1. Record Nr. UNINA9910300642003321

Autore	Brandewinder Mathias
Titolo	Machine Learning Projects for .NET Developers / / by Mathias Brandewinder
Pubbl/distr/stampa	Berkeley, CA : , : Apress : , : Imprint : Apress, , 2015
ISBN	9781430267669 1430267666
Edizione	[1st ed. 2015.]
Descrizione fisica	1 online resource (290 p.)
Collana	Expert's Voice in .NET
Disciplina	004
Soggetti	Artificial intelligence Software engineering Artificial Intelligence Software Engineering/Programming and Operating Systems
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	Includes index.
Nota di contenuto	Contents at a Glance; Contents; About the Author; About the Technical Reviewer; Acknowledgments; Introduction; Chapter 1: 256 Shades of Gray; What Is Machine Learning?; A Classic Machine Learning Problem: Classifying Images; Our Challenge: Build a Digit Recognizer; Distance Functions in Machine Learning; Start with Something Simple; Our First Model, C# Version; Dataset Organization; Reading the Data; Computing Distance between Images; Writing a Classifier; So, How Do We Know It Works?; Cross-validation; Evaluating the Quality of Our Model; Improving Your Model Introducing F# for Machine Learning Live Scripting and Data Exploration with F# Interactive; Creating our First F# Script; Dissecting Our First F# Script; Creating Pipelines of Functions; Manipulating Data with Tuples and Pattern Matching; Training and Evaluating a Classifier Function; Improving Our Model; Experimenting with Another Definition of Distance; Factoring Out the Distance Function; So, What Have We Learned?; What to Look for in a Good Distance Function; Models Don't Have to Be Complicated; Why F#?; Going Further; Chapter 2: Spam or Ham? Our Challenge: Build a Spam-Detection Engine Getting to Know Our

Dataset; Using Discriminated Unions to Model Labels; Reading Our Dataset; Deciding on a Single Word; Using Words as Clues; Putting a Number on How Certain We Are; Bayes' Theorem; Dealing with Rare Words; Combining Multiple Words; Breaking Text into Tokens; Naively Combining Scores; Simplified Document Score; Implementing the Classifier; Extracting Code into Modules; Scoring and Classifying a Document; Introducing Sets and Sequences; Learning from a Corpus of Documents; Training Our First Classifier
Implementing Our First Tokenizer Validating Our Design Interactively; Establishing a Baseline with Cross-validation; Improving Our Classifier; Using Every Single Word; Does Capitalization Matter?; Less Is more; Choosing Our Words Carefully; Creating New Features; Dealing with Numeric Values; Understanding Errors; So What Have We Learned?; Chapter 3: The Joy of Type Providers; Exploring StackOverflow data; The StackExchange API; Using the JSON Type Provider; Building a Minimal DSL to Query Questions; All the Data in the World; The World Bank Type Provider; The R Type Provider
Analyzing Data Together with R Data Frames Deedle, a .NET Data Frame; Data of the World, Unite!; So, What Have We Learned?; Going Further; Chapter 4: Of Bikes and Men; Getting to Know the Data; What's in the Dataset?; Inspecting the Data with FSharp.Charting; Spotting Trends with Moving Averages; Fitting a Model to the Data; Defining a Basic Straight-Line Model; Finding the Lowest-Cost Model; Finding the Minimum of a Function with Gradient Descent; Using Gradient Descent to Fit a Curve; A More General Model Formulation; Implementing Gradient Descent
Stochastic Gradient Descent

Sommario/riassunto

Machine Learning Projects for .NET Developers shows you how to build smarter .NET applications that learn from data, using simple algorithms and techniques that can be applied to a wide range of real-world problems. You'll code each project in the familiar setting of Visual Studio, while the machine learning logic uses F#, a language ideally suited to machine learning applications in .NET. If you're new to F#, this book will give you everything you need to get started. If you're already familiar with F#, this is your chance to put the language into action in an exciting new context. In a series of fascinating projects, you'll learn how to: Build an optical character recognition (OCR) system from scratch Code a spam filter that learns by example Use F#'s powerful type providers to interface with external resources (in this case, data analysis tools from the R programming language) Transform your data into informative features, and use them to make accurate predictions Find patterns in data when you don't know what you're looking for Predict numerical values using regression models Implement an intelligent game that learns how to play from experience Along the way, you'll learn fundamental ideas that can be applied in all kinds of real-world contexts and industries, from advertising to finance, medicine, and scientific research. While some machine learning algorithms use fairly advanced mathematics, this book focuses on simple but effective approaches. If you enjoy hacking code and data, this book is for you.