

# Deep Learning Basics Github Pages

Thank you very much for downloading **Deep Learning Basics Github Pages** . Maybe you have knowledge that, people have look hundreds times for their favorite readings like this Deep Learning Basics Github Pages , but end up in infectious downloads.

Rather than reading a good book with a cup of coffee in the afternoon, instead they are facing with some malicious bugs inside their laptop.

Deep Learning Basics Github Pages is available in our book collection an online access to it is set as public so you can get it instantly.

Our digital library spans in multiple locations, allowing you to get the most less latency time to download any of our books like this one.

Kindly say, the Deep Learning Basics Github Pages is universally compatible with any devices to read

**Computational Genomics with R** - Altuna Akalin 2020-12-16  
Computational Genomics with R provides a starting point for beginners in genomic data analysis and also guides more advanced practitioners to sophisticated data analysis techniques in genomics. The book covers topics from R programming, to machine learning and statistics, to the latest genomic data analysis techniques. The text provides accessible information and explanations, always with the genomics context in the background. This also contains practical and well-documented examples in R so readers can analyze their data by simply reusing the code presented. As the field of computational genomics is interdisciplinary, it requires different starting points for people with different backgrounds. For example, a biologist might skip sections on basic genome biology and start with R programming, whereas a computer scientist might want to start with genome biology. After

reading: You will have the basics of R and be able to dive right into specialized uses of R for computational genomics such as using Bioconductor packages. You will be familiar with statistics, supervised and unsupervised learning techniques that are important in data modeling, and exploratory analysis of high-dimensional data. You will understand genomic intervals and operations on them that are used for tasks such as aligned read counting and genomic feature annotation. You will know the basics of processing and quality checking high-throughput sequencing data. You will be able to do sequence analysis, such as calculating GC content for parts of a genome or finding transcription factor binding sites. You will know about visualization techniques used in genomics, such as heatmaps, meta-gene plots, and genomic track visualization. You will be familiar with analysis of different high-throughput sequencing data sets, such

as RNA-seq, ChIP-seq, and BS-seq. You will know basic techniques for integrating and interpreting multi-omics datasets. Altuna Akalin is a group leader and head of the Bioinformatics and Omics Data Science Platform at the Berlin Institute of Medical Systems Biology, Max Delbrück Center, Berlin. He has been developing computational methods for analyzing and integrating large-scale genomics data sets since 2002. He has published an extensive body of work in this area. The framework for this book grew out of the yearly computational genomics courses he has been organizing and teaching since 2015.

R for Data Science - Hadley Wickham 2016-12-12

Learn how to use R to turn raw data into insight, knowledge, and understanding. This book introduces you to R, RStudio, and the tidyverse, a collection of R packages designed to work together to make data science fast, fluent, and fun. Suitable for readers with no previous programming experience, R for Data Science is designed to get you doing data science as quickly as possible. Authors Hadley Wickham and Garrett Grolemund guide you through the steps of importing, wrangling, exploring, and modeling your data and communicating the results. You'll get a complete, big-picture understanding of the data science cycle, along with basic tools you need to manage the details. Each section of the book is paired with exercises to help you practice what you've learned along the way. You'll learn how to: Wrangle—transform your datasets into a form convenient for analysis Program—learn powerful R tools for solving data problems with greater clarity and ease Explore—examine your data, generate hypotheses, and quickly test them Model—provide a low-dimensional summary that captures true "signals" in your dataset

Communicate—learn R Markdown for integrating prose, code, and results

**Deep Learning from the Basics** - Koki Saitoh 2021-03-08

Discover ways to implement various deep learning algorithms by leveraging Python and other technologies Key Features Learn deep learning models through several activities Begin with simple machine learning problems, and finish by building a complex system of your own Teach your machines to see by mastering the technologies required for image recognition Book Description Deep learning is rapidly becoming the most preferred way of solving data problems. This is thanks, in part, to its huge variety of mathematical algorithms and their ability to find patterns that are otherwise invisible to us. Deep Learning from the Basics begins with a fast-paced introduction to deep learning with Python, its definition, characteristics, and applications. You'll learn how to use the Python interpreter and the script files in your applications, and utilize NumPy and Matplotlib in your deep learning models. As you progress through the book, you'll discover backpropagation—an efficient way to calculate the gradients of weight parameters—and study multilayer perceptrons and their limitations, before, finally, implementing a three-layer neural network and calculating multidimensional arrays. By the end of the book, you'll have the knowledge to apply the relevant technologies in deep learning. What you will learn Use Python with minimum external sources to implement deep learning programs Study the various deep learning and neural network theories Learn how to determine learning coefficients and the initial values of weights Implement trends such as Batch Normalization, Dropout, and Adam Explore applications like automatic driving, image generation, and reinforcement learning Who

this book is for Deep Learning from the Basics is designed for data scientists, data analysts, and developers who want to use deep learning techniques to develop efficient solutions. This book is ideal for those who want a deeper understanding as well as an overview of the technologies. Some working knowledge of Python is a must. Knowledge of NumPy and pandas will be beneficial, but not essential.

**Artificial Intelligence** - David L. Poole 2017-09-25  
Artificial Intelligence presents a practical guide to AI, including agents, machine learning and problem-solving simple and complex domains.

*Deep Learning* - Andrew Glassner 2021-06-22  
A richly-illustrated, full-color introduction to deep learning that offers visual and conceptual explanations instead of equations. You'll learn how to use key deep learning algorithms without the need for complex math. Ever since computers began beating us at chess, they've been getting better at a wide range of human activities, from writing songs and generating news articles to helping doctors provide healthcare. Deep learning is the source of many of these breakthroughs, and its remarkable ability to find patterns hiding in data has made it the fastest growing field in artificial intelligence (AI). Digital assistants on our phones use deep learning to understand and respond intelligently to voice commands; automotive systems use it to safely navigate road hazards; online platforms use it to deliver personalized suggestions for movies and books - the possibilities are endless. *Deep Learning: A Visual Approach* is for anyone who wants to understand this fascinating field in depth, but without any of the advanced math and programming usually required to grasp its internals. If you want to know how these tools work,

and use them yourself, the answers are all within these pages. And, if you're ready to write your own programs, there are also plenty of supplemental Python notebooks in the accompanying Github repository to get you going. The book's conversational style, extensive color illustrations, illuminating analogies, and real-world examples expertly explain the key concepts in deep learning, including:

- How text generators create novel stories and articles
- How deep learning systems learn to play and win at human games
- How image classification systems identify objects or people in a photo
- How to think about probabilities in a way that's useful to everyday life
- How to use the machine learning techniques that form the core of modern AI

Intellectual adventurers of all kinds can use the powerful ideas covered in *Deep Learning: A Visual Approach* to build intelligent systems that help us better understand the world and everyone who lives in it. It's the future of AI, and this book allows you to fully envision it. Full Color Illustrations

[Grokking Deep Learning](#) - Andrew W. Trask 2019-01-23  
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

**Deep Learning with Python** - Francois Chollet 2017-11-30 Summary Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning has made remarkable progress in recent years. We went from near-unusable speech and image recognition, to near-human accuracy. We went from machines that couldn't beat a serious Go player, to defeating a world champion. Behind this progress is deep learning—a combination of engineering advances, best practices, and theory that enables a wealth of previously impossible smart applications. About the Book Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. You'll explore challenging concepts and practice with applications in computer vision, natural-language processing, and generative models. By the time you

finish, you'll have the knowledge and hands-on skills to apply deep learning in your own projects. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image-classification models Deep learning for text and sequences Neural style transfer, text generation, and image generation About the Reader Readers need intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the Author François Chollet works on deep learning at Google in Mountain View, CA. He is the creator of the Keras deep-learning library, as well as a contributor to the TensorFlow machine-learning framework. He also does deep-learning research, with a focus on computer vision and the application of machine learning to formal reasoning. His papers have been published at major conferences in the field, including the Conference on Computer Vision and Pattern Recognition (CVPR), the Conference and Workshop on Neural Information Processing Systems (NIPS), the International Conference on Learning Representations (ICLR), and others. Table of Contents PART 1 - FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of neural networks Getting started with neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE Deep learning for computer vision Deep learning for text and sequences Advanced deep-learning best practices Generative deep learning Conclusions appendix A - Installing Keras and its dependencies on Ubuntu appendix B - Running Jupyter notebooks on an EC2 GPU instance Dive Into Deep Learning - Joanne Quinn 2019-07-15 The leading experts in system change and learning, with their school-based partners around the world, have

created this essential companion to their runaway best-seller, *Deep Learning: Engage the World Change the World*. This hands-on guide provides a roadmap for building capacity in teachers, schools, districts, and systems to design deep learning, measure progress, and assess conditions needed to activate and sustain innovation. *Dive Into Deep Learning: Tools for Engagement* is rich with resources educators need to construct and drive meaningful deep learning experiences in order to develop the kind of mindset and know-how that is crucial to becoming a problem-solving change agent in our global society. Designed in full color, this easy-to-use guide is loaded with tools, tips, protocols, and real-world examples. It includes:

- A framework for deep learning that provides a pathway to develop the six global competencies needed to flourish in a complex world – character, citizenship, collaboration, communication, creativity, and critical thinking.
- Learning progressions to help educators analyze student work and measure progress.
- Learning design rubrics, templates and examples for incorporating the four elements of learning design: learning partnerships, pedagogical practices, learning environments, and leveraging digital.
- Conditions rubrics, teacher self-assessment tools, and planning guides to help educators build, mobilize, and sustain deep learning in schools and districts. Learn about, improve, and expand your world of learning. Put the joy back into learning for students and adults alike. Dive into deep learning to create learning experiences that give purpose, unleash student potential, and transform not only learning, but life itself.

Probabilistic Machine Learning - Kevin P. Murphy 2022-03-01

A detailed and up-to-date introduction to machine learning, presented through the unifying lens of probabilistic modeling and Bayesian decision theory. This book offers a detailed and up-to-date introduction to machine learning (including deep learning) through the unifying lens of probabilistic modeling and Bayesian decision theory. The book covers mathematical background (including linear algebra and optimization), basic supervised learning (including linear and logistic regression and deep neural networks), as well as more advanced topics (including transfer learning and unsupervised learning). End-of-chapter exercises allow students to apply what they have learned, and an appendix covers notation. Probabilistic Machine Learning grew out of the author's 2012 book, *Machine Learning: A Probabilistic Perspective*. More than just a simple update, this is a completely new book that reflects the dramatic developments in the field since 2012, most notably deep learning. In addition, the new book is accompanied by online Python code, using libraries such as scikit-learn, JAX, PyTorch, and Tensorflow, which can be used to reproduce nearly all the figures; this code can be run inside a web browser using cloud-based notebooks, and provides a practical complement to the theoretical topics discussed in the book. This introductory text will be followed by a sequel that covers more advanced topics, taking the same probabilistic approach.

*Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow* - Aurélien Géron 2019-09-05

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

**Deep Learning Illustrated** - Jon Krohn 2019-08-05

"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.

**Foundations of Machine Learning, second edition** -

Mehryar Mohri 2018-12-25

A new edition of a graduate-level machine learning

textbook that focuses on the analysis and theory of algorithms. This book is a general introduction to machine learning that can serve as a textbook for graduate students and a reference for researchers. It covers fundamental modern topics in machine learning while providing the theoretical basis and conceptual tools needed for the discussion and justification of algorithms. It also describes several key aspects of the application of these algorithms. The authors aim to present novel theoretical tools and concepts while giving concise proofs even for relatively advanced topics. Foundations of Machine Learning is unique in its focus on the analysis and theory of algorithms. The first four chapters lay the theoretical foundation for what follows; subsequent chapters are mostly self-contained. Topics covered include the Probably Approximately Correct (PAC) learning framework; generalization bounds based on Rademacher complexity and VC-dimension; Support Vector Machines (SVMs); kernel methods; boosting; on-line learning; multi-class classification; ranking; regression; algorithmic stability; dimensionality reduction; learning automata and languages; and reinforcement learning. Each chapter ends with a set of exercises. Appendixes provide additional material including concise probability review. This second edition offers three new chapters, on model selection, maximum entropy models, and conditional entropy models. New material in the appendixes includes a major section on Fenchel duality, expanded coverage of concentration inequalities, and an entirely new entry on information theory. More than half of the exercises are new to this edition.

**Deep Learning with JavaScript** - Stanley Bileschi

2020-01-24

Summary Deep learning has transformed the fields of computer vision, image processing, and natural language applications. Thanks to TensorFlow.js, now JavaScript developers can build deep learning apps without relying on Python or R. Deep Learning with JavaScript shows developers how they can bring DL technology to the web. Written by the main authors of the TensorFlow library, this new book provides fascinating use cases and in-depth instruction for deep learning apps in JavaScript in your browser or on Node. Foreword by Nikhil Thorat and Daniel Smilkov. About the technology Running deep learning applications in the browser or on Node-based backends opens up exciting possibilities for smart web applications. With the TensorFlow.js library, you build and train deep learning models with JavaScript. Offering uncompromising production-quality scalability, modularity, and responsiveness, TensorFlow.js really shines for its portability. Its models run anywhere JavaScript runs, pushing ML farther up the application stack. About the book In Deep Learning with JavaScript, you'll learn to use TensorFlow.js to build deep learning models that run directly in the browser. This fast-paced book, written by Google engineers, is practical, engaging, and easy to follow. Through diverse examples featuring text analysis, speech processing, image recognition, and self-learning game AI, you'll master all the basics of deep learning and explore advanced concepts, like retraining existing models for transfer learning and image generation. What's inside - Image and language processing in the browser - Tuning ML models with client-side data - Text and image creation with generative deep learning - Source code samples to test and modify About the reader For JavaScript programmers interested in deep learning. About the author Shanging

Cai, Stanley Bileschi and Eric D. Nielsen are software engineers with experience on the Google Brain team, and were crucial to the development of the high-level API of TensorFlow.js. This book is based in part on the classic, Deep Learning with Python by François Chollet. TOC: PART 1 - MOTIVATION AND BASIC CONCEPTS 1 • Deep learning and JavaScript PART 2 - A GENTLE INTRODUCTION TO TENSORFLOW.JS 2 • Getting started: Simple linear regression in TensorFlow.js 3 • Adding nonlinearity: Beyond weighted sums 4 • Recognizing images and sounds using convnets 5 • Transfer learning: Reusing pretrained neural networks PART 3 - ADVANCED DEEP LEARNING WITH TENSORFLOW.JS 6 • Working with data 7 • Visualizing data and models 8 • Underfitting, overfitting, and the universal workflow of machine learning 9 • Deep learning for sequences and text 10 • Generative deep learning 11 • Basics of deep reinforcement learning PART 4 - SUMMARY AND CLOSING WORDS 12 • Testing, optimizing, and deploying models 13 • Summary, conclusions, and beyond **Deep Learning with PyTorch** - Luca Pietro Giovanni Antiga 2020-07-01

"We finally have the definitive treatise on PyTorch! It covers the basics and abstractions in great detail. I hope this book becomes your extended reference document." –Soumith Chintala, co-creator of PyTorch Key Features Written by PyTorch's creator and key contributors Develop deep learning models in a familiar Pythonic way Use PyTorch to build an image classifier for cancer detection Diagnose problems with your neural network and improve training with data augmentation Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About The Book Every other day we hear about new ways to put deep learning to good use: improved medical imaging,



accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands. Instantly familiar to anyone who knows Python data tools like NumPy and Scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It's great for building quick models, and it scales smoothly from laptop to enterprise. Deep Learning with PyTorch teaches you to create deep learning and neural network systems with PyTorch. This practical book gets you to work right away building a tumor image classifier from scratch. After covering the basics, you'll learn best practices for the entire deep learning pipeline, tackling advanced projects as your PyTorch skills become more sophisticated. All code samples are easy to explore in downloadable Jupyter notebooks. What You Will Learn Understanding deep learning data structures such as tensors and neural networks Best practices for the PyTorch Tensor API, loading data in Python, and visualizing results Implementing modules and loss functions Utilizing pretrained models from PyTorch Hub Methods for training networks with limited inputs Sifting through unreliable results to diagnose and fix problems in your neural network Improve your results with augmented data, better model architecture, and fine tuning This Book Is Written For For Python programmers with an interest in machine learning. No experience with PyTorch or other deep learning frameworks is required. About The Authors Eli Stevens has worked in Silicon Valley for the past 15 years as a software engineer, and the past 7 years as Chief Technical Officer of a startup making medical device software. Luca Antiga is co-founder and CEO of an AI engineering company located in Bergamo, Italy, and a regular contributor to PyTorch. Thomas Viehmann is a Machine Learning and PyTorch

speciality trainer and consultant based in Munich, Germany and a PyTorch core developer. Table of Contents PART 1 - CORE PYTORCH 1 Introducing deep learning and the PyTorch Library 2 Pretrained networks 3 It starts with a tensor 4 Real-world data representation using tensors 5 The mechanics of learning 6 Using a neural network to fit the data 7 Telling birds from airplanes: Learning from images 8 Using convolutions to generalize PART 2 - LEARNING FROM IMAGES IN THE REAL WORLD: EARLY DETECTION OF LUNG CANCER 9 Using PyTorch to fight cancer 10 Combining data sources into a unified dataset 11 Training a classification model to detect suspected tumors 12 Improving training with metrics and augmentation 13 Using segmentation to find suspected nodules 14 End-to-end nodule analysis, and where to go next PART 3 - DEPLOYMENT 15 Deploying to production **Introduction to Deep Learning** - Eugene Charniak 2019-01-29

A project-based guide to the basics of deep learning. This concise, project-driven guide to deep learning takes readers through a series of program-writing tasks that introduce them to the use of deep learning in such areas of artificial intelligence as computer vision, natural-language processing, and reinforcement learning. The author, a longtime artificial intelligence researcher specializing in natural-language processing, covers feed-forward neural nets, convolutional neural nets, word embeddings, recurrent neural nets, sequence-to-sequence learning, deep reinforcement learning, unsupervised models, and other fundamental concepts and techniques. Students and practitioners learn the basics of deep learning by working through programs in Tensorflow, an open-source machine learning framework. "I find I learn computer science material best by

sitting down and writing programs," the author writes, and the book reflects this approach. Each chapter includes a programming project, exercises, and references for further reading. An early chapter is devoted to Tensorflow and its interface with Python, the widely used programming language. Familiarity with linear algebra, multivariate calculus, and probability and statistics is required, as is a rudimentary knowledge of programming in Python. The book can be used in both undergraduate and graduate courses; practitioners will find it an essential reference.

IPython Interactive Computing and Visualization Cookbook  
- Cyrille Rossant 2014-09-25

Intended to anyone interested in numerical computing and data science: students, researchers, teachers, engineers, analysts, hobbyists... Basic knowledge of Python/NumPy is recommended. Some skills in mathematics will help you understand the theory behind the computational methods.

**Neural Networks and Deep Learning** - Charu C. Aggarwal  
2018-08-25

This book covers both classical and modern models in deep learning. The primary focus is on the theory and algorithms of deep learning. The theory and algorithms of neural networks are particularly important for understanding important concepts, so that one can understand the important design concepts of neural architectures in different applications. Why do neural networks work? When do they work better than off-the-shelf machine-learning models? When is depth useful? Why is training neural networks so hard? What are the pitfalls? The book is also rich in discussing different applications in order to give the practitioner a flavor of how neural architectures are designed for different

types of problems. Applications associated with many different areas like recommender systems, machine translation, image captioning, image classification, reinforcement-learning based gaming, and text analytics are covered. The chapters of this book span three categories: The basics of neural networks: Many traditional machine learning models can be understood as special cases of neural networks. An emphasis is placed in the first two chapters on understanding the relationship between traditional machine learning and neural networks. Support vector machines, linear/logistic regression, singular value decomposition, matrix factorization, and recommender systems are shown to be special cases of neural networks. These methods are studied together with recent feature engineering methods like word2vec. Fundamentals of neural networks: A detailed discussion of training and regularization is provided in Chapters 3 and 4. Chapters 5 and 6 present radial-basis function (RBF) networks and restricted Boltzmann machines. Advanced topics in neural networks: Chapters 7 and 8 discuss recurrent neural networks and convolutional neural networks. Several advanced topics like deep reinforcement learning, neural Turing machines, Kohonen self-organizing maps, and generative adversarial networks are introduced in Chapters 9 and 10. The book is written for graduate students, researchers, and practitioners. Numerous exercises are available along with a solution manual to aid in classroom teaching. Where possible, an application-centric view is highlighted in order to provide an understanding of the practical uses of each class of techniques.

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

Whether you're a software engineer aspiring to enter the world of deep learning, a veteran data scientist, or a hobbyist with a simple dream of making the next viral AI app, you might have wondered where to begin. This step-by-step guide teaches you how to build practical deep learning applications for the cloud, mobile, browsers, and edge devices using a hands-on approach. Relying on years of industry experience transforming deep learning research into award-winning applications, Anirudh Koul, Siddha Ganju, and Meher Kasam guide you through the process of converting an idea into something that people in the real world can use. Train, tune, and deploy computer vision models with Keras, TensorFlow, Core ML, and TensorFlow Lite Develop AI for a range of devices including Raspberry Pi, Jetson Nano, and Google Coral Explore fun projects, from Silicon Valley's Not Hotdog app to 40+ industry case studies Simulate an autonomous car in a video game environment and build a miniature version with reinforcement learning Use transfer learning to train models in minutes Discover 50+ practical tips for maximizing model accuracy and speed, debugging, and scaling to millions of users

Hands-On Machine Learning with R - Brad Boehmke  
2019-11-07

Hands-on Machine Learning with R provides a practical and applied approach to learning and developing intuition into today's most popular machine learning methods. This book serves as a practitioner's guide to the machine learning process and is meant to help the reader learn to apply the machine learning stack within R, which includes using various R packages such as glmnet, h2o, ranger, xgboost, keras, and others to effectively model and gain insight from their data. The book favors a hands-on approach, providing an intuitive

understanding of machine learning concepts through concrete examples and just a little bit of theory. Throughout this book, the reader will be exposed to the entire machine learning process including feature engineering, resampling, hyperparameter tuning, model evaluation, and interpretation. The reader will be exposed to powerful algorithms such as regularized regression, random forests, gradient boosting machines, deep learning, generalized low rank models, and more! By favoring a hands-on approach and using real world data, the reader will gain an intuitive understanding of the architectures and engines that drive these algorithms and packages, understand when and how to tune the various hyperparameters, and be able to interpret model results. By the end of this book, the reader should have a firm grasp of R's machine learning stack and be able to implement a systematic approach for producing high quality modeling results. Features:

- Offers a practical and applied introduction to the most popular machine learning methods.
- Topics covered include feature engineering, resampling, deep learning and more.
- Uses a hands-on approach and real world data.

Foundations of Deep Reinforcement Learning - Laura Graesser  
2019-11-20

The Contemporary Introduction to Deep Reinforcement Learning that Combines Theory and Practice Deep reinforcement learning (deep RL) combines deep learning and reinforcement learning, in which artificial agents learn to solve sequential decision-making problems. In the past decade deep RL has achieved remarkable results on a range of problems, from single and multiplayer games—such as Go, Atari games, and DotA 2—to robotics. Foundations of Deep Reinforcement Learning is an introduction to deep RL that uniquely combines both

theory and implementation. It starts with intuition, then carefully explains the theory of deep RL algorithms, discusses implementations in its companion software library SLM Lab, and finishes with the practical details of getting deep RL to work. This guide is ideal for both computer science students and software engineers who are familiar with basic machine learning concepts and have a working understanding of Python. Understand each key aspect of a deep RL problem Explore policy- and value-based algorithms, including REINFORCE, SARSA, DQN, Double DQN, and Prioritized Experience Replay (PER) Delve into combined algorithms, including Actor-Critic and Proximal Policy Optimization (PPO) Understand how algorithms can be parallelized synchronously and asynchronously Run algorithms in SLM Lab and learn the practical implementation details for getting deep RL to work Explore algorithm benchmark results with tuned hyperparameters Understand how deep RL environments are designed Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

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

Master the essential skills needed to recognize and solve complex problems with machine learning and deep learning. Using real-world examples that leverage the popular Python machine learning ecosystem, this book is your perfect companion for learning the art and science of machine learning to become a successful practitioner. The concepts, techniques, tools, frameworks, and methodologies used in this book will teach you how to think, design, build, and execute machine learning systems and projects successfully. Practical Machine

Learning with Python follows a structured and comprehensive three-tiered approach packed with hands-on examples and code. Part 1 focuses on understanding machine learning concepts and tools. This includes machine learning basics with a broad overview of algorithms, techniques, concepts and applications, followed by a tour of the entire Python machine learning ecosystem. Brief guides for useful machine learning tools, libraries and frameworks are also covered. Part 2 details standard machine learning pipelines, with an emphasis on data processing analysis, feature engineering, and modeling. You will learn how to process, wrangle, summarize and visualize data in its various forms. Feature engineering and selection methodologies will be covered in detail with real-world datasets followed by model building, tuning, interpretation and deployment. Part 3 explores multiple real-world case studies spanning diverse domains and industries like retail, transportation, movies, music, marketing, computer vision and finance. For each case study, you will learn the application of various machine learning techniques and methods. The hands-on examples will help you become familiar with state-of-the-art machine learning tools and techniques and understand what algorithms are best suited for any problem. Practical Machine Learning with Python will empower you to start solving your own problems with machine learning today! What You'll Learn Execute end-to-end machine learning projects and systems Implement hands-on examples with industry standard, open source, robust machine learning tools and frameworks Review case studies depicting applications of machine learning and deep learning on diverse domains and industries Apply a wide range of machine learning models including

regression, classification, and clustering. Understand and apply the latest models and methodologies from deep learning including CNNs, RNNs, LSTMs and transfer learning. Who This Book Is For IT professionals, analysts, developers, data scientists, engineers, graduate students

**Machine Learning Refined** - Jeremy Watt 2016-09-08

Providing a unique approach to machine learning, this text contains fresh and intuitive, yet rigorous, descriptions of all fundamental concepts necessary to conduct research, build products, tinker, and play. By prioritizing geometric intuition, algorithmic thinking, and practical real world applications in disciplines including computer vision, natural language processing, economics, neuroscience, recommender systems, physics, and biology, this text provides readers with both a lucid understanding of foundational material as well as the practical tools needed to solve real-world problems. With in-depth Python and MATLAB/OCTAVE-based computational exercises and a complete treatment of cutting edge numerical optimization techniques, this is an essential resource for students and an ideal reference for researchers and practitioners working in machine learning, computer science, electrical engineering, signal processing, and numerical optimization.

Machine Learning with PyTorch and Scikit-Learn -

Sebastian Raschka 2022-02-25

This book of the bestselling and widely acclaimed Python Machine Learning series is a comprehensive guide to machine and deep learning using PyTorch's simple to code framework. Purchase of the print or Kindle book includes a free eBook in PDF format. Key Features Learn applied machine learning with a solid foundation in theory

Clear, intuitive explanations take you deep into the theory and practice of Python machine learning Fully updated and expanded to cover PyTorch, transformers, XGBoost, graph neural networks, and best practices Book Description Machine Learning with PyTorch and Scikit-Learn is a comprehensive guide to machine learning and deep learning with PyTorch. It acts as both a step-by-step tutorial and a reference you'll keep coming back to as you build your machine learning systems. Packed with clear explanations, visualizations, and examples, the book covers all the essential machine learning techniques in depth. While some books teach you only to follow instructions, with this machine learning book, we teach the principles allowing you to build models and applications for yourself. Why PyTorch? PyTorch is the Pythonic way to learn machine learning, making it easier to learn and simpler to code with. This book explains the essential parts of PyTorch and how to create models using popular libraries, such as PyTorch Lightning and PyTorch Geometric. You will also learn about generative adversarial networks (GANs) for generating new data and training intelligent agents with reinforcement learning. Finally, this new edition is expanded to cover the latest trends in deep learning, including graph neural networks and large-scale transformers used for natural language processing (NLP). This PyTorch book is your companion to machine learning with Python, whether you're a Python developer new to machine learning or want to deepen your knowledge of the latest developments. What you will learn Explore frameworks, models, and techniques for machines to 'learn' from data Use scikit-learn for machine learning and PyTorch for deep learning Train machine learning classifiers on images, text, and more Build and train neural networks,

transformers, and boosting algorithms Discover best practices for evaluating and tuning models Predict continuous target outcomes using regression analysis Dig deeper into textual and social media data using sentiment analysis Who this book is for If you have a good grasp of Python basics and want to start learning about machine learning and deep learning, then this is the book for you. This is an essential resource written for developers and data scientists who want to create practical machine learning and deep learning applications using scikit-learn and PyTorch. Before you get started with this book, you'll need a good understanding of calculus, as well as linear algebra.

**Python Data Science Handbook** - Jake VanderPlas  
2016-11-21

For many researchers, Python is a first-class tool mainly because of its libraries for storing, manipulating, and gaining insight from data. Several resources exist for individual pieces of this data science stack, but only with the Python Data Science Handbook do you get them all—IPython, NumPy, Pandas, Matplotlib, Scikit-Learn, and other related tools. Working scientists and data crunchers familiar with reading and writing Python code will find this comprehensive desk reference ideal for tackling day-to-day issues: manipulating, transforming, and cleaning data; visualizing different types of data; and using data to build statistical or machine learning models. Quite simply, this is the must-have reference for scientific computing in Python. With this handbook, you'll learn how to use: IPython and Jupyter: provide computational environments for data scientists using Python NumPy: includes the ndarray for efficient storage and manipulation of dense data arrays in Python Pandas:

features the DataFrame for efficient storage and manipulation of labeled/columnar data in Python Matplotlib: includes capabilities for a flexible range of data visualizations in Python Scikit-Learn: for efficient and clean Python implementations of the most important and established machine learning algorithms

*Practical Machine Learning with Rust* - Joydeep Bhattacharjee 2019-12-10

Explore machine learning in Rust and learn about the intricacies of creating machine learning applications. This book begins by covering the important concepts of machine learning such as supervised, unsupervised, and reinforcement learning, and the basics of Rust. Further, you'll dive into the more specific fields of machine learning, such as computer vision and natural language processing, and look at the Rust libraries that help create applications for those domains. We will also look at how to deploy these applications either on site or over the cloud. After reading Practical Machine Learning with Rust, you will have a solid understanding of creating high computation libraries using Rust. Armed with the knowledge of this amazing language, you will be able to create applications that are more performant, memory safe, and less resource heavy. What You Will Learn Write machine learning algorithms in Rust Use Rust libraries for different tasks in machine learning Create concise Rust packages for your machine learning applications Implement NLP and computer vision in Rust Deploy your code in the cloud and on bare metal servers Who This Book Is For Machine learning engineers and software engineers interested in building machine learning applications in Rust.

**Deep Learning For Physics Research** - Martin Erdmann  
2021-06-25

A core principle of physics is knowledge gained from data. Thus, deep learning has instantly entered physics and may become a new paradigm in basic and applied research. This textbook addresses physics students and physicists who want to understand what deep learning actually means, and what is the potential for their own scientific projects. Being familiar with linear algebra and parameter optimization is sufficient to jump-start deep learning. Adopting a pragmatic approach, basic and advanced applications in physics research are described. Also offered are simple hands-on exercises for implementing deep networks for which python code and training data can be downloaded.

*Deep Learning and the Game of Go* - Kevin Ferguson  
2019-01-06

Summary Deep Learning and the Game of Go teaches you how to apply the power of deep learning to complex reasoning tasks by building a Go-playing AI. After exposing you to the foundations of machine and deep learning, you'll use Python to build a bot and then teach it the rules of the game. Foreword by Thore Graepel, DeepMind Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology The ancient strategy game of Go is an incredible case study for AI. In 2016, a deep learning-based system shocked the Go world by defeating a world champion. Shortly after that, the upgraded AlphaGo Zero crushed the original bot by using deep reinforcement learning to master the game. Now, you can learn those same deep learning techniques by building your own Go bot! About the Book Deep Learning and the Game of Go introduces deep learning by teaching you to build a Go-winning bot. As you progress, you'll apply increasingly complex training techniques and strategies using the

Python deep learning library Keras. You'll enjoy watching your bot master the game of Go, and along the way, you'll discover how to apply your new deep learning skills to a wide range of other scenarios! What's inside Build and teach a self-improving game AI Enhance classical game AI systems with deep learning Implement neural networks for deep learning About the Reader All you need are basic Python skills and high school-level math. No deep learning experience required. About the Author Max Pumperla and Kevin Ferguson are experienced deep learning specialists skilled in distributed systems and data science. Together, Max and Kevin built the open source bot BetaGo. Table of Contents PART 1 - FOUNDATIONS Toward deep learning: a machine-learning introduction Go as a machine-learning problem Implementing your first Go bot PART 2 - MACHINE LEARNING AND GAME AI Playing games with tree search Getting started with neural networks Designing a neural network for Go data Learning from data: a deep-learning bot Deploying bots in the wild Learning by practice: reinforcement learning Reinforcement learning with policy gradients Reinforcement learning with value methods Reinforcement learning with actor-critic methods PART 3 - GREATER THAN THE SUM OF ITS PARTS AlphaGo: Bringing it all together AlphaGo Zero: Integrating tree search with reinforcement learning Interpretable Machine Learning - Christoph Molnar 2020 This book is about making machine learning models and their decisions interpretable. After exploring the concepts of interpretability, you will learn about simple, interpretable models such as decision trees, decision rules and linear regression. Later chapters focus on general model-agnostic methods for interpreting black box models like feature importance and accumulated

local effects and explaining individual predictions with Shapley values and LIME. All interpretation methods are explained in depth and discussed critically. How do they work under the hood? What are their strengths and weaknesses? How can their outputs be interpreted? This book will enable you to select and correctly apply the interpretation method that is most suitable for your machine learning project.

High-Dimensional Data Analysis with Low-Dimensional Models - John Wright 2022-01-13

Connecting theory with practice, this systematic and rigorous introduction covers the fundamental principles, algorithms and applications of key mathematical models for high-dimensional data analysis. Comprehensive in its approach, it provides unified coverage of many different low-dimensional models and analytical techniques, including sparse and low-rank models, and both convex and non-convex formulations. Readers will learn how to develop efficient and scalable algorithms for solving real-world problems, supported by numerous examples and exercises throughout, and how to use the computational tools learnt in several application contexts.

Applications presented include scientific imaging, communication, face recognition, 3D vision, and deep networks for classification. With code available online, this is an ideal textbook for senior and graduate students in computer science, data science, and electrical engineering, as well as for those taking courses on sparsity, low-dimensional structures, and high-dimensional data. Foreword by Emmanuel Candès.

**Deep Learning** - Ian Goodfellow 2016-11-10

An introduction to a broad range of topics in deep learning, covering mathematical and conceptual background, deep learning techniques used in industry,

and research perspectives. "Written by three experts in the field, Deep Learning is the only comprehensive book on the subject." –Elon Musk, cochair of OpenAI; cofounder and CEO of Tesla and SpaceX Deep learning is a form of machine learning that enables computers to learn from experience and understand the world in terms of a hierarchy of concepts. Because the computer gathers knowledge from experience, there is no need for a human computer operator to formally specify all the knowledge that the computer needs. The hierarchy of concepts allows the computer to learn complicated concepts by building them out of simpler ones; a graph of these hierarchies would be many layers deep. This book introduces a broad range of topics in deep learning. The text offers mathematical and conceptual background, covering relevant concepts in linear algebra, probability theory and information theory, numerical computation, and machine learning. It describes deep learning techniques used by practitioners in industry, including deep feedforward networks, regularization, optimization algorithms, convolutional networks, sequence modeling, and practical methodology; and it surveys such applications as natural language processing, speech recognition, computer vision, online recommendation systems, bioinformatics, and videogames. Finally, the book offers research perspectives, covering such theoretical topics as linear factor models, autoencoders, representation learning, structured probabilistic models, Monte Carlo methods, the partition function, approximate inference, and deep generative models. Deep Learning can be used by undergraduate or graduate students planning careers in either industry or research, and by software engineers who want to begin using deep learning in their products or platforms. A



website offers supplementary material for both readers and instructors.

Machine Learning in Action - Peter Harrington 2012-04-03

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 Deep Learning for Coders with fastai and PyTorch - Jeremy Howard 2020-06-29

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

**Text Analytics with Python** - Dipanjan Sarkar 2016-11-30  
Derive useful insights from your data using Python. You will learn both basic and advanced concepts, including text and language syntax, structure, and semantics. You will focus on algorithms and techniques, such as text classification, clustering, topic modeling, and text summarization. Text Analytics with Python teaches you the techniques related to natural language processing and text analytics, and you will gain the skills to know which technique is best suited to solve a particular problem. You will look at each technique and algorithm with both a bird's eye view to understand how it can be used as well as with a microscopic view to understand the mathematical concepts and to implement them to solve your own problems. What You Will Learn: Understand the major concepts and techniques of natural language processing (NLP) and text analytics, including syntax and structure Build a text classification system to categorize news articles, analyze app or game reviews using topic modeling and text summarization, and cluster popular movie synopses and analyze the sentiment of movie reviews Implement Python and popular open source libraries in NLP and text analytics, such as the natural language toolkit (nltk), gensim, scikit-learn, spaCy and Pattern Who This Book Is For : IT professionals, analysts, developers, linguistic experts, data scientists, and anyone with a keen interest in linguistics, analytics, and generating insights from textual data

**Machine Learning** - Kevin P. Murphy 2012-08-24

A comprehensive introduction to machine learning that uses probabilistic models and inference as a unifying approach. Today's Web-enabled deluge of electronic data calls for automated methods of data analysis. Machine learning provides these, developing methods that can automatically detect patterns in data and then use the uncovered patterns to predict future data. This textbook offers a comprehensive and self-contained introduction to the field of machine learning, based on a unified, probabilistic approach. The coverage combines breadth and depth, offering necessary background material on such topics as probability, optimization, and linear algebra as well as discussion of recent developments in the field, including conditional random fields, L1 regularization, and deep learning. The book is written in an informal, accessible style, complete with pseudo-code for the most important algorithms. All topics are copiously illustrated with color images and worked examples drawn from such application domains as biology, text processing, computer vision, and robotics. Rather than providing a cookbook of different heuristic methods, the book stresses a principled model-based approach, often using the language of graphical models to specify models in a concise and intuitive way. Almost all the models described have been implemented in a MATLAB software package—PMTK (probabilistic modeling toolkit)—that is freely available online. The book is suitable for upper-level undergraduates with an introductory-level college math background and beginning graduate students.

*The Principles of Deep Learning Theory* - Daniel A. Roberts 2022-05-26

This volume develops an effective theory approach to understanding deep neural networks of practical

relevance.

**Graph Neural Networks: Foundations, Frontiers, and Applications** - Lingfei Wu 2022-01-03

Deep Learning models are at the core of artificial intelligence research today. It is well known that deep learning techniques are disruptive for Euclidean data, such as images or sequence data, and not immediately applicable to graph-structured data such as text. This gap has driven a wave of research for deep learning on graphs, including graph representation learning, graph generation, and graph classification. The new neural network architectures on graph-structured data (graph neural networks, GNNs in short) have performed remarkably on these tasks, demonstrated by applications in social networks, bioinformatics, and medical informatics. Despite these successes, GNNs still face many challenges ranging from the foundational methodologies to the theoretical understandings of the power of the graph representation learning. This book provides a comprehensive introduction of GNNs. It first discusses the goals of graph representation learning and then reviews the history, current developments, and future directions of GNNs. The second part presents and reviews fundamental methods and theories concerning GNNs while the third part describes various frontiers that are built on the GNNs. The book concludes with an overview of recent developments in a number of applications using GNNs. This book is suitable for a wide audience including undergraduate and graduate students, postdoctoral researchers, professors and lecturers, as well as industrial and government practitioners who are new to this area or who already have some basic background but want to learn more about advanced and promising techniques and applications.

**Mathematics for Machine Learning** - Marc Peter Deisenroth 2020-04-23

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.

**Introduction to Data Science** - Rafael A. Irizarry 2019-11-20

Introduction to Data Science: Data Analysis and Prediction Algorithms with R introduces concepts and skills that can help you tackle real-world data analysis challenges. It covers concepts from probability, statistical inference, linear regression, and machine learning. It also helps you develop skills such as R programming, data wrangling, data visualization, predictive algorithm building, file organization with

UNIX/Linux shell, version control with Git and GitHub, and reproducible document preparation. This book is a textbook for a first course in data science. No previous knowledge of R is necessary, although some experience with programming may be helpful. The book is divided into six parts: R, data visualization, statistics with R, data wrangling, machine learning, and productivity tools. Each part has several chapters meant to be presented as one lecture. The author uses motivating case studies that realistically mimic a data scientist's experience. He starts by asking specific questions and answers these through data analysis so concepts are learned as a means to answering the questions. Examples of the case studies included are: US murder rates by state, self-reported student heights, trends in world health and economics, the impact of vaccines on infectious disease rates, the financial crisis of 2007-2008, election forecasting, building a baseball team, image processing of hand-written digits, and movie recommendation systems. The statistical concepts used to answer the case study questions are only briefly introduced, so complementing with a probability and statistics textbook is highly recommended for in-depth understanding of these concepts. If you read and understand the chapters and complete the exercises, you will be prepared to learn the more advanced concepts and skills needed to become an expert.

*Introduction to Machine Learning with Python* - Andreas C. Müller 2016-09-26

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

*Deep Learning through Sparse and Low-Rank Modeling* - Zhangyang Wang 2019-04-26

Deep Learning through Sparse Representation and Low-Rank Modeling bridges classical sparse and low rank models- those that emphasize problem-specific Interpretability- with recent deep network models that have enabled a larger learning capacity and better utilization of Big Data. It shows how the toolkit of deep learning is closely tied with the sparse/low rank methods and algorithms, providing a rich variety of theoretical and analytic tools to guide the design and interpretation of deep learning models. The development of the theory and models is supported by a wide variety of applications in computer vision, machine learning, signal processing,

and data mining. This book will be highly useful for researchers, graduate students and practitioners working in the fields of computer vision, machine learning, signal processing, optimization and statistics. Combines classical sparse and low-rank models and algorithms with the latest advances in deep learning networks Shows how

the structure and algorithms of sparse and low-rank methods improves the performance and interpretability of Deep Learning models Provides tactics on how to build and apply customized deep learning models for various applications