

# Book R Deep Learning Essentials

[Deep Learning with Python](#)  
[Deep Learning](#)  
 Implement advanced deep learning and neural network models using TensorFlow and Keras  
[Machine Learning with R](#)  
[TensorFlow for Deep Learning](#)  
[with Applications in R](#)  
[Machine Learning Refined](#)  
[Deep Learning with R for Beginners](#)  
[Python Data Science Handbook](#)  
[Introduction to Machine Learning](#)  
[Deep Learning from Scratch](#)  
[R Machine Learning Essentials](#)  
[Hands-On Machine Learning with R](#)  
[Introduction to Data Science](#)  
[Machine Learning Essentials](#)  
 Your hands-on guide to the fundamentals of deep learning and neural network modeling  
[Essential Techniques for Predictive Analysis](#)  
[A Practitioner's Approach](#)  
[Import, Tidy, Transform, Visualize, and Model Data](#)  
[The Essentials of Data Science: Knowledge Discovery Using R](#)  
[Deep Learning with JavaScript](#)  
[Designing Next-Generation Machine Intelligence Algorithms](#)  
[Machine Learning in Python](#)  
[Deeper Learning](#)  
[F# for Machine Learning Essentials](#)  
[From Linear Regression to Reinforcement Learning](#)  
[R Deep Learning Projects](#)  
[Fundamentals of Machine Learning for Predictive Data Analytics, second edition](#)  
[Advanced Deep Learning with R](#)  
[Data Analysis and Prediction Algorithms with R](#)  
[Microsoft Azure Essentials Azure Machine Learning](#)  
[Building with Python from First Principles](#)  
[Deep Learning](#)  
[Foundations of Machine Learning, second edition](#)  
[Statistical Tools, Machine Learning, and R-Statistical Software Overview](#)  
[Neural networks in TensorFlow.js](#)  
[Deep Learning with R](#)  
[Algorithms, Worked Examples, and Case Studies](#)  
[R Deep Learning Essentials](#)

*Book R Deep Learning Essentials*

Downloaded from [ns1.galaxy.mu](#) by guest

## EVELIN MELTON

[Deep Learning with Python](#) "O'Reilly Media, Inc."

With the reinvigoration of neural networks in the 2000s, deep learning has become an extremely active area of research, one that's paving the way for modern machine learning. In this practical book, author Nikhil Buduma provides examples and clear explanations to guide you through major concepts of this complicated field. Companies such as Google, Microsoft, and Facebook are actively growing in-house deep-learning teams. For the rest of us, however, deep learning is still a pretty complex and difficult subject to grasp. If you're familiar with Python, and have a background in calculus, along with a basic understanding of machine learning, this book will get you started. Examine the foundations of machine learning and neural networks Learn how to train feed-forward neural networks Use TensorFlow to implement your first neural network Manage problems that arise as you begin to make networks deeper Build neural networks that analyze complex images Perform effective dimensionality reduction using autoencoders Dive deep into sequence analysis to examine language Learn the fundamentals of reinforcement learning

[Deep Learning](#) Packt Publishing Ltd

Get to grips with the essentials of deep learning by leveraging the power of Python Key Features Your one-stop solution to get started with the essentials of deep learning and neural network modeling Train different kinds of neural networks to tackle various problems in Natural Language Processing, computer vision, speech recognition, and more Covers popular Python libraries such as Tensorflow, Keras, and more, along with tips on training, deploying and optimizing your deep learning models in the best possible manner Book Description Deep Learning a trending topic in the field of Artificial Intelligence today and can be considered to be an advanced form of machine learning, which is quite tricky to master. This book will help you take your first steps in training efficient deep learning models and applying them in various practical scenarios. You will model, train, and deploy different kinds of neural networks such as Convolutional Neural Network, Recurrent Neural Network, and will see some of their applications in real-world domains including computer vision, natural language processing, speech recognition, and so on. You will build practical projects such as chatbots, implement reinforcement learning to build smart games, and develop expert systems for image captioning and processing. Popular Python library such as TensorFlow is used in this book to build the models. This book also covers solutions for different problems you might come across while training models, such as noisy datasets, small datasets, and more. This book does not assume any prior knowledge of deep learning. By the end of this book, you will have a firm understanding of the basics of deep learning and neural network modeling, along with their practical applications. What you will learn Get to grips with the core concepts of deep learning and neural networks Set up deep learning library such as TensorFlow Fine-tune your deep learning models for NLP and Computer Vision applications Unify different information sources, such as images, text, and speech through deep learning Optimize and fine-tune your deep learning models for better performance Train a deep reinforcement learning model that plays a game better than humans Learn how to make your models get the best out of your GPU or CPU Who this book is for Aspiring data scientists and machine learning experts who have limited or no exposure to deep learning will find this book to be very useful. If you are looking for a resource that gets you up and running with the fundamentals of deep learning and neural networks, this book is for you. As the models in the book are trained using the popular Python-based libraries such as Tensorflow and Keras, it would be useful to have sound programming knowledge of Python.

[Implement advanced deep learning and neural network models using TensorFlow and Keras](#) STHDA

Learn how to solve challenging machine learning problems with TensorFlow, Google's revolutionary new software library for deep learning. If you have some background in basic linear algebra and

calculus, this practical book introduces machine-learning fundamentals by showing you how to design systems capable of detecting objects in images, understanding text, analyzing video, and predicting the properties of potential medicines. TensorFlow for Deep Learning teaches concepts through practical examples and helps you build knowledge of deep learning foundations from the ground up. It's ideal for practicing developers with experience designing software systems, and useful for scientists and other professionals familiar with scripting but not necessarily with designing learning algorithms. Learn TensorFlow fundamentals, including how to perform basic computation Build simple learning systems to understand their mathematical foundations Dive into fully connected deep networks used in thousands of applications Turn prototypes into high-quality models with hyperparameter optimization Process images with convolutional neural networks Handle natural language datasets with recurrent neural networks Use reinforcement learning to solve games such as tic-tac-toe Train deep networks with hardware including GPUs and tensor processing units

[Machine Learning with R](#) Business Expert Press

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

[TensorFlow for Deep Learning](#) John Wiley & Sons

An introduction to a broad range of topics in deep learning, covering mathematical and conceptual background, deep learning techniques used in industry, and research perspectives. "Written by three experts in the field, Deep Learning is the only comprehensive book on the subject." —Elon Musk, cochair of OpenAI; cofounder and CEO of Tesla and SpaceX Deep learning is a form of machine learning that enables computers to learn from experience and understand the world in terms of a hierarchy of concepts. Because the computer gathers knowledge from experience, there is no need for a human computer operator to formally specify all the knowledge that the computer needs. The hierarchy of concepts allows the computer to learn complicated concepts by building them out of simpler ones; a graph of these hierarchies would be many layers deep. This book introduces a broad range of topics in deep learning. The text offers mathematical and conceptual background, covering relevant concepts in linear algebra, probability theory and information theory, numerical computation, and machine learning. It describes deep learning techniques used by practitioners in industry, including deep feedforward networks, regularization, optimization algorithms, convolutional networks, sequence modeling, and practical methodology; and it surveys such applications as natural language processing, speech recognition, computer vision, online recommendation systems, bioinformatics, and videogames. Finally, the book offers research perspectives, covering such theoretical topics as linear factor models, autoencoders, representation learning, structured probabilistic models, Monte Carlo methods, the partition function, approximate inference, and deep generative models. Deep Learning can be used by undergraduate or graduate students planning careers in either industry or research, and by software engineers who want to begin using deep learning in their products or platforms. A website offers supplementary material for both readers and instructors.

*with Applications in R* Simon and Schuster

An Introduction to Statistical Learning provides an accessible overview of the field of statistical learning, an essential toolset for making sense of the vast and complex data sets that have emerged in fields ranging from biology to finance to marketing to astrophysics in the past twenty years. This book presents some of the most important modeling and prediction techniques, along with relevant applications. Topics include linear regression, classification, resampling methods, shrinkage approaches, tree-based methods, support vector machines, clustering, and more. Color graphics and real-world examples are used to illustrate the methods presented. Since the goal of this textbook is

to facilitate the use of these statistical learning techniques by practitioners in science, industry, and other fields, each chapter contains a tutorial on implementing the analyses and methods presented in R, an extremely popular open source statistical software platform. Two of the authors co-wrote *The Elements of Statistical Learning* (Hastie, Tibshirani and Friedman, 2nd edition 2009), a popular reference book for statistics and machine learning researchers. An *Introduction to Statistical Learning* covers many of the same topics, but at a level accessible to a much broader audience. This book is targeted at statisticians and non-statisticians alike who wish to use cutting-edge statistical learning techniques to analyze their data. The text assumes only a previous course in linear regression and no knowledge of matrix algebra.

[Machine Learning Refined](#) Packt Publishing Ltd

*Introduction to Data Science: Data Analysis and Prediction Algorithms with R* introduces concepts and skills that can help you tackle real-world data analysis challenges. It covers concepts from probability, statistical inference, linear regression, and machine learning. It also helps you develop skills such as R programming, data wrangling, data visualization, predictive algorithm building, file organization with UNIX/Linux shell, version control with Git and GitHub, and reproducible document preparation. This book is a textbook for a first course in data science. No previous knowledge of R is necessary, although some experience with programming may be helpful. The book is divided into six parts: R, data visualization, statistics with R, data wrangling, machine learning, and productivity tools. Each part has several chapters meant to be presented as one lecture. The author uses motivating case studies that realistically mimic a data scientist's experience. He starts by asking specific questions and answers these through data analysis so concepts are learned as a means to answering the questions. Examples of the case studies included are: US murder rates by state, self-reported student heights, trends in world health and economics, the impact of vaccines on infectious disease rates, the financial crisis of 2007-2008, election forecasting, building a baseball team, image processing of hand-written digits, and movie recommendation systems. The statistical concepts used to answer the case study questions are only briefly introduced, so complementing with a probability and statistics textbook is highly recommended for in-depth understanding of these concepts. If you read and understand the chapters and complete the exercises, you will be prepared to learn the more advanced concepts and skills needed to become an expert.

*Deep Learning with R for Beginners* CRC Press

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

*Python Data Science Handbook* John Wiley & Sons

Summary Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning has made remarkable progress in recent years. We went from near-unusable speech and image recognition, to near-human accuracy. We went from machines that couldn't beat a serious Go player, to defeating a world champion. Behind this progress is deep learning—a combination of engineering advances, best practices, and theory that enables a wealth of previously impossible smart applications. About the Book Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. You'll explore challenging concepts and practice with applications in computer vision, natural-language processing, and generative models. By the time you finish, you'll have the knowledge and hands-on skills to apply deep learning in your own projects. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image-classification models Deep learning for text and sequences Neural style transfer, text generation, and image generation About the Reader Readers need intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the Author François Chollet works on deep learning at Google in Mountain View, CA. He is the creator of the Keras deep-learning library, as well as a contributor to the TensorFlow machine-learning framework. He also does deep-learning research, with a focus on computer vision and the application of machine learning to formal reasoning. His papers have been published at major conferences in the field, including the Conference on Computer Vision and Pattern Recognition (CVPR), the Conference and Workshop on Neural Information Processing Systems (NIPS), the International Conference on Learning Representations (ICLR), and others. Table of Contents PART 1 - FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of neural networks Getting started with neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE Deep learning for computer vision Deep learning for text and sequences Advanced deep-learning best practices Generative deep learning Conclusions appendix A - Installing Keras and its dependencies on Ubuntu appendix B - Running Jupyter notebooks on an EC2 GPU instance

[Introduction to Machine Learning](#) Cambridge University Press

[R Deep Learning Essentials](#)

[Deep Learning from Scratch](#) O'Reilly Media

Although interest in machine learning has reached a high point, lofty expectations often scuttle projects before they get very far. How can machine learning—especially deep neural networks—make a real difference in your organization? This hands-on guide not only provides the most practical information available on the subject, but also helps you get started building efficient deep learning networks. Authors Adam Gibson and Josh Patterson provide theory on deep learning before introducing their open-source *Deeplearning4j* (DL4J) library for developing production-class workflows. Through real-world examples, you'll learn methods and strategies for training deep network architectures and running deep learning workflows on Spark and Hadoop with DL4J. Dive into machine learning concepts in general, as well as deep learning in particular Understand how deep networks evolved from neural network fundamentals Explore the major deep network architectures, including Convolutional and Recurrent Learn how to map specific deep networks to the right problem Walk through the fundamentals of tuning general neural networks and specific deep network architectures Use vectorization techniques for different data types with DataVec, DL4J's workflow tool Learn how to use DL4J natively on Spark and Hadoop

[R Machine Learning Essentials](#) MIT Press

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

**Hands-On Machine Learning with R** Springer Science & Business Media

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

*Introduction to Data Science* MIT Press

Get up and running with machine learning with F# in a fun and functional way About This Book Design algorithms in F# to tackle complex computing problems Be a proficient F# data scientist using this simple-to-follow guide Solve real-world, data-related problems with robust statistical models, built for a range of datasets Who This Book Is For If you are a C# or an F# developer who now wants to explore the area of machine learning, then this book is for you. Familiarity with theoretical concepts and notation of mathematics and statistics would be an added advantage. What You Will Learn Use F# to find patterns through raw data Build a set of classification systems using Accord.NET, Weka, and F# Run machine learning jobs on the Cloud with MBrace Perform mathematical operations on matrices and vectors using Math.NET Use a recommender system for your own problem domain Identify tourist spots across the globe using inputs from the user with decision tree algorithms In Detail The F# functional programming language enables developers to write simple code to solve complex problems. With F#, developers create consistent and predictable programs that are easier to test and reuse, simpler to parallelize, and are less prone to bugs. If you want to learn how to use F# to build machine learning systems, then this is the book you want. Starting with an introduction to the several categories on machine learning, you will quickly learn to implement time-tested, supervised learning algorithms. You will gradually move on to solving problems on predicting housing pricing using Regression Analysis. You will then learn to use Accord.NET to implement SVM techniques and clustering. You will also learn to build a recommender system for your e-commerce site from scratch. Finally, you will dive into advanced topics such as implementing neural network algorithms while performing sentiment analysis on your data. Style and approach This book is a fast-paced tutorial guide that uses hands-on examples to explain real-world applications of machine learning. Using practical examples, the book will explore several machine learning techniques and also describe how you can use F# to build machine learning systems.

MIT Press

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

*Machine Learning Essentials* "O'Reilly Media, Inc."

The *Essentials of Data Science: Knowledge Discovery Using R* presents the concepts of data science through a hands-on approach using free and open source software. It systematically drives an accessible journey through data analysis and machine learning to discover and share knowledge

from data. Building on over thirty years' experience in teaching and practising data science, the author encourages a programming-by-example approach to ensure students and practitioners attune to the practise of data science while building their data skills. Proven frameworks are provided as reusable templates. Real world case studies then provide insight for the data scientist to swiftly adapt the templates to new tasks and datasets. The book begins by introducing data science. It then reviews R's capabilities for analysing data by writing computer programs. These programs are developed and explained step by step. From analysing and visualising data, the framework moves on to tried and tested machine learning techniques for predictive modelling and knowledge discovery. Literate programming and a consistent style are a focus throughout the book.

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

An accessible introduction to the artificial intelligence technology that enables computer vision, speech recognition, machine translation, and driverless cars. Deep learning is an artificial intelligence technology that enables computer vision, speech recognition in mobile phones, machine translation, AI games, driverless cars, and other applications. When we use consumer products from Google, Microsoft, Facebook, Apple, or Baidu, we are often interacting with a deep learning system. In this volume in the MIT Press Essential Knowledge series, computer scientist John Kelleher offers an accessible and concise but comprehensive introduction to the fundamental technology at the heart of the artificial intelligence revolution. Kelleher explains that deep learning enables data-driven decisions by identifying and extracting patterns from large datasets; its ability to learn from complex data makes deep learning ideally suited to take advantage of the rapid growth in big data and computational power. Kelleher also explains some of the basic concepts in deep learning, presents a history of advances in the field, and discusses the current state of the art. He describes the most important deep learning architectures, including autoencoders, recurrent neural networks, and long short-term networks, as well as such recent developments as Generative Adversarial Networks and capsule networks. He also provides a comprehensive (and comprehensible) introduction to the two fundamental algorithms in deep learning: gradient descent and backpropagation. Finally, Kelleher considers the future of deep learning—major trends, possible developments, and significant challenges.

*Essential Techniques for Predictive Analysis* Packt Publishing Ltd

Hands-on Machine Learning with R provides a practical and applied approach to learning and developing intuition into today's most popular machine learning methods. This book serves as a practitioner's guide to the machine learning process and is meant to help the reader learn to apply the machine learning stack within R, which includes using various R packages such as glmnet, h2o, ranger, xgboost, keras, and others to effectively model and gain insight from their data. The book favors a hands-on approach, providing an intuitive understanding of machine learning concepts through concrete examples and just a little bit of theory. Throughout this book, the reader will be exposed to the entire machine learning process including feature engineering, resampling, hyperparameter tuning, model evaluation, and interpretation. The reader will be exposed to powerful algorithms such as regularized regression, random forests, gradient boosting machines, deep learning, generalized low rank models, and more! By favoring a hands-on approach and using real 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.

*A Practitioner's Approach* Simon and Schuster

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.

*Import, Tidy, Transform, Visualize, and Model Data* CRC Press

The goal of machine learning is to program computers to use example data or past experience to solve a given problem. Many successful applications of machine learning exist already, including systems that analyze past sales data to predict customer behavior, optimize robot behavior so that a task can be completed using minimum resources, and extract knowledge from bioinformatics data. Introduction to Machine Learning is a comprehensive textbook on the subject, covering a broad array of topics not usually included in introductory machine learning texts. Subjects include supervised learning; Bayesian decision theory; parametric, semi-parametric, and nonparametric methods; multivariate analysis; hidden Markov models; reinforcement learning; kernel machines; graphical models; Bayesian estimation; and statistical testing. Machine learning is rapidly becoming a skill that computer science students must master before graduation. The third edition of Introduction to Machine Learning reflects this shift, with added support for beginners, including selected solutions for exercises and additional example data sets (with code available online). Other substantial changes include discussions of outlier detection; ranking algorithms for perceptrons and support vector machines; matrix decomposition and spectral methods; distance estimation; new kernel algorithms; deep learning in multilayered perceptrons; and the nonparametric approach to Bayesian methods. All learning algorithms are explained so that students can easily move from the equations in the book to a computer program. The book can be used by both advanced undergraduates and graduate students. It will also be of interest to professionals who are concerned with the application of machine learning methods.