
Building Machine Learning Systems With Python Willi Richert

Eventually, you will extremely discover a supplementary experience and triumph by spending more cash. still when? complete you recognize that you require to get those every needs behind having significantly cash? Why dont you try to acquire something basic in the beginning? Thats something that will lead you to comprehend even more roughly speaking the globe, experience, some places, behind history, amusement, and a lot more?

It is your certainly own become old to play-act reviewing habit. along with guides you could enjoy now is Building Machine Learning Systems With Python Willi Richert below.



Building Machine Learning Powered Applications
Packt Publishing Ltd

Learn how to solve challenging machine learning problems with TensorFlow, Google's revolutionary new software library for deep learning. If you have some background in basic linear algebra and calculus, this practical book introduces machine-learning fundamentals by showing you how to design systems capable of detecting objects in images, understanding text, analyzing video, and predicting the properties of potential medicines. TensorFlow

for Deep Learning teaches concepts through practical examples and helps you build knowledge of deep learning foundations from the ground up. It's ideal for practicing developers with experience designing software systems, and useful for scientists and other professionals familiar with scripting but not necessarily with designing learning algorithms. Learn TensorFlow fundamentals, including how to perform basic computation Build simple learning systems to understand their mathematical foundations Dive into fully connected deep networks used in thousands of applications Turn prototypes into high-quality models with hyperparameter optimization Process images with convolutional neural networks Handle natural language datasets with recurrent neural networks Use reinforcement learning to solve games such as tic-tac-toe Train deep networks with hardware including GPUs and tensor processing units
Machine Learning in Action "O'Reilly Media, Inc."

Foundational Hands-On Skills for Succeeding with Real Data Science Projects This pragmatic book introduces both machine learning and data science, bridging gaps between data scientist and engineer, and helping you bring these techniques into production. It helps ensure that your efforts actually solve your problem, and offers unique coverage of real-world optimization in production settings. – From the Foreword by Paul Dix, series editor Machine Learning in Production is a crash course in data science and machine learning for people who need to solve real-world problems in production environments. Written for technically competent “accidental data scientists” with more curiosity and ambition than formal training, this complete and rigorous introduction stresses practice, not theory. Building on agile principles, Andrew and Adam Kelleher show how to quickly deliver significant value in production, resisting overhyped tools and unnecessary complexity. Drawing on their extensive experience, they help you ask useful questions and then execute production projects from start to finish. The authors show just how much information you can glean with straightforward queries, aggregations, and visualizations, and they teach indispensable error analysis methods to avoid costly mistakes. They turn to workhorse machine learning techniques such as linear regression, classification, clustering, and Bayesian inference, helping you choose the right algorithm for each production problem. Their concluding section on hardware, infrastructure, and distributed systems offers unique and invaluable guidance on optimization in production environments. Andrew and Adam always focus on what matters in production: solving the problems that offer the highest return on investment, using the simplest, lowest-risk approaches that work. Leverage agile principles to maximize development efficiency in production projects Learn from practical Python code examples and visualizations that bring essential

algorithmic concepts to life Start with simple heuristics and improve them as your data pipeline matures Avoid bad conclusions by implementing foundational error analysis techniques Communicate your results with basic data visualization techniques Master basic machine learning techniques, starting with linear regression and random forests Perform classification and clustering on both vector and graph data Learn the basics of graphical models and Bayesian inference Understand correlation and causation in machine learning models Explore overfitting, model capacity, and other advanced machine learning techniques Make informed architectural decisions about storage, data transfer, computation, and communication Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

Building Machine Learning Systems with Python - Second Edition Packt Publishing Ltd

Automate data and model pipelines for faster machine learning applications Key Features Build automated modules for different machine learning components Understand each component of a machine learning pipeline in depth Learn to use different open source AutoML and feature engineering platforms Book Description AutoML is designed to automate parts of Machine Learning. Readily available AutoML tools are making data science practitioners' work easy and are received well in the advanced analytics community. Automated Machine Learning covers the necessary foundation needed to create automated machine learning modules and helps you get up to speed with them in the most practical way possible. In this book, you'll learn how to automate different tasks in the machine learning pipeline such as data preprocessing, feature selection, model training, model optimization, and much more. In addition to this, it demonstrates how you can use the available automation libraries, such as auto-sklearn and MLBox, and create and extend your own custom AutoML

components for Machine Learning. By the end of this book, you will have a clearer understanding of the different aspects of automated Machine Learning, and you'll be able to incorporate automation tasks using practical datasets. You can leverage your learning from this book to implement Machine Learning in your projects and get a step closer to winning various machine learning competitions. What you will learn

- Understand the fundamentals of Automated Machine Learning systems
- Explore auto-sklearn and MLBox for AutoML tasks
- Automate your preprocessing methods along with feature transformation
- Enhance feature selection and generation using the Python stack
- Assemble individual components of ML into a complete AutoML framework
- Demystify hyperparameter tuning to optimize your ML models
- Dive into Machine Learning concepts such as neural networks and autoencoders
- Understand the information costs and trade-offs associated with AutoML

Who this book is for If you're a budding data scientist, data analyst, or Machine Learning enthusiast and are new to the concept of automated machine learning, this book is ideal for you. You'll also find this book useful if you're an ML engineer or data professional interested in developing quick machine learning pipelines for your projects. Prior exposure to Python programming will help you get the best out of this book.

Mathematics for Machine Learning "O'Reilly Media, Inc."

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

TensorFlow for Deep Learning Morgan & Claypool Publishers

Take a systematic approach to understanding the fundamentals of machine learning and deep learning from the ground up and how they are applied in practice. You will use this comprehensive guide for building and deploying learning models to address complex use cases while leveraging the computational resources of Google Cloud Platform. Author Ekaba Bisong shows you how machine learning tools and techniques are used to predict or classify events based on a set of interactions between variables known as features or attributes in a particular dataset. He teaches you how deep learning extends the machine learning algorithm of neural networks to learn complex tasks that are difficult for computers to perform, such as recognizing faces

and understanding languages. And you will know how to leverage cloud computing to accelerate data science and machine learning deployments. Building Machine Learning and Deep Learning Models on Google Cloud Platform is divided into eight parts that cover the fundamentals of machine learning and deep learning, the concept of data science and cloud services, programming for data science using the Python stack, Google Cloud Platform (GCP) infrastructure and products, advanced analytics on GCP, and deploying end-to-end machine learning solution pipelines on GCP. What You'll Learn Understand the principles and fundamentals of machine learning and deep learning, the algorithms, how to use them, when to use them, and how to interpret your results Know the programming concepts relevant to machine and deep learning design and development using the Python stack Build and interpret machine and deep learning models Use Google Cloud Platform tools and services to develop and deploy large-scale machine learning and deep learning products Be aware of the different facets and design choices to consider when modeling a learning problem Productionalize machine learning models into software products Who This Book Is For Beginners to the practice of data science and applied machine learning, data scientists at all levels, machine learning engineers, Google Cloud Platform data engineers/architects, and software developers

Packt Publishing Ltd

"This video, with the help of practical projects, highlights how TensorFlow can be used in different scenarios--this includes projects for training models, machine learning, deep learning, and working with various neural networks. Each project provides exciting and insightful exercises that will teach you how to use TensorFlow and show you how layers of data can be explored by working with tensors"--Resource description page.

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

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

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

Building Machine Learning Pipelines "O'Reilly Media, Inc."

Build machine and deep learning systems with the newly released TensorFlow 2 and Keras for the lab, production, and mobile devices
Key Features
Introduces and then uses TensorFlow 2 and Keras right from the start
Teaches key machine and deep learning techniques
Understand the fundamentals of deep learning and machine learning through clear explanations and extensive code samples
Book Description
Deep Learning with TensorFlow 2 and Keras, Second Edition teaches neural networks and deep learning techniques alongside TensorFlow (TF) and Keras. You'll learn how to write deep

learning applications in the most powerful, popular, and scalable machine learning stack available. TensorFlow is the machine learning library of choice for professional applications, while Keras offers a simple and powerful Python API for accessing TensorFlow. TensorFlow 2 provides full Keras integration, making advanced machine learning easier and more convenient than ever before. This book also introduces neural networks with TensorFlow, runs through the main applications (regression, ConvNets (CNNs), GANs, RNNs, NLP), covers two working example apps, and then dives into TF in production, TF mobile, and using TensorFlow with AutoML. What you will learn
Build machine learning and deep learning systems with TensorFlow 2 and the Keras API
Use Regression analysis, the most popular approach to machine learning
Understand ConvNets (convolutional neural networks) and how they are essential for deep learning systems such as image classifiers
Use GANs (generative adversarial networks) to create new data that fits with existing patterns
Discover RNNs (recurrent neural networks) that can process sequences of input intelligently, using one part of a sequence to correctly interpret another
Apply deep learning to natural human language and interpret natural language texts to produce an appropriate response
Train your models on the cloud and put TF to work in real environments
Explore how Google tools can automate simple ML workflows without the need for complex modeling
Who this book is for
This book is for Python developers and data scientists who want to build machine learning and deep learning systems with TensorFlow. This book gives you the theory and practice required to use Keras, TensorFlow 2, and AutoML to build machine learning systems. Some knowledge of machine learning is expected.

Building Machine Learning Systems with Python - Third Edition BPB Publications

Deep learning is often viewed as the exclusive domain of math PhDs and big tech companies. But as this hands-on guide demonstrates, programmers comfortable with Python can achieve impressive results in deep learning with little math background, small amounts of data, and minimal code. How? With fastai, the first library to provide a consistent interface to the most frequently used deep learning applications. Authors Jeremy Howard and Sylvain Gugger, the creators of fastai, show you how to train a model on a wide range of tasks using fastai and PyTorch. You'll also dive progressively further into deep learning theory to gain a complete understanding of the algorithms behind the scenes. Train models in computer vision, natural language processing, tabular data, and collaborative filtering Learn the latest deep learning techniques that matter most in practice Improve accuracy, speed, and reliability by understanding how deep learning models work Discover how to turn your models into web applications Implement deep learning algorithms from scratch Consider the ethical implications of your work Gain insight from the foreword by PyTorch cofounder, Soumith Chintala

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

Machine learning has become an integral part of many commercial applications and research projects, but this field is not exclusive to large companies with extensive research teams. If you use Python, even as a beginner, this book will teach you practical ways to build your own machine learning solutions. With all the data available today, machine learning applications are limited only by your imagination. You'll learn the steps necessary to create a successful machine-learning application with Python and the scikit-learn library. Authors Andreas Müller and Sarah Guido focus on the practical aspects of using machine learning algorithms, rather than the math behind them. Familiarity with the NumPy and matplotlib libraries will help you get even more from this book. With this book, you'll learn: Fundamental concepts and applications of machine learning Advantages and

shortcomings of widely used machine learning algorithms How to represent data processed by machine learning, including which data aspects to focus on Advanced methods for model evaluation and parameter tuning The concept of pipelines for chaining models and encapsulating your workflow Methods for working with text data, including text-specific processing techniques Suggestions for improving your machine learning and data science skills

Machine Learning for Beginners Apress

Summary Machine Learning in Action is unique book that blends the foundational theories of machine learning with the practical realities of building tools for everyday data analysis. You'll use the flexible Python programming language to build programs that implement algorithms for data classification, forecasting, recommendations, and higher-level features like summarization and simplification. About the Book A machine is said to learn when its performance improves with experience. Learning requires algorithms and programs that capture data and ferret out the interesting or useful patterns. Once the specialized domain of analysts and mathematicians, machine learning is becoming a skill needed by many. Machine Learning in Action is a clearly written tutorial for developers. It avoids academic language and takes you straight to the techniques you'll use in your day-to-day work. Many (Python) examples present the core algorithms of statistical data processing, data analysis, and data visualization in code you can reuse. You'll understand the concepts and how they fit in with tactical tasks like classification, forecasting, recommendations, and higher-level features like summarization and simplification. Readers need no prior experience with machine learning or statistical processing. Familiarity with Python is helpful. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. What's Inside A no-nonsense introduction Examples showing common ML tasks Everyday data analysis Implementing classic algorithms like Apriori and Adaboos Table of Contents PART 1 CLASSIFICATION Machine learning basics Classifying with k-Nearest Neighbors Splitting datasets one

feature at a time: decision trees Classifying with probability theory: naïve Bayes Logistic regression Support vector machines Improving classification with the AdaBoost meta algorithm PART 2 FORECASTING NUMERIC VALUES WITH REGRESSION Predicting numeric values: regression Tree-based regression PART 3 UNSUPERVISED LEARNING Grouping unlabeled items using k-means clustering Association analysis with the Apriori algorithm Efficiently finding frequent itemsets with FP-growth PART 4 ADDITIONAL TOOLS Using principal component analysis to simplify data Simplifying data with the singular value decomposition Big data and MapReduce

Up and Running Google AutoML and AI Platform: Building Machine Learning and NLP Models Using AutoML and AI Platform for Production Environment (English Edition) Packt Publishing Ltd

The fundamental mathematical tools needed to understand machine learning include linear algebra, analytic geometry, matrix decompositions, vector calculus, optimization, probability and statistics. These topics are traditionally taught in disparate courses, making it hard for data science or computer science students, or professionals, to efficiently learn the mathematics. This self-contained textbook bridges the gap between mathematical and machine learning texts, introducing the mathematical concepts with a minimum of prerequisites. It uses these concepts to derive four central machine learning methods: linear regression, principal component analysis, Gaussian mixture models and support vector machines. For students and others with a mathematical background, these derivations provide a starting point to machine learning texts. For those learning the mathematics for the first time, the methods help build intuition

and practical experience with applying mathematical concepts. Every chapter includes worked examples and exercises to test understanding. Programming tutorials are offered on the book's web site.

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

Machine Learning Systems Addison-Wesley Professional "The authors' clear visual style provides a comprehensive look at what's currently possible with artificial neural networks as well as a glimpse of the magic that's to come." –Tim Urban, author of Wait But Why Fully Practical, Insightful Guide to Modern Deep Learning Deep learning is transforming software, facilitating powerful new artificial intelligence capabilities, and driving unprecedented algorithm performance. Deep Learning Illustrated is uniquely intuitive and offers a complete introduction to the discipline's techniques. Packed with full-color figures and easy-to-follow code, it sweeps away the complexity of building deep learning models, making the subject approachable and fun to learn. World-class instructor and practitioner Jon Krohn—with visionary content from Grant Beyleveld and beautiful illustrations by Aglaé Bassens—presents straightforward analogies to explain

what deep learning is, why it has become so popular, and how it relates to other machine learning approaches. Krohn has created a practical reference and tutorial for developers, data scientists, researchers, analysts, and students who want to start applying it. He illuminates theory with hands-on Python code in accompanying Jupyter notebooks. To help you progress quickly, he focuses on the versatile deep learning library Keras to nimbly construct efficient TensorFlow models; PyTorch, the leading alternative library, is also covered. You'll gain a pragmatic understanding of all major deep learning approaches and their uses in applications ranging from machine vision and natural language processing to image generation and game-playing algorithms. Discover what makes deep learning systems unique, and the implications for practitioners Explore new tools that make deep learning models easier to build, use, and improve Master essential theory: artificial neurons, training, optimization, convolutional nets, recurrent nets, generative adversarial networks (GANs), deep reinforcement learning, and more Walk through building interactive deep learning applications, and move forward with your own artificial intelligence projects Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

Building Machine Learning Systems with Python Apress

Build cutting edge machine and deep learning systems for the lab, production, and mobile devices Key Features Understand the fundamentals of deep learning and machine learning through clear explanations and extensive code samples Implement graph neural networks, transformers using Hugging Face and TensorFlow Hub,

and joint and contrastive learning Learn cutting-edge machine and deep learning techniques Book Description Deep Learning with TensorFlow and Keras teaches you neural networks and deep learning techniques using TensorFlow (TF) and Keras. You'll learn how to write deep learning applications in the most powerful, popular, and scalable machine learning stack available. TensorFlow 2.x focuses on simplicity and ease of use, with updates like eager execution, intuitive higher-level APIs based on Keras, and flexible model building on any platform. This book uses the latest TF 2.0 features and libraries to present an overview of supervised and unsupervised machine learning models and provides a comprehensive analysis of deep learning and reinforcement learning models using practical examples for the cloud, mobile, and large production environments. This book also shows you how to create neural networks with TensorFlow, runs through popular algorithms (regression, convolutional neural networks (CNNs), transformers, generative adversarial networks (GANs), recurrent neural networks (RNNs), natural language processing (NLP), and graph neural networks (GNNs)), covers working example apps, and then dives into TF in production, TF mobile, and TensorFlow with AutoML. What you will learn Learn how to use the popular GNNs with TensorFlow to carry out graph mining tasks Discover the world of transformers, from pretraining to fine-tuning to evaluating them Apply self-supervised learning to natural language processing, computer vision, and audio signal processing Combine probabilistic and deep learning models using TensorFlow Probability Train your models on the cloud and put TF to work in real environments Build machine learning and deep

learning systems with TensorFlow 2.x and the Keras API Who this book is for This hands-on machine learning book is for Python developers and data scientists who want to build machine learning and deep learning systems with TensorFlow. This book gives you the theory and practice required to use Keras, TensorFlow, and AutoML to build machine learning systems. Some machine learning knowledge would be useful. We don't assume TF knowledge.

Machine Learning Engineering Apress

Through a series of recent breakthroughs, deep learning has boosted the entire field of machine learning. Now, even programmers who know close to nothing about this technology can use simple, efficient tools to implement programs capable of learning from data. This practical book shows you how. By using concrete examples, minimal theory, and two production-ready Python frameworks—Scikit-Learn and TensorFlow—author Aurélien Géron helps you gain an intuitive understanding of the concepts and tools for building intelligent systems. You'll learn a range of techniques, starting with simple linear regression and progressing to deep neural networks. With exercises in each chapter to help you apply what you've learned, all you need is programming experience to get started. Explore the machine learning landscape, particularly neural nets Use Scikit-Learn to track an example machine-learning project end-to-end Explore several training models, including support vector machines, decision trees, random forests, and ensemble methods Use the TensorFlow library to build and train neural nets Dive into neural net architectures, including convolutional nets, recurrent nets, and deep reinforcement learning Learn techniques for training and scaling deep neural nets

Introducing Machine Learning Simon and Schuster

This book describes deep learning systems: the algorithms, compilers, and processor components to efficiently train and deploy deep learning models for commercial applications. The exponential growth in computational power is slowing at a time when the amount of compute consumed by state-of-the-art deep learning (DL) workloads is rapidly growing. Model size, serving latency, and power constraints are a significant challenge in the deployment of DL models for many applications. Therefore, it is imperative to codesign algorithms, compilers, and hardware to accelerate advances in this field with holistic system-level and algorithm solutions that improve performance, power, and efficiency. Advancing DL systems generally involves three types of engineers: (1) data scientists that utilize and develop DL algorithms in partnership with domain experts, such as medical, economic, or climate scientists; (2) hardware designers that develop specialized hardware to accelerate the components in the DL models; and (3) performance and compiler engineers that optimize software to run more efficiently on a given hardware. Hardware engineers should be aware of the characteristics and components of production and academic models likely to be adopted by industry to guide design decisions impacting future hardware. Data scientists should be aware of deployment platform constraints when designing models. Performance engineers should support optimizations across diverse models, libraries, and hardware targets. The purpose of this book is to provide a solid understanding of (1) the design, training, and applications of DL algorithms in industry; (2) the compiler techniques to map deep learning code to hardware targets; and (3) the critical hardware

features that accelerate DL systems. This book aims to facilitate innovation for the advancement of DL systems. It is written for engineers working in one or more of these areas who seek to understand the entire system stack in order to better collaborate with engineers working in other parts of the system stack. The book details advancements and adoption of DL models in industry, explains the training and deployment process, describes the essential hardware architectural features needed for today's and future models, and details advances in DL compilers to efficiently execute algorithms across various hardware targets. Unique in this book is the holistic exposition of the entire DL system stack, the emphasis on commercial applications, and the practical techniques to design models and accelerate their performance. The author is fortunate to work with hardware, software, data scientist, and research teams across many high-technology companies with hyperscale data centers. These companies employ many of the examples and methods provided throughout the book.

Building Machine Learning Systems with Python Packt Publishing Ltd

Design efficient machine learning systems that give you more accurate results About This Book Gain an understanding of the machine learning design process Optimize machine learning systems for improved accuracy Understand common programming tools and techniques for machine learning Develop techniques and strategies for dealing with large amounts of data from a variety of sources Build models to solve unique tasks Who This Book Is For This book is for data scientists, scientists, or just the curious. To get the most out of this book, you will need to know some linear algebra and some Python, and have a basic knowledge of machine learning concepts. What You Will Learn Gain an understanding of the machine learning design process Optimize the error function of your machine learning system Understand the common programming patterns used in machine learning

co-Discover optimizing techniques that will help you get the most from your data Find out how to design models uniquely suited to your task In Detail Machine learning is one of the fastest growing trends in modern computing. It has applications in a wide range of fields, including economics, the natural sciences, web development, and business modeling. In order to harness the power of these systems, it is essential that the practitioner develops a solid understanding of the underlying design principles. There are many reasons why machine learning models may not give accurate results. By looking at these systems from a design perspective, we gain a deeper understanding of the underlying algorithms and the optimisation methods that are available. This book will give you a solid foundation in the machine learning design process, and enable you to build customised machine learning models to solve unique problems. You may already know about, or have worked with, some of the off-the-shelf machine learning models for solving common problems such as spam detection or movie classification, but to begin solving more complex problems, it is important to adapt these models to your own specific needs. This book will give you this understanding and more. Style and approach This easy-to-follow, step-by-step guide covers the most important machine learning models and techniques from a design perspective.

Practical Machine Learning with Python Packt Publishing Ltd

Learn to build powerful machine learning models quickly and deploy large-scale predictive applications About This Book Design, engineer and deploy scalable machine learning solutions with the power of Python Take command of Hadoop and Spark with Python for effective machine learning on a map reduce framework Build state-of-the-art models and develop personalized recommendations to perform machine learning at scale Who This Book Is For This book is for anyone who intends to work with large and complex data sets. Familiarity with basic Python and machine learning concepts is recommended. Working knowledge in statistics and computational mathematics would also be helpful. What You Will Learn Apply the most scalable machine learning algorithms Work with modern state-of-the-art large-scale machine learning techniques Increase predictive accuracy with deep learning and scalable data-handling techniques Improve

your work by combining the MapReduce framework with Spark Build powerful ensembles at scale Use data streams to train linear and non-linear predictive models from extremely large datasets using a single machine In Detail Large Python machine learning projects involve new problems associated with specialized machine learning architectures and designs that many data scientists have yet to tackle. But finding algorithms and designing and building platforms that deal with large sets of data is a growing need. Data scientists have to manage and maintain increasingly complex data projects, and with the rise of big data comes an increasing demand for computational and algorithmic efficiency. Large Scale Machine Learning with Python uncovers a new wave of machine learning algorithms that meet scalability demands together with a high predictive accuracy. Dive into scalable machine learning and the three forms of scalability. Speed up algorithms that can be used on a desktop computer with tips on parallelization and memory allocation. Get to grips with new algorithms that are specifically designed for large projects and can handle bigger files, and learn about machine learning in big data environments. We will also cover the most effective machine learning techniques on a map reduce framework in Hadoop and Spark in Python. Style and Approach This efficient and practical title is stuffed full of the techniques, tips and tools you need to ensure your large scale Python machine learning runs swiftly and seamlessly. Large-scale machine learning tackles a different issue to what is currently on the market. Those working with Hadoop clusters and in data intensive environments can now learn effective ways of building powerful machine learning models from prototype to production. This book is written in a style that programmers from other languages (R, Julia, Java, Matlab) can follow.

Deep Learning Illustrated Manning

Companies are spending billions on machine learning projects, but it's money wasted if the models can't be deployed effectively. In this practical guide, Hannes Hapke and Catherine Nelson walk you through the steps of automating a machine learning pipeline using the TensorFlow ecosystem. You'll learn

the techniques and tools that will cut deployment time from days to minutes, so that you can focus on developing new models rather than maintaining legacy systems. Data scientists, machine learning engineers, and DevOps engineers will discover how to go beyond model development to successfully productize their data science projects, while managers will better understand the role they play in helping to accelerate these projects. Understand the steps to build a machine learning pipeline Build your pipeline using components from TensorFlow Extended Orchestrate your machine learning pipeline with Apache Beam, Apache Airflow, and Kubeflow Pipelines Work with data using TensorFlow Data Validation and TensorFlow Transform Analyze a model in detail using TensorFlow Model Analysis Examine fairness and bias in your model performance Deploy models with TensorFlow Serving or TensorFlow Lite for mobile devices Learn privacy-preserving machine learning techniques