

Classification And Regression Trees By Leo Breiman

Data Mining and Knowledge Discovery Handbook organizes all major concepts, theories, methodologies, trends, challenges and applications of data mining (DM) and knowledge discovery in databases (KDD) into a coherent and unified repository. This book first surveys, then provides comprehensive yet concise algorithmic descriptions of methods, including classic methods plus the extensions and novel methods developed recently. This volume concludes with in-depth descriptions of data mining applications in various interdisciplinary industries including finance, marketing, medicine, biology, engineering, telecommunications, software, and security. Data Mining and Knowledge Discovery Handbook is designed for research scientists and graduate-level students in computer science and engineering. This book is also suitable for professionals in fields such as computing applications, information systems management, and strategic research management.

The Complete Beginner's Guide to Understanding and Building Machine Learning Systems with Python Machine Learning with Python for Everyone will help you master the processes, patterns, and strategies you need to build effective learning systems, even if you're an absolute beginner. If you can write some Python code, this book is for you, no matter how little college-level math you know. Principal instructor Mark E. Fenner relies on plain-English stories, pictures, and Python examples to communicate the ideas of machine learning. Mark begins by discussing machine learning and what it can do; introducing key mathematical and computational topics in an approachable manner; and walking you through the first steps in building, training, and evaluating learning systems. Step by step, you'll fill out the components of a practical learning system, broaden your toolbox, and explore some of the field's most sophisticated and exciting techniques. Whether you're a student, analyst, scientist, or hobbyist, this guide's insights will be applicable to every learning system you ever build or use. Understand machine learning algorithms, models, and core machine learning concepts Classify examples with classifiers, and quantify examples with regressors Realistically assess performance of machine learning systems Use feature engineering to smooth rough data into useful forms Chain multiple components into one system and tune its performance Apply machine learning techniques to images and text Connect the core concepts to neural networks and graphical models Leverage the Python scikit-learn library and other powerful tools Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details. Leverage the power of machine learning and Swift programming to build intelligent iOS applications with ease Key Features Implement effective machine learning solutions for your iOS applications Use Swift and Core ML to build and deploy popular machine learning models Develop neural networks for natural language processing and computer vision Book Description Machine learning as a field promises to bring increased intelligence to the software by helping us learn and analyse information efficiently and discover certain patterns that humans cannot. This book will be your guide as you embark on an exciting journey in machine learning using the popular Swift language. We'll start with machine learning basics in the first part of the book to develop a lasting intuition about fundamental machine learning concepts. We explore various supervised and unsupervised statistical learning techniques and how to implement them in Swift, while the third section walks you through deep learning techniques with the help of typical real-world cases. In the last section, we will dive into some hard core topics such as model compression, GPU acceleration and provide some recommendations to avoid common mistakes during machine learning application development. By the end of the book, you'll be able to develop intelligent applications written in Swift that can learn for themselves. What you will learn Learn rapid model prototyping with Python and Swift Deploy pre-trained models to iOS using Core ML Find hidden patterns in the data using unsupervised learning Get a deeper understanding of the clustering techniques Learn modern compact architectures of neural networks for iOS devices Train neural networks for image processing and natural language processing Who this book is for iOS developers who wish to create smarter iOS applications using the power of machine learning will find this book to be useful. This book will also benefit data science professionals who are interested in performing machine learning on mobile devices. Familiarity with Swift programming is all you need to get started with this book.

A Visual Introduction for Beginners

Predictive Accuracy of Classification and Regression Trees (CART) Versus Conjoint Analysis

Discrete Data Analysis with R

Effective strategies to manage data science projects and build a sustainable team

Essential Tools for Working with Data

Machine Learning with Python for Everyone

International Federation of Classification Societies *The International Federation of Classification Societies (IFCS) is an agency for the dissemination of technical and scientific information concerning classification and multivariate data analysis in the broad sense and in as wide a range of applications as possible; founded in 1985 in Cambridge (UK) by the following Scientific Societies and Groups: - British Classification Society - BCS - Classification Society of North America - CSNA - Gesellschaft für Klassifikation - GfKI - Japanese Classification Society - JCS - Classification Group of Italian Statistical Society - CGSIS - Societe Francophone de Classification - SFC Now the IFCS includes also the following Societies: - Dutch-Belgian Classification Society - VOC - Polish Classification Section - SKAD - Portuguese Classification Association - CLAD - Group at Large - Korean Classification Society - KCS IFCS-98, the Sixth Conference of the International Federation of Classification Societies, was held in Rome, from July 21 to 24, 1998. Five preceding conferences were held in Aachen (Germany), Charlottesville (USA), Edinburgh (UK), Paris (France), Kobe (Japan).*

An Applied Treatment of Modern Graphical Methods for Analyzing Categorical Data *Discrete Data Analysis with R: Visualization and Modeling Techniques for Categorical and Count Data presents an applied treatment of modern methods for the analysis of categorical data, both discrete response data and frequency data. It explains how to use graphical meth*

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.*

Springer Handbook of Engineering Statistics

Tree-based Machine Learning Algorithms

Data Mining and Knowledge Discovery Handbook

Nature Inspired Computing

Evolutionary Decision Trees in Large-Scale Data Mining

Machine Learning Algorithms From Scratch with Python

Discovering knowledge from big multivariate data, recorded every days, requires specialized machine learning techniques. This book presents an easy to use practical guide in R to compute the most popular machine learning methods for exploring real word data sets, as well as, for building predictive models. The main parts of the book include: A) Unsupervised learning methods, to explore and discover knowledge from a large multivariate data set using clustering and principal component methods. You will learn hierarchical clustering, k-means, principal component analysis and correspondence analysis methods. B) Regression analysis, to predict a quantitative outcome value using linear regression and non-linear regression strategies. C) Classification techniques, to predict a qualitative outcome value using logistic regression, discriminant analysis, naive bayes classifier and support vector machines. D) Advanced machine learning methods, to build robust regression and classification models using k-nearest neighbors methods, decision tree models, ensemble methods (bagging, random forest and boosting). E) Model selection methods, to select automatically the best combination of predictor variables for building an optimal predictive model. These include, best subsets selection methods, stepwise regression and penalized regression (ridge, lasso and elastic net regression models). We also present principal component-based regression methods, which are useful when the data contain multiple correlated predictor variables. F) Model validation and evaluation techniques for measuring the performance of a predictive model. G) Model diagnostics for detecting and fixing a potential problems in a predictive model. The book presents the basic principles of these tasks and provide many examples in R. This book offers solid guidance in data mining for students and researchers. Key features: - Covers machine learning algorithm and implementation - Key mathematical concepts are presented - Short, self-contained chapters with practical examples.

This book presents a unified framework, based on specialized evolutionary algorithms, for the global induction of various types of classification and regression trees from data. The resulting univariate or oblique trees are significantly smaller than those produced by standard top-down methods, an aspect that is critical for the interpretation of mined patterns by domain analysts. The approach presented here is extremely flexible and can easily be adapted to specific data mining applications, e.g. cost-sensitive model trees for financial data or multi-test trees for gene expression data. The global induction can be efficiently applied to large-scale data without the need for extraordinary resources. With a simple GPU-based acceleration, datasets composed of millions of instances can be mined in minutes. In the event that the size of the datasets makes the fastest memory computing impossible, the Spark-based implementation on computer clusters, which offers impressive fault tolerance and scalability potential, can be applied.

For many researchers, Python is a first-class tool mainly because of its libraries for storing, manipulating, and gaining insight from data. Several resources exist for individual pieces of this data science stack, but only with the Python Data Science Handbook do you get them all—IPython, NumPy, Pandas, Matplotlib, Scikit-Learn, and other related tools. Working scientists and data crunchers familiar with reading and writing Python code will find this comprehensive desk reference ideal for tackling day-to-day issues: manipulating, transforming, and cleaning data; visualizing different types of data; and using data to build statistical or machine learning models. Quite simply, this is the must-have reference for scientific computing in Python. With this handbook, you'll learn how to use: IPython and Jupyter: provide computational environments for data scientists using Python NumPy: includes the ndarray for efficient storage and manipulation of dense data arrays in Python Pandas: features the DataFrame for efficient storage and manipulation of labeled/columnar data in Python Matplotlib: includes capabilities for a flexible range of data visualizations in Python Scikit-Learn: for efficient and clean Python implementations of the most important and established machine learning algorithms

Interpretable Machine Learning

Using Classification and Regression Trees

Visualization and Modeling Techniques for Categorical and Count Data

A User Manual for Identifying Indicators of Vulnerability to Famine and Chronic Food Insecurity

Decision Trees and Random Forests

Delve into decision trees, which are graphs that use a branching method to determine all possible outcomes of a decision. Trees for continuous outcomes are called regression trees, while those for categorical outcomes are called classification trees. Learn how and when to use each, producing inferences that are easily understood by non-statisticians.

Classification and regression trees (CART) is one of the several contemporary statistical techniques with good promise for research in many academic fields. There are very few books on CART, especially on applied CART. This book, as a good practical primer with a focus on applications, introduces the relatively new statistical technique of CART as a powerful analytical tool. The easy-to-understand (non-technical) language and illustrative graphs (tables) as well as the use of the popular statistical software program (SPSS) appeal to readers without strong statistical background. This book helps readers understand the foundation, the operation, and the interpretation of CART analysis, thus becoming knowledgeable consumers and skillful users of CART. The chapter on advanced CART procedures not yet well-discussed in the literature allows readers to effectively seek further empowerment of their research designs by extending the analytical power of CART to a whole new level. This highly practical book is specifically written for academic researchers, data analysts, and graduate students in many disciplines such as economics, social sciences, medical sciences, and sport sciences who do not have strong statistical background but still strive to take full advantage of CART as a powerful analytical tool for research in their fields.

This dissertation consists of two parts. In the first part, we introduce the algorithm Multinomial Logistic Regression Tree. Logistic regression tree recursively partitions the data and ts a logistic regression at each partition. It combines logistic regression and tree model. The tree structure of the model can handle the nonlinear features automatically and the node logistic regression model can provide prediction of response class probabilities. Previous logistic regression algorithms are designed for binary response data only. Here we extend the model to multinomial response data. Our algorithm also supports exhaustive search for numerical cut point selection. In the second part, we compare the prediction accuracy of nine popular regression algorithms on data sets with missing values. To handle the missing data, we consider two approaches, the default method of the algorithm and missing data imputation. At low missing rate, regression tree GUIDE and M5 achieve the best and at high missing rate, tree ensemble GUIDE Forest and Random Forest have the best performance. A new method of GUIDE imputation is the best imputation method for most of the regression algorithms in our experiment.

Learning Statistics

Practical Solutions from Preprocessing to Deep Learning

Theory and Applications

Hands-On Machine Learning with R

Decision Trees and Their Application for Classification and Regression Problems

Classification and Regression Trees, CART

This is the first comprehensive book dedicated entirely to the field of decision trees in data mining and covers all aspects of this important technique. Decision trees have become one of the most powerful and popular approaches in knowledge discovery and data mining, the science and technology of exploring large and complex bodies of data in order to discover useful patterns. The area is of great importance because it enables modeling and knowledge extraction from the abundance of data available.

Both theoreticians and practitioners are continually seeking techniques to make the process more efficient, cost-effective and accurate. Decision trees, originally implemented in decision theory and statistics, are highly effective tools in other areas such as data mining, text mining, information extraction, machine learning, and pattern recognition. This book invites readers to explore the many benefits in data mining that decision trees offer: Self-explanatory and easy to follow when compacted Able to handle a variety of input data: nominal, numeric and textual Able to process datasets that may have errors or missing values High predictive performance for a relatively small computational effort Available in many data mining packages over a variety of platforms Useful for various tasks, such as classification, regression, clustering and feature selection

Presents a unified, efficient model of random decision forests which can be used in a number of applications such as scene recognition from photographs, object recognition in images, automatic diagnosis from radiological scans and document analysis.

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

A Practical Primer

Decision Trees, Random Forests, and Boosting

Episode 19: Regression Trees and Classification Trees

Practical Guide in R

Making Sense of Data

Artificial Intelligence for iOS

In this practical book, four Cloudera data scientists present a set of self-contained patterns for performing large-scale data analysis with Spark. The authors bring Spark, statistical methods, and real-world data sets together to teach you how to approach analytics problems by example. You'll start with an introduction to Spark and its ecosystem, and then dive into patterns that apply common techniques—classification, collaborative filtering, and anomaly detection among others—to fields such as genomics, security, and finance. If you have an entry-level understanding of machine learning and statistics, and you program in Java, Python, or Scala, you'll find these patterns useful for working on your own data applications. Patterns include: Recommending music and the Audioscrobbler data set Predicting forest cover with decision trees Anomaly detection in network traffic with K-means clustering Understanding Wikipedia with Latent Semantic Analysis Analyzing co-occurrence networks with GraphX Geospatial and temporal data analysis on the New York City Taxi Trips data Estimating financial risk through Monte Carlo simulation Analyzing genomics data and the BDG project Analyzing neuroimaging data with PySpark and Thunder

Missing data pose challenges to real-life data analysis. Simple ad-hoc fixes, like deletion or mean imputation, only work under highly restrictive conditions, which are often not met in practice. Multiple imputation replaces each missing value by multiple plausible values. The variability between these replacements reflects our ignorance of the true (but missing) value. Each of the completed data set is then analyzed by standard methods, and the results are pooled to obtain unbiased estimates with correct confidence intervals. Multiple imputation is a general approach that also inspires novel solutions to old problems by reformulating the task at hand as a missing-data problem. This is the second edition of a popular book on multiple imputation, focused on explaining the application of methods through

detailed worked examples using the MICE package as developed by the author. This new edition incorporates the recent developments in this fast-moving field. This class-tested book avoids mathematical and technical details as much as possible: formulas are accompanied by verbal statements that explain the formula in accessible terms. The book sharpens the reader's intuition on how to think about missing data, and provides all the tools needed to execute a well-grounded quantitative analysis in the presence of missing data.

This is the first comprehensive book dedicated entirely to the field of decision trees in data mining and covers all aspects of this important technique. Decision trees have become one of the most powerful and popular approaches in knowledge discovery and data mining, the science and technology of exploring large and complex bodies of data in order to discover useful patterns. The area is of great importance because it enables modeling and knowledge extraction from the abundance of data available. Both theoreticians and practitioners are continually seeking techniques to make the process more efficient, cost-effective and accurate. Decision trees, originally implemented in decision theory and statistics, are highly effective tools in other areas such as data mining, text mining, information extraction, machine learning, and pattern recognition. This book invites readers to explore the many benefits in data mining that decision trees offer:: Self-explanatory and easy to follow when compacted; Able to handle a variety of input data: nominal, numeric and textual; Able to process datasets that may have errors or missing values; High predictive performance for a relatively small computational effort; Available in many data mining packages over a variety of platforms; Useful for various tasks, such as classification, regression, clustering and feature selection . Sample Chapter(s). Chapter 1: Introduction to Decision Trees (245 KB). Chapter 6: Advanced Decision Trees (409 KB). Chapter 10: Fuzzy Decision Trees (220 KB). Contents: Introduction to Decision Trees; Growing Decision Trees; Evaluation of Classification Trees; Splitting Criteria; Pruning Trees; Advanced Decision Trees; Decision Forests; Incremental Learning of Decision Trees; Feature Selection; Fuzzy Decision Trees; Hybridization of Decision Trees with Other Techniques; Sequence Classification Using Decision Trees. Readership: Researchers, graduate and undergraduate students in information systems, engineering, computer science, statistics and management.

Python Data Science Handbook

Trading Based on Classification and Regression Trees

Patterns for Learning from Data at Scale

Proceedings of a Workshop held in Heidelberg, April 2-4, 1979

Machine Learning with Python Cookbook

Flexible Imputation of Missing Data, Second Edition

The methodology used to construct tree structured rules is the focus of this monograph. Unlike many other statistical procedures, which moved from pencil and paper to calculators, this text's use of trees was unthinkable before computers. Both the practical and theoretical sides have been developed in the authors' study of tree methods. Classification and Regression Trees reflects these two sides, covering the use of trees as a data analysis method, and in a more mathematical framework, proving some of their fundamental properties.

In today's global and highly competitive environment, continuous improvement in the processes and products of any field of engineering is essential for survival. This book gathers together the full range of statistical techniques required by engineers from all fields. It will assist them to gain sensible statistical feedback on how their processes or products are functioning and to give them realistic predictions of how these could be improved. The handbook will be essential reading for all engineers and engineering-connected managers who are serious about keeping their methods and products at the cutting edge of quality and competitiveness.

As the first book devoted to relational data mining, this coherently written multi-author monograph provides a thorough introduction and systematic overview of the area. The first part introduces the reader to the basics and principles of classical knowledge discovery in databases and inductive logic programming; subsequent chapters by leading experts assess the techniques in relational data mining in a principled and comprehensive way; finally, three chapters deal with advanced applications in various fields and refer the reader to resources for relational data mining. This book will become a valuable source of reference for R&D professionals active in relational data mining. Students as well as IT professionals and ambitious practitioners interested in learning about relational data mining will appreciate the book as a useful text and gentle introduction to this exciting new field.

Managing Data Science

Data Mining with Decision Trees

Meta-CART

On Using Classification and Regression Trees in Machine Learning

Proceedings of the 6th Conference of the International Federation of Classification Societies (IFCS-98) Università "La Sapienza", Rome, 21-24 July, 1998

Contributions to Classification and Regression Trees

Tree methods are some of the best and most commonly used methods in the field of statistical learning. They are widely used in classification and regression modeling. This thesis introduces the concept and focuses more on decision trees such as Classification and Regression Trees (CART) used for classification and regression predictive modeling problems. We also introduced some ensemble methods such as bagging, random forest and boosting. These methods were introduced to improve the performance and accuracy of the models constructed by classification and regression tree models. This work also provides an in-depth understanding of how the CART models are constructed, the algorithm behind the construction and also using cost-complexity approaching in tree pruning for regression trees and classification error rate approach used for pruning classification trees. We took two real-life examples, which we used to solve classification problem such as classifying the type of cancer based on tumor type, size and other parameters present in the dataset and regression problem such as predicting the first year GPA of a college student based on high school GPA, SAT scores and other parameters present in the dataset.

This book presents machine learning models and algorithms to address big data classification problems. Existing machine learning techniques like the decision tree (a hierarchical approach), random forest (an ensemble hierarchical approach), and deep learning (a layered approach) are highly suitable for the system that can handle such problems. This book helps readers, especially students and newcomers to the field of big data and machine learning, to gain a quick understanding of the techniques and technologies; therefore, the theory, examples, and programs (Matlab and R) presented in this book have been simplified, hardcoded, repeated, or spaced for improvements. They provide vehicles to test and understand the complicated concepts of various topics in the field. It is expected that the readers adopt these programs to experiment with the examples, and then modify or write their own programs toward advancing their knowledge for solving more complex and challenging problems. The presentation format of this book focuses on simplicity, readability, and dependability so that both undergraduate and graduate students as well as new researchers, developers, and practitioners in this field can easily trust and grasp the concepts, and learn them effectively. It has been written to reduce the mathematical complexity and help the vast majority of readers to understand the topics and get interested in the field. This book consists of four parts, with the total of 14 chapters. The first part mainly focuses on the topics that are needed to help analyze and understand data and big data. The second part covers the topics that can explain the systems required for processing big data. The third part presents the topics required to understand and select machine learning techniques to classify big data. Finally, the fourth part concentrates on the topics that explain the scaling-up machine learning, an important solution for modern big data problems.

This volume comprises the select proceedings of the annual convention of the Computer Society of India. Divided into 10 topical volumes, the proceedings present papers on state-of-the-art research, surveys, and succinct reviews. The volumes cover diverse topics ranging from communications networks to big data analytics, and from system architecture to cyber security. This volume focuses on Nature Inspired Computing. The contents of this book will be useful to researchers and students alike.

Machine Learning Models and Algorithms for Big Data Classification

Advances in Data Science and Classification

Growing Classification and Regression Trees on Network Data

Applied Predictive Modeling

Smoothing Techniques for Curve Estimation

The Integration of Meta-analysis and Classification & Regression Trees

You must understand algorithms to get good at machine learning. The problem is that they are only ever explained using Math. No longer. In this Ebook, finally cut through the math and learn exactly how machine learning algorithms work. Using clear explanations, simple pure Python code (no libraries!) and step-by-step tutorials you will discover how to load and prepare data, evaluate model skill, and implement a suite of linear, nonlinear and ensemble machine learning algorithms from scratch.

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.

Classification and Regression TreesRoutledge

Machine Learning with Swift

Classification and Regression Trees

A Practical Guide to Exploratory Data Analysis and Data Mining

Classification and Regression Trees Regression Trees

Machine Learning Essentials

A Unified Framework for Classification, Regression, Density Estimation, Manifold, Learning and Semi-Supervised Learning

If you want to learn how decision trees and random forests work, plus create your own, this visual book is for you. The fact is, decision tree and random forest algorithms are powerful and likely touch your life everyday. From online search to product development and credit scoring, both types of algorithms are at work behind the scenes in many modern applications and services. They are also used in countless industries such as medicine, manufacturing and finance to help companies make better decisions and reduce risk. Whether coded or scratched out by hand, both algorithms are powerful tools that can make a significant impact. This book is a visual introduction for beginners that unpacks the fundamentals of decision trees and random forests. If you want to dig into the basics with a visual twist plus create your own algorithms in Python, this book is for you.

Understand data science concepts and methodologies to manage and deliver top-notch solutions for your organization Key Features Learn the basics of data science and explore its possibilities and limitations Manage data science projects and assemble teams effectively even in the most challenging situations Understand management principles and approaches for data science projects to streamline the innovation process Book Description Data science and machine learning can transform any organization and unlock new opportunities. However, employing the right management strategies is crucial to guide the solution from prototype to production. Traditional approaches often fail as they don't entirely meet the conditions and requirements necessary for current data science projects. In this book, you'll explore the right approach to data science project management, along with useful tips and best practices to guide you along the way. After understanding the practical applications of data science and artificial intelligence, you'll see how to incorporate them into your solutions. Next, you will go through the data science project life cycle, explore the common pitfalls encountered at each step, and learn how to avoid them. Any data science project requires a skilled team, and this book will offer the right advice for hiring and growing a data science team for your organization. Later, you'll be shown how to efficiently manage and improve your data science projects through the use of DevOps and ModelOps. By the end of this book, you will be well versed with various data science solutions and have gained practical insights into tackling the different challenges that you'll encounter on a daily basis. What you will learn Understand the underlying problems of building a strong data science pipeline Explore the different tools for building and deploying data science solutions Hire, grow, and sustain a data science team Manage data science projects through all stages, from prototype to production Learn how to use ModelOps to improve your data science pipelines Get up to speed with the model testing techniques used in both development and production stages Who this book is for This book is for data scientists, analysts, and program managers who want to use data science for business productivity by incorporating data science workflows efficiently. Some understanding of basic data science concepts will be useful to get the most out of this book.

"Learn how to use decision trees and random forests for classification and regression, their respective limitations, and how the algorithms that build them work. Each chapter introduces a new data concern and then walks you through modifying the code, thus building the engine just-in-time. Along the way you will gain experience making decision trees and random forests work for you."--Back cover.

Thinking with Examples for Effective Learning

Proceedings of CSI 2015

Advanced Analytics with Spark

Relational Data Mining

Decision Forests