Math Behind Machine Learning

MCDOWELL LANG

Mathematical Foundations of Infinite-Dimensional Statistical Models Springer
Machine Learning: An Applied Mathematics Introduction covers the essential mathematics behind all of the following topics - K Nearest Neighbours; K Means Clustering; Naïve Bayes Classifier; Regression Methods; Support Vector Machines; Self-Organizing Maps; Decision Trees; Neural Networks; Reinforcement Learning
Deep Learning Pragmatic Bookshelf
Just like electricity, Machine Learning will revolutionize our life in many ways – some of which are not even conceivable today. This book provides a thorough conceptual understanding of Machine Learning techniques and algorithms. Many of the mathematical concepts are explained in an intuitive manner. The book starts with an overview of machine learning and the underlying Mathematical and Statistical concepts before moving onto machine learning topics. It gradually builds up the depth, covering many of the present day machine learning algorithms, ending in Deep Learning and Reinforcement Learning algorithms. The book also covers some of the popular Machine Learning applications. The material in this book is agnostic to any specific programming language or hardware so that readers can try these concepts on whichever platforms they are already familiar with. Offers a comprehensive introduction to Machine Learning, while not assuming any prior knowledge of the topic; Provides a complete overview of available techniques and algorithms in conceptual terms, covering various application domains of machine learning; Not tied to any specific software language or hardware implementation.
Linear Algebra and Learning from Data Machine Learning Mastery
Develop intelligent machine learning systems with SparkAbout This Book*Get to the grips with the latest version of Apache Spark*Utilize Spark's machine learning library to implement predictive analytics*Leverage Spark's powerful tools to load, analyze, clean, and transform your dataWho This Book Is ForIf you have a basic knowledge of machine learning and want to implement various machine-learning concepts in the context of Spark ML, this book is for you. You should be well versed with the Scala and Python languages.What You Will Learn*Get hands-on with the latest version of Spark ML*Create your first Spark program with Scala and Python*Set up and configure a development environment for Spark on your own computer, as well as on Amazon EC2*Access public machine learning datasets and use Spark to load, process, clean, and transform data*Use Spark's machine learning library to implement programs by utilizing well-known machine learning models*Deal with large-scale text data, including feature extraction and using text data as input to your machine learning models*Write Spark functions to evaluate the performance of your machine learning modelsIn DetailSpark ML is the machine learning module of Spark. It uses in-memory RDDs to process machine learning models faster for clustering, classification, and regression.This book will teach you about popular machine learning algorithms and their implementation. You will learn how various machine learning concepts are implemented in the context of Spark ML. You will start by installing Spark in a single and multinode cluster. Next you'll see how to execute Scala and Python based programs for Spark ML. Then we will take a few datasets and go deeper into clustering, classification, and regression. Toward the end, we will also cover text processing using Spark ML.Once you have learned the concepts, they can be applied to implement algorithms in either green-field implementations or to migrate existing systems to this new platform. You can migrate from Mahout or Scikit to use Spark ML.

Probability Samuel Hack
Introduces machine learning and its algorithmic paradigms, explaining the principles behind automated learning approaches and the considerations underlying their usage.
Introduction to Probability Bloomsbury Publishing
Preface -- Combinatorics -- Probability -- Expectation values -- Distributions -- Gaussian approximations -- Correlation and regression -- Appendices.
Linear Algebra Done Right MIT Press
From machine learning and data science to engineering and finance, linear algebra is an important prerequisite for the careers of today and of the future. There aren't many resources out there that give simple detailed examples and that walk you through the topics step by step. Many resources out there are either too dry or too difficult. This book aims to teach linear algebra step-by-step with examples that are simple but concrete.
Linear Algebra and Optimization for Machine Learning MIT Press
Summary Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Deep learning, a branch of artificial intelligence, teaches computers to learn by using neural networks, technology inspired by the human brain. Online text translation, self-driving cars, personalized product recommendations, and virtual voice assistants are just a few of the exciting modern advancements possible thanks to deep learning. About the Book Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Using only Python and its math-supporting library, NumPy, you'll train your own neural networks to see and understand images, translate text into different languages, and even write like Shakespeare! When you're done, you'll be fully prepared to move on to mastering deep learning frameworks. What's inside The science behind deep learning Building and training your own neural networks Privacy concepts, including federated learning Tips for continuing your pursuit of deep learning About the Reader For readers with high school-level math and intermediate programming skills. About the Author Andrew Trask is a PhD student at Oxford University and a research scientist at DeepMind. Previously, Andrew was a researcher and analytics product manager at Digital Reasoning, where he trained the world's largest artificial neural network and helped guide the analytics roadmap for the Synthesys cognitive computing platform. Table of Contents Introducing deep learning: why you should learn it Fundamental concepts: how do machines learn? Introduction to neural prediction: forward propagation Introduction to neural learning: gradient descent Learning multiple weights at a time: generalizing gradient descent Building your first deep neural network: introduction to backpropagation How to picture neural networks: in your head and on paper Learning signal and ignoring noise:introduction to regularization and batching Modeling probabilities and nonlinearities: activation functions Neural learning about edges and corners: intro to convolutional neural networks Neural networks that understand language: king - man + woman == ? Neural networks that write like Shakespeare: recurrent layers for variable-length data Introducing automatic optimization: let's build a deep learning framework Learning to write like Shakespeare: long short-term memory Deep learning on unseen data: introducing federated learning Where to go from here: a brief guide
Math for Machine Learning Linear Algebra for Beginners
Math for Deep Learning provides the essential math you need to understand deep learning discussions, explore more complex implementations, and better use the deep learning toolkits. With Math for Deep Learning, you'll learn the essential mathematics used by and as a background for deep learning. You'll work through Python examples to learn key deep learning related topics in probability, statistics, linear algebra, differential calculus, and matrix calculus as well as how to implement data flow in a neural network, backpropagation, and gradient descent. You'll also use Python to work through the mathematics that underlies those algorithms and even build a fully-functional neural network. In addition you'll find coverage of gradient descent including variations commonly used by the deep learning community: SGD, Adam, RMSprop, and Adagrad/Adadelta.
Machine Learning for Hackers Springer Nature
An intuitive, yet precise introduction to probability theory, stochastic processes, statistical inference, and probabilistic models used in science, engineering, economics, and related fields. This is the currently used textbook for an introductory probability course at the Massachusetts Institute of Technology, attended by a large number of undergraduate and graduate students, and for a leading online class on the subject. The book covers the fundamentals of probability theory (probabilistic models, discrete and continuous random variables, multiple random variables, and limit theorems), which are typically part of a first course on the subject. It also contains a number of more advanced topics, including transforms, sums of random variables, a fairly detailed introduction to Bernoulli, Poisson, and Markov processes, Bayesian inference, and an introduction to classical statistics. The book strikes a balance between simplicity in exposition and sophistication in analytical reasoning. Some of the more mathematically rigorous analysis is explained intuitively in the main text, and then developed in detail (at the level of advanced calculus) in the numerous solved theoretical problems.
No Starch Press
Based on the author's experience in teaching data science for more than 10 years, Mathematics and Programming for Machine Learning with R: From the Ground Up reveals how machine learning algorithms do their magic and explains how these algorithms can be implemented in code. It is designed to provide readers with an understanding of the reasoning behind machine learning algorithms as well as how to program them. Written for novice programmers, the book progresses step-by-step, providing the coding skills needed to implement machine learning algorithms in R. The book begins with simple implementations and fundamental concepts of logic, sets, and probability before moving to the coverage of powerful deep learning algorithms. The first eight chapters deal with probability-based machine learning algorithms, and the last eight chapters deal with machine learning based on artificial neural networks. The first half of the book does not require mathematical sophistication, although familiarity with probability and statistics would be helpful. The second half assumes the reader is familiar with at least one semester of calculus. The text guides novice R programmers through algorithms and their application and along the way; the reader gains programming confidence in tackling advanced R programming challenges. Highlights of the book include: More than 400 exercises A strong emphasis on improving programming skills and guiding beginners to the implementation of full-fledged algorithms Coverage of fundamental computer and mathematical concepts including logic, sets, and probability In-depth explanations of machine learning algorithms
Machine Learning Mathematics Createspace Independent Publishing Platform
This practical guide provides nearly 200 self-contained recipes to help you solve machine learning challenges you may encounter in your daily work. If you're comfortable with Python and its libraries, including pandas and scikit-learn, you'll be able to address specific problems such as loading data, handling text or numerical data, model selection, and dimensionality reduction and many other topics. Each recipe includes code that you can copy and paste into a toy dataset to ensure that it actually works. From there, you can insert, combine, or adapt the code to help construct your application. Recipes also include a discussion that explains the solution and provides meaningful context. This cookbook takes you beyond theory and concepts by providing the nuts and bolts you need to construct working machine learning applications. You'll find recipes for: Vectors, matrices, and arrays Handling numerical and categorical data, text, images, and dates and times Dimensionality reduction using feature extraction or feature selection Model evaluation and selection Linear and logical regression, trees and forests, and k-nearest neighbors Support vector machines (SVM), naïve Bayes, clustering, and neural networks Saving and loading trained models
Math and Architectures of Deep Learning Cambridge University Press
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.
Data Science and Machine Learning Cambridge University Press
TODAY ONLY 55% OFF for Bookstores! Are you an aspiring entrepreneur? Are you an amateur software developer looking for a break in the world of machine learning? Then this is the book for you. Machine learning is the way of the future - and breaking into this highly lucrative and ever-evolving field is a great way for your career, or business, to prosper. Inside this guide, you'll find simple, easy-to-follow explanations of the fundamental concepts behind machine learning, from the mathematical and statistical concepts to the programming behind them. With a wide range of comprehensive advice including machine learning models, neural networks, statistics, and much more, this guide is a highly effective tool for mastering this incredible technology. Inside, you will: Learn the Fundamental Concepts of Machine Learning Algorithms, and Their Impact in Resolving Modern Day Business Problems Understand The Four Fundamental Types of Machine Learning Algorithm Master the Concept of "Statistical Learning", a Descriptive Statistics-Based Machine Learning Algorithm Dive into the Development and Application of Six of the Most Popular Supervised and Unsupervised Machine Learning Algorithms, With Details on Linear Regression, Logistic Regression And More Learn Everything You Need to Know about Neural Networks and Data Pipelines Master the Concept of "General Setting of Learning", a Fundamental of Machine Learning Development Overview The Basics, Importance, and Applications of Data Science With Details on the "Team Data Science Process" Lifecycle And a Free Bonus! Covering everything you need to know about machine learning, now you can master the mathematics and statistics behind this field and develop your very own neural networks! Whether you want to use machine learning to help your business, or you're a programmer looking to expand your skills, this book is a must-read for anyone interested in the world of machine learning. Buy Now to Discover How You Can Master Machine Learning Today!
Grokking Machine Learning Springer Science & Business Media
Covering all the main approaches in state-of-the-art machine learning research, this will set a new standard as an introductory textbook.
Linear Algebra for Beginners: Open Doors to Great Careers Mathematics for Machine Learning
Math and Architectures of Deep Learning sets out the foundations of DL usefully and accessibly to working practitioners. Math and Architectures of Deep Learning bridges the gap between theory and practice, laying out the math of deep learning side by side with practical implementations in Python and PyTorch. You'll peer inside the "black box" to understand how your code is working, and learn to comprehend cutting-edge research you can turn into practical applications. Math and Architectures of Deep Learning sets out the foundations of DL usefully and accessibly to working practitioners. Each chapter explores a new fundamental DL concept or architectural pattern, explaining the underpinning mathematics and demonstrating how they work in practice with well-annotated Python code. You'll start with a primer of basic algebra, calculus, and statistics, working your way up to state-of-the-art DL paradigms taken from the latest research. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.
Mathematics and Programming for Machine Learning with R Manning
This book explains the math behind machine learning using simple but concrete examples. This book will get you started in machine learning in a smooth and natural way, preparing you for more advanced topics and dispelling the belief that machine learning is complicated, difficult, and intimidating.
Machine Learning with Python Cookbook "O'Reilly Media, Inc."
Linear algebra and the foundations of deep learning, together at last! From Professor Gilbert Strang, acclaimed author of Introduction to Linear Algebra, comes Linear Algebra and Learning from Data, the first textbook that teaches linear algebra together with deep learning and neural nets. This readable yet rigorous textbook contains a complete course in the linear algebra and related mathematics that students need to know to get to grips with learning from data. Included are: the four fundamental subspaces, singular value decompositions, special matrices, large matrix computation techniques, compressed sensing, probability and statistics, optimization, the architecture of neural nets, stochastic gradient descent and backpropagation.
Linear Algebra "O'Reilly Media, Inc."
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
Deep Learning for Coders with fastai and PyTorch Simon and Schuster
Grokking Machine Learning presents machine learning algorithms and techniques in a way that anyone can understand. This book skips the confused academic jargon and offers clear explanations that require only basic algebra. As you go, you'll build interesting projects with Python, including models for spam detection and image recognition. You'll also pick up practical skills for cleaning and preparing data.
Probabilistic Graphical Models MIT Press
Distills key concepts from linear algebra, geometry, matrices, calculus, optimization, probability and statistics that are used in machine learning.