

Book R Deep Learning Essentials

This book helps readers understand the mathematics of machine learning, and apply them in different situations. It is divided into two basic parts, the first of which introduces readers to the theory of linear algebra, probability, and data distributions and its applications to machine learning. It also includes a detailed introduction to the concepts and constraints of machine learning and what is involved in designing a learning algorithm. This part helps readers understand the mathematical and statistical aspects of machine learning. In turn, the second part discusses the algorithms used in supervised and unsupervised learning. It works out each learning algorithm mathematically and encodes it in R to produce customized learning applications. In the process, it touches upon the specifics of each algorithm and the science behind its formulation. The book includes a wealth of worked-out examples along with R codes. It explains the code for each algorithm, and readers can modify the code to suit their own needs. The book will be of interest to all researchers who intend to use R for machine learning, and those who are interested in the practical aspects of implementing learning algorithms for data analysis. Further, it will be particularly useful and informative for anyone who has struggled to relate the concepts of mathematics and statistics to machine learning.

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.

Build automatic classification and prediction models using unsupervised learning
About This Book- Harness the ability to build algorithms for unsupervised data using deep learning concepts with R- Master the common problems faced such as overfitting of data, anomalous datasets, image recognition, and performance tuning while building the models- Build models relating to neural networks, prediction and deep prediction
Who This Book Is For
This book caters to aspiring data scientists who are well versed with machine learning concepts with R and are looking to explore the deep learning paradigm using the packages available in R. You should have a fundamental understanding of the R language and be comfortable with statistical algorithms and machine learning techniques, but you do not need to be well versed with deep learning concepts.
What You Will Learn- Set up the R package H2O to train deep learning models- Understand the core concepts behind deep learning models- Use Autoencoders to identify anomalous data or outliers- Predict or classify data automatically using deep neural networks- Build generalizable models using regularization to avoid overfitting the training data
In Detail
Deep learning is a branch of machine learning based on a set of algorithms that attempt to model high-level abstractions in data by using model architectures. With the superb memory management and the full integration with multi-node big data platforms, the H2O engine has become more and more popular among data scientists in the field of deep learning. This book will introduce you to the deep learning package H2O with R and help you understand the concepts of deep learning. We will start by setting up important deep learning packages available in R and then move towards building models related to neural networks, prediction, and deep prediction, all of this with the help of real-life examples. After installing the H2O package, you will learn about prediction algorithms. Moving ahead, concepts such as overfitting data, anomalous data, and deep prediction models are explained. Finally, the book will cover concepts relating to tuning and optimizing models.
Style and approach
This book takes a practical approach to showing you the concepts of deep learning with the R programming language. We will start with setting up important deep learning packages available in R and then move towards building models related to neural network, prediction, and deep prediction -

Read Book Book R Deep Learning Essentials

and all of this with the help of real-life examples.

Explore the world of neural networks by building powerful deep learning models using the R ecosystem Key Features Get to grips with the fundamentals of deep learning and neural networks Use R 3.5 and its libraries and APIs to build deep learning models for computer vision and text processing Implement effective deep learning systems in R with the help of end-to-end projects Book Description Deep learning finds practical applications in several domains, while R is the preferred language for designing and deploying deep learning models. This Learning Path introduces you to the basics of deep learning and even teaches you to build a neural network model from scratch. As you make your way through the chapters, you'll explore deep learning libraries and understand how to create deep learning models for a variety of challenges, right from anomaly detection to recommendation systems. The book will then help you cover advanced topics, such as generative adversarial networks (GANs), transfer learning, and large-scale deep learning in the cloud, in addition to model optimization, overfitting, and data augmentation. Through real-world projects, you'll also get up to speed with training convolutional neural networks (CNNs), recurrent neural networks (RNNs), and long short-term memory networks (LSTMs) in R. By the end of this Learning Path, you'll be well versed with deep learning and have the skills you need to implement a number of deep learning concepts in your research work or projects. This Learning Path includes content from the following Packt products: R Deep Learning Essentials - Second Edition by Joshua F. Wiley and Mark Hodnett R Deep Learning Projects by Yuxi (Hayden) Liu and Pablo Maldonado What you will learn Implement credit card fraud detection with autoencoders Train neural networks to perform handwritten digit recognition using MXNet Reconstruct images using variational autoencoders Explore the applications of autoencoder neural networks in clustering and dimensionality reduction Create natural language processing (NLP) models using Keras and TensorFlow in R Prevent models from overfitting the data to improve generalizability Build shallow neural network prediction models Who this book is for This Learning Path is for aspiring data scientists, data analysts, machine learning developers, and deep learning enthusiasts who are well versed in machine learning concepts and are looking to explore the deep learning paradigm using R. A fundamental understanding of R programming and familiarity with the basic concepts of deep learning are necessary to get the most out of this Learning Path.

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

Summary Deep learning has transformed the fields of computer vision, image processing, and natural language applications. Thanks to

Read Book Book R Deep Learning Essentials

TensorFlow.js, now JavaScript developers can build deep learning apps without relying on Python or R. Deep Learning with JavaScript shows developers how they can bring DL technology to the web. Written by the main authors of the TensorFlow library, this new book provides fascinating use cases and in-depth instruction for deep learning apps in JavaScript in your browser or on Node. Foreword by Nikhil Thorat and Daniel Smilkov. About the technology Running deep learning applications in the browser or on Node-based backends opens up exciting possibilities for smart web applications. With the TensorFlow.js library, you build and train deep learning models with JavaScript. Offering uncompromising production-quality scalability, modularity, and responsiveness, TensorFlow.js really shines for its portability. Its models run anywhere JavaScript runs, pushing ML farther up the application stack. About the book In Deep Learning with JavaScript, you'll learn to use TensorFlow.js to build deep learning models that run directly in the browser. This fast-paced book, written by Google engineers, is practical, engaging, and easy to follow. Through diverse examples featuring text analysis, speech processing, image recognition, and self-learning game AI, you'll master all the basics of deep learning and explore advanced concepts, like retraining existing models for transfer learning and image generation. What's inside - Image and language processing in the browser - Tuning ML models with client-side data - Text and image creation with generative deep learning - Source code samples to test and modify About the reader For JavaScript programmers interested in deep learning. About the author Shangqing Cai, Stanley Bileschi and Eric D. Nielsen are software engineers with experience on the Google Brain team, and were crucial to the development of the high-level API of TensorFlow.js. This book is based in part on the classic, Deep Learning with Python by François Chollet. TOC: PART 1 - MOTIVATION AND BASIC CONCEPTS 1 • Deep learning and JavaScript PART 2 - A GENTLE INTRODUCTION TO TENSORFLOW.JS 2 • Getting started: Simple linear regression in TensorFlow.js 3 • Adding nonlinearity: Beyond weighted sums 4 • Recognizing images and sounds using convnets 5 • Transfer learning: Reusing pretrained neural networks PART 3 - ADVANCED DEEP LEARNING WITH TENSORFLOW.JS 6 • Working with data 7 • Visualizing data and models 8 • Underfitting, overfitting, and the universal workflow of machine learning 9 • Deep learning for sequences and text 10 • Generative deep learning 11 • Basics of deep reinforcement learning PART 4 - SUMMARY AND CLOSING WORDS 12 • Testing, optimizing, and deploying models 13 • Summary, conclusions, and beyond

Take a hands-on approach to understanding deep learning and build smart applications that can recognize images and interpret text Key Features Understand how to implement deep learning with TensorFlow and Keras Learn the fundamentals of computer vision and image recognition Study the architecture of different neural networks Book Description Are you fascinated by how deep learning powers intelligent applications such as self-driving cars, virtual assistants, facial recognition devices, and chatbots to process data and solve complex problems? Whether you are familiar with machine learning or are new to this domain, The Deep Learning Workshop will make it easy for you to understand deep learning with the help of interesting examples and exercises throughout. The book starts by highlighting the relationship between deep learning, machine learning, and artificial intelligence and helps you get comfortable with the TensorFlow 2.0 programming structure using hands-on exercises. You'll understand neural networks, the structure of a perceptron, and how to use TensorFlow to create and train models. The book will then let you explore the fundamentals of computer vision by performing image recognition exercises with convolutional neural networks (CNNs) using Keras. As you advance, you'll be able to make your model more powerful by implementing text embedding and sequencing the data using popular deep learning solutions. Finally, you'll get to grips with bidirectional recurrent neural networks (RNNs) and build generative adversarial networks (GANs) for image synthesis. By the end of this deep learning book, you'll have learned the skills essential for building deep learning models with TensorFlow and Keras. What you will learn Understand how deep learning,

Read Book Book R Deep Learning Essentials

machine learning, and artificial intelligence are different Develop multilayer deep neural networks with TensorFlow Implement deep neural networks for multiclass classification using Keras Train CNN models for image recognition Handle sequence data and use it in conjunction with RNNs Build a GAN to generate high-quality synthesized images Who this book is for If you are interested in machine learning and want to create and train deep learning models using TensorFlow and Keras, this workshop is for you. A solid understanding of Python and its packages, along with basic machine learning concepts, will help you to learn the topics quickly.

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.

Build machine learning algorithms, prepare data, and dig deep into data prediction techniques with R About This Book Harness the power of R for statistical computing and data science Explore, forecast, and classify data with R Use R to apply common machine learning algorithms to real-world scenarios Who This Book Is For 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. What You Will Learn Harness the power of R to build common machine learning algorithms with real-world data science applications Get to grips with techniques in R to clean and prepare your data for analysis and visualize your results Discover the different types of machine learning models and learn what is best to meet your data needs and solve data analysis problems Classify your data with Bayesian and nearest neighbour methods Predict values using R to build decision trees, rules, and support vector machines Forecast numeric values with linear regression and model your data with neural networks Evaluate and improve the performance of machine learning models Learn specialized machine learning techniques for text mining, social network data, and big data In Detail Machine learning, at its core, is concerned with transforming data into actionable knowledge. This makes machine learning well suited to the present-day era of big data. Given the growing prominence of R's cross-platform, zero-cost statistical programming environment, there has never been a better time to start applying machine learning to your data. Machine learning with R offers a powerful set of methods to quickly and easily gain insight from your data to both, veterans and beginners in data analytics. Want to turn your data into actionable knowledge, predict outcomes that make real impact, and have constantly developing insights? R gives you access to all the power you need to master exceptional machine learning techniques. The second edition of Machine Learning with R provides you with an introduction to the essential skills required in data science. Without shying away from technical theory, it is written to provide focused and practical knowledge to get you building algorithms and crunching your data, with minimal previous experience. With this book, you'll

discover all the analytical tools you need to gain insights from complex data and learn to choose the correct algorithm for your specific needs. Through full engagement with the sort of real-world problems data-wrangers face, you'll learn to apply machine learning methods to deal with common tasks, including classification, prediction, forecasting, market analysis, and clustering. Transform the way you think about data; discover machine learning with R. Style and approach How can we use machine learning to transform data into action? This book uses a series of simple steps to show you. Using practical examples, the book illustrates how to prepare data for analysis, choose a machine learning method, and measure its success.

Discover best practices for choosing, building, training, and improving deep learning models using Keras-R, and TensorFlow-R libraries Key Features Implement deep learning algorithms to build AI models with the help of tips and tricks Understand how deep learning models operate using expert techniques Apply reinforcement learning, computer vision, GANs, and NLP using a range of datasets Book Description Deep learning is a branch of machine learning based on a set of algorithms that attempt to model high-level abstractions in data. Advanced Deep Learning with R will help you understand popular deep learning architectures and their variants in R, along with providing real-life examples for them. This deep learning book starts by covering the essential deep learning techniques and concepts for prediction and classification. You will learn about neural networks, deep learning architectures, and the fundamentals for implementing deep learning with R. The book will also take you through using important deep learning libraries such as Keras-R and TensorFlow-R to implement deep learning algorithms within applications. You will get up to speed with artificial neural networks, recurrent neural networks, convolutional neural networks, long short-term memory networks, and more using advanced examples. Later, you'll discover how to apply generative adversarial networks (GANs) to generate new images; autoencoder neural networks for image dimension reduction, image de-noising and image correction and transfer learning to prepare, define, train, and model a deep neural network. By the end of this book, you will be ready to implement your knowledge and newly acquired skills for applying deep learning algorithms in R through real-world examples. What you will learn Learn how to create binary and multi-class deep neural network models Implement GANs for generating new images Create autoencoder neural networks for image dimension reduction, image de-noising and image correction Implement deep neural networks for performing efficient text classification Learn to define a recurrent convolutional network model for classification in Keras Explore best practices and tips for performance optimization of various deep learning models Who this book is for This book is for data scientists, machine learning practitioners, deep learning researchers and AI enthusiasts who want to develop their skills and knowledge to implement deep learning techniques and algorithms using the power of R. A solid understanding of machine learning and working knowledge of the R programming language are required.

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

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

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

The Essentials of Data Science: Knowledge Discovery Using R presents the concepts of data science through a hands-on approach using free and open source software. It systematically drives an accessible journey through data analysis and machine learning to discover and share knowledge from data. Building on over thirty years' experience in teaching and practising data science, the author encourages a programming-by-example approach to ensure students and practitioners attune to the practise of data science while building their data skills. Proven frameworks are provided as reusable templates. Real world case studies then provide insight for the data scientist to swiftly adapt the templates to new tasks and datasets. The book begins by introducing data science. It then reviews R's capabilities for analysing data by writing computer programs. These programs are developed and explained step by step. From analysing and visualising data, the framework moves on to tried and tested machine learning techniques for predictive modelling and knowledge discovery. Literate programming and a consistent style are a focus throughout the book.

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

5 real-world projects to help you master deep learning concepts Key Features Master the different deep learning paradigms and build real-world projects related to text generation, sentiment analysis, fraud detection, and more Get to grips with R's impressive range of Deep Learning libraries and frameworks such as deepnet, MXNetR, Tensorflow, H2O, Keras, and text2vec Practical projects that show you how to implement different neural networks with helpful tips, tricks, and best practices Book Description R is a popular programming language used by statisticians and mathematicians for statistical analysis, and is popularly used for deep learning. Deep Learning, as we all know, is one of the trending topics today, and is finding practical applications in a lot of domains. This book demonstrates end-to-end implementations of five real-world projects on popular topics in deep learning such as handwritten digit recognition, traffic light detection, fraud detection, text generation, and sentiment analysis. You'll learn how to train effective neural networks in R—including convolutional neural networks, recurrent neural networks, and LSTMs—and apply them in practical scenarios. The book also highlights how neural networks can be trained using GPU capabilities. You will use popular R libraries and packages—such as MXNetR, H2O, deepnet, and more—to implement the projects. By the end of this book, you will have a better understanding of deep learning concepts and techniques and how to use them in a practical setting. What you will learn Instrument Deep Learning models with packages such as deepnet, MXNetR, Tensorflow, H2O, Keras, and text2vec Apply neural networks to perform handwritten digit recognition using MXNet Get the knack of CNN models, Neural Network API, Keras, and TensorFlow for traffic sign classification -Implement credit card fraud detection with Autoencoders Master reconstructing images using variational autoencoders Wade through sentiment analysis from movie reviews Run from past to future and vice versa with bidirectional Long Short-Term Memory (LSTM) networks Understand the applications of Autoencoder Neural Networks in clustering and dimensionality reduction Who this book is for Machine learning professionals and data scientists looking to master deep learning by implementing practical projects in R will find this book a useful resource. A knowledge of R programming and the basic concepts of deep learning is required to get the best out of this book.

The second edition of a comprehensive introduction to machine learning approaches used in predictive data analytics, covering both theory and practice. Machine learning is often used to build predictive models by extracting patterns from large datasets. These models are used in predictive data analytics applications including price prediction, risk assessment, predicting customer behavior, and document classification. This introductory textbook offers a detailed and focused treatment of the most important machine learning approaches used in predictive data analytics, covering both theoretical concepts and practical applications. Technical and mathematical material is augmented with explanatory worked examples, and case studies illustrate the application of these models in the broader business context. This second edition covers recent developments in machine learning, especially in a new chapter on deep learning, and two new chapters that go beyond predictive analytics to cover unsupervised learning and reinforcement learning.

Tackle the complex challenges faced while building end-to-end deep learning models using modern R libraries Key Features Understand the intricacies of R deep learning packages to perform a range of deep learning tasks Implement deep learning

techniques and algorithms for real-world use cases Explore various state-of-the-art techniques for fine-tuning neural network models Book Description Deep learning (DL) has evolved in recent years with developments such as generative adversarial networks (GANs), variational autoencoders (VAEs), and deep reinforcement learning. This book will get you up and running with R 3.5.x to help you implement DL techniques. The book starts with the various DL techniques that you can implement in your apps. A unique set of recipes will help you solve binomial and multinomial classification problems, and perform regression and hyperparameter optimization. To help you gain hands-on experience of concepts, the book features recipes for implementing convolutional neural networks (CNNs), recurrent neural networks (RNNs), and Long short-term memory (LSTMs) networks, as well as sequence-to-sequence models and reinforcement learning. You'll then learn about high-performance computation using GPUs, along with learning about parallel computation capabilities in R. Later, you'll explore libraries, such as MXNet, that are designed for GPU computing and state-of-the-art DL. Finally, you'll discover how to solve different problems in NLP, object detection, and action identification, before understanding how to use pre-trained models in DL apps. By the end of this book, you'll have comprehensive knowledge of DL and DL packages, and be able to develop effective solutions for different DL problems. What you will learn Work with different datasets for image classification using CNNs Apply transfer learning to solve complex computer vision problems Use RNNs and their variants such as LSTMs and Gated Recurrent Units (GRUs) for sequence data generation and classification Implement autoencoders for DL tasks such as dimensionality reduction, denoising, and image colorization Build deep generative models to create photorealistic images using GANs and VAEs Use MXNet to accelerate the training of DL models through distributed computing Who this book is for This deep learning book is for data scientists, machine learning practitioners, deep learning researchers and AI enthusiasts who want to learn key tasks in deep learning domains using a recipe-based approach. A strong understanding of machine learning and working knowledge of the R programming language is mandatory. 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.

Data science and analytics have emerged as the most desired fields in driving business decisions. Using the techniques and methods of data science, decision makers can uncover hidden patterns in their data, develop algorithms and models that help improve processes and make key business decisions. Data science is a data driven decision making approach that uses several different areas and disciplines with a purpose of extracting insights and knowledge from structured and unstructured data. The algorithms and models of data science along with machine learning and predictive modeling are widely used in solving business problems and predicting future outcomes. This book combines the key concepts of data science and analytics to help you gain a practical understanding of these fields. The four different sections of the book are divided into chapters that explain the core of data science. Given the booming interest in data science, this book is timely and informative.

Get to grips with the essentials of deep learning by leveraging the power of Python Key Features Your one-stop solution to get started with the essentials of deep learning and neural network modeling Train different kinds of neural networks to tackle various problems in Natural Language Processing, computer vision, speech recognition, and more Covers popular Python libraries such as Tensorflow, Keras, and more, along with tips on training, deploying and optimizing your deep learning models in the best possible manner Book Description Deep Learning a trending topic in the field of Artificial Intelligence today and can be considered to be an advanced form of machine learning, which is quite tricky to master. This book will help you take your first steps in training efficient deep learning models and applying them in various practical scenarios. You will model, train, and deploy different kinds of neural networks such as Convolutional Neural Network, Recurrent Neural Network, and will see some of their applications in real-world domains including computer vision, natural language processing, speech recognition, and so on. You will build practical projects such as chatbots, implement reinforcement learning to build smart games, and develop expert systems for image captioning and processing. Popular Python library such as TensorFlow is used in this book to build the models. This book also covers solutions for different problems you might come across while training models, such as noisy datasets, small datasets, and more. This book does not assume any prior knowledge of deep learning. By the end of this book, you will have a firm understanding of the basics of deep learning and neural network modeling, along with their practical applications. What you will learn Get to grips with the core concepts of deep learning and neural networks Set up deep learning library such as TensorFlow Fine-tune your deep learning models for NLP and Computer Vision applications Unify different information

sources, such as images, text, and speech through deep learning Optimize and fine-tune your deep learning models for better performance Train a deep reinforcement learning model that plays a game better than humans Learn how to make your models get the best out of your GPU or CPU Who this book is for Aspiring data scientists and machine learning experts who have limited or no exposure to deep learning will find this book to be very useful. If you are looking for a resource that gets you up and running with the fundamentals of deep learning and neural networks, this book is for you. As the models in the book are trained using the popular Python-based libraries such as Tensorflow and Keras, it would be useful to have sound programming knowledge of Python.

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

Machine learning is an intimidating subject until you know the fundamentals. If you understand basic coding concepts, this introductory guide will help you gain a solid foundation in machine learning principles. Using the R programming language, you'll first start to learn with regression modelling and then move into more advanced topics such as neural networks and tree-based methods. Finally, you'll delve into the frontier of machine learning, using the caret package in R. Once you develop a familiarity with topics such as the difference between regression and classification models, you'll be able to solve an array of machine learning problems. Author Scott V. Burger provides several examples to help you build a working knowledge of machine learning. Explore machine learning models, algorithms, and data training Understand machine learning algorithms for supervised and unsupervised cases Examine statistical concepts for designing data for use in models Dive into linear regression models used in business and science Use single-layer and multilayer neural networks for calculating outcomes Look at how tree-based models work, including popular decision trees Get a comprehensive view of the machine learning ecosystem in R Explore the powerhouse of tools available in R's caret package

Deep learning doesn't have to be intimidating. Until recently, this machine-learning method required years of study, but with frameworks such as Keras and Tensorflow, software engineers without a background in machine learning can quickly enter the field. With the recipes in this cookbook, you'll learn how to solve deep-learning problems for classifying and generating text, images, and music. Each chapter consists of several recipes needed to complete a single project, such as training a music recommending system. Author Douwe Osinga also provides a chapter with half a dozen techniques to help you if you're stuck. Examples are written in Python with code available on GitHub as a set of Python notebooks. You'll learn how to: Create applications that will serve real users Use word embeddings to calculate text similarity Build a movie recommender system based on Wikipedia links Learn how AIs see the world by visualizing their internal state Build a model to suggest emojis for pieces of text Reuse pretrained networks to build an inverse image search service Compare how GANs, autoencoders and LSTMs generate icons Detect music styles and index song collections

If you want to learn how to develop effective machine learning solutions to your business problems in R, this book is for you. It would be helpful to have a bit of familiarity with basic object-oriented programming concepts, but no prior

experience is required.

Data visualization is one of the most important part of data science. Many books and courses present a catalogue of graphics but they don't teach you which charts to use according to the type of the data. In this book, we start by presenting the key graphic systems and packages available in R, including R base graphs, lattice and ggplot2 plotting systems. Next, we provide more than 200 practical examples to create great graphics for the right data using either the ggplot2 package and extensions or the traditional R graphics. With this book, you 'll learn: - How to quickly create beautiful graphics using ggplot2 packages - How to properly customize and annotate the plots - Type of graphics for visualizing categorical and continuous variables - How to add automatically p-values to box plots, bar plots and alternatives - How to add marginal density plots and correlation coefficients to scatter plots - Key methods for analyzing and visualizing multivariate data - R functions and packages for plotting time series data - How to combine multiple plots on one page to create production-quality figures.

This practical guide will teach you how deep learning (DL) can be used to solve complex real-world problems. Key Features Explore deep reinforcement learning (RL), from the first principles to the latest algorithms Evaluate high-profile RL methods, including value iteration, deep Q-networks, policy gradients, TRPO, PPO, DDPG, D4PG, evolution strategies and genetic algorithms Keep up with the very latest industry developments, including AI-driven chatbots Book Description Recent developments in reinforcement learning (RL), combined with deep learning (DL), have seen unprecedented progress made towards training agents to solve complex problems in a human-like way. Google's use of algorithms to play and defeat the well-known Atari arcade games has propelled the field to prominence, and researchers are generating new ideas at a rapid pace. Deep Reinforcement Learning Hands-On is a comprehensive guide to the very latest DL tools and their limitations. You will evaluate methods including Cross-entropy and policy gradients, before applying them to real-world environments. Take on both the Atari set of virtual games and family favorites such as Connect4. The book provides an introduction to the basics of RL, giving you the know-how to code intelligent learning agents to take on a formidable array of practical tasks. Discover how to implement Q-learning on 'grid world' environments, teach your agent to buy and trade stocks, and find out how natural language models are driving the boom in chatbots. What you will learn Understand the DL context of RL and implement complex DL models Learn the foundation of RL: Markov decision processes Evaluate RL methods including Cross-entropy, DQN, Actor-Critic, TRPO, PPO, DDPG, D4PG and others Discover how to deal with discrete and continuous action spaces in various environments Defeat Atari arcade games using the value iteration method Create your own OpenAI Gym environment to train a stock trading agent Teach your agent to play Connect4 using AlphaGo Zero Explore the very latest deep RL research on topics

including AI-driven chatbots Who this book is for Some fluency in Python is assumed. Basic deep learning (DL) approaches should be familiar to readers and some practical experience in DL will be helpful. This book is an introduction to deep reinforcement learning (RL) and requires no background in RL.

Solve real-world data problems with R and machine learning Key Features Third edition of the bestselling, widely acclaimed R machine learning book, updated and improved for R 3.6 and beyond Harness the power of R to build flexible, effective, and transparent machine learning models Learn quickly with a clear, hands-on guide by experienced machine learning teacher and practitioner, Brett Lantz Book Description Machine learning, at its core, is concerned with transforming data into actionable knowledge. R offers a powerful set of machine learning methods to quickly and easily gain insight from your data. Machine Learning with R, Third Edition provides a hands-on, readable guide to applying machine learning to real-world problems. Whether you are an experienced R user or new to the language, Brett Lantz teaches you everything you need to uncover key insights, make new predictions, and visualize your findings. This new 3rd edition updates the classic R data science book to R 3.6 with newer and better libraries, advice on ethical and bias issues in machine learning, and an introduction to deep learning. Find powerful new insights in your data; discover machine learning with R. What you will learn Discover the origins of machine learning and how exactly a computer learns by example Prepare your data for machine learning work with the R programming language Classify important outcomes using nearest neighbor and Bayesian methods Predict future events using decision trees, rules, and support vector machines Forecast numeric data and estimate financial values using regression methods Model complex processes with artificial neural networks — the basis of deep learning Avoid bias in machine learning models Evaluate your models and improve their performance Connect R to SQL databases and emerging big data technologies such as Spark, H2O, and TensorFlow Who this book is for Data scientists, students, and other practitioners who want a clear, accessible guide to machine learning with R.

An accessible introduction to the artificial intelligence technology that enables computer vision, speech recognition, machine translation, and driverless cars. Deep learning is an artificial intelligence technology that enables computer vision, speech recognition in mobile phones, machine translation, AI games, driverless cars, and other applications. When we use consumer products from Google, Microsoft, Facebook, Apple, or Baidu, we are often interacting with a deep learning system. In this volume in the MIT Press Essential Knowledge series, computer scientist John Kelleher offers an accessible and concise but comprehensive introduction to the fundamental technology at the heart of the artificial intelligence revolution. Kelleher explains that deep learning enables data-driven decisions by identifying and extracting patterns from large datasets; its ability to learn from complex data makes deep learning ideally suited to take

advantage of the rapid growth in big data and computational power. Kelleher also explains some of the basic concepts in deep learning, presents a history of advances in the field, and discusses the current state of the art. He describes the most important deep learning architectures, including autoencoders, recurrent neural networks, and long short-term networks, as well as such recent developments as Generative Adversarial Networks and capsule networks. He also provides a comprehensive (and comprehensible) introduction to the two fundamental algorithms in deep learning: gradient descent and backpropagation. Finally, Kelleher considers the future of deep learning—major trends, possible developments, and significant challenges.

Deep Learning Essentials Your hands-on guide to the fundamentals of deep learning and neural network modeling Packt Publishing Ltd

The fundamental mathematical tools needed to understand machine learning include linear algebra, analytic geometry, matrix decompositions, vector calculus, optimization, probability and statistics. These topics are traditionally taught in disparate courses, making it hard for data science or computer science students, or professionals, to efficiently learn the mathematics. This self-contained textbook bridges the gap between mathematical and machine learning texts, introducing the mathematical concepts with a minimum of prerequisites. It uses these concepts to derive four central machine learning methods: linear regression, principal component analysis, Gaussian mixture models and support vector machines. For students and others with a mathematical background, these derivations provide a starting point to machine learning texts. For those learning the mathematics for the first time, the methods help build intuition and practical experience with applying mathematical concepts. Every chapter includes worked examples and exercises to test understanding. Programming tutorials are offered on the book's web site.

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.

Deep learning is often viewed as the exclusive domain of math PhDs and big tech companies. But as this hands-on guide demonstrates, programmers comfortable with Python can achieve impressive results in deep learning with little math background, small amounts of data, and minimal code. How? With *fastai*, the first library to provide a consistent interface to the most frequently used deep learning applications. Authors Jeremy Howard and Sylvain Gugger, the creators of *fastai*, show you how to train a model on a wide range of tasks using *fastai* and PyTorch. You'll also dive progressively further into deep learning theory to gain a complete understanding of the algorithms behind the scenes. Train models in computer vision, natural language processing, tabular data, and collaborative filtering Learn the latest deep learning techniques that matter most in practice Improve accuracy, speed, and reliability by understanding how deep learning models work Discover how to turn your models into web applications Implement deep learning algorithms from scratch Consider the ethical implications of your work Gain insight from the foreword by PyTorch cofounder, Soumith Chintala

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.

Build machine and deep learning systems with the newly released TensorFlow 2 and Keras for the lab, production, and mobile devices Key Features Introduces and then uses TensorFlow 2 and Keras right from the start Teaches key machine and deep learning techniques Understand the fundamentals of deep learning and machine learning through clear explanations and extensive code samples Book Description Deep Learning with TensorFlow 2 and Keras, Second Edition teaches neural networks and deep learning techniques alongside TensorFlow (TF) and Keras. You'll learn how to write deep learning applications in the most powerful, popular, and scalable machine learning stack available. TensorFlow is the machine learning library of choice for professional applications, while Keras offers a simple and powerful Python API for accessing TensorFlow. TensorFlow 2 provides full Keras integration, making advanced machine learning easier and more convenient than ever before. This book also introduces neural networks with TensorFlow, runs through the main applications (regression, ConvNets (CNNs), GANs, RNNs, NLP), covers

two working example apps, and then dives into TF in production, TF mobile, and using TensorFlow with AutoML. What you will learn Build machine learning and deep learning systems with TensorFlow 2 and the Keras API Use Regression analysis, the most popular approach to machine learning Understand ConvNets (convolutional neural networks) and how they are essential for deep learning systems such as image classifiers Use GANs (generative adversarial networks) to create new data that fits with existing patterns Discover RNNs (recurrent neural networks) that can process sequences of input intelligently, using one part of a sequence to correctly interpret another Apply deep learning to natural human language and interpret natural language texts to produce an appropriate response Train your models on the cloud and put TF to work in real environments Explore how Google tools can automate simple ML workflows without the need for complex modeling Who this book is for This book is for Python developers and data scientists who want to build machine learning and deep learning systems with TensorFlow. Whether or not you have done machine learning before, this book gives you the theory and practice required to use Keras, TensorFlow 2, and AutoML to build machine learning systems.

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 the groundbreaking advances of deep learning with this extensively revised new edition of the bestselling original. Learn directly from the creator of Keras and master practical Python deep learning techniques that are easy to apply in the real world. In *Deep Learning with Python, Second Edition* you will learn: Deep learning from first principles Image classification and image segmentation Timeseries forecasting Text classification and machine translation Text generation, neural style transfer, and image generation *Deep Learning with Python* has taught thousands of readers how to put the full capabilities of deep learning into action. This extensively revised second edition introduces deep learning using Python and Keras, and is loaded with insights for both novice and experienced ML practitioners. You'll learn practical techniques that are easy to apply in the real world, and important theory for perfecting neural networks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Recent innovations in deep learning unlock exciting new software capabilities like

automated language translation, image recognition, and more. Deep learning is quickly becoming essential knowledge for every software developer, and modern tools like Keras and TensorFlow put it within your reach—even if you have no background in mathematics or data science. This book shows you how to get started. About the book *Deep Learning with Python, Second Edition* introduces the field of deep learning using Python and the powerful Keras library. In this revised and expanded new edition, Keras creator François Chollet offers insights for both novice and experienced machine learning practitioners. As you move through this book, you'll build your understanding through intuitive explanations, crisp illustrations, and clear examples. You'll quickly pick up the skills you need to start developing deep-learning applications. What's inside Deep learning from first principles Image classification and image segmentation Time series forecasting Text classification and machine translation Text generation, neural style transfer, and image generation About the reader For readers with intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the author François Chollet is a software engineer at Google and creator of the Keras deep-learning library. Table of Contents 1 What is deep learning? 2 The mathematical building blocks of neural networks 3 Introduction to Keras and TensorFlow 4 Getting started with neural networks: Classification and regression 5 Fundamentals of machine learning 6 The universal workflow of machine learning 7 Working with Keras: A deep dive 8 Introduction to deep learning for computer vision 9 Advanced deep learning for computer vision 10 Deep learning for timeseries 11 Deep learning for text 12 Generative deep learning 13 Best practices for the real world 14 Conclusions

Machine learning is a field of AI where we build systems that learn from data. This book explains complicated concepts with real-world applications. It demonstrates the power of R and machine learning extensively while highlighting the constraints. Finally, it will walk you through topics such as text analysis, time series, and deep learning.

This book demonstrates how to use deep Learning in R for machine learning, image classification, and natural language processing. It covers topics such as convolutional networks, recurrent neural networks, transfer learning and deep learning in the cloud. By the end of this book, you will be able to apply deep learning to real-world projects.

Get up and running with machine learning with F# in a fun and functional way About This Book- Design algorithms in F# to tackle complex computing problems- Be a proficient F# data scientist using this simple-to-follow guide- Solve real-world, data-related problems with robust statistical models, built for a range of datasets Who This Book Is For If you are a C# or an F# developer who now wants to explore the area of machine learning, then this book is for you. Familiarity with theoretical concepts and notation of mathematics and statistics would be an added advantage. What You Will Learn- Use F# to find patterns through raw data- Build a set of classification systems using Accord.NET, Weka, and F#- Run machine learning jobs on the Cloud with MBrace- Perform mathematical operations on matrices and vectors using Math.NET- Use a recommender system for your own problem domain- Identify tourist spots across the globe using inputs from the user with decision tree algorithms In Detail The F# functional programming language enables developers to write simple code to solve complex problems. With F#, developers create consistent and predictable programs that are easier to test and reuse, simpler to parallelize, and are less prone to bugs. If you want to learn how to use F# to build machine learning systems, then this is the book you want. Starting with an introduction to the several categories on machine learning, you will quickly learn to implement time-tested, supervised learning algorithms.

Read Book Book R Deep Learning Essentials

You will gradually move on to solving problems on predicting housing pricing using Regression Analysis. You will then learn to use Accord.NET to implement SVM techniques and clustering. You will also learn to build a recommender system for your e-commerce site from scratch. Finally, you will dive into advanced topics such as implementing neural network algorithms while performing sentiment analysis on your data. Style and approach This book is a fast-paced tutorial guide that uses hands-on examples to explain real-world applications of machine learning. Using practical examples, the book will explore several machine learning techniques and also describe how you can use F# to build machine learning systems.

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). 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

“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 Python programmers with an interest in machine learning. No experience with PyTorch or other deep learning frameworks is required.

About The Authors Eli Stevens has worked in Silicon Valley for the past 15 years as a software engineer, and the past 7 years as Chief Technical Officer of a startup making medical device software. Luca Antiga is co-founder and CEO of an AI engineering company located in Bergamo, Italy, and a regular contributor to PyTorch. Thomas Viehmann is a Machine Learning and PyTorch speciality trainer and consultant based in Munich, Germany and a PyTorch core developer.

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

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

Read Book Book R Deep Learning Essentials

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

[Copyright: fc2eeb67a387a85b12716bcd7bc78ceb](https://www.folio3.com/author/fc2eeb67a387a85b12716bcd7bc78ceb/)