
Machine Learning With R Expert Techniques For Pre

Machine Learning for Hackers
Machine Learning with R Cookbook
Learning Deep Architectures for AI
Machine Learning with R
Machine Learning for Email
The Elements of Statistical Learning
Mastering Machine Learning with R
The Big R-Book
Mathematics and Programming for Machine
Learning with R
Machine Learning Mastery With R
Advanced Deep Learning with R
Machine Learning Using R
R for Data Science
Mathematics for Machine Learning
Deep Learning for Coders with fastai and PyTorch
Mastering Machine Learning with R
Interpretable Machine Learning
Neural Networks with R
R Machine Learning By Example
Applied Predictive Modeling
Machine Learning Algorithms
Data Science in Education Using R
R in a Nutshell

Learning Quantitative Finance with R
Introduction to Machine Learning with R
Machine Learning with R, the tidyverse, and mlr
Learning Bayesian Models with R
Machine Learning with R
Machine Learning For Dummies
Machine Learning with R
Hands-On Ensemble Learning with R
Deep Learning with R
Machine Learning for Absolute Beginners
Practical Machine Learning in R
Hands-On Machine Learning with R
Machine Learning in Action
Practical Big Data Analytics
Supervised Machine Learning for Text Analysis in R
Understanding Machine Learning
Hands-On Deep Learning with R

*Machine
Learning
With R
Expert
Techniques
For Pre*

*Downloaded
from
worldimpex.com
by guest*

PORTER AUDRINA

Machine Learning for
Hackers John Wiley &
Sons

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 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.

Machine Learning with R Cookbook

Packt Publishing Ltd
One of Mark Cuban's top reads for better understanding A.I. (inc.com, 2021) Your comprehensive entry-

level guide to machine learning While machine learning expertise doesn't quite mean you can create your own Turing Test-proof android—as in the movie *Ex Machina*—it is a form of artificial intelligence and one of the most exciting technological means of identifying opportunities and solving problems fast and on a large scale. Anyone who masters the principles of machine learning is mastering a big part of our tech future and opening up incredible new directions in careers that include fraud detection, optimizing search results, serving real-time ads, credit-scoring, building accurate and sophisticated pricing models—and way, way

more. Unlike most machine learning books, the fully updated 2nd Edition of *Machine Learning For Dummies* doesn't assume you have years of experience using programming languages such as Python (R source is also included in a downloadable form with comments and explanations), but lets you in on the ground floor, covering the entry-level materials that will get you up and running building models you need to perform practical tasks. It takes a look at the underlying—and fascinating—math principles that power machine learning but also shows that you don't need to be a math whiz to build fun new tools and apply them to your work and

study. Understand the history of AI and machine learning Work with Python 3.8 and TensorFlow 2.x (and R as a download) Build and test your own models Use the latest datasets, rather than the worn out data found in other books Apply machine learning to real problems Whether you want to learn for college or to enhance your business or career performance, this friendly beginner's guide is your best introduction to machine learning, allowing you to become quickly confident using this amazing and fast-developing technology that's impacting lives for the better all over the world.

Learning Deep Architectures for AI
Packt Publishing Ltd

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 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
Machine Learning with
R Packt Publishing Ltd
 Explore powerful R
 packages to create
 predictive models

using ensemble
 methods Key Features
 Implement machine
 learning algorithms to
 build ensemble-
 efficient models
 Explore powerful R
 packages to create
 predictive models
 using ensemble
 methods Learn to build
 ensemble models on
 large datasets using a
 practical approach
 Book Description
 Ensemble techniques
 are used for combining
 two or more similar or
 dissimilar machine
 learning algorithms to
 create a stronger
 model. Such a model
 delivers superior
 prediction power and
 can give your datasets
 a boost in accuracy.
 Hands-On Ensemble
 Learning with R begins
 with the important
 statistical resampling
 methods. You will then
 walk through the

central trilogy of ensemble techniques – bagging, random forest, and boosting – then you'll learn how they can be used to provide greater accuracy on large datasets using popular R packages. You will learn how to combine model predictions using different machine learning algorithms to build ensemble models. In addition to this, you will explore how to improve the performance of your ensemble models. By the end of this book, you will have learned how machine learning algorithms can be combined to reduce common problems and build simple efficient ensemble models with the help of real-world examples. What you will learn Carry out an

essential review of re-sampling methods, bootstrap, and jackknife Explore the key ensemble methods: bagging, random forests, and boosting Use multiple algorithms to make strong predictive models Enjoy a comprehensive treatment of boosting methods Supplement methods with statistical tests, such as ROC Walk through data structures in classification, regression, survival, and time series data Use the supplied R code to implement ensemble methods Learn stacking method to combine heterogeneous machine learning models Who this book is for This book is for you if you are a data scientist or machine

learning developer who wants to implement machine learning techniques by building ensemble models with the power of R. You will learn how to combine different machine learning algorithms to perform efficient data processing. Basic knowledge of machine learning techniques and programming knowledge of R would be an added advantage.

Machine Learning for Email Simon and Schuster

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 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.

The Elements of Statistical Learning

Packt Publishing Ltd

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

Mastering Machine Learning with R

Cambridge University Press

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/liv-evideo/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 *The Big R-Book* John Wiley & Sons Summary Machine Learning in Action is unique book that blends the foundational theories of machine learning with the practical realities of building tools for everyday data analysis. You'll use the flexible Python programming language to build programs that implement algorithms for data classification, forecasting, recommendations, and higher-level features like summarization and simplification. About the Book A machine is said to learn when its performance improves

with experience. Learning requires algorithms and programs that capture data and ferret out the interesting or useful patterns. Once the specialized domain of analysts and mathematicians, machine learning is becoming a skill needed by many. Machine Learning in Action is a clearly written tutorial for developers. It avoids academic language and takes you straight to the techniques you'll use in your day-to-day work. Many (Python) examples present the core algorithms of statistical data processing, data analysis, and data visualization in code you can reuse. You'll understand the concepts and how they fit in with tactical tasks

like classification, forecasting, recommendations, and higher-level features like summarization and simplification. Readers need no prior experience with machine learning or statistical processing. Familiarity with Python is helpful. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book.

What's Inside A no-nonsense introduction

Examples showing common ML tasks

Everyday data analysis

Implementing classic algorithms like Apriori and Adaboos

Table of Contents PART 1

CLASSIFICATION

Machine learning basics

Classifying with k-Nearest Neighbors

Splitting datasets one

feature at a time:

decision trees

Classifying with probability theory:

naïve Bayes

Logistic regression

Support vector machines

Improving classification with the AdaBoost meta algorithm

PART 2

FORECASTING

NUMERIC VALUES WITH REGRESSION

Predicting numeric values:

regression

Tree-based regression

PART 3

UNSUPERVISED LEARNING

Grouping unlabeled items using k-means clustering

Association analysis with the Apriori algorithm

Efficiently finding frequent itemsets with FP-growth

PART 4

ADDITIONAL TOOLS

Using principal component analysis to simplify data

Simplifying data with the singular value

decomposition Big data and MapReduce
Mathematics and Programming for Machine Learning with R CRC Press
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 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.

Machine Learning

Mastery With R Simon and Schuster

Learn how to use R to turn raw data into insight, knowledge, and understanding. This book introduces you to R, RStudio, and the tidyverse, a

collection of R packages designed to work together to make data science fast, fluent, and fun.

Suitable for readers with no previous programming experience, *R for Data Science* is designed to get you doing data science as quickly as possible. Authors Hadley Wickham and Garrett Grolemund guide you through the steps of importing, wrangling, exploring, and modeling your data and communicating the results. You'll get a complete, big-picture understanding of the data science cycle, along with basic tools you need to manage the details. Each section of the book is paired with exercises to help you practice what you've learned

along the way. You'll learn how to:

Wrangle—transform your datasets into a form convenient for analysis

Program—learn powerful R tools for solving data problems with greater clarity and ease

Explore—examine your data, generate hypotheses, and quickly test them

Model—provide a low-dimensional summary that captures true "signals" in your dataset

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

Advanced Deep Learning with R CRC Press

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.
Machine Learning Using R Machine Learning Mastery
Build strong foundation

for entering the world of Machine Learning and data science with the help of this comprehensive guide
About This Book Get started in the field of Machine Learning with the help of this solid, concept-rich, yet highly practical guide. Your one-stop solution for everything that matters in mastering the whats and whys of Machine Learning algorithms and their implementation. Get a solid foundation for your entry into Machine Learning by strengthening your roots (algorithms) with this comprehensive guide. Who This Book Is For This book is for IT professionals who want to enter the field of data science and are very new to Machine Learning. Familiarity with languages such as

R and Python will be invaluable here. What You Will Learn Acquaint yourself with important elements of Machine Learning Understand the feature selection and feature engineering process Assess performance and error trade-offs for Linear Regression Build a data model and understand how it works by using different types of algorithm Learn to tune the parameters of Support Vector machines Implement clusters to a dataset Explore the concept of Natural Processing Language and Recommendation Systems Create a ML architecture from scratch. In Detail As the amount of data continues to grow at an almost incomprehensible rate,

being able to understand and process data is becoming a key differentiator for competitive organizations. Machine learning applications are everywhere, from self-driving cars, spam detection, document search, and trading strategies, to speech recognition. This makes machine learning well-suited to the present-day era of Big Data and Data Science. The main challenge is how to transform data into actionable knowledge. In this book you will learn all the important Machine Learning algorithms that are commonly used in the field of data science. These algorithms can be used for supervised as well as unsupervised learning,

reinforcement learning, and semi-supervised learning. A few famous algorithms that are covered in this book are Linear regression, Logistic Regression, SVM, Naive Bayes, K-Means, Random Forest, TensorFlow, and Feature engineering. In this book you will also learn how these algorithms work and their practical implementation to resolve your problems. This book will also introduce you to the Natural Processing Language and Recommendation systems, which help you run multiple algorithms simultaneously. On completion of the book you will have mastered selecting Machine Learning algorithms for clustering, classification, or

regression based on for your problem. Style and approach An easy-to-follow, step-by-step guide that will help you get to grips with real - world applications of Algorithms for Machine Learning.

R for Data Science Now Publishers Inc

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. *Mathematics for Machine Learning* Packt Publishing Ltd

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.

Deep Learning for Coders with fastai and PyTorch

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

Mastering Machine Learning with R
Lulu.com

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.
Interpretable Machine Learning "O'Reilly Media, Inc."
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.

Neural Networks with R Packt Publishing Ltd
Theoretical results suggest that in order to

learn the kind of complicated functions that can represent high-level abstractions (e.g. in vision, language, and other AI-level tasks), one may need deep architectures. Deep architectures are composed of multiple levels of non-linear operations, such as in neural nets with many hidden layers or in complicated propositional formulae re-using many sub-formulae. Searching the parameter space of deep architectures is a difficult task, but learning algorithms such as those for Deep Belief Networks have recently been proposed to tackle this problem with notable success, beating the state-of-the-art in certain areas. This paper discusses the

motivations and principles regarding learning algorithms for deep architectures, in particular those exploiting as building blocks unsupervised learning of single-layer models such as Restricted Boltzmann Machines, used to construct deeper models such as Deep Belief Networks.

R Machine Learning By Example Packt Publishing Ltd

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.

Applied Predictive 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.