


Programming ML.NET



Professional



Dino Esposito
Francesco Esposito

FREE SAMPLE CHAPTER

SHARE WITH OTHERS



Programming ML.Net

Dino Esposito
Francesco Esposito

Programming ML.Net

**Published with the authorization of Microsoft Corporation by:
Pearson Education, Inc.**

Copyright © 2022 by Pearson Education, Inc.

All rights reserved. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. For information regarding permissions, request forms, and the appropriate contacts within the Pearson Education Global Rights & Permissions Department, please visit www.pearson.com/permissions.

No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

ISBN-13: 978-0-13-738365-8

ISBN-10: 0-13-738365-7

Library of Congress Control Number: 2021952995

ScoutAutomatedPrintCode

Trademarks

Microsoft and the trademarks listed at <http://www.microsoft.com> on the “Trademarks” webpage are trademarks of the Microsoft group of companies. All other marks are property of their respective owners.

Warning and Disclaimer

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an “as is” basis. The author, the publisher, and Microsoft Corporation shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the programs accompanying it.

Special Sales

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at corpsales@pearsoned.com or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact intlcs@pearson.com.

Editor-in-Chief

Brett Bartow

Executive Editor

Loretta Yates

Sponsoring Editor

Charvi Arora

Development Editor

Rick Kughen

Managing Editor

Sandra Schroeder

Senior Project Editor

Tracey Croom

Copy Editor

Rick Kughen

Indexer

Timothy Wright

Proofreader

Abigail Manheim

Technical Editor

Bri Achtmann

Cover Designer

Twist Creative, Seattle

Composer

codeMantra

To Silvia, Michela and new dreams
— *Dino Esposito*

To my loved ones, to whom I couldn't help but dedicate a book
— *Francesco Esposito*

This page intentionally left blank

Contents at a Glance

	<i>Acknowledgments</i>	xv
	<i>Introduction</i>	xvii
CHAPTER 1	Artificially Intelligent Software	1
CHAPTER 2	An Architectural Perspective of ML.NET	9
CHAPTER 3	The Foundation of ML.NET	27
CHAPTER 4	Prediction Tasks	45
CHAPTER 5	Classification Tasks	73
CHAPTER 6	Clustering Tasks	99
CHAPTER 7	Anomaly Detection Tasks	119
CHAPTER 8	Forecasting Tasks	141
CHAPTER 9	Recommendation Tasks	159
CHAPTER 10	Image Classification Tasks	175
CHAPTER 11	Overview of Neural Networks	189
CHAPTER 12	A Neural Network to Recognize Passports	205
APPENDIX A	Model Explainability	219
	<i>Index</i>	225

This page intentionally left blank

Contents

Acknowledgments.....	xv
Introduction	xvii
Chapter 1 Artificially Intelligent Software	1
How We Ended Up with Software	2
The Formalization of Computing Machines	2
The Engineering of Computing Machines	3
The Birth of Artificial Intelligence	3
Software as a Side Effect	4
The Role of Software Today	4
Automating Tasks	5
Mirroring the Real World.....	5
Empowering People	6
AI Is Just Software	7
Chapter 2 An Architectural Perspective of ML.NET	9
Life Beyond Python.....	10
Why Is Python So Popular in Machine Learning?.....	10
Taxonomy of Python Machine Learning Libraries.....	11
End-to-End Solutions on Top of Python Models.....	13
Introducing ML.NET	13
The Learning Pipeline in ML.NET	14
Model Training Executive Summary	19
Consuming a Trained Model	23
Making the Model Callable from the Outside	23
Other Deployment Scenarios.....	23
From Data Science to Programming	24
Summary	25

Chapter 3	The Foundation of ML.NET	27
	On the Way to Data Engineering	27
	The Role of a Data Scientist.	28
	The Role of a Data Engineer	29
	The Role of an ML Engineer	30
	The Data to Start From.	30
	Making Sense of the Available Data.	31
	Building a Data Processing Pipeline.	32
	The Training Step	35
	Picking an Algorithm	36
	Measuring the Actual Value of an Algorithm	36
	Planning the Testing Phase.	37
	A Look at the Metrics	38
	Consuming the Model from a Client Application	39
	Getting the Model File	39
	The Overall Project	39
	Making a Taxi Fare Prediction	40
	Scalability Concerns	42
	Devising an Adequate User Interface	42
	Summary	43
Chapter 4	Prediction Tasks	45
	The Pipeline and the Chain of Estimators.	46
	Data Views	46
	Transformers	47
	Estimators.	47
	Pipelines	48
	The Regression ML Task.	48
	General Aspects of ML Tasks	49
	Supported Regression Algorithms.	49
	Supported Validation Techniques.	52

Using the Regression Task	54
A Look at the Available Training Data	54
Feature Engineering	58
Accessing the Content of Datasets	60
Composing the Training Pipeline	62
The ML Devil's Advocate	70
Simple and Linear Regression	70
Nonlinear Regression	71
Summary	71

Chapter 5 Classification Tasks 73

The Binary Classification ML Task	73
Supported Algorithms	73
Supported Validation Techniques	75
Binary Classification for Sentiment Analysis	75
A Look at the Available Training Data	75
Feature Engineering	79
Composing the Training Pipeline	81
The Multiclass Classification ML Task	85
Supported Algorithms	85
Using the Multiclass Classification Task	87
A Look at the Available Data	88
Composing the Training Pipeline	90
The ML Devil's Advocate	96
The Many Faces of Classification	97
Another Perspective on Sentiment Analysis	98
Summary	98

Chapter 6 Clustering Tasks 99

The Clustering ML Task	99
Unsupervised Learning	99
A Look at the Available Training Data	100
Feature Engineering	104
Clustering Algorithms	105

	Composing the Training Pipeline	109
	Setting Up a Client Application	111
	The ML Devil's Advocate	114
	Clustering Is Always the First Step	114
	Unsupervised Reduction of the Dataset.....	115
	Summary	117
Chapter 7	Anomaly Detection Tasks	119
	What Is an Anomaly?	119
	General Approaches to Detect Anomalies.....	120
	Time Series Data.....	120
	Statistical Techniques	122
	Machine Learning Approaches	123
	The Anomaly Detection ML Task	125
	A Look at the Available Training Data	125
	Composing the Training Pipeline	128
	Setting Up a Client Application	134
	The ML Devil's Advocate	137
	Predictive Maintenance.....	137
	Fraudulent Financial Operations	139
	Summary	140
Chapter 8	Forecasting Tasks	141
	Predicting the Future	141
	Simple Forecasting Methods	141
	Mathematical Foundation of Forecasting	142
	Common Decomposition Algorithms.....	144
	The SSA Algorithm.....	144
	The Forecast ML Task	146
	A Look at the Available Data	146
	Composing the Training Pipeline	148
	Setting Up a Client Application	151
	The ML Devil's Advocate	154

Not A Random Walk in the Park?	154
Other Approaches to Time Series	155
Energy Production Prediction	156
Summary	158
Chapter 9 Recommendation Tasks	159
Inside Information Retrieval Systems	159
The Basic Art of Ranking	160
The Flexible Art of Recommendation	161
The Delicate Art of Collaborative Filtering	162
The ML Recommendation Task	163
A Look at the Available Data	163
Composing the Training Pipeline	166
Setting Up a Client Application	170
ML Devil's Advocate	172
If You're Like Netflix	172
What If You're Not Like Netflix?	173
Summary	173
Chapter 10 Image Classification Tasks	175
Transfer Learning	175
Popular Image Processing Neural Networks	176
Other Image Neural Networks	176
Transfer Learning via Composition	176
The Transfer Learning Pattern in ML.NET	177
Overall Purpose of the New Image Classifier	178
A Look at the Available Data	178
Composing the Training Pipeline	180
Setting Up a Client Application	182
The ML Image Classification Task	184
The Image Classification API	184
Using the Image Classification API	185

The ML Devil's Advocate	186
The Magic of the Human Brain	186
Handcrafted Neural Networks	187
Retraining	188
Summary	188
Chapter 11 Overview of Neural Networks	189
Feed-forward Neural Networks	189
Artificial Neurons	190
Layers of the Network	191
The Logistic Neuron	193
Training a Neural Network	194
More Sophisticated Neural Networks	197
Stateful Neural Networks	197
Convolutional Neural Networks	199
Auto-Encoders	202
Summary	203
Chapter 12 A Neural Network to Recognize Passports	205
Using Azure Cognitive Services	205
Anatomy and Solution of the Problem	206
Working with the ID Form Recognizer	207
Crafting Your Own Neural Network	210
Topology of the Neural Network	211
Training Pains	215
The ML Devil's Advocate	216
Commodity Versus Vertical Solutions	216
When Are Custom Solutions Inevitable?	217
Summary	217

Appendix A Model Explainability	219
Software Intelligence	219
The Super Theory of Artificial Intelligence	220
Machine Learning Black Boxes	221
Interpretability and Explainability	221
Explainability Techniques	223
Conclusion	224
<i>Index</i>	225

This page intentionally left blank

Acknowledgments

FROM DINO:

It's the second time that the two of us, father and son, have written a machine learning book and a lot has changed since our last one, two years ago. In this book, we really joined forces—I put my software experience on the table, and Francesco gave his freshness, energy, and mathematical skills. We learned both how tricky it can be to put machine learning solutions in production and how “easy” it can be to hide those little gems in the folds of “normal” ASP.NET applications.

In the past two years, we achieved other results and, for example, we solidified our grasp of software for professional tennis and expanded to healthcare, agriculture, and customer care. The common denominator is always that one: intelligent software that ends up doing intelligent things. It's not about replacing humans and killing jobs—quite the reverse. It's about replacing tasks with automated procedures, thus freeing humans from boring, automatable tasks and keeping them engaged in more interesting activities.

With Giorgio Garcia-Agreda and Gaetano Guarino, and the entire Crionet crew, we're making our tennis fanatic dreams bigger every day. We are changing the games. With Vito Lanzotti and the KBMS Data Force team, we're making silent history by turning doctors' operating dreams into concrete and applicable artifacts, thus smoothing the way patients receive care. With Salvo Intilisano and Daniel Intilisano of Karma Enterprise, it was really a matter of technical karma. Same mindset, same vision, and same father-and-son business model! Agriculture won't be the same after the project ends, and the bees will be grateful!

The Youbiquitous team is growing, and the business is now spread over multiple pairs of strong shoulders—mainly those of Matteo, Luciano, Martina, Filippo, and Gabriele. Thank you all for taking the time to keep the business running while we were having fun with ML.NET.

Finally, any book is teamwork, and it is our pleasure to call out the names of those who made it ultimately possible. Last but not certainly least! A monumental thank-you goes to Loretta Yates as the acquisition editor, Charvi Arora as the sponsoring editor, Rick Kughen as the development and copy editor, and Bri Achtman as the technical editor.

FROM FRANCESCO:

I'm 23, grown enough to live alone but young enough to feel my heart beating for my grandparents. Sadly enough, the number is smaller than a book ago. A warm thought goes to Grandpa Salvatore and a hug to Grandma Concetta and Grandma Leda: I love you. On an even more personal side, this book is for Gianfranco—friend, business partner, second father, and grandfather. He taught me how to do things right and forgot to teach me how to do it wrong. And this book is also for Michela, who is strong enough to pursue her own way—no matter what—and smart enough to choose a good path!

Introduction

*We need men who can dream of things that never were, and ask,
“why not?”*

—John F. Kennedy, *Speech to the Irish Parliament, June 1963*

Today, the quest for data scientists is continuous, the data seems to be abundant, and cloud computing power is available. Is it the perfect world for the definitive triumph of machine learning? As we see things, we have all the necessary ingredients to cook up the “applied AI,” but we still lack a clear and effective method for combining them.

The purpose of data science is, like the purpose of science, to show that something is possible. Data science, though, doesn’t productionize solutions. That’s the purpose of another branch of the machine learning universe—data engineering.

Companies are wildly looking for data scientists, but the outcome of a good data science team is typically a runnable model whose software quality is often that of a prototype rather than of a production-ready artifact. Algorithms are tightly bound to data, and data must be complete, clean, and balanced. Who’s in charge of this part of the job is often unclear, and as a result, the job is often partially done at best. Yet, a data science team disconnected from the rest of the applied AI pipeline that makes it to production is still a due investment for a large organization whose business produces large quantities of data (such as energy utilities, financial institutions, and manufacturing farms). For smaller companies with significantly more limited budgets, the outcome of some applied data science can be cheaper to buy as a service.

From data science to production, there is usually a long way in between and a lot of work on data. Following are a few points to consider:

- How is data stored? On a daily or hourly basis?
- Should data be temporarily copied in some intermediate format?
- What kind of transformation is required for the model to work? How can you automate that?
- How is the performance of the model once deployed to production?
- How often is the model expected to need retrained to stay adherent to live data?
- If retraining is a frequent operation, how to automate any related tasks?

- What about collecting updated datasets, running the training, and deploying the up-to-date model?

The biggest issue we experience with machine learning models traces back to the data employed. In July 2021, the *MIT Technology Review* published an article about the impact of AI in facing the COVID pandemic. The article's bottom line is that many of the problems uncovered by a large review of aptly developed models are linked to the poor data quality that researchers used to develop their tools. Hence, nearly all tools were of nearly no effective use. This leads to a better understanding of the role of data engineering and data quality. Treating data via CSV sparse files is sufficient for probing an idea, but in order to build a robust infrastructure, you need to switch to some database (relational, NoSQL or graph) and some serious query language, and for this purpose, likely move beyond Python and enter into classic programming. Data science is not enough without serious programming and database skills. On the other hand, isn't searching for business-specific insights in data all that you do?

AI in general, and machine learning in particular, are now played as a tradeoff between commodity and direct solutions for vertical problems. Commoditized cloud services offer security, stability, and acceptable quality. They don't cover every possible scenario, though. But they're expanding, and more will expand in the near future.

All this is creating the environment for building the same old software but with more powerful tools. We're not just talking about the primitives of programming languages and classes from some frameworks. We're also talking about intelligent and predictive tools backed by machine learning algorithms and commoditized cloud (or containerized) services.

In this scenario, ML.NET acts as a perfect bridge between data engineering and commoditized data science, and it is fully integrated with the .NET framework. ML.NET commodities come in different forms: built-in algorithms for shallow learning, facilitated access to Azure cloud services, and integration with pre-trained models, such as Keras or TensorFlow networks.

Who Should Read This Book?

In our vision, if you adopt the .NET stack, then ML.NET is the perfect tool to do machine learning, whatever that ultimately means in terms of the internal gears of your chosen algorithms and models.

Hence, this book is for .NET developers willing (or needing) to approach the world of machine learning. It's ideal if you're a software developer adding data science and

machine learning skills to your arsenal. It's ideal if you're a data scientist willing to learn more about software beyond Python. Both categories, though, need to learn more and more about the other.

Who Should Not Read This Book?

This book discusses machine learning through the lenses of ML.NET, which is a platform-specific library. It is tailored more to data engineers and ML engineers than plain data scientists. To clarify, the core responsibilities of an ML engineer are to physically incorporate an externally trained model into client applications and perform the much more delicate task of supervising the building and training of a model based on data science specs. The book discusses the tools for doing this.

If you're not much interested in the actual productionizing of the machine learning solution, this book is probably not the best you can get. It doesn't open your mind to cutting-edge data science techniques, but it teaches you how to start leveraging what the ML.NET team has been doing for years—to integrate simple but effective machine learning solutions in .NET.

Organization of This Book

This book is divided into three sections.

- Chapters 1-3 provide a foundational overview of the library.
- Chapters 4-10 outline the dedicated tasks for data processing, training, and evaluation for common problems, such as regression, classification, ranking, anomaly detection, and more.
- Chapters 11-13 are dedicated to neural networks that might come into play when none of the shallow learning tasks is found to be suitable. Also, we include an overview of neural networks and an example of passport recognition that uses both commoditized Azure cognitive services and a handmade custom Keras network.

Lastly, Appendix A discusses model explainability.

System Requirements

You will need the following hardware and software to complete the practice exercises in this book:

- A computer running Windows 10, Linux, or macOS
- Visual Studio 2019, any edition, or superior; Visual Studio Code
- Internet connection to download software or chapter examples

Code Samples

All the code included in the book, including possible errata and extensions, can be found at MicrosoftPressStore.com/ProgrammingMLNET/downloads.

Errata, updates, & book support

We've made every effort to ensure the accuracy of this book and its companion content. You can access updates to this book—in the form of a list of submitted errata and their related corrections—at:

MicrosoftPressStore.com/ProgrammingMLNET/errata

If you discover an error that is not already listed, please submit it to us at the same page.

For additional book support and information, please visit <http://www.MicrosoftPressStore.com/Support>.

Please note that product support for Microsoft software and hardware is not offered through the previous addresses. For help with Microsoft software or hardware, go to <http://support.microsoft.com>.

Stay in Touch

Let's keep the conversation going! We're on Twitter: <http://twitter.com/MicrosoftPress>.

An Architectural Perspective of ML.NET

“In mathematics, the art of asking questions is more valuable than solving problems.”

—Georg Cantor, *Doctoral thesis, 1867*

As consumers, we continuously experience the pleasant effects of cognitive AI (for example, Amazon, Google, Apple, Microsoft, and Netflix). As people, we hope to see the same power applied to healthcare. In the general enterprise sector, where companies much smaller and rich than the web giants strive, the adoption of AI is slow and steady. This is precisely the feeling that descends from the point of intelligent software. Very few running businesses need the same level of cognitive AI we find in, say, Alexa or Cortana. All running businesses, though, would benefit from smarter features.

What is intelligent software, then?

Any software is statically designed to be aware of the context, but only intelligent software is designed to run while being dynamically aware of the business context. On the other hand, isn't this just what intelligence is supposed to be—the ability to acquire knowledge and turn it into expertise? In a nutshell, intelligence combines cognitive capabilities, including perception, memory, language, and reasoning and uses a specific learning approach to extract, transform, and store information. Turning all this into code requires ad hoc tools that are different from the basic logical equipment provided in any programming language or core framework.

Marketing departments love to generally identify these tools as artificial intelligence, specifically in the machine learning (ML) section. How do we do ML?

Most ML solutions today are built using tools from the Python ecosystem. It's a matter of convenience, however, rather than a matter of technological merit. In this chapter, we introduce the ML.NET platform—the .NET way to machine learning and the core topic of the book. More than just that, though, we will aim at providing an architectural perspective of a generic ML solution and presenting the reasons that, in our opinion, make ML.NET the right thing at the right time.

Life Beyond Python

In the collective imagination, ML is tightly coupled to using the Python programming language. Even from a surface-level look at job descriptions in tons of recruitment posts, this is clear. There are both historical and convenience reasons for languages like Python and C++ to be at the forefront of ML. However, there's no strict business reason or technical argument that prevents .NET and related languages (C# and F#) from being effectively used to build ML models.

Why Is Python So Popular in Machine Learning?

Python is an interpreted and object-oriented programming language created by Guido Van Rossum in the late 1980s with the declared goals of syntax minimalism and readability. The vision of Python as a programming language is that of a small core language engine with a large standard library and an easily extensible interpreter.

Python, which was born in a scientific environment, has become the de facto standard programming language for scientists to practice, explore, and experiment with numbers. In a way, it took the place that Fortran held in the 1960s and '70s. In the beginning, using Python in a hot new scientific field such as machine learning was a natural choice, and over time—given the natural extensibility of the language—it led to the creation of a vast ecosystem of dedicated libraries and tools. In turn, that reinforced the belief that using Python for building computational models was the best option.

Today, most data scientists find Python comfortable to use for machine learning projects, and that is probably because of a combination of the language's simplicity, the available tools, and plenty of examples. As developers, we also find Python comfortable to use to reshape data quickly to find the most appropriate format, test algorithms quickly, and explore different directions.

Once a clear path is outlined, the ML model must be trained and integrated into a runtime environment, its performance with live data must be monitored, and due changes must be applied and redeployed. It's the ML life cycle that is also known as MLOps. When you move away from experimenting with tools and libraries and look just for what enterprises need—working and maintainable code—Python shows structural limits. At the very minimum, it's yet another stack to integrate into .NET or Java stacks, which is how most business applications are written.



NOTE It's difficult to go from ML experimentation (often done in Python and via Notebooks) to deployment. In fact, according to a 2020 State of Enterprise Machine Learning report, only 22 percent of companies using machine learning have successfully deployed a machine learning model into production. See <https://bit.ly/3y8BxOH>

That's one of the advantages of ML.NET—.NET makes it super easy to bring projects to production.

Taxonomy of Python Machine Learning Libraries

The ecosystem of tools and libraries available in Python can be divided into five main areas: data manipulation, data visualization, numerical computing, model training, and neural networks. It's probably not an exhaustive list because many other libraries exist that perform other tasks and focus on some specific areas of machine learning, such as natural language processing and image recognition.

When using Python, the steps to build a machine learning pipeline are typically performed within the boundaries of a notebook. A notebook is a document created in a specific web or local interactive environment called Jupyter Notebook. (See <https://jupyter.org>.) Each notebook contains a combination of executable Python code, richly formatted text, data grids, charts, and pictures through which you build and share your development story. In some way, a notebook is comparable to a Visual Studio project solution.

In a notebook, you perform tasks such as data manipulation, plotting, and training, and you can use a number of predefined and battle-tested libraries.

Data Manipulation and Analysis

Pandas (<https://pandas.pydata.org>) is a library centered around the DataFrame object through which developers can load and manipulate in-memory tabular data. The object can import content from CSV files, text files, and SQL databases, and it provides core capabilities, such as conditional search, filtering, indexing and sorting, data slicing, grouping, and column operations (such as add, remove, and rename). The DataFrame object has built-in capability to flexibly reshape and pivot data and merge multiple frames. It also works well with time-series data.

The Pandas library is ideal for data preparation operations. Its integration with interactive notebooks enables you to perform on-the-fly testing of different configurations and data groupings.

Data Visualization

Matplotlib (<https://matplotlib.org>) is a helper library that isn't directly related to any of the common tasks of a machine learning pipeline, but it comes very handy to visually represent data during the various phases of the data preparation step or metrics obtained after evaluating trained models.

In general terms, it's a mere data visualization library built for Python code. It includes a 2D/3D rendering engine and supports common types of graphs, such as histograms, pie charts, and bar charts. Graphs are fully customizable in terms of line styles, font properties, axes, legends, and the like.

Numeric Computing

Because Python is a language that is largely used in scientific environments, it can't be without a bunch of extensions specifically designed for numerical computation. In this area, NumPy and SciPy are popular libraries, though they have slightly different capabilities.

NumPy (<https://www.numpy.org>) focuses on array operations and provides facilities to create, manipulate, and reshape one-dimensional and multidimensional arrays. Also, the library supplies linear algebra, Fourier transform, and random number operations.

SciPy (<https://scipy.org>) extends NumPy with polynomials, file I/O, image and signal processing, and more advanced features such as integration, interpolation, optimization, and statistics.

In the area of scientific computation, another Python library that is worth mentioning is Theano (<https://github.com/Theano/Theano>). Theano evaluates mathematical expressions based on multidimensional arrays, very efficiently making transparent use of the GPU. It also does symbolic differentiation for functions with one or more inputs.

Model Training

Though it was originally designed for data mining, today, scikit-learn (<https://scikit-learn.org>) is a library mainly focused on model training. It provides implementations of popular algorithms for regression, classification, and clustering. Also, scikit-learn provides methods for data preprocessing, such as dimensionality reduction, feature extraction, and normalization.

In a nutshell, scikit-learn is the Python foundation for shallow learning.

Neural Networks

Shallow learning is an area of machine learning that covers a broad array of fundamental problems such as regression and classification. Outside the realm of shallow learning, there are deep learning and neural networks. For building neural networks in Python, more specialized libraries exist.

TensorFlow (<https://www.tensorflow.org>) is probably the most popular library for training deep neural networks. It is part of a comprehensive framework that can be programmed at various levels. For example, you can use the high-level Keras API to build neural networks or manually build the desired topology and specify via code forward and activation steps, custom layers, and training loops. Overall, TensorFlow is an end-to-end machine learning platform providing facilities also to train and deploy.

Keras (<https://keras.io>) is probably the easiest way to get into the dazzling world of deep learning. It offers a very straightforward programming interface that at least comes in handy for quick prototyping. As mentioned, Keras can be used from within TensorFlow.

Yet another option is PyTorch, available at <https://pytorch.org>. PyTorch is the Python adaptation of an existing C-based library specialized in natural language processing and computer vision. Of the three neural network options, Keras is, by far, the ideal entry point and the tool of choice as long as it can deliver what you're looking for. PyTorch and TensorFlow do the same job of building sophisticated neural networks, but they use different approaches to the task. TensorFlow requires you to define the entire topology of the network before you can train it. In contrast, PyTorch follows a more agile approach and provides a more dynamic method for making changes to the graph. In some ways, their differences can be summarized as "waterfall versus agile." PyTorch is younger and doesn't have TensorFlow's huge community behind it.

End-to-End Solutions on Top of Python Models

With Python, you can easily find a way to build and train a machine learning model. Ultimately, a model is a binary file that must be loaded into a client application and invoked. Usually, a Java or .NET application serves as the client application for an ML model.

There are three main ways to consume a trained model:

- Hosting the trained model in a web service and making it accessible via a REST or gRPC API.
- Importing the trained model as a serialized file in the application and interacting with it through the programming interface provided by the infrastructure it is built upon (for example, TensorFlow or *scikit-learn*). This is possible only if the founding infrastructure provides bindings for the language to which the client application.
- The trained model is exposed via the new universal ONNX format, and the client application incorporates a wrapper for consuming ONNX binaries.

While the web service option is the most commonly used, a direct API that is specific for the client language of choice might seem the fastest way to consume a trained model. There are a couple of aspects to review, however:

- Using a direct API can prevent you from taking advantage of hardware acceleration and network distribution. In fact, if the API is hosted locally, any dedicated hardware (such as a GPU) is up to you. For this reason, if you want to invoke a graph at a very high rate in real-time, then you should consider using an ad hoc, hardware-accelerated cloud host.
- A binding for the specific trained model might not exist for the language of your choice. For example, TensorFlow natively supports Python, C, C++, Go, Swift, and Java.

Invoking a Python (or C++) library from within .NET code is not an unsurmountable technical issue. However, invoking a specific library, such as a trained machine learning model, is usually harder than calling a plain Python or C++ class.

In summary, an ML solution doesn't live on its own and must be framed in the context of an end-to-end business solution. Because many business solutions are based on the .NET stack, it was about time that a platform for training ML models natively in .NET came out. Using ML.NET, you can stay in the .NET ecosystem and don't have to deal with integrating Python into .NET applications.

Introducing ML.NET

First released in the spring of 2019, ML.NET is a free cross-platform and open-source .NET framework designed to build and train machine learning models and host them within .NET applications. See <https://dotnet.microsoft.com/apps/machinelearning-ai/ml-dotnet>.

ML.NET aims to provide the same set of capabilities that data scientists and developers can find in the Python ecosystem, as described earlier. Specifically built for .NET developers (for example, the API reflects common patterns of .NET frameworks and related development practices), ML.NET is built around the concept of the classic ML pipeline: collect data, set the algorithm, train, and deploy. In addition, any required programming steps sound familiar to anybody using the .NET framework and C# and F# programming languages.

The most interesting aspect of ML.NET is that it offers a quite pragmatic programming platform arranged around the idea of predefined *learning tasks*. The library comes equipped to make it relatively easy—even for machine learning newbies—to tackle common machine learning scenarios such as sentiment analysis, fraud detection, or price prediction as if they were just plain programming.

Compared to the pillars of the Python ecosystem presented earlier, ML.NET can be seen primarily as the counterpart of the scikit-learn model building library. The framework, however, also includes some basic facilities for data preparation and analysis that you can find in Pandas or NumPy. ML.NET also allows for the consumption of deep-learning models (specifically, TensorFlow and ONNX). Also, developers can train image classification and object detection models via Model Builder. It is remarkable, though, that the whole ML.NET library is built atop the tremendous power of the whole .NET Core framework.

The ML.NET framework is available as a set of NuGet packages. To start building models, you don't need more than that. However, as of version 16.6.1, Visual Studio also ships the Model Builder wizard that analyzes your input data and chooses the best available algorithm. We return to Model Builder in Chapter 3, "The Foundation of ML.NET."

The Learning Pipeline in ML.NET

A typical ML.NET solution is commonly articulated in three distinct projects:

- An application that orchestrates the steps of any machine learning pipeline: data collection, feature engineering, model selection, training, evaluation, and storage of the trained model
- A class library to contain the data types necessary to have the final model make a prediction once hosted in a client application. Note though that the input and output schemas do not strictly require its own project since these classes can be defined in the same project where training or consumption occurs
- A client application (website or a mobile or desktop application)

The orchestrator can be any type of .NET application, but the most natural choice is to have it coded as a console application.

It is worth noting, though, that this particular piece of code is not a one-off application that stops existing once it has given birth to the model. More often than not, the model has to be re-created many times before production and especially after being run in production. For this reason, the trainer application must be devised to be reusable and easy to configure and maintain.

Getting Started

As simple as it sounds, you can start by creating three such projects manually in Visual Studio and make them look like Figure 2-1. The figure presents three embryonal projects with many files and references missing but with sufficient details to deliver the big picture.

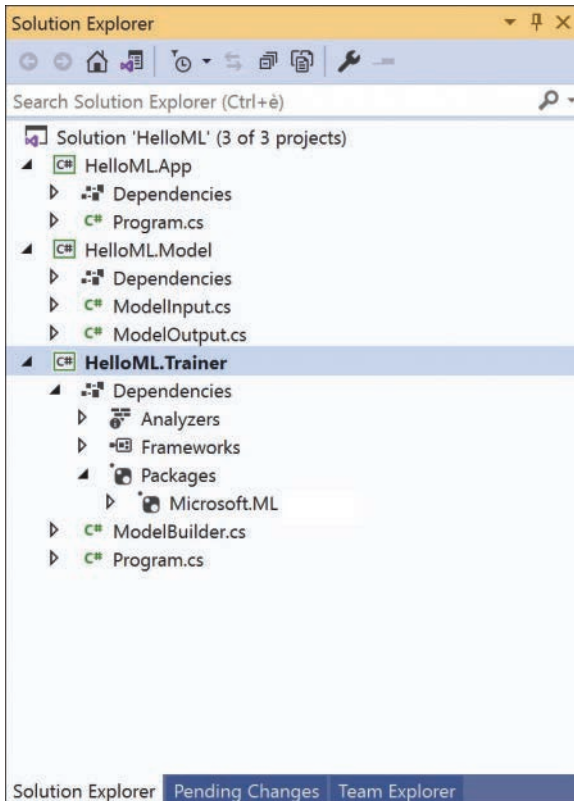


FIGURE 2-1 The skeleton of an ML.NET project in Visual Studio

Aside from the core references to the .NET framework of choice (whether 3.x or 5), the only additional piece you need to bring in is the `Microsoft.ML` NuGet package.

The package is not comprehensive, meaning that depending on what you intend to do, installing more packages might be necessary. However, the package is sufficient to get you started and enables you to experiment with the library. Let's focus on the trainer application and see what it takes to interact with the ML.NET library.

The Pipeline Entry Point

The entry point in the ML.NET pipeline is the `MLContext` object. You use it in much the same way you use the Entity Framework `DbContext` object or the connection object to a database library. You need to have an instance of this class shared across the various objects that participate in the building of the

model. A common practice employed in most tutorials—including the sample code generated by the aforementioned Model Builder wizard—is wrapping the model building workflow in a dedicated class, often just named `ModelBuilder`.

```
public static class ModelBuilder
{
    private static MLContext Context = new MLContext();

    // Main method
    public static void CreateModel(string inputDataFileName, string outputModelFileName)
    {
        // Load data

        // Build training pipeline

        // Train Model

        // QUICK evaluation of the model

        // Save the output model
    }
}
```

The instance of the `MLContext` class is global to the class methods, and the name of the files containing train data and the final output file are passed as arguments. The body of the `CreateModel` method (or whatever name you choose for it) develops around a few steps that involve more specific classes of the ML.NET library for activities such as data transformation, feature engineering, model selection, training, evaluation, and persistence.

Data Preparation

The ML.NET framework can read data from a variety of data sources (for example, a CSV-style text file, a binary file, or any `IEnumerable`-based object) and does that through the services of a few specialized loaders built around a specific interface—the `IDataView` interface, a flexible and efficient way of describing tabular data.

An `IDataView`-based loader works as a database cursor and supplies methods to navigate around the data set at any acceptable pace. It also provides an in-memory cache and methods to write the modified content back to disk. Here's a quick example:

```
// Create the context for the pipeline
Context = new MLContext();

// Load data into the pipeline via the DataView object
var dataView = Context.Data.LoadFromTextFile<ModelInput>(INPUT_DATA_FILE);
```

The sample code loads training data from the specified file and manages it as a collection of `ModelInput` types. Needless to say, the `ModelInput` type is a custom class that reflects the rows of

data loaded from the text file. The snippet below shows a sample `ModelInput` class. The `LoadColumn` attribute refers to the position of the CSV column to which the property binds.

```
public class ModelInput
{
    [LoadColumn(0)]
    public string Month { get; set; }

    [LoadColumn(1)]
    public float Sales { get; set; }
}
```

In the code, there's an even more relevant point to emphasize. The code makes a key assumption that the loaded data is already in a format that is acceptable for machine learning operations. More often than not, instead, some transformation is necessary to perform on top of available data—the most important of which is that all data must be numeric because algorithms can only read numbers. Here's a realistic scenario:

Your customer has a lot of data coming from a variety of data sources, whether timeline series, sparse Office documents, or database tables populated by online web frontends. In its raw format, the data—no matter the amount—might not be very usable. The appropriate format of the data depends on the desired result and the selected training algorithm. Therefore before mounting the final pipeline, a number of data transformation actions might be necessary, such as rendering data in columns, adding ad hoc feature columns, removing columns, aggregating and normalizing values, and adding density wherever possible. Depending on the context, these steps might be accomplished only once or every time the model is trained.



NOTE At first sight, it might seem that integrating data processing in the pipeline is a waste of time and that there's no value in doing it every time the model is built. More in general, instead, it is a matter of trade-off. We're usually talking about large quantities of data, and processing it to some intermediate format may be expensive. On the other hand, if the gap between raw and cleaned data is not that big, transforming data on each build of the model can deliver tremendous flexibility because you can change the transformation parameters as is convenient. It's a pure speed versus flexibility trade-off.

Trainers and Their Categorization

Training is the crucial phase of a machine learning pipeline. The training consists of picking an algorithm, setting in some way its configuration parameters, and running it repeatedly on a given (training) data set. The output of the training phase is the set of parameters that lead the algorithm to generate the best results. In the ML.NET jargon, the algorithm is called the *trainer*. More precisely, in ML.NET, a trainer is an algorithm plus a task. The same algorithm (say, L-BFGS) can be used for different tasks, such as regression or multiclass classification.

Table 2-1 lists some of the supported trainers grouped in a few tasks. We cover ML.NET tasks in more depth throughout the rest of the book and investigate their programming interface.

TABLE 2-1 ML.NET Tasks Related to Training

Task	Description
AnomalyDetection	Aims at detecting unexpected or unusual events or behaviors compared to the received training
BinaryClassification	Aims at classifying data in one of two categories
Clustering	Aims at splitting data in a number of possibly correlated groups without knowing which aspects could possibly make data items related
Forecasting	Addresses forecasting problems
MulticlassClassification	Aims at classifying data in three or more categories
Ranking	Addresses ranking problems
Regression	Aims at predicting the value of a data item

Figure 2-2 presents the list of ML.NET task objects as they show up through IntelliSense from the MLContext pipeline entry point.

```
public static void CreateModel(string inputDataFileName, string outputModelFileName)
{
    // Load data
    Context.Regression..
    AnomalyDetection AnomalyDetectionCatalog
    // Bl BinaryClassification BinaryClassificationCatalog
    Clustering ClusteringCatalog
    // Tr ComponentCatalog ComponentCatalog
    Data DataOperationsCatalog
    // QL Forecasting ForecastingCatalog
    Model ModelOperationsCatalog
    // Sa MulticlassClassification MulticlassClassificationCatalog
    Ranking RankingCatalog
    Regression RegressionCatalog
    Transforms TransformsCatalog
}
```

FIGURE 2-2 The list of ML.NET task objects

Each of the task objects listed in Table 2-1 has a Trainers property that lists the predefined algorithms supported by the framework. For example, for a prediction task, a good algorithm is the Online Gradient Descent algorithm.

```
var dataProcessPipeline = mlContext.Transforms.Text.FeaturizeText(...);
var trainer = mlContext.Regression.Trainers.OnlineGradientDescent(...);
var trainingPipeline = dataProcessPipeline.Append(trainer);
```

The code snippet selects an instance of the algorithm and then appends it to the data processing pipeline at the end of which the compiled model will come out. This short code contains the essence of the whole ML.NET programming model. This is the way in which the whole pipeline is built step by step and then run.

It is also worth noting that ML.NET supports several specific algorithms for each predefined task. In particular, for the regression task the ML.NET framework also supports the “Poisson Regression” and the “Stochastic Dual Coordinate Ascent” algorithms, and many more can be added to the project at any time through new NuGet packages.

Once the pipeline is built and fully configured, it is ready to run on the provided data. In this regard, the pipeline is a sort of abstract workflow that processes data in a way comparable to how in .NET LINQ queryable objects work on collections and data sets.

Once trained, a model is nothing more than the serialization of a computation graph that represents some sort of mathematical expression or, in some cases, a decision tree. The exact details of the expression depend on the algorithm and to some extent, the nature of the problem’s category.

Model Training Executive Summary

Explaining the mechanics of model training is well beyond the scope of this book, which remains strictly focused on the ML.NET framework. However, at least a brief recap of what it means and how it works is in order. For an in-depth analysis, you can easily find online resources as well as books. In particular, you can take a look at our book *Introducing Machine Learning* (Microsoft Press, 2020), in which we mainly focus on the mathematics behind most problems and the algorithmic solutions discovered so far for each class.

Purpose of the Training Phase

Abstractly speaking, an algorithm is the sequence of steps that lead to the solution of a problem. In artificial intelligence, there are two main classes of problems: classification of entities and predictions. In each of these classes, we find several subclasses, such as ranking, forecasting, regression, anomaly detection, image and text analysis, and so forth.

In fact, the output of a machine learning pipeline is a software artifact made of an algorithm (or a chain of algorithms) whose parametric parts (settings and configurable elements) have been adjusted based on the provided training data. In other words, the output of a machine learning pipeline is the instance of an algorithm that, much like the instance of an object-oriented language class, has been initialized to hold a given configuration. The configuration to use for the instance of the algorithm is discovered during the training phase. The schema is outlined in Figure 2-3.




FIGURE 2-3 Overall schema of the training phase in a machine learning process

The Computation Graph

As mentioned, the model is a mathematical black-box—a computation graph—that takes input and computes an output. Input and output are lists of numbers, and in an object-oriented context—as in .NET—they are modeled using classes.

Figure 2-4 presents an abstract and concrete view of how a client application ultimately uses a trained model. Input data flows in, and the gears of the graph crunch numbers and produce a response for the application to deal with.




FIGURE 2-4 Overall schema of a trained model being used in production

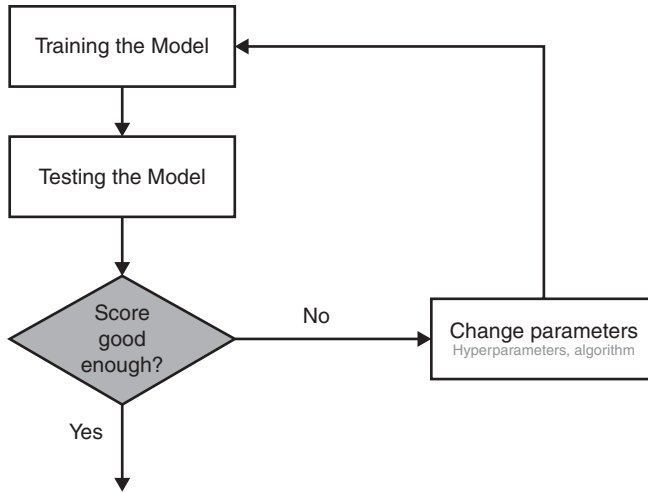
For example, if you have a model trained to detect possible fraudulent transactions in a financial application, the graph in the model will be called to process a numeric representation of the transaction and produce some values that could be interpreted as to whether the transaction should be approved, denied, or just flagged for further investigation.

Performance of the Model

The term *machine learning* sounds fascinating, but it is not always fully representative of what really happens in a low-level ML framework such as ML.NET or Python’s scikit-learn. At this level, the training phase just iteratively processes records in the training data set to minimize an error function.

- The function that produces values—the computation graph—is defined by the selected algorithm.
- The error function is yet another element added to the processing pipeline and also depends, to some extent, on the selected algorithm.
- The error function measures the distance between values produced by the graph on testing data and expected data embedded in the training data set.
- The graph enters the training phase with a default configuration that is updated if the measured error is too large for the desired goal.
- When a good compromise between speed and accuracy is reached, the training ends, and the current configuration of the graph is frozen and serialized for production.

The whole process develops iteratively in the training phase within the ML.NET framework. The steps are summarized in Figure 2-5.



Deploy to production

FIGURE 2-5 The generation of an ML model

It should be noted that the evaluation phase depicted here happens within the ML.NET framework and, more in general, within the boundaries of the ML framework of choice. The actual performance evaluated is obtainable on training data with a given set of algorithm parameters (referred to as hyperparameters) and internally computed coefficients.

This is not the same as measuring the performance of the model in production. At this stage, you only measure how good the model is at work on sample data. However, sample data is only expected to be a realistic snapshot of data the model will face once deployed to production. A more crucial evaluation phase will take place later and might even take to rebuilding the model based on different hyperparameters and—why not?—a different algorithm.

In ML.NET, the quality of the model during the training phase is measured using special components called *evaluators*.

A Look at Evaluators

An evaluator is a component that implements a given metric. Evaluation metrics are specific to the class of the algorithm and, in ML.NET, to the ongoing ML task. A good introduction on which evaluator is deemed appropriate for each ML.NET task can be found at

<https://docs.microsoft.com/en-us/dotnet/machine-learning/resources/metrics>.

A more in-depth discussion about the mathematical reasons that make each metric qualified for a given task can be found in the aforementioned book *Introducing Machine Learning*.

For example, for a prediction problem such as estimating the cost of a taxi ride (and in general for regression and ranking/recommendation problems), a key metric to consider is Squared Loss, also known as Mean Squared Error (MSE). This metric works by measuring how close a regression line is to

test expected values. For each input test value, the evaluator takes the distance between the computed and the expected response, squares it, and then calculates the mean. The squaring is applied to increase the relevance of larger differences.

Interestingly, Model Builder, which is embedded in Visual Studio, does some of the work for you. It first lets you choose the class of the problem (the ML task) and indicate the training data set. Based on that, it automatically selects a few matching algorithms, trains them, and measures the performance according to automatically selected metrics. Then it makes a final call and recommends how you should start coding your machine learning solution.

In general, there are a few things that could possibly go bad in a machine learning project:

- The selected algorithm (or algorithm hyperparameters) might not be the most appropriate to explore the given data set.
- The original data set needs more (or less) column transformation.
- The original data set is too small for the intended purpose.

As an example, Table 2-2 summarizes the scores of various algorithms selected by Model Builder for a prediction (regression) task.

TABLE 2-2 Multiple algorithms (and scores) for a sample regression task

Algorithm	Squared Loss	Absolute Loss	RSquared	RMS Loss
LightGbmRegression	4.49	0.38	0.9513	2.12
FastTreeTweedieRegression	4.70	0.44	0.9491	2.17
FastTreeRegression	4.83	0.41	0.9486	2.19
SdcaRegression	10.52	0.87	0.8845	3.27

After a test run of Model Builder, it shows that all featured algorithms ended up with good marks, but Model Builder ranked them in the order shown, thus suggesting we use the `LightGbmRegression` algorithm based on the evidence provided by metrics. Look in particular at the Squared Loss column. The score is acceptably good for the first three ranked algorithms and significantly worse for `SdcaRegression`. On the other hand, `SdcaRegression` is much faster to train. The golden rule of machine learning is that everything is a trade-off.



NOTE Another aspect to consider is that once the model goes to production—even with the best metrics—there might still be chances that things go wrong and the predictions are not in line with business expectations. When this happens, the odds are that an inadequate set of data rows was used for training. Inadequate, at least, in comparison to the real data the model was called to manage in production.

Consuming a Trained Model

At the end of the training phase, you have a model that contains instructions on which algorithm to run and which configuration to use. The model file is a zipped file in some serialization format. Note that a universal, interoperable format exists and is the ONNX format. ML.NET supports it.

As is, however, the model is a dead thing. To bring it to life, you need to load it in a runtime environment so that an API can be exposed to invoke the computation from the outside.

Making the Model Callable from the Outside

Once saved to a file—typically a ZIP file—the model is simply the flat description of a computation to be done on some input data. The first step is wrapping it into a framework engine that knows how to deserialize the graph and execute it on some input data.

ML.NET has a tailor-made set of methods ready to use. Here's the skeleton of the code you need to invoke a previously trained model in ML.NET.

```
public ModelOutput RunModel(string modelName, ModelInput input)
{
    var ml = new MLContext();
    var model = ml.Model.Load(modelName, out var schema);
    var engine = ml.Model.CreatePredictionEngine<ModelInput, ModelOutput>(model);
    return engine.Predict(input);
}
```

The sample function takes the path to the serialized model file and the input data to which a prediction is made. If the model estimates the cost of a taxi ride, then the class `ModelInput` describes the ride for which a quote is required. Typically, you will find that the model uses details such as distance, time of day, type of service requested, traffic conditions, area of the city involved, and whatever else is established. The `ModelOutput` class describes the output of the algorithm used for training. Usually, it's a simple C# class with just a few numeric properties. Here's an example:

```
public class ModelOutput
{
    public double Prediction { get; set; }
}
```

The ML.NET shell code creates an instance of a prediction engine that will carry the task of deserializing and executing the graph and return the calculated value. From the software developer's perspective, invoking an ML model is in no way different from calling a class library method.

Other Deployment Scenarios

Direct embedding of a trained model in the client application is one—and by far the simplest—deployment scenario. There are a couple of potentially sore points to emphasize.

One is the cost of deserializing the model and turning it into an executable computation graph for the runtime environment of choice—in this case, the .NET framework. The other is the (related) cost of setting up a prediction engine. Both operations can be quite expensive to perform if the client application is, say, a web application subject to thousands of calls per second. This is where an artifact like `PredictionEnginePool` comes in handy.

Therefore, the code snippet shown earlier is great for understanding the process but not necessarily good for production. More realistically, a company trains a model to expose a business-critical process as a service to various running software applications. This means that the model should be incorporated in a kind of web service, and proper layers of caching and balancing should be used to ensure proper performance.

In a nutshell, a trained model can be seen as a business black box to be used as a local class library, as a web service, or even as a microservice with its own storage and micro frontend. No option is favorable over the others, but all are feasible options for the architect to choose.

From Data Science to Programming

If you look at the trained model as an autonomous, black-boxed artifact integrated in a given type of software application, you should be able to see also the frontier between data science and programming. Data science contributes the model; programming makes it usable. Both aspects are strictly needed and unavoidable.

A trained model is nothing if not surrounded by a decent programming interface, whether in the form of a class library or a web service. To build an effective model, specific skills are required. First, you need domain expertise. Second, statistics and mathematics and the ability to discern between algorithms and metrics and interpret numbers are required. In extreme cases, the ability to develop new algorithms (including neural networks) or customize existing ones are also required. These skills very rarely belong to developers.

In much the same way, exposing a functional model requires due attention to the overall performance and scalability of the host application and care of the user experience. A taxi ride predictor model ultimately needs numbers to represent any sort of information. But you can hardly expect that people using the app on the go enter their destination through numbers. This is programming work.

In this scenario, ML.NET takes an interesting challenge: enabling developers to code their own machine learning tasks autonomously at least for relatively simple instances of problems and where a sharp precision is not the goal. This is just the ultimate purpose of ML tasks and AutoML—the engine that lies behind Model Builder. In this book, we deeply cover ML tasks but also dedicate a few final chapters to give problems a more real-world perspective. High precision, if necessary, comes at a cost!

Summary

ML.NET is now slated to become the reference platform for machine learning in the .NET space. It is mainly limited to shallow learning and doesn't offer direct support for building neural networks and deep learning (the support is only for consuming existing networks). On the other hand, also in the Python space there are libraries for shallow learning (scikit-learn) and libraries for building neural networks.

It is, instead, much more interesting and promising is the overall approach aimed at making machine learning easy to consume and relatively easy to design for developers. No developers will turn into expert data scientists overnight—not even after digesting the content of this book—but any savvy developer passionate about newer technologies and artificial intelligence would be incredibly comfortable with getting into the dazzling world of machine learning through ML.NET.

We've already mentioned something, but it helps to reinforce the concept: Although Python is quite popular among data scientists, there's no strict reason why machine learning models can't be developed and tested in .NET (or other languages, including Java and Go). It's all about the ecosystem and ease of use. ML.NET relies on the .NET Core infrastructure and Visual Studio.

Let's now go with a simple but not-so-trivial and complete example: taxi fare prediction. The next chapter includes a bit of feature engineering, feature selection, and, more importantly, a client web application.

Index

A

- accessing, datasets, 60
- accuracy, recommender systems and, 173
- activation function, 190–191, 193–194
- AI (artificial intelligence), 3–7, 186, 205. *See also* ML (machine learning)
 - automation and, 5
 - cognitive services and, 6
 - insights and, 6
 - machine learning, 7
 - software and, 7
 - thinking machines, 3
- algorithms, 15–17, 18–19, 33, 35. *See also* trainers
 - backpropagation, 195–196
 - binary classification, 74
 - choosing, 36, 62–63, 81
 - clustering
 - DBSCAN, 107–108, 124
 - K-Means, 105–106, 109–111
 - OPTICS, 106–108
 - computational complexity, 37
 - decomposition, 144
 - density, 107
 - K-Modes, 106
 - KNN (K-nearest neighbor), 123
 - LightGbmRegression, 36
 - linear, 91–92
 - LOF (Local Outlier Factor), 124
 - measuring the value of, 36–37
 - MF (Matrix Factorization), 166–168
 - multiclass classification, 85–86, 90–91
 - OCSVM (One Class Support Vector Machine), 125
 - prediction, 22, 40–41
 - Quicksort, 37
 - random forest, 155
 - ranking, 160–161
 - regression, 38, 48, 49–50
 - SdcaMaximumEntropy, 92
 - SR-CNN (Spectral Residual and Convolutional Neural Network), 132
 - SSA (Singular Spectrum Analysis), 132, 144, 149–150, 154
 - decomposition step, 144–145
 - reconstruction step, 145
 - anomaly detection, 119–120. *See also* time series
 - composing the training pipeline, 128
 - fraud and, 139–140
 - IoT and, 137–138
 - loading data and feature engineering, 127–128
 - semi-supervised learning, 124–125
 - setting up a client application, 134–136
 - supervised classification, 123
 - in a time series
 - change points, detecting, 130–131
 - outliers, 121–122
 - randomized PCA (principle component analysis), 132–134
 - spikes, detecting, 128–130
 - SR-CNN (Spectral Residual and Convolutional Neural Network), 132
 - SSA (Singular Spectrum Analysis), 132
 - statistical techniques, 122–123
 - unsupervised clustering, 124
- Apache Spark, 48
- APIs
 - Azure Cognitive Services, 205–206, 207–210
 - Image Classification, 184, 185–186
 - bottleneck phase, 185
 - training phase, 185
- Append method, 47
- AppendCacheCheckpoint method, 90–91
- artificial neurons, 190
 - logistic neuron, 193–194
 - perceptron, 1, 190–191
- ASP.NET, 68

auto-encoders

- auto-encoders, 202
 - applications, 202–203
 - schema, 202
- Automata theory, 3
- automation, AI and, 5
- AutoML, 24
- Azure Cognitive Services, 205–206
 - form recognizer service, 207–210
 - processing the input image, 211

B

- backpropagation algorithm, 195–196
- binary classification, 73
 - calibrators, 74–75
 - choosing the algorithm, 81
 - dataset, 75–77
 - partitioning, 77–78
 - evaluating the model, 82
 - metrics, 83
 - programmatically holdout, 78–79
 - for sentiment analysis, 75
 - setting up a client application, 84
 - supported algorithms, 73–74
 - text featurization, 79–81
 - trainers, 74
 - validation techniques, 75
- binning, 60
- Bombe, the, 2
- bulk rating, 173

C

- calculus ratorator*, 1
- calibrators, binary classification, 74–75
- change points, 125, 130–131
- choosing, algorithms, 36, 62–63, 81, 90–91
- Church, Alonzo, 2–3
- Church-Turing thesis, 2
- classes
 - binding output columns to, 112
 - datasets and, 31–32
 - helper, 147
 - ImageData, 179
 - ModelBuilder, 15–16
 - modeling data to, 103
 - ModelInput, 16–17
- classification, 19, 45, 73, 97
 - binary, 73

- algorithms, 74
- calibrators, 74–75
- choosing the algorithm, 81
- dataset, 75–77
- evaluating the model, 82
- metrics, 83
- partitioning the dataset, 77–78
- programmatically holdout, 78–79
- for sentiment analysis, 75
- setting up a client application, 84
- supported algorithms, 73–74
- text featurization, 79–81
- validation techniques, 75
- modeling data to classes, 103
- multiclass, 85, 87–88
 - confusion matrix, 94–95
 - dataset, 88
 - metrics, 93–94
 - object pooling, 96
 - setting up a client application, 95–96
 - supported algorithms, 85
 - text featurization, 89–90
 - trainers, 85–86, 90–91
- naïve Bayesian, 86
- One-Versus-All, 86–87
- One-Versus-One, 87
- supervised, 123
- client application, 67–68
 - anomaly detection, 134–136
 - binary classification, 84
 - clustering, 111
 - forecasting, 151–152
 - creating a time series engine, 152
 - creating checkpoints, 152–154
 - discount factor, 154
 - image classification
 - skeleton of the app, 182–183
 - user interface, 183–184
 - multiclass classification, 95–96
 - recommender system, 170
 - skeleton of the app, 170–171
 - user interface, 171–172
- cloud computing, 5, 217
- clustering, 99, 114–115
 - algorithms
 - DBSCAN, 107–108, 124
 - K-Means, 105–106, 109–111
 - K-Modes, 106
 - OPTICS, 106–108
 - applying persistent transformations, 101–102

- binding output columns to a C# class, 112
- composing the training pipeline, 109
- customer segmentation, 100
- elbow method, 106
- feature engineering, 104
- inspecting the transformed dataset, 111
 - purpose of data in, 103–104
 - reducing the number of features, 115–116
 - reducing the number of rows, 116
 - saving clusters to separate files, 113–114
 - setting up a client application, 111
 - silhouette method, 106
 - unsupervised, 124
 - unsupervised learning and, 99–100, 115
- CNN (convolutional neural networks), 199–202
- cognitive services, 6
- collaborative filtering, 162, 172
- collective outliers, 121
- Colossus, 2
- combined metrics, 83
- commoditization of solutions, 216–217
- computing machines, 2. *See also* AI (artificial intelligence)
 - Bombe, the, 2
 - calculus ratiocinator*, 1
 - Church-Turing thesis, 2
 - engineering of, 3
 - ENIAC computer, 2
 - Enigma, 2
 - Godel's theorems of incompleteness, 1
 - human brain and, 3
 - Lorenz machine, 2
 - software and, 2
 - Turing machine, 1–2
 - Von Neumann architecture, 4
- confusion matrix, 94–95
- consuming a trained model, 23, 39
 - deployment options, 23–24
 - getting the model file, 39
 - making the model callable from the outside, 23
- contextual outliers, 121
- CopyColumns method, 59
- CPU (central processing unit), 4
- credit cards, fraud detection, 139–140
- CrossValidate method, 65–66, 75
- CrossValidateNonCalibrated method, 75
- cross-validation, 38, 65–66
 - binary classification and, 75
 - holdout, 52
 - k-fold, 52–53

- custom solutions, 217
- cybernetics, 3
- cycles, 142

D

- data engineering, 29
- data science, 24–28, 30–31
- data trend, 123
- data views, 34, 40–41, 46–47, 61
- data visualization, Matplotlib library, 11
- datasets, 33–35, 54–55. *See also* ETL (Extract-Transform-Load) pipeline; time series
 - accessing, 60
 - augmenting, 55–56
 - binary classification, 75–78
 - classes and, 31–32
 - clustering, 99, 100–101
 - applying persistent transformations, 101–102
 - customer segmentation, 100
 - modeling data to classes, 103
 - reducing the number of features, 115–116
 - reducing the number of rows, 116
 - data science and, 31
 - feature engineering, 58
 - clustering and, 104
 - preliminary physical operations, 58–59
 - text featurization, 79–81
 - image classification, 178–180
 - loading, 34, 56–57
 - memory and, 61–62
 - multiclass classification, 88
 - mapping target classes to numerical values, 90
 - text featurization, 89
 - normalization, 33
 - outliers, 58–59
 - preparing, 51
 - recommender system, 163
 - feature engineering, 165
 - schema, 163–164
 - selecting columns of data, 164–165
 - schema, 57
 - splitting, 78–79
 - supported data sources, 57–58
 - test, 77–78
 - time series, univariate and multivariate, 126–127
 - training, 77
 - transformers, 47
 - validation, 77–78

DBSCAN algorithm

- DBSCAN algorithm, 124
- decomposition algorithms, 144
 - SSA (Singular Spectrum Analysis), 144
 - decomposition step, 144–145
 - reconstruction step, 145
- deep learning, 12, 192
- density algorithms, 107, 123
- DetectIdChangePoint method, 131
- DetectIdSpike method, 128–129
- DetectSpikeBySsa method, 132

E

- elbow method, 106
- energy production prediction, 156
 - forecasting pipeline, 157
 - power plant data, 156–157
 - weather forecasts, 156
- ENIAC computer, 2
- Enigma machine, 2
- error function, 63–64
- estimators, 47–48
- ETL (Extract-Transform-Load) pipeline, 29, 32, 34–35
 - data views, 46–47
 - loading data, 34, 56–57
 - scalability concerns, 42
 - supported data sources, 57–58
 - transformations, 33, 35
 - algorithms, 33
 - feature engineering, 58–60
 - normalization, 33
 - transformers, 47
- Evaluate method, 38, 64, 82
- evaluators, 21–22
- expert systems, 7

F

- F1-score, 83
- feature engineering, 58
 - anomaly detection and, 127–128
 - binning, 60
 - building the transformation pipeline, 59
 - clustering and, 104
 - key-value mapping, 60
 - normalization, 59–60
 - one hot encoding, 60
 - preliminary physical operations, 58–59
 - recommendation and, 165

- text featurization, 79–80
 - mapping target classes to numerical values, 90
 - multiclass classification and, 89
- feature importance, 53–54
- feed-forward neural networks, 189, 192, 197
- FilterRowsByColumn method, 149
- Fit method, 36, 48, 63, 129
- ForecastBySsa method, 149–150
- forecasting, 141–142, 146
 - applying the algorithm, 149–150
 - dataset, 146–147
 - helper classes, 147
 - loading data from a database source, 148
 - saving and evaluating the model, 150–151
 - separating training and testing data, 148–149
 - setting up a client application, 151–154
 - creating a time series engine, 152
 - creating checkpoints, 152–154
 - discount factor, 154
 - in a time series, 154–155
 - weather, 156
- Fortran, 10
- fraud detection, 139–140
- Frege, Gottlob, 1

G

- general recursive function, 2
- GetOutputSchema method, 47, 48
- GetRowCursor method, 34
- Godel's theorems of incompleteness, 1, 2

H

- hashing, 60
- helper classes, 147
- Hoare, Tony, 37
- holdout cross-validation, 37, 52
- human brain, the
 - computing machines and, 3
 - object detection, 187
- Hyndman, Rob, 126

I

- IDataView loader, 16
- IEstimator interface, 47
- image classification, 175. *See also* neural networks;
- passport recognition system

- composing the training pipeline, 180
 - adding the TensorFlow model to the pipeline, 181
 - retraining the TensorFlow model, 181–182
- dataset, 178–180
- mapping to a canonical classification problem, 178
- NuGet packages, 177
- retraining and, 188
- setting up a client application, 182
 - skeleton of the app, 182–183
 - user interface, 183–184
- transfer learning and, 175
- Image Classification API, 184, 185–186
 - bottleneck phase, 185
 - training phase, 185
- ImageData class, 179
- Inception v3, 176
- information retrieval systems, 159–160
- instantiation, 67, 68–70
- intelligent software, 9
- interface
 - IDataView, 60
 - IEstimator, 47
 - ITransformer, 47
- Internet, 4
- Introducing Machine Learning*, 19, 21
- IoT (Internet of Things), anomaly detection and, 137–138
- isolation forest, 124
- ITransformer interface, 47

J-K

- Jupyter Notebook, 11
- Keras, 12
- kernel matrix, 200–201
- key-value mapping, 60
- k-fold cross-validation, 52–53
- K-Means algorithm, 105–106, 109–111
- K-Modes algorithm, 106
- KNN (K-nearest neighbor) algorithm, 123, 169–170
- KPI (key performance indicators), 120

L

- lambda calculus, 2
- learning tasks, ML.NET, 14
- Leibniz, Gottfried, 1

- libraries
 - NimbusML, 109
 - OpenCV, 211–212
 - Python, 11
 - Matplotlib, 11
 - model training, 12
 - neural network, 12
 - numeric computing, 11–12
 - Pandas, 11
 - scikit-learn, 12, 14
- LightGbmRegression algorithm, 36
- linear algorithms, 91–92
- linear regression, 71
- LoadFromTextFile method, 179
- loading data, 34, 56–58
- LOF (Local Outlier Factor) algorithm, 124
- logistic neuron, 193
- logistic regression, 81–82
- Lorenz machine, 2
- LSTM (Long Short-Term Memory) neural network, 155, 157, 199

M

- MakePrediction method, 41
- Matplotlib, 11
- McCarthy, John, 3
- McCulloch, Warren, 189
- memory
 - context, 197–198
 - datasets and, 61–62
 - LSTM, 155
- methods
 - Append, 47
 - AppendCacheCheckpoint, 90–91
 - CopyColumns, 59
 - CrossValidate, 65–66, 75
 - CrossValidateNonCalibrated, 75
 - DetectLidChangePoint, 131
 - DetectLidSpike, 128–129
 - DetectSpikeBySsa, 132
 - Evaluate, 38, 64, 82
 - FilterRowsByColumn, 149
 - Fit, 36, 48, 63, 129
 - ForecastBySsa, 149–150
 - GetOutputSchema, 47–48
 - GetRowCursor, 34
 - LoadFromTextFile, 179
 - MakePrediction, 41
 - Plain, 134

methods

- Predict, 68
 - Preview, 47
 - Transform, 47
 - Zip, 151
 - metrics, 21–22, 38
 - binary classification, 83
 - combined, 83
 - multiclass classification, 93–94
 - regression, 64
 - MF (Matrix Factorization) algorithm, 166–168
 - ML (machine learning), 7, 9, 20–21. *See also* Python
 - consuming a trained model, 13
 - engineer, 29–30
 - models, 13
 - neural networks, 45
 - prediction, 48
 - Python and, 9–10
 - roles, 35
 - data engineer, 29
 - data scientist, 28
 - ML engineer, 29–30
 - shallow learning, 12
 - training, 12, 17–19. *See also* training
 - performance and, 20–21
 - purpose of, 19
 - ML.NET, 7, 10, 13–14. *See also* tasks
 - algorithms, 15–17, 18–19
 - cross-validation, 65–66
 - data views, 34, 46–47, 61
 - datasets, accessing, 60
 - estimators, 47–48
 - evaluators, 21–22
 - GetRowCursor method, 34
 - learning pipeline, 14
 - data preparation, 16–17
 - entry point, 15–16
 - error function, 20
 - evaluation phase, 21
 - trainers, 17–19
 - learning tasks, 14
 - Model Builder, 14
 - NuGet package, 15
 - orchestrator, 14
 - pipelines, 48
 - projects, 15, 39
 - tasks, 49
 - trainers, 49–51
 - transfer learning, 177
 - transformers, 47
 - MLOps, 10
 - Model Builder, 22, 24
 - ModelBuilder class, 15–16
 - ModelInput class, 16–17
 - model(s). *See also* Inception v3; testing; training
 - composition, 176, 178
 - computational graph, 20
 - cross-validation, 65–66
 - error function, 63–64
 - explainability, 222–223
 - feature importance, 53–54
 - fitting, 63
 - instantiation, 68–70
 - interpretability, 221–222
 - packaging, 66–67
 - regularization, 53
 - TensorFlow, 181–182
 - training, 12, 17–19
 - performance and, 20–21
 - purpose of, 19
 - user interface, 42–43
 - validation, 64–65
 - moving average, 122–123
 - MRZ (Machine-Readable Zone), 212–214
 - multiclass classification, 85, 87–88
 - composing the training pipeline, 90
 - evaluating the model, 92–94
 - selecting the algorithm, 90–91
 - switching back to text, 92
 - confusion matrix, 94–95
 - dataset, 88
 - metrics, 93–94
 - object pooling, 96
 - setting up a client application, 95–96
 - supported algorithms, 85
 - text featurization, 89
 - trainers, 85–86, 90–91
 - multivariate time series, 126–127
- ## N
- naïve Bayesian classification, 86
 - NAND gate, 191
 - Netflix prize challenge, 167, 169, 172
 - .NET, 9
 - neural networks, 7, 12, 45, 165, 178, 187, 189, 191, 211
 - activation function, 190–191
 - artificial neurons, 190
 - auto-encoders, 202
 - applications, 202–203

- schema, 202
- building, 178
- checkpoints, 152–154
- convolutional, 199–202
- crafting your own, 210
- custom, 217
- enabling to learn, 193
- feed-forward, 189
- handcrafting, 187
- hidden layers, 192
- image processing, 176
- Inception v3, 176
- logistic neuron, 193–194
- LSTM (Long Short-Term Memory), 155, 157, 199
- NAND gate, 191
- perceptron, 190
- predictive maintenance, 137–138
- Python libraries, 12
- recurrent, 197
 - architecture, 198
 - “deep”, 199
 - memory context, 197–198
- stateful, 197
- topology, 211
- training, 194, 215
 - backpropagation algorithm, 195–196
 - using an MRZ generator, 215–216
 - using TensorFlow generators, 216

Newton, Isaac, 1

NimbusML library, 109

noise removal, 124

nonlinear regression, 71

normalization, 59–60

notebook, 11

NuGet packages, image classification, 177

numeric computing libraries, 11–12

NumPy, 12

O

object

- detection, 175, 187
- pooling, 42, 96

OCR (optical character recognition), 205, 206–207

OCSVM (One Class Support Vector Machine) algorithm, 125

one hot encoding, 60

One-Versus-All classification, 86–87

One-Versus-One classification, 87

OpenCV library, 211–212

OPTICS algorithm, 106–108

orchestrator, 14

outliers

- change points and, 125
- collective, 121
- contextual, 121
- point, 121
- in a time series, 121–122

overfitting, 53

P

Pandas, 11

partitioning a dataset, 77–78. *See also* clustering

passport recognition system, 205–206

- extracting the MRZ region, 212–214
- form recognizer service, 207–210
- image input, 206
- MRZ (Machine-Readable Zone), 206
- OCR (optical character recognition), 205–207
- processing the input image, 211
- recognizing the MRZ content, 214–215
- text detection, 206
- training
 - using an MRZ generator, 215–216
 - using TensorFlow generators, 216

PCA (principle component analysis), 132–134

perceptron, 1, 190–191

personalization, 161–162

pipelines, 48

- chain of estimators, 66
- multiple, 65–66

Pitts, Walter, 189

Plain method, 134

point outliers, 121

Predict method, 68

prediction, 19, 21–22, 45, 141. *See also* forecasting

- algorithms, 22, 35, 36
- data science and, 31
- energy production, 156
 - forecasting pipeline, 157
 - power plant data, 156–157
 - weather forecasts, 156
- regression, 48, 54
 - augmenting the dataset, 55–56
 - datasets, 54–55
 - taxi fare, 40–41

predictive maintenance, 137–138

Preview method

- Preview method, 47
- programmatically holdout, 78–79
- programming languages, 4, 9–10. *See also* Python projects, 39
- Python, 9–10
 - libraries, 11, 14
 - Matplotlib, 11
 - model training, 12
 - neural network, 12
 - NimbusML, 109
 - numeric computing, 11–12
 - Pandas, 11
 - scikit-learn, 12
 - notebook, 11
 - popularity of, 10
- PyTorch, 12

- logistic, 81–82
- metrics, 38, 64
- nonlinear, 71
- time series, 71
- trainers, 49–50
- regularization, 53
- retraining, 188. *See also* transfer learning
- RNN (recurrent neural network), 197
 - architecture, 198
 - “deep”, 199
 - memory context, 197–198
- roles
 - data engineer, 29
 - data scientist, 28
 - ML engineer, 29–30
- Russell, Bertrand, 1

Q-R

- Quicksort algorithm, 37
- random forest algorithm, 155
- random walk, 154–155
- randomized PCA (principle component analysis), 132–134
- ranking, 159–161, 162
 - information retrieval systems, 159–160
 - versus recommendation, 161
- recommender system, 161, 163
 - accuracy and, 173
 - bulk rating, 173
 - collaborative filtering, 162, 172
 - composing the training pipeline, 166
 - dataset, 163
 - schema, 163–164
 - selecting columns of data, 164–165
 - evaluating the model, 168–169
 - feature engineering, 165
 - information retrieval systems, 159–160
 - KNN (K-nearest neighbor) algorithm, 169–170
 - MF (Matrix Factorization) algorithm, 166–168
 - personalization and, 161–162
 - versus ranking, 161
 - relevance by date, 166
 - setting up a client application, 170
 - skeleton of the app, 170–171
 - user interface, 171–172
- regression, 48–49, 54, 70
 - dataset, 54–56
 - linear, 71

S

- schema, 57
 - auto-encoder, 202
 - time series data, 127
- scikit-learn, 12, 14
- SciPy, 12
- SDCA (Stochastic Dual Coordinate Ascent) trainer, 91–92, 99
- SdcaMaximumEntropy trainer, 92
- seasonality, 142
- semi-supervised learning, 124–125
- sentiment analysis, 75, 98
- shallow learning, 12
- sigmoid function, 193–194
- silhouette method, 106
- software, 1–2, 4
 - AI and, 7
 - in cars, 5–6
 - cognitive services, 6
 - empowerment and, 6
 - goals of, 5
 - Godel’s theorems of incompleteness, 1
 - intelligent, 9
 - programming languages, 4
 - role of, 4–5
- solutions
 - commoditization of, 216–217
 - custom, 217
- spikes, detecting, 128–130
- SR-CNN (Spectral Residual and Convolutional Neural Network), 132

SSA (Singular Spectrum Analysis) algorithm, 132, 144, 149–150, 154
 decomposition step, 144–145
 reconstruction step, 145
 stateful neural networks, 197. *See also* RNN (recurrent neural network)
 stationarity, 143
 statistical market analysis, 31
 supervised classification, 123
 SVD (Singular Value Decomposition), 132–133
 SVM (Support Vector Machine) algorithm, 81–82, 99
 symbolic calculation, 1

T

tasks, 49

- anomaly detection, 119–120
 - change points, detecting, 130–131
 - composing the training pipeline, 128
 - fraud and, 139–140
 - loading data and feature engineering, 127–128
 - outliers in time series data, 121–122
 - semi-supervised learning, 124–125
 - setting up a client application, 134–136
 - spikes, detecting, 128–130
 - SR-CNN (Spectral Residual and Convolutional Neural Network), 132
 - SSA (Singular Spectrum Analysis), 132
 - statistical techniques, 122–123
 - supervised classification, 123
 - in a time series, 120–121
 - unsupervised clustering, 124
- binary classification, 73
 - calibrators, 74–75
 - choosing the algorithm, 81
 - dataset, 75–77
 - evaluating the model, 82
 - metrics, 83
 - partitioning the dataset, 77–78
 - programmatic holdout, 78–79
 - for sentiment analysis, 75
 - setting up a client application, 84
 - supported algorithms, 73–74
 - text featurization, 79–81
 - trainers, 74
 - validation techniques, 75
- classification, unsupervised learning, 99–100
- clustering, 114–115
 - applying persistent transformations, 101–102
 - binding output columns to a C# class, 112
 - customer segmentation, 100
 - DBSCAN algorithm, 107–108
 - elbow method, 106
 - feature engineering, 104
 - inspecting the transformed dataset, 111
 - K-Means algorithm, 109–111
 - K-Modes algorithm, 106
 - modeling data to classes, 103
 - OPTICS algorithm, 106–108
 - purpose of data in, 103–104
 - reducing the number of features, 115–116
 - reducing the number of rows, 116
 - saving clusters to separate files, 113–114
 - setting up a client application, 111
 - silhouette method, 106
- forecasting, 146
 - applying the algorithm, 149–150
 - dataset, 146–147
 - helper classes, 147
 - loading data from a database source, 148
 - saving and evaluating the model, 150–151
 - separating training and testing data, 148–149
 - weather, 156
- image classification, 175
 - composing the training pipeline, 180, 181–182
 - dataset, 178–179
 - mapping to a canonical classification problem, 178
 - setting up a client application, 182–184
 - transfer learning and, 175
- multiclass classification, 85, 87–88
 - confusion matrix, 94–95
 - dataset, 88
 - metrics, 93–94
 - object pooling, 96
 - setting up a client application, 95–96
 - supported algorithms, 85
 - trainers, 85–86
- naïve Bayesian classification, 86
- One-Versus-All classification, 86–87
- One-Versus-One classification, 87
- ranking, 160–162
 - information retrieval systems, 159–160
 - versus recommendation, 161
- recommendation, 159, 161, 163
 - accuracy and, 173
 - bulk rating, 173
 - client application, 170–172
 - collaborative filtering, 162, 172

tasks

- composing the training pipeline, 166
- dataset, 163–165
- evaluating the model, 168–169
- feature engineering, 165
- information retrieval systems, 159–160
- KNN (K-nearest neighbor) algorithm, 169–170
- MF (Matrix Factorization) algorithm, 166–168
 - personalization and, 161–162
 - versus ranking, 161
 - relevance by date, 166
- regression, 49–50, 54
- TensorFlow model, 12, 181
 - neural network training, 216
 - retraining, 181–182
- testing, 37, 52
 - cross-validation, 38, 65–66
 - holdout, 37, 52
 - k-fold, 52–53
 - datasets, 77–78
 - metrics, 38
 - validation, 64–65
- text featurization, 79–81
 - mapping target classes to numerical values, 90
 - multiclass classification and, 89
- Theano, 12
- thinking machines, 3. *See also* AI (artificial intelligence)
- time series
 - anomaly detection
 - change points, detecting, 130–131
 - randomized PCA (principle component analysis), 132–134
 - spikes, detecting, 128–130
 - SR-CNN (Spectral Residual and Convolutional Neural Network), 132
 - SSA (Singular Spectrum Analysis), 132
 - statistical techniques, 122–123
 - cycles, 142
 - multivariate, 126–127
 - outliers, 121–122
 - random walk, 154–155
 - regression, 71
 - seasonality, 142
 - stationarity, 143
 - trends, 142
 - univariate, 126–127
 - usual pattern, 120–121
- trainers
 - binary classification, 74
 - configuration, 50–51
 - multiclass classification, 85–86, 90–91
 - regression, 49–50
 - SDCA (Stochastic Dual Coordinate Ascent), 91
 - training, 12, 17–19. *See also* consuming a trained model; ETL (Extract-Transform-Load) pipeline;
- testing
 - algorithms, 35
 - choosing, 36
 - measuring the value of, 36–37
 - anomaly detection, 128
 - detecting change points, 130–132
 - detecting spikes, 128–130
 - setting up a client application, 134–136
 - using randomized PCA, 132–134
 - using the SR-CNN service, 132
 - classification, 90–92
 - confusion matrix, 94–95
 - evaluating the model, 92–94
 - managing multiple prediction engine pools, 96
 - selecting the algorithm, 90–91
 - setting up a client application, 95–96
 - switching back to text, 92
 - clustering, 109
 - inspecting the transformed dataset, 111
 - running the K-Means algorithm, 109
 - setting up a client application, 111
 - data transformation, 17
 - dataset, 77
 - forecasting, 148
 - applying the algorithm, 149–150
 - loading data from a database source, 148
 - saving and evaluating the model, 150–151
 - separating training and testing data, 148–149
 - image classification, 180
 - adding the TensorFlow model to the pipeline, 181
 - retraining the TensorFlow model, 181–182
 - multiclass classification, 90–91
 - neural networks, 194
 - backpropagation algorithm, 195–196
 - using an MRZ generator, 215–216
 - using TensorFlow generators, 216
 - overfitting, 53
 - performance and, 20–21
 - prediction, 62
 - cross-validation of the model, 65–66
 - fitting the model, 63
 - identifying the training algorithm, 62–63
 - invoking a trained model, 68–70

- loss function of the model, 63–64
- packaging the trained model, 66–67
- setting up a client application, 67–68
- validation of the model, 64–65
- preparing the dataset, 51
- purpose of, 19
- recommender system, 166
 - evaluating the model, 168–169
 - MF (Matrix Factorization) algorithm, 166–168
- regularization, 53
- transfer learning, 175, 178
 - dataset, 178–179
 - in ML.NET, 177
 - via composition, 176
- Transform method, 47
- trends, 142
- T-shaped skill set, 30
- Turing, Alan, 1, 3, 29
- Turing machine, 1–2

U

- univariate time series, 126–127
- unsupervised clustering, 124

- unsupervised learning, 99–100, 115
- user interface, 42–43

V

- validation. *See also* cross-validation
 - binary classification and, 75
 - dataset, 77–78
- Van Rossum, Guido, 10
- variance, 52, 58
- vertical solutions, 216–217
- Visual Studio, 15, 22, 62–63
- Von Neumann architecture, 4
- von Neumann, John, 2–3

W-X-Y-Z

- weather forecasting, 156
- Y-shaped skill set, 30
- Zip method, 151