

Master Machine Learning Algorithms

If you're an experienced programmer interested in crunching data, this book will get you started with machine learning—a toolkit of algorithms that enables computers to train themselves to automate useful tasks. Authors Drew Conway and John Myles White help you understand machine learning and statistics tools through a series of hands-on case studies, instead of a traditional math-heavy presentation. Each chapter focuses on a specific problem in machine learning, such as classification, prediction, optimization, and recommendation. Using the R programming language, you'll learn how to analyze sample datasets and write simple machine learning algorithms. Machine Learning for Hackers is ideal for programmers from any background, including business, government, and academic research. Develop a naïve Bayesian classifier to determine if an email is spam, based only on its text Use linear regression to predict the number of page views for the top 1,000 websites Learn optimization techniques by attempting to break a simple letter cipher Compare and contrast U.S. Senators statistically, based on their voting records Build a “whom to follow” recommendation system from Twitter data

Master Machine Learning Algorithms Discover How They Work and Implement Them From Scratch Machine Learning Mastery

Under the aegis of machine learning in our data-driven machine age, computers are programming themselves and learning about and solving an extraordinary range of problems, from the mundane to the most daunting. Today it is machine learning programs that enable Amazon and Netflix to predict what users will like, Apple to power Siri's ability to understand voices, and Google to pilot cars. These programs are already helping us fight the war on cancer and predict the movements of the stock market, and they are making great headway with instant language translation and discovering new laws of nature. But machine learning is incomplete, and its practitioners across the globe are seeking the most powerful algorithm of all. The Master Algorithm will not be limited to solving particular problems but will be able to learn anything and solve any problem, however difficult, and Pedro Domingos, a trailblazing computer scientist, is at the very forefront of the search for it. With the Master Algorithm in hand and data as its fuel, machine learning essentially the automation of discovery, a kind of scientific method on steroids will become the most powerful technology humanity has ever devised. And The Master Algorithm will be its bible.

You must understand the algorithms to get good (and be recognized as being good) at machine learning. In this Ebook, finally cut through the math and learn exactly how machine learning algorithms work, then implement them from scratch, step-by-step.

A practical guide to implementing supervised and unsupervised machine learning algorithms in Python

Reinforcement Learning Algorithms with Python

Machine Learning Mastery With Python

The Master Algorithm

From Theory to Algorithms

An Introduction

Study Deep Learning Through Data Science. How to Build Artificial Intelligence Through Concepts of Statistics, Algorithms, Analysis and Data Mining

Are you looking forward to learning Machine Learning? No matter your skill level, this powerful blueprint 3 Books in 1 will get you started to help you begin and expand your knowledge. Inside you'll find: 1.) Master Machine Learning Fundamental For Beginners Framing - work out the scope and look at the project as a whole Get insights from your data Prepare your data so the algorithms can see the data patterns Explore models and compare the best ones Fine-tune and put your models together in a fantastic solution 2) Master Machine Learning For Business Leaders Machine Learning for your Business Data Mining Techniques Business Optimization Machine Learning for Marketing Machine Learning for Finance Day Trading with Machine Learning 3) Master Machine Learning For Aspiring Data Scientists The different types of machine learning algorithms, including regression and classification Data cleansing The downsides of some algorithms The ins and outs of machine learning iteration Training models Exploratory analysis And many more... So pick up your copy with this Powerful 3 books in 1 and start building your confidence, skills and begin your journey to Master the Concepts of Machine Learning.

Concepts of Machine Learning with Practical Approaches. KEY FEATURES ● Includes real-scenario examples to explain the working of Machine Learning algorithms. ● Includes graphical and statistical representation to simplify modeling Machine Learning and Neural Networks. ● Full of Python codes, numerous exercises, and model question papers for data science students. DESCRIPTION The book offers the readers the fundamental concepts of Machine Learning techniques in a user-friendly language. The book aims to give in-depth knowledge of the different Machine Learning (ML) algorithms and the practical implementation of the various ML approaches. This book covers different Supervised Machine Learning algorithms such as Linear Regression Model, Naïve Bayes classifier Decision Tree, K-nearest neighbor, Logistic Regression, Support Vector Machine, Random forest algorithms, Unsupervised Machine Learning algorithms such as k-means clustering, Hierarchical Clustering, Probabilistic clustering, Association rule mining, Apriori Algorithm, f-p growth algorithm, Gaussian mixture model and Reinforcement Learning algorithm such as Markov Decision Process (MDP), Bellman equations, policy evaluation using Monte Carlo, Policy iteration and Value iteration, Q-Learning, State-Action-Reward-State-Action (SARSA). It also includes various feature extraction and feature selection techniques, the Recommender System, and a brief overview of Deep Learning. By the end of this book, the reader can understand Machine Learning concepts and easily implement various ML algorithms to real-world problems. WHAT YOU WILL LEARN ● Perform feature extraction and feature selection techniques. ● Learn to select the best Machine Learning algorithm for a given problem. ● Get a stronghold in using popular Python libraries like Scikit-learn, pandas, and matplotlib. ● Practice how to implement different types of Machine Learning techniques. ● Learn about Artificial Neural Network along with the Back Propagation Algorithm. ● Make use of various recommended systems with powerful algorithms. WHO THIS BOOK IS FOR This book is designed for data science and analytics students, academicians, and researchers who want to explore the concepts of machine learning and practice the understanding of real cases. Knowing basic statistical and programming concepts would be good, although not mandatory. TABLE OF CONTENTS 1. Introduction 2. Supervised Learning Algorithms 3. Unsupervised Learning 4. Introduction to the Statistical Learning Theory 5. Semi-Supervised Learning and Reinforcement Learning 6. Recommended Systems

Machine learning is not just for professors. Weka is a top machine learning platform that provides an easy-to-use graphical interface and state-of-the-art algorithms. In this Ebook, learn exactly how to get started with applied machine learning using the Weka platform.

A thought-provoking and wide-ranging exploration of machine learning and the race to build computer intelligences as flexible as our own. In the world's top research labs and universities, the race is on to invent the ultimate learning algorithm: one capable of discovering any knowledge from data, and doing anything we want, before we even ask. In *The Master Algorithm*, Pedro Domingos lifts the veil to give us a peek inside the learning machines that power Google, Amazon, and your smartphone. He assembles a blueprint for the future universal learner--the Master Algorithm--and discusses what it will mean for business, science, and society. If data-ism is today's philosophy, this book is its bible.

Ensemble Learning Algorithms With Python

A Problem-Solver's Guide to Building Real-World Intelligent Systems

Data Science and Machine Learning

Supervised, Unsupervised, and Advanced Learning

Step-By-Step Guide to Implement Machine Learning Algorithms with Python

Practical Machine Learning with Python

"This textbook is a well-rounded, rigorous, and informative work presenting the mathematics behind modern machine learning techniques. It hits all the right notes: the choice of topics is up-to-date and perfect for a course on data science for mathematics students at the advanced undergraduate or early graduate level. This book fills a sorely-needed gap in the existing literature by not sacrificing depth for breadth, presenting proofs of major theorems and subsequent derivations, as well as providing a copious amount of Python code. I only wish a book like this had been around when I first began my journey!" -Nicholas Hoell, University of Toronto "This is a well-written book that provides a deeper dive into data-scientific methods than many introductory texts. The writing is clear, and the text logically builds up regularization, classification, and decision trees. Compared to its probable competitors, it carves out a unique niche. -Adam Loy, Carleton College The purpose of *Data Science and Machine Learning: Mathematical and Statistical Methods* is to provide an accessible, yet comprehensive textbook intended for students interested in gaining a better understanding of the mathematics and statistics that underpin the rich variety of ideas and machine learning algorithms in data science. Key Features: Focuses on mathematical understanding. Presentation is self-contained, accessible, and comprehensive. Extensive list of exercises and worked-out examples. Many concrete algorithms with Python code. Full color throughout. The Authors: Dirk P. Kroese, PhD, is a Professor of Mathematics and Statistics at The University of Queensland. He has published over 120 articles and five books in a wide range of areas in mathematics, statistics, data science, machine learning, and Monte Carlo methods. He is a pioneer of the well-known Cross-Entropy method—an adaptive Monte Carlo technique, which is being used around the world to help solve difficult estimation and optimization problems in science, engineering, and finance. Zdravko Botev, PhD, is an Australian Mathematical Science Institute Lecturer in Data Science and Machine Learning with an appointment at the University of New South Wales in Sydney, Australia. He is the recipient of the 2018 Christopher Heyde Medal of the Australian Academy of Science for distinguished research in the Mathematical Sciences. Thomas Taimre, PhD, is a Senior Lecturer of Mathematics and Statistics at The University of Queensland. His research interests range from applied probability and Monte Carlo methods to applied physics and the remarkably universal self-mixing effect in lasers. He has published over 100 articles, holds a patent, and is the coauthor of *Handbook of Monte Carlo Methods* (Wiley). Radislav Vaisman, PhD, is a Lecturer of Mathematics and Statistics at The University of Queensland. His research interests lie at the intersection of applied probability, machine learning, and computer science. He has published over 20 articles and two books. This book provides conceptual understanding of machine learning algorithms through supervised, unsupervised, and advanced learning techniques. The book consists of four parts: foundation, supervised learning, unsupervised learning, and advanced learning. The first part provides the fundamental materials, background, and simple machine learning algorithms, as the preparation for studying machine learning algorithms. The second and the third parts provide understanding of the supervised learning algorithms and the unsupervised learning algorithms as the core parts. The last part provides advanced machine learning algorithms: ensemble learning, semi-supervised learning, temporal learning, and reinforced learning. Provides comprehensive coverage of both learning algorithms: supervised and unsupervised learning; Outlines the computation paradigm for solving classification, regression, and clustering; Features essential techniques for building the a new generation of machine learning.

Forecasting is required in many situations. Stocking an inventory may require forecasts of demand months in advance. Telecommunication routing requires traffic forecasts a few minutes ahead. Whatever the circumstances or time horizons involved, forecasting is an important aid in effective and efficient planning. This textbook provides a comprehensive introduction to forecasting methods and presents enough information about each method for readers to use them sensibly.

Deep learning neural networks have become easy to define and fit, but are still hard to configure. Discover exactly how to improve the performance of deep learning neural network models on your predictive modeling projects. With clear explanations, standard Python libraries, and step-by-step tutorial lessons, you'll discover how to better train your models, reduce overfitting, and make more accurate predictions.

Data Science for Business 2019 (2 BOOKS IN 1)

Master Supervised and Unsupervised Learning Algorithms with Real Examples (English Edition)

Expert techniques for implementing popular machine learning algorithms, fine-tuning your models, and understanding how they work, 2nd Edition

Machine Learning in Action

Machine Learning

Hands-On Reinforcement Learning with Python

Master Data Analytics & Machine Learning with Optimized Marketing Strategies (Artificial Intelligence, Neural Networks, Algorithms & Predictive Modelling)

The Python ecosystem with scikit-learn and pandas is required for operational machine learning. Python is the rising platform for professional machine learning because you can use the same code to explore different models in R&D then deploy it directly to production. In this Ebook, learn exactly how to get started and apply machine learning using the Python ecosystem.

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

This book introduces basic-to-advanced deep learning algorithms used in a production environment by AI researchers and principal data scientists; it explains algorithms intuitively, including the underlying math, and shows how to implement them using popular Python-based deep learning libraries such as TensorFlow.

?This book includes 2 Manuscripts? Are you looking for new ways to grow your business, with resources you already have? Do you want to know how the big players like Netflix, Amazon, or Shopify use data analytics to MULTIPLY their growth? Keep listening to learn how to use data analytics to maximize YOUR business.

Hands-On Machine Learning with scikit-learn and Scientific Python Toolkits

Better Deep Learning

Machine Learning Mastery With Weka

Machine Learning Algorithms From Scratch with Python

The Ultimate Guide for Beginners to Master Machine Learning with Practical Applications to Artificial Intelligence and Deep Learning with Python

Master Machine Learning Fundamentals for Beginners, Business Leaders and Aspiring Data Scientists

Discover the Mathematical Language of Data in Python

Optimization happens everywhere. Machine learning is one example of such and gradient descent is probably the most famous algorithm for performing optimization. Optimization means to find the best value of some function or model. That can be the maximum or the minimum according to some metric. Using clear explanations, standard Python libraries, and step-by-step tutorial lessons, you will learn how to find the optimum point to numerical functions confidently using modern optimization algorithms.

It's time to dispel the myth that machine learning is difficult. Grokking Machine Learning teaches you how to apply ML to your projects using only standard Python code and high school-level math. No specialist knowledge is required to tackle the hands-on exercises using readily available machine learning tools! In Grokking Machine Learning, expert machine learning engineer Luis Serrano introduces the most valuable ML techniques and teaches you how to make them work for you. Practical examples illustrate each new concept to ensure you're grokking as you go. You'll build models for spam detection, language analysis, and image recognition as you lock in each carefully-selected skill. Packed with easy-to-follow Python-based exercises and mini-projects, this book sets you on the path to becoming a machine learning expert. Key Features · Different types of machine learning, including supervised and unsupervised learning · Algorithms for simplifying, classifying, and splitting data · Machine learning packages and tools · Hands-on exercises with fully-explained Python code samples For readers with intermediate programming knowledge in Python or a similar language. About the technology Machine learning is a collection of mathematically-based techniques and algorithms that enable computers to identify patterns and generate predictions from data. This revolutionary data analysis approach is behind everything from recommendation systems to self-driving cars, and is transforming industries from finance to art.

Comprehensive recipes to give you valuable insights on Transformers, Reinforcement Learning, and more Key Features Deep Learning solutions from Kaggle Masters and Google Developer Experts Get to grips with the fundamentals including variables, matrices, and data sources Learn advanced techniques to make your algorithms faster and more accurate Book Description The independent recipes in Machine Learning Using TensorFlow Cookbook will teach you how to perform complex data computations and gain valuable insights into your data. Dive into recipes on training models, model evaluation, sentiment analysis, regression analysis, artificial neural networks, and deep learning - each using Google ' s machine learning library, TensorFlow. This cookbook covers the fundamentals of the TensorFlow library, including variables, matrices, and various data sources. You ' ll discover real-world implementations of Keras and TensorFlow and learn how to use estimators to train linear models and boosted trees, both for classification and regression. Explore the practical applications of a variety of deep learning architectures, such as recurrent neural networks and Transformers, and see how they can be used to solve computer vision and natural language processing (NLP) problems. With the help of this book, you will be proficient in using TensorFlow, understand deep learning from the basics, and be able to implement machine learning algorithms in real-world scenarios. What you will learn Take TensorFlow into production Implement and fine-tune Transformer models for various NLP tasks Apply reinforcement learning algorithms using the TF-Agents framework Understand linear regression techniques and use Estimators to train linear models Execute neural networks and improve predictions on tabular data Master convolutional neural networks and recurrent neural networks through practical recipes Who this book is for If you are a data scientist or a machine learning engineer, and you want to skip detailed theoretical explanations in favor of building production-ready machine learning models using TensorFlow, this book is for you. Basic familiarity with Python, linear algebra, statistics, and machine learning is necessary to make the most out of this book.

Reinforcement learning is a self-evolving type of machine learning that takes us closer to achieving true artificial intelligence. This easy-to-follow guide explains everything from scratch using rich examples written in Python.

Get Started, Build Accurate Models and Work Through Projects Step-by-Step

Python Machine Learning

Grokking Machine Learning

Machine Learning Mastery With R

Machine Learning Mathematics

Reinforcement Learning, second edition

Mathematical and Statistical Methods

MACHINE LEARNING - PYTHON Buy the Paperback version of this book, and get the Kindle eBook version included for FREE! Do You Want to Become An Expert Of Machine Learning?? Start Getting this Book and Follow My Step by Step Explanations! Click Add To Cart Now! This book is for anyone who would like to learn how to develop machine-learning systems. We will cover the most important concepts about machine learning algorithms, in both a theoretical and a practical way, and we'll implement many machine-learning algorithms using the Scikit-learn library in the Python programming language. In the first chapter, you'll learn the most important concepts of machine learning, and, in the next chapter, you'll work mainly with the classification. In the last chapter you'll learn how to train your model. I assume that you've knowledge of the basics of programming This book contains illustrations and step-by-step explanations with bullet points and exercises for easy and enjoyable learning. Benefits of reading this book that you're not going to find anywhere else: Introduction to Machine Learning Classification How to train a Model Different Models Combinations Don't miss out on this new step by step guide to Machine Learning. All you need to do is scroll up and click on the BUY NOW button to learn all about it!

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.

Develop self-learning algorithms and agents using TensorFlow and other Python tools, frameworks, and libraries Key Features Learn, develop, and deploy advanced reinforcement learning algorithms to solve a variety of tasks Understand and develop model-free and model-based algorithms for building self-learning agents Work with advanced Reinforcement Learning concepts and algorithms such as imitation learning and evolution strategies Book Description Reinforcement Learning (RL) is a popular and promising branch of AI that involves making smarter models and agents that can automatically determine ideal behavior based on changing requirements. This book will help you master RL algorithms and understand their implementation as you build self-learning agents. Starting with an introduction to the tools, libraries, and setup needed to work in the RL environment, this book covers the building blocks of RL and delves into value-based methods, such as the application of Q-learning and SARSA algorithms. You'll learn how to use a combination of Q-learning and neural networks to solve complex problems. Furthermore, you'll study the policy gradient methods, TRPO, and PPO, to improve performance and stability, before moving on to the DDPG and TD3 deterministic algorithms. This book also covers how imitation learning techniques work and how Dagger can teach an agent to drive. You'll discover evolutionary strategies and black-box optimization techniques, and see how they can improve RL algorithms. Finally, you'll get to grips with exploration approaches, such as UCB and UCB1, and develop a meta-algorithm called ESBS. By the end of the book, you'll have worked with key RL algorithms to overcome challenges in real-world applications, and be part of the RL research community. What you will learn Develop an agent to play CartPole using the OpenAI Gym interface Discover the model-based reinforcement learning paradigm Solve the Frozen Lake problem with dynamic programming Explore Q-learning and SARSA with a view to playing a taxi game Apply Deep Q-Networks (DQNs) to Atari games using Gym Study policy gradient algorithms, including Actor-Critic and REINFORCE Understand and apply PPO and TRPO in continuous locomotion environments Get to

grips with evolution strategies for solving the lunar lander problem Who this book is for If you are an AI researcher, deep learning user, or anyone who wants to learn reinforcement learning from scratch, this book is for you. You'll also find this reinforcement learning book useful if you want to learn about the advancements in the field. Working knowledge of Python is necessary.

R has been the gold standard in applied machine learning for a long time. Surveys show that it is the most popular platform used by professional data scientists. It is also preferred by the best data scientists in the world. In this Ebook, learn how to get started, practice and apply machine learning using the R platform.

Interpretable Machine Learning

Machine Learning Foundations

Optimization for Machine Learning

Hands-On Deep Learning Algorithms with Python

Case Studies and Algorithms to Get You Started

Master deep learning algorithms with extensive math by implementing them using TensorFlow

Discover How They Work and Implement Them From Scratch

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.

Updated and revised second edition of the bestselling guide to exploring and mastering the most important algorithms for solving complex machine learning problems Key Features Updated to include new algorithms and techniques Code updated to Python 3.8 & TensorFlow 2.x New coverage of regression analysis, time series analysis, deep learning models, and cutting-edge applications Book Description Mastering Machine Learning Algorithms, Second Edition helps you harness the real power of machine learning algorithms in order to implement smarter ways of meeting today's overwhelming data needs. This newly updated and revised guide will help you master algorithms used widely in semi-supervised learning, reinforcement learning, supervised learning, and unsupervised learning domains. You will use all the modern libraries from the Python ecosystem – including NumPy and Keras – to extract features from varied complexities of data. Ranging from Bayesian models to the Markov chain Monte Carlo algorithm to Hidden Markov models, this machine learning book teaches you how to extract features from your dataset, perform complex dimensionality reduction, and train supervised and semi-supervised models by making use of Python-based libraries such as scikit-learn. You will also discover practical applications for complex techniques such as maximum likelihood estimation, Hebbian learning, and ensemble learning, and how to use TensorFlow 2.x to train effective deep neural networks. By the end of this book, you will be ready to implement and solve end-to-end machine learning problems and use case scenarios. What you will learn Understand the characteristics of a machine learning algorithm Implement algorithms from supervised, semi-supervised, unsupervised, and RL domains Learn how regression works in time-series analysis and risk prediction Create, model, and train complex probabilistic models Cluster high-dimensional data and evaluate model accuracy Discover how artificial neural networks work – train, optimize, and validate them Work with autoencoders, Hebbian networks, and GANs Who this book is for This book is for data science professionals who want to delve into complex ML algorithms to understand how various machine learning models can be built. Knowledge of Python programming is required.

Linear algebra is a pillar of machine learning. You cannot develop a deep understanding and application of machine learning without it. In this laser-focused Ebook, you will finally cut through the equations, Greek letters, and confusion, and discover the topics in linear algebra that you need to know. Using clear explanations, standard Python libraries, and step-by-step tutorial lessons, you will discover what linear algebra is, the importance of linear algebra to machine learning, vector, and matrix operations, matrix factorization, principal component analysis, and much more.

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

Learn, understand, and develop smart algorithms for addressing AI challenges

Develop Deep Learning Models on Theano and TensorFlow Using Keras

Deep Learning With Python

Machine Learning for Hackers

Mastering Machine Learning Algorithms

Master reinforcement and deep reinforcement learning using OpenAI Gym and TensorFlow

The Art and Science of Algorithms that Make Sense of Data

Predictive performance is the most important concern on many classification and regression problems. Ensemble learning algorithms combine the predictions from multiple models and are designed to perform better than any contributing ensemble member. Using clear explanations, standard Python libraries, and step-by-step tutorial lessons, you will discover how to confidently and effectively improve predictive modeling performance using ensemble algorithms.

? 55% OFF for Bookstores! NOW at \$ 13.49 instead of \$ 29.97! LAST DAYS! ? Do you want to learn how to design and master different Machine Learning algorithms quickly and easily? Your Customers Will Love This Amazing Guide! Today, we live in the era of Artificial Intelligence. Self-driving cars, customized product recommendations, real-time pricing, speech and facial recognition are just a few examples proving this truth. Also, think about medical diagnostics or automation of mundane and repetitive labor tasks; all these highlight the fact that we live in interesting times. From research topics to projects and applications in different stages of production, there is a lot going on in the world of Machine Learning. Machines and automation represent a huge part of our daily life. They are becoming part of our experience and existence. This is Machine Learning. Artificial Intelligence is currently one of the most thriving fields any programmer would wish to delve into, and for a good reason: this is the future! Simply put, Machine Learning is about teaching machines to think and make decisions as we would. The difference between the way machines learn and the way we do is that while for the most part we learn from experiences, machines learn from data. Starting from scratch, Python Machine Learning explains how this happens, how machines build their experience and compounding knowledge. Data forms the core of Machine Learning because within data lie truths whose depths exceed our imagination. The computations machines can perform on data are incredible, beyond anything a human brain could do. Once we introduce data to a machine learning model, we must create an environment where we update the data stream frequently. This builds the machine's learning ability. The more data Machine Learning models are exposed to, the easier it is for these models to expand their potential. Some of the topics that we will discuss inside include: What is Machine Learning and how it is applied in real-world situations Understanding the differences between Machine Learning, Deep Learning, and Artificial Intelligence Supervised learning, unsupervised learning, and semi-supervised learning The place of Regression techniques in Machine Learning, including Linear Regression in Python Machine learning training models How to use Lists and Modules in Python The 12 essential libraries for Machine Learning in Python What is the Tensorflow library Artificial Neural Networks And Much More! While most books only focus on widespread details without going deeper into the different models and techniques, Python Machine Learning explains how to master the concepts of Machine Learning technology and helps you to understand how researchers are breaking the boundaries of Data Science to mimic human intelligence in machines using various Machine Learning algorithms. Even if some concepts of Machine Learning algorithms can appear complex to most computer programming beginners, this book takes the time to explain them in a simple and concise way. Would You Like To Know More? Buy It NOW And Let Your Customers Get Addicted To This Amazing Book!

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

Explore and master the most important algorithms for solving complex machine learning problems. Key Features Discover high-performing machine learning algorithms and understand how they work in depth. One-stop solution to mastering supervised, unsupervised, and semi-supervised machine learning algorithms and their implementation. Master concepts related to algorithm tuning, parameter optimization, and more Book Description Machine learning is a subset of AI that aims to make modern-day computer systems smarter and more intelligent. The real power of machine learning resides in its algorithms, which make even the most difficult things capable of being handled by machines. However, with the advancement in the technology and requirements of data, machines will have to be smarter than they are today to meet the overwhelming data needs; mastering these algorithms and using them optimally is the need of the hour. Mastering Machine Learning Algorithms is your complete guide to quickly getting to grips with popular machine learning algorithms. You will be introduced to the most widely used algorithms in supervised, unsupervised, and semi-supervised machine learning, and will learn how to use them in the best possible manner. Ranging from Bayesian models to the MCMC algorithm to Hidden Markov models, this book will teach you how to extract features from your dataset and perform dimensionality reduction by making use of Python-based libraries such as scikit-learn. You will also learn how to use Keras and TensorFlow to train effective neural networks. If you are looking for a single resource to study, implement,

and solve end-to-end machine learning problems and use-cases, this is the book you need. What you will learn Explore how a ML model can be trained, optimized, and evaluated Understand how to create and learn static and dynamic probabilistic models Successfully cluster high-dimensional data and evaluate model accuracy Discover how artificial neural networks work and how to train, optimize, and validate them Work with Autoencoders and Generative Adversarial Networks Apply label spreading and propagation to large datasets Explore the most important Reinforcement Learning techniques Who this book is for This book is an ideal and relevant source of content for data science professionals who want to delve into complex machine learning algorithms, calibrate models, and improve the predictions of the trained model. A basic knowledge of machine learning is preferred to get the best out of this guide.

Forecasting: principles and practice

Basics of Linear Algebra for Machine Learning

Machine Learning Algorithms

Machine Learning Using TensorFlow Cookbook

Statistical Methods for Machine Learning

Master Machine Learning Algorithms

Expert techniques to implement popular machine learning algorithms and fine-tune your models

The significantly expanded and updated new edition of a widely used text on reinforcement learning, one of the most active research areas in artificial intelligence. Reinforcement learning, one of the most active research areas in artificial intelligence, is a computational approach to learning whereby an agent tries to maximize the total amount of reward it receives while interacting with a complex, uncertain environment. In Reinforcement Learning, Richard Sutton and Andrew Barto provide a clear and simple account of the field's key ideas and algorithms. This second edition has been significantly expanded and updated, presenting new topics and updating coverage of other topics. Like the first edition, this second edition focuses on core online learning algorithms, with the more mathematical material set off in shaded boxes. Part I covers as much of reinforcement learning as possible without going beyond the tabular case for which exact solutions can be found. Many algorithms presented in this part are new to the second edition, including UCB, Expected Sarsa, and Double Learning. Part II extends these ideas to function approximation, with new sections on such topics as artificial neural networks and the Fourier basis, and offers expanded treatment of off-policy learning and policy-gradient methods. Part III has new chapters on reinforcement learning's relationships to psychology and neuroscience, as well as an updated case-studies chapter including AlphaGo and AlphaGo Zero, Atari game playing, and IBM Watson's wagering strategy. The final chapter discusses the future societal impacts of reinforcement learning.

Introduces machine learning and its algorithmic paradigms, explaining the principles behind automated learning approaches and the considerations underlying their usage.

You must understand algorithms to get good at machine learning. The problem is that they are only ever explained using Math. No longer. In this Ebook, finally cut through the math and learn exactly how machine learning algorithms work. Using clear explanations, simple pure Python code (no libraries!) and step-by-step tutorials you will discover how to load and prepare data, evaluate model skill, and implement a suite of linear, nonlinear and ensemble machine learning algorithms from scratch.

Statistics is a pillar of machine learning. You cannot develop a deep understanding and application of machine learning without it. Cut through the equations, Greek letters, and confusion, and discover the topics in statistics that you need to know. Using clear explanations, standard Python libraries, and step-by-step tutorial lessons, you will discover the importance of statistical methods to machine learning, summary stats, hypothesis testing, nonparametric stats, resampling methods, and much more.

Make Better Predictions with Bagging, Boosting, and Stacking

Analyze Data, Develop Models, and Work Through Projects

Discover how to Transform Data into Knowledge with Python

Create powerful machine learning algorithms with TensorFlow

How the Quest for the Ultimate Learning Machine Will Remake Our World

Understand Your Data, Create Accurate Models, and Work Projects End-to-End

Train Faster, Reduce Overfitting, and Make Better Predictions

Covering all the main approaches in state-of-the-art machine learning research, this will set a new standard as an introductory textbook.

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.

Integrate scikit-learn with various tools such as NumPy, pandas, imbalanced-learn, and scikit-surprise and use it to solve real-world machine learning problems Key Features Delve into machine learning with this comprehensive guide to scikit-learn and scientific Python Master the art of data-driven problem-solving with hands-on examples Foster your theoretical and practical knowledge of supervised and unsupervised machine learning algorithms Book Description Machine learning is applied everywhere, from business to research and academia, while scikit-learn is a versatile library that is popular among machine learning practitioners. This book serves as a practical guide for anyone looking to provide hands-on machine learning solutions with scikit-learn and Python toolkits. The book begins with an explanation of machine learning concepts and fundamentals, and strikes a balance between theoretical concepts and their applications. Each chapter covers a different set of algorithms, and shows you how to use them to solve real-life problems. You'll also learn about various key supervised and unsupervised machine learning algorithms using practical examples. Whether it is an instance-based learning algorithm, Bayesian estimation, a deep neural network, a tree-based ensemble, or a recommendation system, you'll gain a thorough understanding of its theory and learn when to apply it. As you advance, you'll learn how to deal with unlabeled data and when to use different clustering and anomaly detection algorithms. By the end of this machine learning book, you'll have learned how to take a data-driven approach to provide end-to-end machine learning solutions. You'll also have discovered how to formulate the problem at hand, prepare required data, and evaluate and deploy models in production. What you will learn Understand when to use supervised, unsupervised, or reinforcement learning algorithms Find out how to collect and prepare your data for machine learning tasks Tackle imbalanced data and optimize your algorithm for a bias or variance tradeoff Apply supervised and unsupervised algorithms to overcome various machine learning challenges Employ best practices for tuning your algorithm's hyper parameters Discover how to use neural networks for classification and regression Build, evaluate, and deploy your machine learning solutions to production Who this book is for This book is for data scientists, machine learning practitioners, and anyone who wants to learn how machine learning algorithms work and to build different machine learning models using the Python ecosystem. The book will help you take your knowledge of machine learning to the next level by grasping its ins and outs and tailoring it to your needs. Working knowledge of Python and a basic understanding of underlying mathematical and statistical concepts is required.

Mathematics for Machine Learning

Understanding Machine Learning