

## Deep Learning With Python

"The first edition of Deep Learning with Python is one of the best books on the subject. The second edition made it even better." - Todd Cook

The bestseller revised! Deep Learning with Python, Second Edition is a comprehensive introduction to the field of deep learning using Python and the powerful Keras library. Written by Google AI researcher François Chollet, the creator of Keras, this revised edition has been updated with new chapters, new tools, and cutting-edge techniques drawn from the latest research. You'll build your understanding through practical examples and intuitive explanations that make the complexities of deep learning accessible and understandable.

Machine learning has made remarkable progress in recent years. We've gone from near-unusable speech recognition, to near-human accuracy. From machines that couldn't beat a serious Go player, to defeating a world champion. Medical imaging diagnostics, weather forecasting, and natural language question answering have suddenly become tractable problems. Behind this progress is deep learning--a combination of engineering advances, best practices, and theory that enables a wealth of previously impossible smart applications across every industry sector.

about the book Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. You'll learn directly from the creator of Keras, François Chollet, building your understanding through intuitive explanations and practical examples. Updated from the original bestseller with over 50% new content, this second edition includes new chapters, cutting-edge innovations, and coverage of the very latest deep learning tools. 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 Image-classification, image segmentation, and object detection Deep learning for natural language processing Timeseries forecasting 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 AI research, with a focus on abstraction and 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.

Build real-world Artificial Intelligence applications with Python to intelligently interact with the world around you

About This Book Step into the amazing world of intelligent apps using this comprehensive guide Enter the world of Artificial Intelligence, explore it, and create your own applications Work through simple yet insightful examples that will get you up and running with Artificial Intelligence in no time Who This Book Is For This book is for Python developers who want to build real-world Artificial Intelligence applications. This book is friendly to Python beginners, but being familiar with Python would be useful to play around with the code. It will also be useful for experienced Python programmers who are looking to use Artificial Intelligence techniques in their existing technology stacks. What You Will Learn Realize different classification and regression techniques Understand the concept of clustering and how to use it to automatically segment data See how to build an intelligent recommender system Understand logic programming and how to use it Build automatic speech recognition systems Understand the basics of heuristic search and genetic programming Develop games using Artificial Intelligence Learn how

reinforcement learning works Discover how to build intelligent applications centered on images, text, and time series data See how to use deep learning algorithms and build applications based on it In Detail Artificial Intelligence is becoming increasingly relevant in the modern world where everything is driven by technology and data. It is used extensively across many fields such as search engines, image recognition, robotics, finance, and so on. We will explore various real-world scenarios in this book and you'll learn about various algorithms that can be used to build Artificial Intelligence applications. During the course of this book, you will find out how to make informed decisions about what algorithms to use in a given context. Starting from the basics of Artificial Intelligence, you will learn how to develop various building blocks using different data mining techniques. You will see how to implement different algorithms to get the best possible results, and will understand how to apply them to real-world scenarios. If you want to add an intelligence layer to any application that's based on images, text, stock market, or some other form of data, this exciting book on Artificial Intelligence will definitely be your guide! Style and approach This highly practical book will show you how to implement Artificial Intelligence. The book provides multiple examples enabling you to create smart applications to meet the needs of your organization. In every chapter, we explain an algorithm, implement it, and then build a smart application.

Summary Deep Learning with R introduces the world of deep learning using the powerful Keras library and its R language interface. The book builds your understanding of deep learning through intuitive explanations and practical examples. Continue your journey into the world of deep learning with Deep Learning with R in Motion, a practical, hands-on video course available exclusively at Manning.com ([www.manning.com/livevideo/deep-?learning-with-r-in-motion](http://www.manning.com/livevideo/deep-?learning-with-r-in-motion)). 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. Deep-learning systems now enable previously impossible smart applications, revolutionizing image recognition and natural-language processing, and identifying complex patterns in data. The Keras deep-learning library provides data scientists and developers working in R a state-of-the-art toolset for tackling deep-learning tasks. About the Book Deep Learning with R introduces the world of deep learning using the powerful Keras library and its R language interface. Initially written for Python as Deep Learning with Python by Keras creator and Google AI researcher François Chollet and adapted for R by RStudio founder J. J. Allaire, this book builds your understanding of deep learning through intuitive explanations and practical examples. You'll practice your new skills with R-based applications in computer vision, natural-language processing, and generative models. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image classification and generation Deep learning for text and sequences About the Reader You'll need intermediate R programming skills. No previous experience with machine learning or deep learning is assumed. About the Authors François Chollet is a deep-learning researcher at Google and the author of the Keras library. J.J. Allaire is the founder of RStudio and the author of the R interfaces to TensorFlow and Keras. 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

Take your machine learning skills to the next level by mastering Deep Learning concepts and algorithms using Python. About This Book Explore and create intelligent systems using cutting-edge deep learning techniques Implement deep learning algorithms and work with revolutionary libraries in Python Get real-world examples and easy-to-follow tutorials on Theano, TensorFlow, H2O and more Who This Book

## Get Free Deep Learning With Python

Is For This book is for Data Science practitioners as well as aspirants who have a basic foundational understanding of Machine Learning concepts and some programming experience with Python. A mathematical background with a conceptual understanding of calculus and statistics is also desired. What You Will Learn Get a practical deep dive into deep learning algorithms Explore deep learning further with Theano, Caffe, Keras, and TensorFlow Learn about two of the most powerful techniques at the core of many practical deep learning implementations: Auto-Encoders and Restricted Boltzmann Machines Dive into Deep Belief Nets and Deep Neural Networks Discover more deep learning algorithms with Dropout and Convolutional Neural Networks Get to know device strategies so you can use deep learning algorithms and libraries in the real world In Detail With an increasing interest in AI around the world, deep learning has attracted a great deal of public attention. Every day, deep learning algorithms are used broadly across different industries. The book will give you all the practical information available on the subject, including the best practices, using real-world use cases. You will learn to recognize and extract information to increase predictive accuracy and optimize results. Starting with a quick recap of important machine learning concepts, the book will delve straight into deep learning principles using Sci-kit learn. Moving ahead, you will learn to use the latest open source libraries such as Theano, Keras, Google's TensorFlow, and H2O. Use this guide to uncover the difficulties of pattern recognition, scaling data with greater accuracy and discussing deep learning algorithms and techniques. Whether you want to dive deeper into Deep Learning, or want to investigate how to get more out of this powerful technology, you'll find everything inside. Style and approach Python Machine Learning by example follows practical hands on approach. It walks you through the key elements of Python and its powerful machine learning libraries with the help of real world projects.

Learn, understand, and implement deep neural networks in a math- and programming-friendly approach using Keras and Python. The book focuses on an end-to-end approach to developing supervised learning algorithms in regression and classification with practical business-centric use-cases implemented in Keras. The overall book comprises three sections with two chapters in each section. The first section prepares you with all the necessary basics to get started in deep learning. Chapter 1 introduces you to the world of deep learning and its difference from machine learning, the choices of frameworks for deep learning, and the Keras ecosystem. You will cover a real-life business problem that can be solved by supervised learning algorithms with deep neural networks. You'll tackle one use case for regression and another for classification leveraging popular Kaggle datasets. Later, you will see an interesting and challenging part of deep learning: hyperparameter tuning; helping you further improve your models when building robust deep learning applications. Finally, you'll further hone your skills in deep learning and cover areas of active development and research in deep learning. At the end of Learn Keras for Deep Neural Networks, you will have a thorough understanding of deep learning principles and have practical hands-on experience in developing enterprise-grade deep learning solutions in Keras. What You'll Learn Master fast-paced practical deep learning concepts with math- and programming-friendly abstractions. Design, develop, train, validate, and deploy deep neural networks using the Keras framework Use best practices for debugging and validating deep learning models Deploy and integrate deep learning as a service into a larger software service or product Extend deep learning principles into other popular frameworks Who This Book Is For Software engineers and data engineers with basic programming skills in any language and who are keen on exploring deep learning for a career move or an enterprise project. Applied machine learning with a solid foundation in theory. Revised and expanded for TensorFlow 2, GANs, and reinforcement learning. Key Features Third edition of the bestselling, widely acclaimed Python machine learning book Clear and intuitive explanations take you deep into the theory and practice of Python machine learning Fully updated and expanded to cover TensorFlow 2, Generative Adversarial Network

models, reinforcement learning, and best practices Book Description Python Machine Learning, Third Edition is a comprehensive guide to machine learning and deep learning with Python. 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 working 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, Raschka and Mirjalili teach the principles behind machine learning, allowing you to build models and applications for yourself. Updated for TensorFlow 2.0, this new third edition introduces readers to its new Keras API features, as well as the latest additions to scikit-learn. It's also expanded to cover cutting-edge reinforcement learning techniques based on deep learning, as well as an introduction to GANs. Finally, this book also explores a subfield of natural language processing (NLP) called sentiment analysis, helping you learn how to use machine learning algorithms to classify documents. This 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 Master the frameworks, models, and techniques that enable machines to 'learn' from data Use scikit-learn for machine learning and TensorFlow for deep learning Apply machine learning to image classification, sentiment analysis, intelligent web applications, and more Build and train neural networks, GANs, and other models 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 know some Python and you want to use machine learning and deep learning, pick up this book. Whether you want to start from scratch or extend your machine learning knowledge, this is an essential resource. Written for developers and data scientists who want to create practical machine learning and deep learning code, this book is ideal for anyone who wants to teach computers how to learn from data.

This practical guide provides nearly 200 self-contained recipes to help you solve machine learning challenges you may encounter in your daily work. If you're comfortable with Python and its libraries, including pandas and scikit-learn, you'll be able to address specific problems such as loading data, handling text or numerical data, model selection, and dimensionality reduction and many other topics. Each recipe includes code that you can copy and paste into a toy dataset to ensure that it actually works. From there, you can insert, combine, or adapt the code to help construct your application. Recipes also include a discussion that explains the solution and provides meaningful context. This cookbook takes you beyond theory and concepts by providing the nuts and bolts you need to construct working machine learning applications. You'll find recipes for: Vectors, matrices, and arrays Handling numerical and categorical data, text, images, and dates and times Dimensionality reduction using feature extraction or feature selection Model evaluation and selection Linear and logical regression, trees and forests, and k-nearest neighbors Support vector machines (SVM), naïve Bayes, clustering, and neural networks Saving and loading trained models

Python Deep Learning Projects book will simplify and ease how deep learning works, and demonstrate how neural networks play a vital role in exploring predictive analytics across different domains. You will explore projects in the field of computational linguistics, computer vision, machine translation, pattern recognition and many more

Python makes machine learning easy for beginners and experienced developers With computing power increasing exponentially and costs decreasing at the same time, there is no better time to learn machine learning using Python. Machine learning tasks that once required enormous processing power are now possible on desktop machines. However, machine learning is not for the faint of heart—it requires a good foundation in statistics, as well as programming knowledge. Python Machine Learning will help coders of all levels master one of the most in-

demand programming skillsets in use today. Readers will get started by following fundamental topics such as an introduction to Machine Learning and Data Science. For each learning algorithm, readers will use a real-life scenario to show how Python is used to solve the problem at hand.

- Python data science—manipulating data and data visualization
- Data cleansing
- Understanding Machine learning algorithms
- Supervised learning algorithms
- Unsupervised learning algorithms
- Deploying machine learning models

Python Machine Learning is essential reading for students, developers, or anyone with a keen interest in taking their coding skills to the next level. 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 Solve different problems in modelling deep neural networks using Python, Tensorflow, and Keras with this practical guide About This Book Practical recipes on training different neural network models and tuning them for optimal performance Use Python frameworks like TensorFlow, Caffe, Keras, Theano for Natural Language Processing, Computer Vision, and more A hands-on guide covering the common as well as the not so common problems in deep learning using Python Who This Book Is For This book is intended for machine learning professionals who are looking to use deep learning algorithms to create real-world applications using Python. Thorough understanding of the machine learning concepts and Python libraries such as NumPy, SciPy and scikit-learn is expected. Additionally, basic knowledge in linear algebra and calculus is desired. What You Will Learn Implement different neural network models in Python Select the best Python framework for deep learning such as PyTorch, Tensorflow, MXNet and Keras Apply tips and tricks related to neural networks internals, to boost learning performances Consolidate machine learning principles and apply them in the deep learning field Reuse and adapt Python code snippets to everyday problems Evaluate the cost/benefits and performance implication of each discussed solution In Detail Deep Learning is revolutionizing a wide range of industries. For many applications, deep learning has proven to outperform humans by making faster and more accurate predictions. This book

provides a top-down and bottom-up approach to demonstrate deep learning solutions to real-world problems in different areas. These applications include Computer Vision, Natural Language Processing, Time Series, and Robotics. The Python Deep Learning Cookbook presents technical solutions to the issues presented, along with a detailed explanation of the solutions. Furthermore, a discussion on corresponding pros and cons of implementing the proposed solution using one of the popular frameworks like TensorFlow, PyTorch, Keras and CNTK is provided. The book includes recipes that are related to the basic concepts of neural networks. All techniques s, as well as classical networks topologies. The main purpose of this book is to provide Python programmers a detailed list of recipes to apply deep learning to common and not-so-common scenarios. Style and approach Unique blend of independent recipes arranged in the most logical manner

Microservices in .NET, Second Edition teaches you to build and deploy microservices using ASP.NET and Azure services. Summary In Microservices in .NET, Second Edition you will learn how to: Build scalable microservices that are reliable in production Optimize microservices for continuous delivery Design event-based collaboration between microservices Deploy microservices to Kubernetes Set up Kubernetes in Azure Microservices in .NET, Second Edition is a comprehensive guide to building microservice applications using the .NET stack. After a crystal-clear introduction to the microservices architectural style, it teaches you practical microservices development skills using ASP.NET. This second edition of the bestselling original has been revised with up-to-date tools for the .NET ecosystem, and more new coverage of scoping microservices and deploying to Kubernetes. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Microservice architectures connect independent components that must work together as a system. Integrating new technologies like Docker and Kubernetes with Microsoft's familiar ASP.NET framework and Azure cloud platform enables .NET developers to create and manage microservices efficiently. About the book Microservices in .NET, Second Edition teaches you to build and deploy microservices using ASP.NET and Azure services. It lays out microservice architecture simply, and then guides you through several real-world projects, such as building an ecommerce shopping cart. In this fully revised edition, you'll learn about scoping microservices, deploying to Kubernetes, and operations concerns like monitoring, logging, and security. What's inside Optimize microservices for continuous delivery Design event-based collaboration between microservices Deploy microservices to Kubernetes Set up Kubernetes in Azure About the reader For C# developers. No experience with microservices required. About the author Christian Horsdal is an independent consultant with more than 20 years of experience building projects from large-scale microservice systems to tiny embedded systems. Table of Contents PART 1 GETTING STARTED WITH MICROSERVICES 1 Microservices at a glance 2 A basic shopping cart microservice 3 Deploying a microservice to Kubernetes PART 2 BUILDING MICROSERVICES 4 Identifying and scoping

microservices 5 Microservice collaboration 6 Data ownership and data storage 7 Designing for robustness 8 Writing tests for microservices PART 3 HANDLING CROSS-CUTTING CONCERNS: BUILDING A REUSABLE MICROSERVICE PLATFORM 9 Cross-cutting concerns: Monitoring and logging 10 Securing microservice-to-microservice communication 11 Building a reusable microservice platform PART 4 BUILDING APPLICATIONS 12 Creating applications over microservices

This hands-on book will help you make your machine learning models fairer, safer, and more reliable and in turn improve business outcomes. Every chapter introduces a new mission where you learn how to apply interpretation methods to realistic use cases with methods that work for any model type as well as methods specific for deep neural networks. This book introduces basic-to-advanced deep learning algorithms used in a production environment by AI researchers and principal data scientists; it explains algorithms intuitively, including the underlying math, and shows how to implement them using popular Python-based deep learning libraries such as TensorFlow.

This book shows readers how they can successfully analyze data using only two core machine learning algorithms---and how to do so using the popular Python programming language. These algorithms deal with common scenarios faced by all data analysts and data scientists. This book focuses on two algorithm families (linear methods and ensemble methods) that effectively predict outcomes. This type of problem covers a multitude of use cases (what ad to place on a web page, predicting prices in securities markets, detecting credit card fraud, etc.). The focus on two families gives enough room for full descriptions of the mechanisms at work in the algorithms. Then the code examples serve to illustrate the workings of the machinery with specific hackable code. The author will explain in simple terms, using no complex math, how these algorithms work, and will then show how to apply them in Python. He will also provide advice on how to select from among these algorithms, and will show how to prepare the data, and how to use the trained models in practice. The author begins with an overview of the two core algorithms, explaining the types of problems solved by each one. He then introduces a core set of Python programming techniques that can be used to apply these algorithms. The author shows various techniques for building predictive models that solve a range of problems, from simple to complex; he also shows how to measure the performance of each model to ensure you use the right one. The following chapters provide a deep dive into each of the two algorithms: penalized linear regression and ensemble methods. Chapters will show how to apply each algorithm in Python. Readers can directly use the sample code to build their own solutions. Master the practical aspects of implementing deep learning solutions with PyTorch, using a hands-on approach to understanding both theory and practice. This updated edition will prepare you for applying deep learning to real world problems with a sound theoretical foundation and practical know-how with PyTorch, a platform developed by Facebook's

Artificial Intelligence Research Group. You'll start with a perspective on how and why deep learning with PyTorch has emerged as a path-breaking framework with a set of tools and techniques to solve real-world problems. Next, the book will ground you with the mathematical fundamentals of linear algebra, vector calculus, probability and optimization. Having established this foundation, you'll move on to key components and functionality of PyTorch including layers, loss functions and optimization algorithms. You'll also gain an understanding of Graphical Processing Unit (GPU) based computation, which is essential for training deep learning models. All the key architectures in deep learning are covered, including feedforward networks, convolution neural networks, recurrent neural networks, long short-term memory networks, autoencoders and generative adversarial networks. Backed by a number of tricks of the trade for training and optimizing deep learning models, this edition of Deep Learning with Python explains the best practices in taking these models to production with PyTorch. What You'll Learn Review machine learning fundamentals such as overfitting, underfitting, and regularization. Understand deep learning fundamentals such as feed-forward networks, convolution neural networks, recurrent neural networks, automatic differentiation, and stochastic gradient descent. Apply in-depth linear algebra with PyTorch Explore PyTorch fundamentals and its building blocks Work with tuning and optimizing models Who This Book Is For Beginners with a working knowledge of Python who want to understand Deep Learning in a practical, hands-on manner.

Machine Learning Guide for Oil and Gas Using Python: A Step-by-Step Breakdown with Data, Algorithms, Codes, and Applications delivers a critical training and resource tool to help engineers understand machine learning theory and practice, specifically referencing use cases in oil and gas. The reference moves from explaining how Python works to step-by-step examples of utilization in various oil and gas scenarios, such as well testing, shale reservoirs and production optimization. Petroleum engineers are quickly applying machine learning techniques to their data challenges, but there is a lack of references beyond the math or heavy theory of machine learning. Machine Learning Guide for Oil and Gas Using Python details the open-source tool Python by explaining how it works at an introductory level then bridging into how to apply the algorithms into different oil and gas scenarios. While similar resources are often too mathematical, this book balances theory with applications, including use cases that help solve different oil and gas data challenges. Helps readers understand how open-source Python can be utilized in practical oil and gas challenges Covers the most commonly used algorithms for both supervised and unsupervised learning Presents a balanced approach of both theory and practicality while progressing from introductory to advanced analytical techniques

Introduction to Deep Learning and Neural Networks with Python™: A Practical Guide is an intensive step-by-step guide for neuroscientists to fully understand, practice, and build neural networks. Providing math and Python™ code examples

to clarify neural network calculations, by book's end readers will fully understand how neural networks work starting from the simplest model  $Y=X$  and building from scratch. Details and explanations are provided on how a generic gradient descent algorithm works based on mathematical and Python™ examples, teaching you how to use the gradient descent algorithm to manually perform all calculations in both the forward and backward passes of training a neural network. Examines the practical side of deep learning and neural networks Provides a problem-based approach to building artificial neural networks using real data Describes Python™ functions and features for neuroscientists Uses a careful tutorial approach to describe implementation of neural networks in Python™ Features math and code examples (via companion website) with helpful instructions for easy implementation

Deep Learning with Python Manning Publications

Curious to discover the revolutionary technology that is shaping our future and changing the world? Deep learning is a part of the field of computer science and a subset of machine learning that involves computer systems being able to "learn" unsupervised with data that is unlabeled or unstructured. In 2017, AlphaGo, which is AI developed by Google DeepMind and started off by only knowing the rules of the game, was eventually able to train itself and beat Ke Jie, the world No.1 ranked player at the time. Although this may not seem that impressive at first, it is important to understand that Go is a very complex game that many programmers were not able to trump with AI in the past. Although Go is an interesting example, the possibilities of using machine learning are limitless. From retail to medicine to finance, machine learning has the ability to change each industry it comes into contact with. In fact, this revolution has already begun and will only continue to get bigger. According to statista.com, the artificial intelligence industry is set to grow exponentially in the next few years from \$7 Billion in 2018 to \$90 Billion in 2025! This isn't something you can afford to miss. Without a doubt it is the future. However, it is as complex as it is revolutionary. If you do not have a background or any experience in the field, it is easy to get bogged down by all the complicated concepts and term. And if you are at a more advanced level, the information you find won't be thorough enough. In this book, you will find the perfect balance between the information being very thorough and being able to understand it. Although tailored for beginners, it won't contain simple and easily accessible information. You will dive deep into the field but will be carefully led through it in a way that will make everything easy to understand even if you do not have a technical background in computer programming. In this Guide, you will discover... What Machine Learning and Deep Learning Is And How You Can Use It To Change The World How The Field Can Be Broken Down And Learned In A Manageable Way Various Applications and Potential of Deep Learning That You Can Utilize - That You May Never Have Even Imagined Supervised And Unsupervised Learning - And Breaking It Down Step By Step How You Can Create And Train Deep Learning Models Where and How To Install the Best Programs So You Can Get Started Today Sample Codes And Datasets To Practice Along With And much more! If you are finally prepared to begin grasping this revolutionary technology at a high level despite what your technical background may be, Click "Add to Cart" Now! \*\*Get the Kindle eBook version for FREE when you buy the Paperback version of this book!\*\*

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

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.

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

Gain expertise in advanced deep learning domains such as neural networks, meta-learning, graph neural networks, and memory augmented neural networks using the Python ecosystem

**Key Features**

- Get to grips with building faster and more robust deep learning architectures
- Investigate and train convolutional neural network (CNN) models with GPU-accelerated libraries such as TensorFlow and PyTorch
- Apply deep neural networks (DNNs) to computer vision problems, NLP, and GANs

**Book Description**

In order to build robust deep learning systems, you'll need to understand everything from how neural networks work to training CNN models. In this book, you'll discover newly developed deep learning models, methodologies used in the domain, and their implementation based on areas of application. You'll start by understanding the building blocks and the math behind neural networks, and then move on to CNNs and their advanced applications in computer vision. You'll also learn to apply the most popular CNN architectures in object detection and image segmentation. Further on, you'll focus on variational autoencoders and GANs. You'll then use neural networks to extract sophisticated vector representations of words, before going on to cover various types of recurrent networks, such as LSTM and GRU. You'll even explore the attention mechanism to process sequential data without the help of recurrent neural networks (RNNs). Later, you'll use graph neural networks for processing structured data, along with covering meta-learning, which allows you to train neural networks with fewer training samples. Finally, you'll understand how to apply deep learning to autonomous vehicles. By the end of this book, you'll have mastered key deep learning concepts and the different applications of deep learning models in the real world.

**What you will learn**

- Cover advanced and state-of-the-art neural network architectures
- Understand the theory and math behind neural networks
- Train DNNs and apply them to modern deep learning problems
- Use CNNs for object detection and image segmentation
- Implement generative adversarial networks (GANs) and variational autoencoders to generate new images
- Solve natural language processing (NLP) tasks, such as machine translation, using sequence-to-sequence models
- Understand DL techniques, such as meta-learning and graph neural networks

**Who this book is for**

This book is for data scientists, deep learning engineers and researchers, and AI developers who want to further their knowledge of deep learning and build innovative and unique deep learning projects. Anyone looking to get to grips with advanced use cases and methodologies adopted in the deep learning domain using real-world examples will also find this book useful. Basic understanding of deep learning concepts and working knowledge of the Python programming language is assumed.

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

## Get Free Deep Learning With Python

Gain insight from the foreword by PyTorch cofounder, Soumith Chintala

Discover the practical aspects of implementing deep-learning solutions using the rich Python ecosystem. This book bridges the gap between the academic state-of-the-art and the industry state-of-the-practice by introducing you to deep learning frameworks such as Keras, Theano, and Caffe. The practicalities of these frameworks is often acquired by practitioners by reading source code, manuals, and posting questions on community forums, which tends to be a slow and a painful process. Deep Learning with Python allows you to ramp up to such practical know-how in a short period of time and focus more on the domain, models, and algorithms. This book briefly covers the mathematical prerequisites and fundamentals of deep learning, making this book a good starting point for software developers who want to get started in deep learning. A brief survey of deep learning architectures is also included. Deep Learning with Python also introduces you to key concepts of automatic differentiation and GPU computation which, while not central to deep learning, are critical when it comes to conducting large scale experiments. What You Will Learn Leverage deep learning frameworks in Python namely, Keras, Theano, and Caffe Gain the fundamentals of deep learning with mathematical prerequisites Discover the practical considerations of large scale experiments Take deep learning models to production Who This Book Is For Software developers who want to try out deep learning as a practical solution to a particular problem. Software developers in a data science team who want to take deep learning models developed by data scientists to production.

Understand the principles and practices of machine learning and deep learning This hands-on guide lays out machine learning and deep learning techniques and technologies in a style that is approachable, using just the basic math required. Written by a pair of experts in the field, Machine Learning and Deep Learning Using Python and TensorFlow contains case studies in several industries, including banking, insurance, e-commerce, retail, and healthcare. The book shows how to utilize machine learning and deep learning functions in today's smart devices and apps. You will get download links for datasets, code, and sample projects referred to in the text. Coverage includes: Machine learning and deep learning concepts Python programming and statistics fundamentals Regression and logistic regression Decision trees Model selection and cross-validation Cluster analysis Random forests and boosting Artificial neural networks TensorFlow and Keras Deep learning hyperparameters Convolutional neural networks Recurrent neural networks and long short-term memory

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 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

Gain the confidence you need to apply machine learning in your daily work. With this practical guide, author Matthew Kirk shows you how to integrate and test machine learning algorithms in your code, without the academic subtext. Featuring graphs and highlighted code examples throughout, the book features tests with Python's Numpy, Pandas, Scikit-Learn, and SciPy data science libraries. If you're a software engineer or business analyst interested in data science, this book will help you: Reference real-world examples to test each algorithm through engaging, hands-on exercises Apply test-driven development (TDD) to write and run tests before you start coding Explore techniques for improving your machine-learning models with data extraction and feature development Watch out for the risks of machine learning, such as underfitting or overfitting data Work with K-Nearest Neighbors, neural networks, clustering, and other algorithms

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 Francois 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 Francois 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 Francois 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

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.

Practical Deep Learning teaches total beginners how to build the datasets and models needed to train neural networks for your own DL projects. If you've been curious about machine learning but didn't know where to start, this is the book you've been waiting for. Focusing on the subfield of machine learning known as deep learning, it explains core concepts and gives you the foundation you need to start building your own models. Rather than simply outlining recipes for using existing toolkits, Practical Deep Learning teaches you the why of deep learning and will inspire you to explore further. All you need is basic familiarity with computer programming and high school math—the book will cover the rest. After an introduction to Python, you'll move through key topics like how to build a good training dataset, work with the

scikit-learn and Keras libraries, and evaluate your models' performance. You'll also learn:

- How to use classic machine learning models like k-Nearest Neighbors, Random Forests, and Support Vector Machines
- How neural networks work and how they're trained
- How to use convolutional neural networks
- How to develop a successful deep learning model from scratch

You'll conduct experiments along the way, building to a final case study that incorporates everything you've learned. The perfect introduction to this dynamic, ever-expanding field, *Practical Deep Learning* will give you the skills and confidence to dive into your own machine learning projects.

*The Complete Beginner's Guide to Understanding and Building Machine Learning Systems with Python* Machine Learning with Python for Everyone will help you master the processes, patterns, and strategies you need to build effective learning systems, even if you're an absolute beginner. If you can write some Python code, this book is for you, no matter how little college-level math you know. Principal instructor Mark E. Fenner relies on plain-English stories, pictures, and Python examples to communicate the ideas of machine learning. Mark begins by discussing machine learning and what it can do; introducing key mathematical and computational topics in an approachable manner; and walking you through the first steps in building, training, and evaluating learning systems. Step by step, you'll fill out the components of a practical learning system, broaden your toolbox, and explore some of the field's most sophisticated and exciting techniques. Whether you're a student, analyst, scientist, or hobbyist, this guide's insights will be applicable to every learning system you ever build or use. Understand machine learning algorithms, models, and core machine learning concepts Classify examples with classifiers, and quantify examples with regressors Realistically assess performance of machine learning systems Use feature engineering to smooth rough data into useful forms Chain multiple components into one system and tune its performance Apply machine learning techniques to images and text Connect the core concepts to neural networks and graphical models Leverage the Python scikit-learn library and other powerful tools Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

"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**

## Get Free Deep Learning With Python

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

The purpose of this book is two-fold, we focus on detailed coverage of deep learning and transfer learning, comparing and contrasting the two with easy-to-follow concepts and examples. The second area of focus is on real-world examples and research problems using TensorFlow, Keras and Python ecosystem with hands-on examples.

55% OFF for bookstores! Get maximum revenue from the sale of this book. Your customers will love this book.

Learn advanced state-of-the-art deep learning techniques and their applications using popular Python libraries Key Features Build a strong foundation in neural networks and deep learning with Python libraries Explore advanced deep learning techniques and their applications across computer vision and NLP Learn how a computer can navigate in complex environments with reinforcement learning Book Description With the surge in artificial intelligence in applications catering to both business and consumer needs, deep learning is more important than ever for meeting current and future market demands. With this book, you'll explore deep learning, and learn how to put machine learning to use in your projects. This second edition of Python Deep Learning will get you up to speed with deep learning, deep neural networks, and how to train them with high-performance algorithms and popular Python frameworks. You'll uncover different neural network architectures, such as convolutional networks, recurrent neural networks, long short-term memory (LSTM) networks, and capsule networks. You'll also learn how to solve problems in the fields of computer vision, natural language processing (NLP), and speech recognition. You'll study generative model approaches such as variational autoencoders and Generative Adversarial Networks (GANs) to generate images. As you delve into newly evolved areas of reinforcement learning, you'll gain an understanding of state-of-the-art algorithms that are the main components behind popular games Go, Atari, and Dota. By the end of the book, you will be well-versed with the theory of deep learning along with its real-world applications. What you will learn Grasp the mathematical theory behind neural networks and deep learning processes Investigate and resolve computer vision challenges using convolutional networks and capsule networks Solve generative tasks using variational autoencoders and Generative Adversarial Networks Implement complex NLP tasks using recurrent networks (LSTM and GRU) and attention models Explore reinforcement learning and understand how agents behave in a complex environment Get up to date with applications of deep learning in autonomous vehicles Who this book is for This book is

for data science practitioners, machine learning engineers, and those interested in deep learning who have a basic foundation in machine learning and some Python programming experience. A background in mathematics and conceptual understanding of calculus and statistics will help you gain maximum benefit from this book.

Practical Machine Learning for Data Analysis Using Python is a problem solver's guide for creating real-world intelligent systems. It provides a comprehensive approach with concepts, practices, hands-on examples, and sample code. The book teaches readers the vital skills required to understand and solve different problems with machine learning. It teaches machine learning techniques necessary to become a successful practitioner, through the presentation of real-world case studies in Python machine learning ecosystems. The book also focuses on building a foundation of machine learning knowledge to solve different real-world case studies across various fields, including biomedical signal analysis, healthcare, security, economics, and finance. Moreover, it covers a wide range of machine learning models, including regression, classification, and forecasting. The goal of the book is to help a broad range of readers, including IT professionals, analysts, developers, data scientists, engineers, and graduate students, to solve their own real-world problems. Offers a comprehensive overview of the application of machine learning tools in data analysis across a wide range of subject areas Teaches readers how to apply machine learning techniques to biomedical signals, financial data, and healthcare data Explores important classification and regression algorithms as well as other machine learning techniques Explains how to use Python to handle data extraction, manipulation, and exploration techniques, as well as how to visualize data spread across multiple dimensions and extract useful features

With the resurgence of neural networks in the 2010s, deep learning has become essential for machine learning practitioners and even many software engineers. This book provides a comprehensive introduction for data scientists and software engineers with machine learning experience. You'll start with deep learning basics and move quickly to the details of important advanced architectures, implementing everything from scratch along the way. Author Seth Weidman shows you how neural networks work using a first principles approach. You'll learn how to apply multilayer neural networks, convolutional neural networks, and recurrent neural networks from the ground up. With a thorough understanding of how neural networks work mathematically, computationally, and conceptually, you'll be set up for success on all future deep learning projects. This book provides: Extremely clear and thorough mental models—accompanied by working code examples and mathematical explanations—for understanding neural networks Methods for implementing multilayer neural networks from scratch, using an easy-to-understand object-oriented framework Working implementations and clear-cut explanations of convolutional and recurrent neural networks Implementation of these neural network concepts using the popular PyTorch framework

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

