. Acquire, prepare, and combine datasets to drive insights. Optimize business throughput using Bayes Classifier. Discover (more) hidden patterns using KNN and Naive Bayes. Discover (even more) hidden patterns using K-Means and PCA. Use Neural Networks to improve business decision making while using the latest ASP.NET technologies. Explore “Big Data”, distributed computing, and how to deploy machine learning models to IoT devices – making machines self-learning and adapting Along the way, learn about Open Data, Bing maps, and MBrace In Detail .Net is one of the widely used platforms for developing applications. With the meteoric rise of Machine learning, developers are now keen on finding out how can they make their .Net applications smarter. Also, .NET developers are interested into moving into the world of devices and how to apply machine learning techniques to, well, machines. This book is packed with real-world examples to easily use machine learning techniques in your business applications. You will begin with introduction to F# and prepare yourselves for machine learning using .NET framework. You will be writing a simple linear regression model using an example which predicts sales of a product. Forming a base with the regression model, you will start using machine learning libraries available in .NET framework such as Math.NET, Numl.NET and Accord.NET with the help of a sample application. You will then move on to writing multiple linear regressions and logistic regressions. You will learn what is open data and the awesomeness of type providers. Next, you are going to address some of the issues that we have been glossing over so far and take a deep dive into obtaining, cleaning, and organizing our data. You will compare the utility of building a KNN and Naive Bayes model to achieve best possible results. Implementation of Kmeans and PCA using Accord.NET and Numl.NET libraries is covered with the help of an example application. We will then look at many of issues confronting creating real-world machine learning models like overfitting and how to combat them using confusion matrixes, scaling, normalization, and feature selection. You will now enter into the world of Neural Networks

and move your line of business application to a hybrid scientific application. After you have covered all the above machine learning models, you will see how to deal with very large datasets using MBrace and how to deploy machine learning models to Internet of Thing (IoT) devices so that the machine can learn and adapt on the fly Style and approach This book will guide you in learning everything about how to tackle the flood of data being encountered these days in your .NET applications with the help of popular machine learning libraries offered by the .NET framework.

Industrial Machine Learning - Andreas François Vermeulen 2019-11-30

Understand the industrialization of machine learning (ML) and take the first steps toward identifying and generating the transformational disruptors of artificial intelligence (AI). You will learn to apply ML to data lakes in various industries, supplying data professionals with the advanced skills required to handle the future of data engineering and data science. Data lakes currently generated by worldwide industrialized business activities are projected to reach 35 zettabytes (ZB) as the Fourth Industrial Revolution produces an exponential increase of volume, velocity, variety, variability, veracity, visualization, and value. Industrialization of ML evolves from AI and studying pattern recognition against the increasingly unstructured resource stored in data lakes. Industrial Machine Learning supplies advanced, yet practical examples in different industries, including finance, public safety, health care, transportation, manufactory, supply chain, 3D printing, education, research, and data science. The book covers: supervised learning, unsupervised learning, reinforcement learning, evolutionary computing principles, soft robotics disruptors, and hard robotics disruptors. What You Will Learn Generate and identify transformational disruptors of artificial intelligence (AI) Understand the field of machine learning (ML) and apply it to handle big data and process the data lakes in your environment Hone the skills required to handle the future of data engineering and data science Who This Book Is For Intermediate to expert level professionals in the fields of data science, data engineering, machine learning, and data management

Practical Deep Learning for Cloud, Mobile, and Edge - Anirudh Koul 2019-10-14

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

Deep Learning on Graphs - Yao Ma 2021-09-23

A comprehensive text on foundations and techniques of graph neural networks with applications in NLP, data mining, vision and healthcare.

Real-World Python - Lee Vaughan 2020-11-10

A project-based approach to learning Python programming for beginners. Intriguing projects teach you how to tackle challenging problems with code. You've mastered the basics. Now you're ready to explore some of Python's more powerful tools. Real-World Python will show you how. Through a series of hands-on projects, you'll investigate and solve real-world problems using sophisticated computer vision, machine learning, data analysis, and language processing tools. You'll be introduced to important modules like OpenCV, NumPy, Pandas, NLTK, Bokeh, Beautiful Soup, Requests, HoloViews, Tkinter, turtle, matplotlib, and more. You'll create complete, working programs and think through

intriguing projects that show you how to: • Save shipwrecked sailors with an algorithm designed to prove the existence of God • Detect asteroids and comets moving against a starfield • Program a sentry gun to shoot your enemies and spare your friends • Select landing sites for a Mars probe using real NASA maps • Send unbreakable messages based on a book code • Survive a zombie outbreak using data science • Discover exoplanets and alien megastructures orbiting distant stars • Test the hypothesis that we're all living in a computer simulation • And more! If you're tired of learning the bare essentials of Python Programming with isolated snippets of code, you'll relish the relevant and geeky fun of Real-World Python!

Human and Machine Learning - Jianlong Zhou 2018-06-07

With an evolutionary advancement of Machine Learning (ML) algorithms, a rapid increase of data volumes and a significant improvement of computation powers, machine learning becomes hot in different applications. However, because of the nature of “black-box” in ML methods, ML still needs to be interpreted to link human and machine learning for transparency and user acceptance of delivered solutions. This edited book addresses such links from the perspectives of visualisation, explanation, trustworthiness and transparency. The book establishes the link between human and machine learning by exploring transparency in machine learning, visual explanation of ML processes, algorithmic explanation of ML models, human cognitive responses in ML-based decision making, human evaluation of machine learning and domain knowledge in transparent ML applications. This is the first book of its kind to systematically understand the current active research activities and outcomes related to human and machine learning. The book will not only inspire researchers to passionately develop new algorithms incorporating human for human-centred ML algorithms, resulting in the overall advancement of ML, but also help ML practitioners proactively use ML outputs for informative and trustworthy decision making. This book is intended for researchers and practitioners involved with machine learning and its applications. The book will especially benefit researchers in areas like artificial intelligence, decision support systems and human-computer interaction.

Machine Learning Algorithms and Applications - Mettu Srinivas 2021-08-10

Machine Learning Algorithms is for current and ambitious machine learning specialists looking to implement solutions to real-world machine learning problems. It talks entirely about the various applications of machine and deep learning techniques, with each chapter dealing with a novel approach of machine learning architecture for a specific application, and then compares the results with previous algorithms. The book discusses many methods based in different fields, including statistics, pattern recognition, neural networks, artificial intelligence, sentiment analysis, control, and data mining, in order to present a unified treatment of machine learning problems and solutions. All learning algorithms are explained so that the user can easily move from the equations in the book to a computer program.

Industrial Applications of Machine Learning - Pedro Larrañaga 2018-12-12

Industrial Applications of Machine Learning shows how machine learning can be applied to address real-world problems in the fourth industrial revolution, and provides the required knowledge and tools to empower readers to build their own solutions based on theory and practice. The book introduces the fourth industrial revolution and its current impact on organizations and society. It explores machine learning fundamentals, and includes four case studies that address a real-world problem in the manufacturing or logistics domains, and approaches machine learning solutions from an application-oriented point of view. The book should be of special interest to researchers interested in real-world industrial problems. Features Describes the opportunities, challenges, issues, and trends offered by the fourth industrial revolution Provides a user-friendly introduction to machine learning with examples of cutting-edge applications in different industrial sectors Includes four case studies addressing real-world industrial problems solved with machine learning techniques A dedicated website for the book contains the datasets of the case studies for the reader's reproduction, enabling the groundwork

for future problem-solving Uses of three of the most widespread software and programming languages within the engineering and data science communities, namely R, Python, and Weka

TensorFlow Deep Learning Projects - Alexey Grigorev 2018-03-28

Leverage the power of Tensorflow to design deep learning systems for a variety of real-world scenarios Key Features Build efficient deep learning pipelines using the popular Tensorflow framework Train neural networks such as ConvNets, generative models, and LSTMs Includes projects related to Computer Vision, stock prediction, chatbots and more Book Description TensorFlow is one of the most popular frameworks used for machine learning and, more recently, deep learning. It provides a fast and efficient framework for training different kinds of deep learning models, with very high accuracy. This book is your guide to master deep learning with TensorFlow with the help of 10 real-world projects. TensorFlow Deep Learning Projects starts with setting up the right TensorFlow environment for deep learning. Learn to train different types of deep learning models using TensorFlow, including Convolutional Neural Networks, Recurrent Neural Networks, LSTMs, and Generative Adversarial Networks. While doing so, you will build end-to-end deep learning solutions to tackle different real-world problems in image processing, recommendation systems, stock prediction, and building chatbots, to name a few. You will also develop systems that perform machine translation, and use reinforcement learning techniques to play games. By the end of this book, you will have mastered all the concepts of deep learning and their implementation with TensorFlow, and will be able to build and train your own deep learning models with TensorFlow confidently. What you will learn Set up the TensorFlow environment for deep learning Construct your own ConvNets for effective image processing Use LSTMs for image caption generation Forecast stock prediction accurately with an LSTM architecture Learn what semantic matching is by detecting duplicate Quora questions Set up an AWS instance with TensorFlow to train GANs Train and set up a chatbot to understand and interpret human input Build an AI capable of playing a video game by itself –and win it! Who this book is for This book is for data scientists, machine learning developers as well as deep learning practitioners, who want to build interesting deep learning projects that leverage the power of Tensorflow. Some understanding of machine learning and deep learning, and familiarity with the TensorFlow framework is all you need to get started with this book.

Human-in-the-Loop Machine Learning - Robert Munro 2021-07-20

Machine learning applications perform better with human feedback. Keeping the right people in the loop improves the accuracy of models, reduces errors in data, lowers costs, and helps you ship models faster. Human-in-the-loop machine learning lays out methods for humans and machines to work together effectively. You'll find best practices on selecting sample data for human feedback, quality control for human annotations, and designing annotation interfaces. You'll learn to dreate training data for labeling, object detection, and semantic segmentation, sequence labeling, and more. The book starts with the basics and progresses to advanced techniques like transfer learning and self-supervision within annotation workflows.

Machine Learning and Deep Learning in Real-Time Applications - Mahrishi, Mehul 2020-04-24

Artificial intelligence and its various components are rapidly engulfing almost every professional industry. Specific features of AI that have proven to be vital solutions to numerous real-world issues are machine learning and deep learning. These intelligent agents unlock higher levels of performance and efficiency, creating a wide span of industrial applications. However, there is a lack of research on the specific uses of machine/deep learning in the professional realm. Machine Learning and Deep Learning in Real-Time Applications provides emerging research exploring the theoretical and practical aspects of machine learning and deep learning and their implementations as well as their ability to solve real-world problems within several professional disciplines including healthcare, business, and computer science. Featuring coverage on a broad range of topics such as image processing, medical improvements, and smart grids, this book is ideally designed for researchers,

academicians, scientists, industry experts, scholars, IT professionals, engineers, and students seeking current research on the multifaceted uses and implementations of machine learning and deep learning across the globe.

Real-World Machine Learning - Henrik Brink 2016-09-15

Summary Real-World Machine Learning is a practical guide designed to teach working developers the art of ML project execution. Without overdosing you on academic theory and complex mathematics, it introduces the day-to-day practice of machine learning, preparing you to successfully build and deploy powerful ML systems. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning systems help you find valuable insights and patterns in data, which you'd never recognize with traditional methods. In the real world, ML techniques give you a way to identify trends, forecast behavior, and make fact-based recommendations. It's a hot and growing field, and up-to-speed ML developers are in demand. About the Book Real-World Machine Learning will teach you the concepts and techniques you need to be a successful machine learning practitioner without overdosing you on abstract theory and complex mathematics. By working through immediately relevant examples in Python, you'll build skills in data acquisition and modeling, classification, and regression. You'll also explore the most important tasks like model validation, optimization, scalability, and real-time streaming. When you're done, you'll be ready to successfully build, deploy, and maintain your own powerful ML systems. What's Inside Predicting future behavior Performance evaluation and optimization Analyzing sentiment and making recommendations About the Reader No prior machine learning experience assumed. Readers should know Python. About the Authors Henrik Brink, Joseph Richards and Mark Fetherolf are experienced data scientists engaged in the daily practice of machine learning. Table of Contents PART 1: THE MACHINE-LEARNING WORKFLOW What is machine learning? Real-world data Modeling and prediction Model evaluation and optimization Basic feature engineering PART 2: PRACTICAL APPLICATION Example: NYC taxi data Advanced feature engineering Advanced NLP example: movie review sentiment Scaling machine-learning workflows Example: digital display advertising

Real World AI - Alyssa Simpson Rochwerger 2021-03-16

How can you successfully deploy AI? When AI works, it's nothing short of brilliant, helping companies make or save tremendous amounts of money while delighting customers on an unprecedented scale. When it fails, the results can be devastating. Most AI models never make it out of testing, but those failures aren't random. This practical guide to deploying AI lays out a human-first, responsible approach that has seen more than three times the success rate when compared to the industry average. In Real World AI, Alyssa Simpson Rochwerger and Wilson Pang share dozens of AI stories from startups and global enterprises alike featuring personal experiences from people who have worked on global AI deployments that impact billions of people every day. AI for business doesn't have to be overwhelming. Real World AI uses plain language to walk you through an AI approach that you can feel confident about-for your business and for your customers.

Python: Real World Machine Learning - Prateek Joshi 2016-11-14

Learn to solve challenging data science problems by building powerful machine learning models using Python About This Book Understand which algorithms to use in a given context with the help of this exciting recipe-based guide This practical tutorial tackles real-world computing problems through a rigorous and effective approach Build state-of-the-art models and develop personalized recommendations to perform machine learning at scale Who This Book Is For This Learning Path is for Python programmers who are looking to use machine learning algorithms to create real-world applications. It is ideal for Python professionals who want to work with large and complex datasets and Python developers and analysts or data scientists who are looking to add to their existing skills by accessing some of the most powerful recent trends in data science. Experience with Python, Jupyter Notebooks, and

command-line execution together with a good level of mathematical knowledge to understand the concepts is expected. Machine learning basic knowledge is also expected. What You Will Learn Use predictive modeling and apply it to real-world problems Understand how to perform market segmentation using unsupervised learning Apply your new-found skills to solve real problems, through clearly-explained code for every technique and test Compete with top data scientists by gaining a practical and theoretical understanding of cutting-edge deep learning algorithms Increase predictive accuracy with deep learning and scalable data-handling techniques Work with modern state-of-the-art large-scale machine learning techniques Learn to use Python code to implement a range of machine learning algorithms and techniques In Detail Machine learning is increasingly spreading in the modern data-driven world. It is used extensively across many fields such as search engines, robotics, self-driving cars, and more. Machine learning is transforming the way we understand and interact with the world around us. In the first module, Python Machine Learning Cookbook, you will learn how to perform various machine learning tasks using a wide variety of machine learning algorithms to solve real-world problems and use Python to implement these algorithms. The second module, Advanced Machine Learning with Python, is designed to take you on a guided tour of the most relevant and powerful machine learning techniques and you'll acquire a broad set of powerful skills in the area of feature selection and feature engineering. The third module in this learning path, Large Scale Machine Learning with Python, dives into scalable machine learning and the three forms of scalability. It covers the most effective machine learning techniques on a map reduce framework in Hadoop and Spark in Python. This Learning Path will teach you Python machine learning for the real world. The machine learning techniques covered in this Learning Path are at the forefront of commercial practice. This Learning Path combines some of the best that Packt has to offer in one complete, curated package. It includes content from the following Packt products: Python Machine Learning Cookbook by Prateek Joshi Advanced Machine Learning with Python by John Hearty Large Scale Machine Learning with Python by Bastiaan Sjardin, Alberto Boschetti, Luca Massaron Style and approach This course is a smooth learning path that will teach you how to get started with Python machine learning for the real world, and develop solutions to real-world problems. Through this comprehensive course, you'll learn to create the most effective machine learning techniques from scratch and more!

Practical Machine Learning with Python - Dipanjan Sarkar 2017-12-20

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

Python Machine Learning Blueprints - Alexander Combs 2019-01-31

Discover a project-based approach to mastering machine learning concepts by applying them to everyday problems using libraries such as scikit-learn, TensorFlow, and Keras Key FeaturesGet to grips with Python's machine learning libraries including scikit-learn, TensorFlow, and KerasImplement advanced concepts and popular machine learning algorithms in real-world projectsBuild analytics, computer vision, and neural network projects Book Description Machine learning is transforming the way we understand and interact with the world around us. This book is the perfect guide for you to put your knowledge and skills into practice and use the Python ecosystem to cover key domains in machine learning. This second edition covers a range of libraries from the Python ecosystem, including TensorFlow and Keras, to help you implement real-world machine learning projects. The book begins by giving you an overview of machine learning with Python. With the help of complex datasets and optimized techniques, you'll go on to understand how to apply advanced concepts and popular machine learning algorithms to real-world projects. Next, you'll cover projects from domains such as predictive analytics to analyze the stock market and recommendation systems for GitHub repositories. In addition to this, you'll also work on projects from the NLP domain to create a custom news feed using frameworks such as scikit-learn, TensorFlow, and Keras. Following this, you'll learn how to build an advanced chatbot, and scale things up using PySpark. In the concluding chapters, you can look forward to exciting insights into deep learning and you'll even create an application using computer vision and neural networks. By the end of this book, you'll be able to analyze data seamlessly and make a powerful impact through your projects. What you will learnUnderstand the Python data science stack and commonly used algorithmsBuild a model to forecast the performance of an Initial Public Offering (IPO) over an initial discrete trading window Understand NLP concepts by creating a custom news feedCreate applications that will recommend GitHub repositories based on ones you've starred, watched, or forkedGain the skills to build a chatbot from scratch using PySparkDevelop a market-prediction app using stock dataDelve into advanced concepts such as computer vision, neural networks, and deep learningWho this book is for This book is for machine learning practitioners, data scientists, and deep learning enthusiasts who want to take their machine learning skills to the next level by building real-world projects. The intermediate-level guide will help you to implement libraries from the Python ecosystem to build a variety of projects addressing various machine learning domains. Knowledge of Python programming and machine learning concepts will be helpful.

Federated AI for Real-World Business Scenarios - Dinesh C. Verma 2021-10-01

This book provides an overview of Federated Learning and how it can be used to build real-world AI-enabled applications. Real-world AI applications frequently have training data distributed in many different locations, with data at different sites having different properties and different formats. In many cases, data movement is not permitted due to security concerns, bandwidth, cost or regulatory restriction. Under these conditions, techniques of federated learning can enable creation of practical applications. Creating practical applications requires implementation of the cycle of learning from data, inferring from data, and acting based on the inference. This book will be the first one

to cover all stages of the Learn-Infer-Act cycle, and presents a set of patterns to apply federation to all stages. Another distinct feature of the book is the use of real-world applications with an approach that discusses all aspects that need to be considered in an operational system, including handling of data issues during federation, maintaining compliance with enterprise security policies, and simplifying the logistics of federated AI in enterprise contexts. The book considers federation from a manner agnostic to the actual AI models, allowing the concepts to be applied to all varieties of AI models. This book is probably the first one to cover the space of enterprise AI-based applications in a holistic manner.

Machine Learning with Python - Abhishek Vijayvargia 2018-03-01

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

Python Machine Learning Cookbook - Giuseppe Ciaburro 2019-03-30

Discover powerful ways to effectively solve real-world machine learning problems using key libraries including scikit-learn, TensorFlow, and PyTorch Key FeaturesLearn and implement machine learning algorithms in a variety of real-life scenariosCover a range of tasks catering to supervised, unsupervised and reinforcement learning techniquesFind easy-to-follow code solutions for tackling common and not-so-common challengesBook 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 learnUse predictive modeling and apply it to real-world problemsExplore data visualization techniques to interact with your dataLearn how to build a recommendation engineUnderstand how to interact with text data and build models to analyze itWork with speech data and recognize spoken words using Hidden Markov ModelsGet well versed with reinforcement learning, automated ML, and transfer learningWork with image data and build systems for image recognition and biometric face recognitionUse deep neural networks to build an optical character recognition systemWho 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.

Machine Learning and AI for Healthcare - Arjun Panesar 2019-02-04

Explore the theory and practical applications of artificial intelligence (AI) and machine learning in healthcare. This book offers a guided tour of machine learning algorithms, architecture design, and applications of learning in healthcare and big data challenges. You'll discover the ethical implications of healthcare data analytics and the future of AI in population and patient health optimization. You'll also create a machine learning model, evaluate performance and operationalize its outcomes within your organization. Machine Learning and AI for Healthcare provides techniques on how to apply machine learning within your organization and evaluate the efficacy, suitability, and efficiency of AI applications. These are illustrated through leading case studies, including how chronic disease is being redefined

through patient-led data learning and the Internet of Things. What You'll LearnGain a deeper understanding of key machine learning algorithms and their use and implementation within wider healthcare Implement machine learning systems, such as speech recognition and enhanced deep learning/AI Select learning methods/algorithms and tuning for use in healthcare Recognize and prepare for the future of artificial intelligence in healthcare through best practices, feedback loops and intelligent agentsWho This Book Is For Health care professionals interested in how machine learning can be used to develop health intelligence – with the aim of improving patient health, population health and facilitating significant care-payer cost savings.

Real World Machine Learning - Henrik Brink 2016

"Real-World Machine Learning is a practical guide designed to teach working developers the art of ML project execution. It will teach you the concepts and techniques you need to be a successful machine learning practitioner without overdosing you on abstract theory and complex mathematics. By working through immediately relevant examples in Python, you'll build skills in data acquisition and modeling, classification, and regression. You'll also explore the most important tasks like model validation, optimization, scalability, and real-time streaming. When you're done, you'll be ready to successfully build, deploy, and maintain your own powerful ML systems. Machine learning systems help you find valuable insights and patterns in data, which you'd never recognize with traditional methods. In the real world, ML techniques give you a way to identify trends, forecast behavior, and make fact-based recommendations. It's a hot and growing field, and up-to-speed ML developers are in demand."--Resource description page.