

# Deep Learning With R

Examine the latest technological advancements in building a scalable machine-learning model with big data using R. This second edition shows you how to work with a machine-learning algorithm and use it to build a ML model from raw data. You will see how to use R programming with TensorFlow, thus avoiding the effort of learning Python if you are only comfortable with R. As in the first edition, the authors have kept the fine balance of theory and application of machine learning through various real-world use-cases which gives you a comprehensive collection of topics in machine learning. New chapters in this edition cover time series models and deep learning.

What You'll Learn Understand machine learning algorithms using R Master the process of building machine-learning models Cover the theoretical foundations of machine-learning algorithms See industry focused real-world use cases Tackle time series modeling in R Apply deep learning using Keras and TensorFlow in R Who This Book is For Data scientists, data science professionals, and researchers in academia who want to understand the nuances of machine-learning approaches/algorithms in practice using R.

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 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

## Read Book Deep Learning With R

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

## Read Book Deep Learning With R

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.

Summary

Machine learning (ML) is a collection of programming techniques for discovering relationships in data. With ML algorithms, you can cluster and classify data for tasks like making recommendations or fraud detection and make predictions for sales trends, risk analysis, and other forecasts. Once the domain of academic data scientists, machine learning has become a mainstream business process, and tools like the

easy-to-learn R programming language put high-quality data analysis in the hands of any programmer. Machine Learning with R, the tidyverse, and mlr teaches you widely used ML techniques and how to apply them to your own datasets using the R programming language and its powerful ecosystem of tools. This book will get you started! Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the book Machine Learning with R, the tidyverse, and mlr gets you started in machine learning using R Studio and the awesome mlr machine learning package. This practical guide simplifies theory and avoids needlessly complicated statistics or math. All core ML techniques are clearly explained through graphics and easy-to-grasp examples. In each engaging chapter, you'll put a new algorithm into action to solve a quirky predictive analysis problem, including Titanic survival odds, spam email filtering, and poisoned wine investigation. What's inside Using the tidyverse packages to process and plot your data Techniques for supervised and unsupervised learning Classification, regression, dimension reduction, and clustering algorithms Statistics primer to fill gaps in your knowledge About the reader For newcomers to machine learning with basic skills in R. About the author Hefin I. Rhys is a senior laboratory research scientist at the Francis Crick

## Read Book Deep Learning With R

Institute. He runs his own YouTube channel of screencast tutorials for R and RStudio. Table of contents: PART 1 - INTRODUCTION 1. Introduction to machine learning 2. Tidying, manipulating, and plotting data with the tidyverse PART 2 - CLASSIFICATION 3. Classifying based on similarities with k-nearest neighbors 4. Classifying based on odds with logistic regression 5. Classifying by maximizing separation with discriminant analysis 6. Classifying with naive Bayes and support vector machines 7. Classifying with decision trees 8. Improving decision trees with random forests and boosting PART 3 - REGRESSION 9. Linear regression 10. Nonlinear regression with generalized additive models 11. Preventing overfitting with ridge regression, LASSO, and elastic net 12. Regression with kNN, random forest, and XGBoost PART 4 - DIMENSION REDUCTION 13. Maximizing variance with principal component analysis 14. Maximizing similarity with t-SNE and UMAP 15. Self-organizing maps and locally linear embedding PART 5 - CLUSTERING 16. Clustering by finding centers with k-means 17. Hierarchical clustering 18. Clustering based on density: DBSCAN and OPTICS 19. Clustering based on distributions with mixture modeling 20. Final notes and further reading Introduces deep learning systems using the powerful Keras library and its R language interface. The book builds your understanding of deep learning through

intuitive explanations and practical examples. Master machine learning techniques with R to deliver insights for complex projects About This Book Get to grips with the application of Machine Learning methods using an extensive set of R packages Understand the benefits and potential pitfalls of using machine learning methods Implement the numerous powerful features offered by R with this comprehensive guide to building an independent R-based ML system Who This Book Is For If you want to learn how to use R's machine learning capabilities to solve complex business problems, then this book is for you. Some experience with R and a working knowledge of basic statistical or machine learning will prove helpful. What You Will Learn Gain deep insights to learn the applications of machine learning tools to the industry Manipulate data in R efficiently to prepare it for analysis Master the skill of recognizing techniques for effective visualization of data Understand why and how to create test and training data sets for analysis Familiarize yourself with fundamental learning methods such as linear and logistic regression Comprehend advanced learning methods such as support vector machines Realize why and how to apply unsupervised learning methods In Detail Machine learning is a field of Artificial Intelligence to build systems that learn from data. Given the growing prominence of R—a cross-platform, zero-cost statistical programming

## Read Book Deep Learning With R

environment—there has never been a better time to start applying machine learning to your data. The book starts with introduction to Cross-Industry Standard Process for Data Mining. It takes you through Multivariate Regression in detail. Moving on, you will also address Classification and Regression trees. You will learn a couple of “Unsupervised techniques”. Finally, the book will walk you through text analysis and time series. The book will deliver practical and real-world solutions to problems and variety of tasks such as complex recommendation systems. By the end of this book, you will gain expertise in performing R machine learning and will be able to build complex ML projects using R and its packages.

**Style and approach** This is a book explains complicated concepts with easy to follow theory and real-world, practical applications. It demonstrates the power of R and machine learning extensively while highlighting the constraints.

**Master machine learning techniques with R to deliver insights in complex projects**

**About This Book**

Understand and apply machine learning methods using an extensive set of R packages such as XGBOOST

Understand the benefits and potential pitfalls of using machine learning methods such as Multi-Class Classification and Unsupervised Learning

Implement advanced concepts in machine learning with this example-rich guide

**Who This Book Is For** This book is for data science professionals,

## Read Book Deep Learning With R

data analysts, or anyone with a working knowledge of machine learning, with R who now want to take their skills to the next level and become an expert in the field. What You Will Learn Gain deep insights into the application of machine learning tools in the industry Manipulate data in R efficiently to prepare it for analysis Master the skill of recognizing techniques for effective visualization of data Understand why and how to create test and training data sets for analysis Master fundamental learning methods such as linear and logistic regression Comprehend advanced learning methods such as support vector machines Learn how to use R in a cloud service such as Amazon In Detail This book will teach you advanced techniques in machine learning with the latest code in R 3.3.2. You will delve into statistical learning theory and supervised learning; design efficient algorithms; learn about creating Recommendation Engines; use multi-class classification and deep learning; and more. You will explore, in depth, topics such as data mining, classification, clustering, regression, predictive modeling, anomaly detection, boosted trees with XGBOOST, and more. More than just knowing the outcome, you'll understand how these concepts work and what they do. With a slow learning curve on topics such as neural networks, you will explore deep learning, and more. By the end of this book, you will be able to perform machine learning with R

## Read Book Deep Learning With R

in the cloud using AWS in various scenarios with different datasets. Style and approach The book delivers practical and real-world solutions to problems and a variety of tasks such as complex recommendation systems. By the end of this book, you will have gained expertise in performing R machine learning and will be able to build complex machine learning projects using R and its packages. 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

## Read Book Deep Learning With R

approach and using real word 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.

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

# Read Book Deep Learning With R

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

# Read Book Deep Learning With R

instance

Understand how machine learning works and get hands-on experience of using R to build algorithms that can solve various real-world problems

**Key Features**

- Gain a comprehensive overview of different machine learning techniques
- Explore various methods for selecting a particular algorithm
- Implement a machine learning project from problem definition through to the final model

**Book Description**

With huge amounts of data being generated every moment, businesses need applications that apply complex mathematical calculations to data repeatedly and at speed. With machine learning techniques and R, you can easily develop these kinds of applications in an efficient way.

Practical Machine Learning with R begins by helping you grasp the basics of machine learning methods, while also highlighting how and why they work. You will understand how to get these algorithms to work in practice, rather than focusing on mathematical derivations. As you progress from one chapter to another, you will gain hands-on experience of building a machine learning solution in R. Next, using R packages such as rpart, random forest, and multiple imputation by chained equations (MICE), you will learn to implement algorithms including neural net classifier, decision trees, and linear and non-linear regression. As you progress through the book, you'll delve into various machine learning techniques for both supervised and unsupervised learning approaches. In addition to this, you'll gain insights into partitioning the datasets and mechanisms to evaluate the results from each model and be able to compare them. By the end of this book, you will have gained expertise in solving your business problems, starting by forming a good problem statement, selecting the most appropriate model to solve your problem, and then ensuring that you do not overtrain it. What you will learn

Define a problem that can be solved by training

## Read Book Deep Learning With R

a machine learning model Obtain, verify and clean data before transforming it into the correct format for use Perform exploratory analysis and extract features from data Build models for neural net, linear and non-linear regression, classification, and clustering Evaluate the performance of a model with the right metrics Implement a classification problem using the neural net package Employ a decision tree using the random forest library Who this book is for If you are a data analyst, data scientist, or a business analyst who wants to understand the process of machine learning and apply it to a real dataset using R, this book is just what you need. Data scientists who use Python and want to implement their machine learning solutions using R will also find this book very useful. The book will also enable novice programmers to start their journey in data science. Basic knowledge of any programming language is all you need to get started.

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.

Summary Deep Learning with R introduces the world of deep learning using the powerful Keras library and its R language

## Read Book Deep Learning With R

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

# Read Book Deep Learning With R

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 Powerful, independent recipes to build deep learning models in different application areas using R libraries About This Book Master intricacies of R deep learning packages such as mxnet & tensorflow Learn application on deep learning in different domains using practical examples from text, image and speech Guide to set-up deep learning models using CPU and GPU Who This Book Is For Data science professionals or analysts who have performed machine learning tasks and now want to explore deep learning and want a quick reference that could address the pain points while implementing deep learning. Those who wish to have an edge over other deep learning professionals will find this book quite useful. What You Will Learn Build deep learning models in different application areas using TensorFlow, H2O, and MXnet. Analyzing a Deep boltzmann machine Setting up and Analysing Deep belief networks Building supervised model using various machine learning algorithms Set up variants of basic convolution function Represent data using Autoencoders. Explore generative models available in Deep Learning. Discover sequence modeling using Recurrent nets Learn fundamentals of Reinforcement Learning Learn the steps involved in applying Deep Learning in text mining Explore application of deep learning in signal processing Utilize Transfer learning for utilizing pre-trained model Train a deep learning model on a GPU In Detail Deep Learning is the next big thing. It is a part of machine learning. It's favorable results in applications with huge and complex data is

## Read Book Deep Learning With R

remarkable. Simultaneously, R programming language is very popular amongst the data miners and statisticians. This book will help you to get through the problems that you face during the execution of different tasks and Understand hacks in deep learning, neural networks, and advanced machine learning techniques. It will also take you through complex deep learning algorithms and various deep learning packages and libraries in R. It will be starting with different packages in Deep Learning to neural networks and structures. You will also encounter the applications in text mining and processing along with a comparison between CPU and GPU performance. By the end of the book, you will have a logical understanding of Deep learning and different deep learning packages to have the most appropriate solutions for your problems. Style and approach Collection of hands-on recipes that would act as your all-time reference for your deep learning needs

The statistics profession is at a unique point in history. The need for valid statistical tools is greater than ever; data sets are massive, often measuring hundreds of thousands of measurements for a single subject. The field is ready to move towards clear objective benchmarks under which tools can be evaluated. Targeted learning allows (1) the full generalization and utilization of cross-validation as an estimator selection tool so that the subjective choices made by humans are now made by the machine, and (2) targeting the fitting of the probability distribution of the data toward the target parameter representing the scientific question of interest. This book is aimed at both statisticians and applied researchers interested in causal inference and general effect estimation for observational and experimental data. Part I is an accessible introduction to super learning and the targeted maximum likelihood estimator, including related concepts necessary to understand and apply these methods. Parts II-IX handle

## Read Book Deep Learning With R

complex data structures and topics applied researchers will immediately recognize from their own research, including time-to-event outcomes, direct and indirect effects, positivity violations, case-control studies, censored data, longitudinal data, and genomic studies.

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

## Read Book Deep Learning With R

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.

AI framework intended to solve a problem of bias-variance tradeoff for supervised learning methods in real-life applications. The AI framework comprises of bootstrapping to create multiple training and testing data sets with various characteristics, design and analysis of statistical experiments to identify optimal feature subsets and optimal hyper-parameters for ML methods, data contamination to test for the robustness of the classifiers. Key Features: Using ML methods by itself doesn't ensure building classifiers that generalize well for new data Identifying optimal feature subsets and hyper-parameters of ML methods can be resolved using design and analysis of statistical experiments Using a bootstrapping approach to massive sampling of training and tests datasets with various data characteristics (e.g.: contaminated training sets) allows dealing with bias Developing of SAS-based table-driven environment allows managing all meta-data related to the proposed AI framework and creating interoperability with R libraries to accomplish variety of statistical and machine-learning tasks Computer programs in R and SAS that create AI framework are

## Read Book Deep Learning With R

available on GitHub

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

## Read Book Deep Learning With R

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.

Deep Learning with R introduces deep learning and neural networks using the R programming language. The book builds on the understanding of the theoretical and mathematical constructs and enables the reader to create applications on computer vision, natural language processing and transfer learning. The book starts with an introduction to machine learning and moves on to describe the basic architecture, different activation functions, forward propagation, cross-entropy loss and backward propagation of a simple neural network. It goes on to create different code segments to construct

## Read Book Deep Learning With R

deep neural networks. It discusses in detail the initialization of network parameters, optimization techniques, and some of the common issues surrounding neural networks such as dealing with NaNs and the vanishing/exploding gradient problem. Advanced variants of multilayered perceptrons namely, convolutional neural networks and sequence models are explained, followed by application to different use cases. The book makes extensive use of the Keras and TensorFlow frameworks.

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.

Uncover the power of artificial neural networks by implementing them through R code. About This Book Develop a strong background in neural networks with R, to implement them in your applications Build smart systems using the power of deep learning Real-world case studies to illustrate the power of neural network models Who This Book Is For This book is intended for anyone who has a statistical background with knowledge in R and wants to work with neural networks to get better results from complex data. If you are interested in artificial intelligence and deep learning and you want to level up, then this book is what you need! What You Will Learn Set up R packages for neural networks and deep learning Understand the core concepts of artificial neural networks Understand neurons, perceptrons, bias, weights, and activation functions Implement supervised and unsupervised machine learning in R for neural networks Predict and classify data automatically using neural networks Evaluate and fine-tune the models you build. In Detail Neural networks are one of the most fascinating machine learning models for solving complex computational problems efficiently. Neural networks are used to solve wide range of problems in different areas of AI and machine learning. This book explains the niche aspects of neural networking and provides you with foundation to get started with advanced topics. The book begins with neural network design using the neural net package, then you'll build a solid foundation knowledge of how a neural network learns from data, and the principles behind it. This book covers various types of

## Read Book Deep Learning With R

neural network including recurrent neural networks and convoluted neural networks. You will not only learn how to train neural networks, but will also explore generalization of these networks. Later we will delve into combining different neural network models and work with the real-world use cases. By the end of this book, you will learn to implement neural network models in your applications with the help of practical examples in the book. Style and approach A step-by-step guide filled with real-world practical examples.

Based on the author's experience in teaching data science for more than 10 years, *Mathematics and Programming for Machine Learning with R: From the Ground Up* reveals how machine learning algorithms do their magic and explains how these algorithms can be implemented in code. It is designed to provide readers with an understanding of the reasoning behind machine learning algorithms as well as how to program them. Written for novice programmers, the book progresses step-by-step, providing the coding skills needed to implement machine learning algorithms in R. The book begins with simple implementations and fundamental concepts of logic, sets, and probability before moving to the coverage of powerful deep learning algorithms. The first eight chapters deal with probability-based machine learning algorithms, and the last eight chapters deal with machine learning based on artificial neural networks. The first half of the book does not require mathematical sophistication, although familiarity with probability and statistics would be helpful. The second half assumes the reader is familiar with at least one semester of calculus.

## Read Book Deep Learning With R

The text guides novice R programmers through algorithms and their application and along the way; the reader gains programming confidence in tackling advanced R programming challenges. Highlights of the book include: More than 400 exercises A strong emphasis on improving programming skills and guiding beginners to the implementation of full-fledged algorithms Coverage of fundamental computer and mathematical concepts including logic, sets, and probability In-depth explanations of machine learning algorithms

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 it's 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.

Applied Predictive Modeling covers the overall predictive modeling process, beginning with the crucial steps of data preprocessing, data splitting and foundations of model tuning. The text then provides intuitive explanations of numerous common and modern regression and classification techniques, always with an emphasis on illustrating and solving real data problems. The text illustrates all parts of the modeling process through many hands-on, real-life examples, and every chapter contains extensive R code for each step of the process. This multi-purpose text can be used as an introduction to predictive models and the overall modeling process, a practitioner's reference handbook, or as a text for advanced undergraduate or graduate level predictive modeling courses. To that end, each chapter contains problem sets to help solidify the covered concepts and uses data available in the book's R package. This text is intended for a broad audience as both an introduction to predictive models as well as a guide to applying them. Non-mathematical readers will appreciate the intuitive explanations of the techniques while an emphasis on problem-solving with real data across a wide variety of applications will aid practitioners who wish to extend their expertise. Readers should have knowledge of basic statistical ideas, such as correlation

## Read Book Deep Learning With R

and linear regression analysis. While the text is biased against complex equations, a mathematical background is needed for advanced topics.

A practitioner's tools have a direct impact on the success of his or her work. This book will provide the data scientist with the tools and techniques required to excel with statistical learning methods in the areas of data access, data munging, exploratory data analysis, supervised machine learning, unsupervised machine learning and model evaluation. Machine learning and data science are large disciplines, requiring years of study in order to gain proficiency. This book can be viewed as a set of essential tools we need for a long-term career in the data science field – recommendations are provided for further study in order to build advanced skills in tackling important data problem domains. The R statistical environment was chosen for use in this book. R is a growing phenomenon worldwide, with many data scientists using it exclusively for their project work. All of the code examples for the book are written in R. In addition, many popular R packages and data sets will be used.

Understand the fundamentals of machine learning with R and build your own dynamic algorithms to tackle complicated real-world problems successfully About This Book Get to grips with the concepts of machine learning through exciting real-world examples Visualize and solve complex problems by using power-packed R constructs and its robust packages for machine learning Learn to build your own machine learning system with this example-based practical guide Who This Book Is For If you are interested in mining useful information from data using state-of-the-art techniques to make data-

## Read Book Deep Learning With R

driven decisions, this is a go-to guide for you. No prior experience with data science is required, although basic knowledge of R is highly desirable. Prior knowledge in machine learning would be helpful but is not necessary. What You Will Learn Utilize the power of R to handle data extraction, manipulation, and exploration techniques Use R to visualize data spread across multiple dimensions and extract useful features Explore the underlying mathematical and logical concepts that drive machine learning algorithms Dive deep into the world of analytics to predict situations correctly Implement R machine learning algorithms from scratch and be amazed to see the algorithms in action Write reusable code and build complete machine learning systems from the ground up Solve interesting real-world problems using machine learning and R as the journey unfolds Harness the power of robust and optimized R packages to work on projects that solve real-world problems in machine learning and data science In Detail Data science and machine learning are some of the top buzzwords in the technical world today. From retail stores to Fortune 500 companies, everyone is working hard to making machine learning give them data-driven insights to grow their business. With powerful data manipulation features, machine learning packages, and an active developer community, R empowers users to build sophisticated machine learning systems to solve real-world data problems. This book takes you on a data-driven journey that starts with the very basics of R and machine learning and gradually builds upon the concepts to work on projects that tackle real-world problems. You'll begin by getting an understanding of the core concepts and definitions required to appreciate machine learning algorithms and concepts. Building upon the basics, you will then work on three different projects to apply the concepts of machine learning, following current trends and cover major algorithms as well as popular

## Read Book Deep Learning With R

R packages in detail. These projects have been neatly divided into six different chapters covering the worlds of e-commerce, finance, and social-media, which are at the very core of this data-driven revolution. Each of the projects will help you to understand, explore, visualize, and derive insights depending upon the domain and algorithms. Through this book, you will learn to apply the concepts of machine learning to deal with data-related problems and solve them using the powerful yet simple language, R. Style and approach The book is an enticing journey that starts from the very basics to gradually pick up pace as the story unfolds. Each concept is first defined in the larger context of things succinctly, followed by a detailed explanation of their application. Each topic is explained with the help of a project that solves a real real-world problem involving hands-on work thus giving you a deep insight into the world of machine learning.

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

# Read Book Deep Learning With R

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.

R has been the gold standard in applied machine learning for a long time. Surveys show that it is the most popular platform used by professional data scientists. It is also preferred by the best data scientists in the world. In this Ebook, learn how to get started, practice and apply machine learning using the R platform.

## Read Book Deep Learning With R

Guides professionals and students through the rapidly growing field of machine learning with hands-on examples in the popular R programming language Machine learning—a branch of Artificial Intelligence (AI) which enables computers to improve their results and learn new approaches without explicit instructions—allows organizations to reveal patterns in their data and incorporate predictive analytics into their decision-making process. Practical Machine Learning in R provides a hands-on approach to solving business problems with intelligent, self-learning computer algorithms. Bestselling author and data analytics experts Fred Nwanganga and Mike Chapple explain what machine learning is, demonstrate its organizational benefits, and provide hands-on examples created in the R programming language. A perfect guide for professional self-taught learners or students in an introductory machine learning course, this reader-friendly book illustrates the numerous real-world business uses of machine learning approaches. Clear and detailed chapters cover data wrangling, R programming with the popular RStudio tool, classification and regression techniques, performance evaluation, and more. Explores data management techniques, including data collection, exploration and dimensionality reduction Covers unsupervised learning, where readers identify and summarize patterns using approaches such as apriori, eclat and clustering Describes the principles behind the Nearest Neighbor, Decision Tree and Naive Bayes classification techniques Explains how to evaluate and choose the right model, as well as how to improve model performance using ensemble methods such as Random Forest and XGBoost Practical Machine Learning in R is a must-have guide for business analysts, data scientists, and other professionals interested in leveraging the power of AI to solve business problems, as well as students and independent learners seeking to enter the field.

# Read Book Deep Learning With R

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

## Read Book Deep Learning With R

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  
Bridge the gap between a high-level understanding of how an  
algorithm works and knowing the nuts and bolts to tune your  
models better. This book will give you the confidence and  
skills when developing all the major machine learning models.  
In Pro Machine Learning Algorithms, you will first develop the  
algorithm in Excel so that you get a practical understanding of  
all the levers that can be tuned in a model, before  
implementing the models in Python/R. You will cover all the  
major algorithms: supervised and unsupervised learning,  
which include linear/logistic regression; k-means clustering;  
PCA; recommender system; decision tree; random forest;  
GBM; and neural networks. You will also be exposed to the  
latest in deep learning through CNNs, RNNs, and word2vec  
for text mining. You will be learning not only the algorithms,  
but also the concepts of feature engineering to maximize the  
performance of a model. You will see the theory along with  
case studies, such as sentiment classification, fraud

## Read Book Deep Learning With R

detection, recommender systems, and image recognition, so that you get the best of both theory and practice for the vast majority of the machine learning algorithms used in industry. Along with learning the algorithms, you will also be exposed to running machine-learning models on all the major cloud service providers. You are expected to have minimal knowledge of statistics/software programming and by the end of this book you should be able to work on a machine learning project with confidence. What You Will Learn Get an in-depth understanding of all the major machine learning and deep learning algorithms Fully appreciate the pitfalls to avoid while building models Implement machine learning algorithms in the cloud Follow a hands-on approach through case studies for each algorithm Gain the tricks of ensemble learning to build more accurate models Discover the basics of programming in R/Python and the Keras framework for deep learning Who This Book Is For Business analysts/ IT professionals who want to transition into data science roles. Data scientists who want to solidify their knowledge in machine learning.

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

## Read Book Deep Learning With R

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.

Text data is important for many domains, from healthcare to marketing to the digital humanities, but specialized approaches are necessary to create features for machine learning from language. Supervised Machine Learning for Text Analysis in R explains how to preprocess text data for modeling, train models, and evaluate model performance using tools from the tidyverse and tidymodels ecosystem. Models like these can be used to make predictions for new observations, to understand what natural language features or characteristics contribute to differences in the output, and more. If you are already familiar with the basics of predictive modeling, use the comprehensive, detailed examples in this book to extend your skills to the domain of natural language

## Read Book Deep Learning With R

processing. This book provides practical guidance and directly applicable knowledge for data scientists and analysts who want to integrate unstructured text data into their modeling pipelines. Learn how to use text data for both regression and classification tasks, and how to apply more straightforward algorithms like regularized regression or support vector machines as well as deep learning approaches. Natural language must be dramatically transformed to be ready for computation, so we explore typical text preprocessing and feature engineering steps like tokenization and word embeddings from the ground up. These steps influence model results in ways we can measure, both in terms of model metrics and other tangible consequences such as how fair or appropriate model results are.

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

## Read Book Deep Learning With R

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

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

## Read Book Deep Learning With R

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.

Explore and implement deep learning to solve various real-world problems using modern R libraries such as TensorFlow, MXNet, H2O, and Deepnet

**Key Features**

- Understand deep learning algorithms and architectures using R and determine which algorithm is best suited for a specific problem
- Improve models using parameter tuning, feature engineering, and ensembling
- Apply advanced neural network models such as deep autoencoders and generative adversarial networks (GANs) across different domains

**Book Description** Deep learning enables efficient and accurate learning from a massive amount of data. This book will help you overcome a number of challenges using various deep learning algorithms and architectures with R programming. This book starts with a brief overview of machine learning and deep learning and how to build your first neural network. You'll understand the architecture of various deep learning algorithms and their applicable fields, learn how to build deep learning models, optimize hyperparameters, and evaluate model performance. Various deep learning applications in image processing, natural language processing (NLP), recommendation systems, and predictive analytics will

also be covered. Later chapters will show you how to tackle recognition problems such as image recognition and signal detection, programmatically summarize documents, conduct topic modeling, and forecast stock market prices. Toward the end of the book, you will learn the common applications of GANs and how to build a face generation model using them. Finally, you'll get to grips with using reinforcement learning and deep reinforcement learning to solve various real-world problems. By the end of this deep learning book, you will be able to build and deploy your own deep learning applications using appropriate frameworks and algorithms. What you will learn

- Design a feedforward neural network to see how the activation function computes an output
- Create an image recognition model using convolutional neural networks (CNNs)
- Prepare data, decide hidden layers and neurons and train your model with the backpropagation algorithm
- Apply text cleaning techniques to remove uninformative text using NLP
- Build, train, and evaluate a GAN model for face generation
- Understand the concept and implementation of reinforcement learning in R

Who this book is for This book is for data scientists, machine learning engineers, and deep learning developers who are familiar with machine learning and are looking to enhance their knowledge of deep learning using practical examples. Anyone interested in increasing the efficiency of their machine learning applications and exploring various options in R will also find this book useful. Basic knowledge of machine learning techniques and working knowledge of the R programming language is expected.

## Read Book Deep Learning With R

Understand deep learning, the nuances of its different models, and where these models can be applied. The abundance of data and demand for superior products/services have driven the development of advanced computer science techniques, among them image and speech recognition. Introduction to Deep Learning Using R provides a theoretical and practical understanding of the models that perform these tasks by building upon the fundamentals of data science through machine learning and deep learning. This step-by-step guide will help you understand the disciplines so that you can apply the methodology in a variety of contexts. All examples are taught in the R statistical language, allowing students and professionals to implement these techniques using open source tools.

**What You'll Learn**

- Understand the intuition and mathematics that power deep learning models
- Utilize various algorithms using the R programming language and its packages
- Use best practices for experimental design and variable selection
- Practice the methodology to approach and effectively solve problems as a data scientist
- Evaluate the effectiveness of algorithmic solutions and enhance their predictive power

**Who This Book Is For** Students, researchers, and data scientists who are familiar with programming using R. This book also is also of use for those who wish to learn how to appropriately deploy these algorithms in applications where they would be most useful.

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

## Read Book Deep Learning With R

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

## Read Book Deep Learning With R

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.

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

## Read Book Deep Learning With R

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-wranglers 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

## Read Book Deep Learning With R

you. Using practical examples, the book illustrates how to prepare data for analysis, choose a machine learning method, and measure its success.

[Copyright: 20a3b72378c4f0579980a49db0a68dfe](#)