

---

# Gradient Boosting Machine Learning Mastery

---

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

Better Deep Learning

Classification and Regression Trees

Practical Machine Learning with Python

Machine Learning Mastery With Python

Methods and Applications

Mastering Machine Learning for Penetration Testing

Principles and Techniques for Data Scientists

An Introduction to Statistical Learning

Deep Learning with Python

Gradient Boosted Trees with XGBoost and scikit-learn

Train Faster, Reduce Overfitting, and Make Better Predictions

Ensemble Methods

Pattern Classification Using Ensemble Methods

Approaching (Almost) Any Machine Learning Problem

Data Preparation for Machine Learning

Machine Learning For Dummies

Machine Learning with R

From Coding to Deep Learning

Mastering Machine Learning Algorithms

Predict the Future with MLPs, CNNs and LSTMs in Python

Context, Methods and Applications

Intermittent Demand Forecasting

Foundations, Algorithms, and Applications

with Applications in R

Make Better Predictions with Bagging, Boosting, and Stacking

Hands-On Gradient Boosting with XGBoost and scikit-learn  
Mastering Machine Learning Algorithms  
Interpretable Machine Learning  
A Practical Implementation Guide to Predictive Data Analytics Using Python  
Understanding Machine Learning  
Discover How They Work and Implement Them From Scratch  
Programming Machine Learning  
The Elements of Statistical Learning  
Machine Learning Mastery With R  
Data Mining, Inference, and Prediction  
Ensemble Learning Algorithms With Python  
A Problem-Solver's Guide to Building Real-World Intelligent Systems  
The Surprising Truth About What Motivates Us  
Handbook of Research on Machine Learning Applications and Trends: Algorithms, Methods, and Techniques

*Gradient Boosting Machine Learning  
Mastery*

Downloaded from [business.itu.edu.tr](https://business.itu.edu.tr)  
by guest

---

## CHASE ALEXIS

---

### **Understand Your Data, Create Accurate Models, and Work Projects End-to-End** Abhishek Thakur

Become a master at penetration testing using machine learning with Python Key Features Identify ambiguities and breach intelligent security systems Perform unique cyber attacks to breach robust systems Learn to leverage machine learning algorithms Book Description Cyber security is crucial for both businesses and individuals. As systems are getting smarter, we now see machine learning interrupting computer security. With the adoption of machine learning in upcoming security products,

it's important for pentesters and security researchers to understand how these systems work, and to breach them for testing purposes. This book begins with the basics of machine learning and the algorithms used to build robust systems. Once you've gained a fair understanding of how security products leverage machine learning, you'll dive into the core concepts of breaching such systems. Through practical use cases, you'll see how to find loopholes and surpass a self-learning security system. As you make your way through the chapters, you'll focus on topics such as network intrusion detection and AV and IDS evasion. We'll also cover the best practices when identifying ambiguities, and extensive techniques to breach an intelligent system. By the end of this book, you will be well-versed with identifying loopholes in a self-learning security system and will be

able to efficiently breach a machine learning system. What you will learn Take an in-depth look at machine learning Get to know natural language processing (NLP) Understand malware feature engineering Build generative adversarial networks using Python libraries Work on threat hunting with machine learning and the ELK stack Explore the best practices for machine learning Who this book is for This book is for pen testers and security professionals who are interested in learning techniques to break an intelligent security system. Basic knowledge of Python is needed, but no prior knowledge of machine learning is necessary.

[Better Deep Learning](#) Machine Learning Mastery

Imbalanced classification are those classification tasks where the distribution of examples across the classes is not equal. Cut through the equations, Greek letters, and confusion, and discover the specialized techniques data preparation techniques, learning algorithms, and performance metrics that you need to know. Using clear explanations, standard Python libraries, and step-by-step tutorial lessons, you will discover how to confidently develop robust models for your own imbalanced classification projects.

[Classification and Regression Trees](#) Morgan & Claypool Publishers

INTERMITTENT DEMAND FORECASTING The first text to focus on the methods and approaches of intermittent, rather than fast, demand forecasting Intermittent Demand Forecasting is for anyone who is interested in improving forecasts of intermittent demand products, and enhancing the management of inventories. Whether you are a practitioner, at the sharp end of demand planning, a software designer, a student, an academic teaching operational research or operations management courses, or a researcher in this field, we hope that the book will

inspire you to rethink demand forecasting. If you do so, then you can contribute towards significant economic and environmental benefits. No prior knowledge of intermittent demand forecasting or inventory management is assumed in this book. The key formulae are accompanied by worked examples to show how they can be implemented in practice. For those wishing to understand the theory in more depth, technical notes are provided at the end of each chapter, as well as an extensive and up-to-date collection of references for further study. Software developments are reviewed, to give an appreciation of the current state of the art in commercial and open source software.

“Intermittent demand forecasting may seem like a specialized area but actually is at the center of sustainability efforts to consume less and to waste less. Boylan and Syntetos have done a superb job in showing how improvements in inventory management are pivotal in achieving this. Their book covers both the theory and practice of intermittent demand forecasting and my prediction is that it will fast become the bible of the field.”

—Spyros Makridakis, Professor, University of Nicosia, and Director, Institute for the Future and the Makridakis Open Forecasting Center (MOFC). “We have been able to support our clients by adopting many of the ideas discussed in this excellent book, and implementing them in our software. I am sure that these ideas will be equally helpful for other supply chain software vendors and for companies wanting to update and upgrade their capabilities in forecasting and inventory management.”

—Suresh Acharya, VP, Research and Development, Blue Yonder. “As product variants proliferate and the pace of business quickens, more and more items have intermittent demand. Boylan and

Syntetos have long been leaders in extending forecasting and inventory methods to accommodate this new reality. Their book gathers and clarifies decades of research in this area, and explains how practitioners can exploit this knowledge to make their operations more efficient and effective.” —Thomas R. Willemain, Professor Emeritus, Rensselaer Polytechnic Institute. [Practical Machine Learning with Python](#) Cambridge University Press

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. [Machine Learning Mastery With Python](#) Packt Publishing Ltd 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.

*Methods and Applications* John Wiley & Sons

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

*Mastering Machine Learning for Penetration Testing* Apress

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.

**Principles and Techniques for Data Scientists** Springer  
Science & Business Media

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

**An Introduction to Statistical Learning** Machine Learning Mastery

The first book of its kind to review the current status and future direction of the exciting new branch of machine learning/data mining called imbalanced learning. Imbalanced learning focuses on how an intelligent system can learn when it is provided with imbalanced data. Solving imbalanced learning problems is critical in numerous data-intensive networked systems, including surveillance, security, Internet, finance, biomedical, defense, and more. Due to the inherent complex characteristics of imbalanced data sets, learning from such data requires new understandings, principles, algorithms, and tools to transform vast amounts of raw data efficiently into information and knowledge representation. The first comprehensive look at this new branch of machine learning, this book offers a critical review of the problem of imbalanced learning, covering the state of the art in techniques, principles, and real-world applications. Featuring contributions from experts in both academia and industry, *Imbalanced Learning: Foundations, Algorithms, and Applications* provides chapter coverage on: Foundations of Imbalanced Learning, Imbalanced Datasets: From Sampling to Classifiers, Ensemble Methods for Class Imbalance, Learning Class Imbalance, Learning Methods for Support Vector Machines, Class Imbalance and Active Learning, Nonstationary Stream Data Learning with Imbalanced Class Distribution, Assessment Metrics for Imbalanced Learning, Imbalanced Learning: Foundations, Algorithms,

and Applications will help scientists and engineers learn how to tackle the problem of learning from imbalanced datasets, and gain insight into current developments in the field as well as future research directions.

*Deep Learning with Python* Packt Publishing Ltd

It is common wisdom that gathering a variety of views and inputs improves the process of decision making, and, indeed, underpins a democratic society. Dubbed “ensemble learning” by researchers in computational intelligence and machine learning, it is known to improve a decision system’s robustness and accuracy. Now, fresh developments are allowing researchers to unleash the power of ensemble learning in an increasing range of real-world applications. Ensemble learning algorithms such as “boosting” and “random forest” facilitate solutions to key computational issues such as face recognition and are now being applied in areas as diverse as object tracking and bioinformatics. Responding to a shortage of literature dedicated to the topic, this volume offers comprehensive coverage of state-of-the-art ensemble learning techniques, including the random forest skeleton tracking algorithm in the Xbox Kinect sensor, which bypasses the need for game controllers. At once a solid theoretical study and a practical guide, the volume is a windfall for researchers and practitioners alike.

*Gradient Boosted Trees with XGBoost and scikit-learn* Machine Learning Mastery

During the past decade there has been an explosion in computation and information technology. With it have come vast amounts of data in a variety of fields such as medicine, biology, finance, and marketing. The challenge of understanding these

data has led to the development of new tools in the field of statistics, and spawned new areas such as data mining, machine learning, and bioinformatics. Many of these tools have common underpinnings but are often expressed with different terminology. This book describes the important ideas in these areas in a common conceptual framework. While the approach is statistical, the emphasis is on concepts rather than mathematics. Many examples are given, with a liberal use of color graphics. It should be a valuable resource for statisticians and anyone interested in data mining in science or industry. The book's coverage is broad, from supervised learning (prediction) to unsupervised learning. The many topics include neural networks, support vector machines, classification trees and boosting---the first comprehensive treatment of this topic in any book. This major new edition features many topics not covered in the original, including graphical models, random forests, ensemble methods, least angle regression & path algorithms for the lasso, non-negative matrix factorization, and spectral clustering. There is also a chapter on methods for "wide" data ( $p$  bigger than  $n$ ), including multiple testing and false discovery rates. Trevor Hastie, Robert Tibshirani, and Jerome Friedman are professors of statistics at Stanford University. They are prominent researchers in this area: Hastie and Tibshirani developed generalized additive models and wrote a popular book of that title. Hastie co-developed much of the statistical modeling software and environment in R/S-PLUS and invented principal curves and surfaces. Tibshirani proposed the lasso and is co-author of the very successful *An Introduction to the Bootstrap*. Friedman is the co-inventor of many data-mining tools including CART, MARS,

projection pursuit and gradient boosting.

*Train Faster, Reduce Overfitting, and Make Better Predictions*

John Wiley & Sons

Ensemble methods have been called the most influential development in Data Mining and Machine Learning in the past decade. They combine multiple models into one usually more accurate than the best of its components. Ensembles can provide a critical boost to industrial challenges -- from investment timing to drug discovery, and fraud detection to recommendation systems -- where predictive accuracy is more vital than model interpretability. Ensembles are useful with all modeling algorithms, but this book focuses on decision trees to explain them most clearly. After describing trees and their strengths and weaknesses, the authors provide an overview of regularization -- today understood to be a key reason for the superior performance of modern ensembling algorithms. The book continues with a clear description of two recent developments: Importance Sampling (IS) and Rule Ensembles (RE). IS reveals classic ensemble methods -- bagging, random forests, and boosting -- to be special cases of a single algorithm, thereby showing how to improve their accuracy and speed. REs are linear rule models derived from decision tree ensembles. They are the most interpretable version of ensembles, which is essential to applications such as credit scoring and fault diagnosis. Lastly, the authors explain the paradox of how ensembles achieve greater accuracy on new data despite their (apparently much greater) complexity. This book is aimed at novice and advanced analytic researchers and practitioners -- especially in Engineering, Statistics, and Computer Science. Those with little exposure to



ensembles will learn why and how to employ this breakthrough method, and advanced practitioners will gain insight into building even more powerful models. Throughout, snippets of code in R are provided to illustrate the algorithms described and to encourage the reader to try the techniques. The authors are industry experts in data mining and machine learning who are also adjunct professors and popular speakers. Although early pioneers in discovering and using ensembles, they here distill and clarify the recent groundbreaking work of leading academics (such as Jerome Friedman) to bring the benefits of ensembles to practitioners. Table of Contents: Ensembles Discovered / Predictive Learning and Decision Trees / Model Complexity, Model Selection and Regularization / Importance Sampling and the Classic Ensemble Methods / Rule Ensembles and Interpretation Statistics / Ensemble Complexity

### **Ensemble Methods** John Wiley & Sons

Written as a tutorial to explore and understand the power of R for machine learning. This practical guide that covers all of the need to know topics in a very systematic way. For each machine learning approach, each step in the process is detailed, from preparing the data for analysis to evaluating the results. These steps will build the knowledge you need to apply them to your own data science tasks. Intended for those who want to learn how to use R's machine learning capabilities and gain insight from your data. Perhaps you already know a bit about machine learning, but have never used R; or perhaps you know a little R but are new to machine learning. In either case, this book will get you up and running quickly. It would be helpful to have a bit of familiarity with basic programming concepts, but no prior

experience is required.

### *Pattern Classification Using Ensemble Methods* IGI Global

One of Mark Cuban's top reads for better understanding A.I.

(inc.com, 2021) Your comprehensive entry-level guide to machine learning

While machine learning expertise doesn't quite mean you can create your own Turing Test-proof android—as in the movie *Ex Machina*—it is a form of artificial intelligence and one of the most exciting technological means of identifying opportunities and solving problems fast and on a large scale.

Anyone who masters the principles of machine learning is mastering a big part of our tech future and opening up incredible new directions in careers that include fraud detection, optimizing search results, serving real-time ads, credit-scoring, building accurate and sophisticated pricing models—and way, way more.

Unlike most machine learning books, the fully updated 2nd Edition of *Machine Learning For Dummies* doesn't assume you have years of experience using programming languages such as Python (R source is also included in a downloadable form with comments and explanations), but lets you in on the ground floor, covering the entry-level materials that will get you up and running building models you need to perform practical tasks. It takes a look at the underlying—and fascinating—math principles that power machine learning but also shows that you don't need to be a math whiz to build fun new tools and apply them to your work and study. Understand the history of AI and machine learning Work with Python 3.8 and TensorFlow 2.x (and R as a download) Build and test your own models Use the latest datasets, rather than the worn out data found in other books Apply machine learning to real problems Whether you want to



learn for college or to enhance your business or career performance, this friendly beginner's guide is your best introduction to machine learning, allowing you to become quickly confident using this amazing and fast-developing technology that's impacting lives for the better all over the world.

### **Approaching (Almost) Any Machine Learning Problem**

Packt Publishing Ltd

An Introduction to Statistical Learning provides an accessible overview of the field of statistical learning, an essential toolset for making sense of the vast and complex data sets that have emerged in fields ranging from biology to finance to marketing to astrophysics in the past twenty years. This book presents some of the most important modeling and prediction techniques, along with relevant applications. Topics include linear regression, classification, resampling methods, shrinkage approaches, tree-based methods, support vector machines, clustering, and more. Color graphics and real-world examples are used to illustrate the methods presented. Since the goal of this textbook is to facilitate the use of these statistical learning techniques by practitioners in science, industry, and other fields, each chapter contains a tutorial on implementing the analyses and methods presented in R, an extremely popular open source statistical software platform. Two of the authors co-wrote *The Elements of Statistical Learning* (Hastie, Tibshirani and Friedman, 2nd edition 2009), a popular reference book for statistics and machine learning researchers. *An Introduction to Statistical Learning* covers many of the same topics, but at a level accessible to a much broader audience. This book is targeted at statisticians and non-statisticians alike who wish to use cutting-edge statistical

learning techniques to analyze their data. The text assumes only a previous course in linear regression and no knowledge of matrix algebra.

*Data Preparation for Machine Learning* Packt Publishing Ltd

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.

*Machine Learning For Dummies* "O'Reilly Media, Inc."

Feature engineering is a crucial step in the machine-learning pipeline, yet this topic is rarely examined on its own. With this practical book, you'll learn techniques for extracting and transforming features—the numeric representations of raw data—into formats for machine-learning models. Each chapter guides you through a single data problem, such as how to represent text or image data. Together, these examples illustrate the main principles of feature engineering. Rather than simply teach these principles, authors Alice Zheng and Amanda Casari focus on practical application with exercises throughout the book. The closing chapter brings everything together by tackling a real-world, structured dataset with several feature-engineering

techniques. Python packages including numpy, Pandas, Scikit-learn, and Matplotlib are used in code examples. You'll examine: Feature engineering for numeric data: filtering, binning, scaling, log transforms, and power transforms Natural text techniques: bag-of-words, n-grams, and phrase detection Frequency-based filtering and feature scaling for eliminating uninformative features Encoding techniques of categorical variables, including feature hashing and bin-counting Model-based feature engineering with principal component analysis The concept of model stacking, using k-means as a featurization technique Image feature extraction with manual and deep-learning techniques

*Machine Learning with R* Packt Publishing Ltd

This is not a traditional book. The book has a lot of code. If you don't like the code first approach do not buy this book. Making code available on Github is not an option. This book is for people who have some theoretical knowledge of machine learning and deep learning and want to dive into applied machine learning. The book doesn't explain the algorithms but is more oriented towards how and what should you use to solve machine learning and deep learning problems. The book is not for you if you are looking for pure basics. The book is for you if you are looking for guidance on approaching machine learning problems. The book is best enjoyed with a cup of coffee and a laptop/workstation where you can code along. Table of contents: - Setting up your working environment - Supervised vs unsupervised learning - Cross-validation - Evaluation metrics - Arranging machine learning projects - Approaching categorical variables - Feature engineering - Feature selection - Hyperparameter optimization - Approaching

image classification & segmentation - Approaching text classification/regression - Approaching ensembling and stacking - Approaching reproducible code & model serving There are no sub-headings. Important terms are written in bold. I will be answering all your queries related to the book and will be making YouTube tutorials to cover what has not been discussed in the book. To ask questions/doubts, visit this link:

<https://bit.ly/aamlquestions> And Subscribe to my youtube channel: <https://bit.ly/abhitubesub>

*From Coding to Deep Learning* Cambridge University Press

"This book investigates machine learning (ML), one of the most fruitful fields of current research, both in the proposal of new techniques and theoretic algorithms and in their application to real-life problems"--Provided by publisher.

*Mastering Machine Learning Algorithms* Machine Learning Mastery

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

Best Sellers - Books :

- [Stone Maidens By Lloyd Devereux Richards](#)
- [Fahrenheit 451 By Ray Bradbury](#)
- [Twisted Lies \(twisted, 4\)](#)
- [Hunting Adeline \(cat And Mouse Duet\)](#)
- [Jackie: Public, Private, Secret](#)
- [A Court Of Mist And Fury \(a Court Of Thorns And Roses, 2\) By Sarah J. Maas](#)
- [My Butt Is So Christmassy! By Dawn Mcmillan](#)
- [Never Lie: An Addictive Psychological Thriller By Freida Mcfadden](#)
- [Love You Forever By Robert Munsch](#)
- [Tomorrow, And Tomorrow, And Tomorrow: A Novel By Gabrielle Zevin](#)