



Hands-on Azure Cognitive Services

Applying AI and Machine Learning for
Richer Applications

Ed Price
Adnan Masood
Gaurav Aroraa



apress®

Hands-on Azure Cognitive Services

**Applying AI and Machine Learning
for Richer Applications**

**Ed Price
Adnan Masood
Gaurav Aroraa**

Apress®

Hands-on Azure Cognitive Services: Applying AI and Machine Learning for Richer Applications

Ed Price
Redmond, WA, USA

Gaurav Aroraa
Noida, India

Adnan Masood
Temple Terrace, FL, USA

ISBN-13 (pbk): 978-1-4842-7248-0
<https://doi.org/10.1007/978-1-4842-7249-7>

ISBN-13 (electronic): 978-1-4842-7249-7

Copyright © 2021 by Ed Price, Adnan Masood, and Gaurav Aroraa

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

Trademarked names, logos, and images may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, logo, or image we use the names, logos, and images only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Managing Director, Apress Media LLC: Welmoed Spahr
Acquisitions Editor: Smriti Srivastava
Development Editor: Laura Berendson
Coordinating Editor: Shrikant Vishwakarma

Cover designed by eStudioCalamar

Cover image designed by Pexels

Distributed to the book trade worldwide by Springer Science+Business Media LLC, 1 New York Plaza, Suite 4600, New York, NY 10004. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail orders-ny@springer-sbm.com, or visit www.springeronline.com. Apress Media, LLC is a California LLC and the sole member (owner) is Springer Science + Business Media Finance Inc (SSBM Finance Inc). SSBM Finance Inc is a **Delaware** corporation.

For information on translations, please e-mail booktranslations@springernature.com; for reprint, paperback, or audio rights, please e-mail bookpermissions@springernature.com, or visit <http://www.apress.com/rights-permissions>.

Apress titles may be purchased in bulk for academic, corporate, or promotional use. eBook versions and licenses are also available for most titles. For more information, reference our Print and eBook Bulk Sales web page at <http://www.apress.com/bulk-sales>.

Any source code or other supplementary material referenced by the author in this book is available to readers on GitHub via the book's product page, located at www.apress.com/978-1-4842-7248-0. For more detailed information, please visit <http://www.apress.com/source-code>.

Printed on acid-free paper

Thank you to Adnan and Gaurav for bringing their expertise, passion, and discipline to this book. I dedicate my portion of this book (it's probably whatever your favorite part is) to my oldest daughter, Eve. I wrote it when her teenage angstiness hit an all-time high, and she still let me hug her and hang out with her.

—Ed

Hi Mom and Dad!

Waves

—Adnan

My sincere thanks to Ed and Adnan. Without them, this book wouldn't be complete. The year was full of surprises and a lot of struggle throughout the globe. I dedicate this write-up to all the warriors of this era who came forward and helped to save humanity from this pandemic. And to my angel (daughter) Aarchi Arora whose smile always encourages and strengthens me to make me live for this project.

—Gaurav

Table of Contents

- About the Authors..... ix
- About the Technical Reviewer xi
- Acknowledgments xiii
- Introduction xv
- Chapter 1: The Power of Cognitive Services..... 1
 - Overview of Azure Cognitive Services 1
 - Why Azure Cognitive Services 2
 - Exploring the Cognitive Services APIs: Vision, Speech, Language, Web Search, and Decision 3
 - Vision APIs 4
 - Speech APIs..... 8
 - Language APIs 10
 - Web Search APIs..... 12
 - Decision APIs 14
 - Overview of Machine Learning 15
 - Understanding the Use Cases..... 18
 - The COVID-19 SmartApp Scenario 20
 - Summary..... 20
- Chapter 2: The Azure Portal for Cognitive Services..... 21
 - Getting Started with Azure Portal and Azure Cognitive Services 21
 - Getting Started with Azure Cognitive Services 24
 - Azure Marketplace: An Overview of AI + Machine Learning 24
 - Understanding Software Development Kits (SDKs): To Get Started with Your Favorite Programing Language..... 26

TABLE OF CONTENTS

Setting Up Your Visual Studio Template 27

Summary..... 35

Chapter 3: Vision – Identify and Analyze Images and Videos..... 37

Understanding the Vision APIs with Computer Vision 37

Analyzing Images..... 41

 Start Diving for Computer Vision 42

 Identifying a Face 59

 Understanding the Working Behavior of Vision APIs for Video Analysis..... 76

Microsoft Azure Video Indexer 89

Summary..... 96

Chapter 4: Language – Understand Unstructured Text and Models..... 97

Creating and Understanding Language Models 97

Enrichment in Progress – JFK Files Demo..... 99

The Muppet Models – Transformers for NLP..... 115

 Named Entity Recognition with Fine-Tuned BERT 116

Summary of the Language API..... 126

Chapter 5: Speech – Talk to Your Application..... 127

Understanding Speech and Speech Services 128

 Comprehensive Privacy and Security 128

 Getting Started 129

 Translating Speech Real Time into Your Application..... 130

 Speech to Text – Converting Spoken Audio to Text for Interaction 131

 Cognitive Speech Search with LUIS and Speech Studio..... 144

Summary of the Speech API 155

Chapter 6: Decision – Make Smarter Decisions in Your Applications 157

Decision Services and APIs..... 158

Content Moderator Service 159

 Trying It Out – Building Content Moderators 161

 Moderating Text Using Content Moderator Service 167

Moderating Text – Invoking the API	170
Personalizer Service	184
Trying It Out – Building a Movie Personalizer	187
Anomaly Detector Service.....	199
Try It Out – Anomaly Detector Demo.....	201
Metrics Advisor (Preview)	214
Summary and Conclusion	219
References and Further Reading	220
Chapter 7: Search – Add Search Capabilities to Your Application.....	221
The Search Ecosystem.....	222
Azure Cognitive Search.....	223
Searching with Azure Cognitive Search.....	223
Testing Azure Cognitive Search	232
Integrating Azure Cognitive Search in Your Notebooks	236
Searching with Bing Web Search.....	240
Summary and Conclusion	249
References and Further Reading	250
Chapter 8: Deploy and Host Services Using Containers	251
Cognitive Services Containers	252
Hosting Cognitive Services Containers	254
Running a Language Service Container	254
Running an Anomaly Detector Service Container	264
Working with Azure Container Instances	270
Deploying a Cognitive Services Container with Azure Kubernetes Service	281
Summary and Conclusion	289
References and Further Reading	290
Chapter 9: Azure Bot Services	291
The Azure Bot Services Ecosystem.....	292
Building Azure Service COVID-19 Bot.....	292

TABLE OF CONTENTS

Augmenting the Bot Knowledge Base from the Web 310

Conclusion and Summary 319

References and Further Reading 319

Chapter 10: Azure Machine Learning 321

 Azure Machine Learning Stack 322

 Hello World with Azure Machine Learning 330

 Creating a Machine Learning Model for Classification 338

 Classification Model Using Azure Machine Learning Designer 343

 Conclusion and Summary 353

 References and Further Reading 354

Index..... 355

About the Authors



Ed Price is Senior Program Manager in Engineering at Microsoft, with an MBA degree in technology management. Previously, he led Azure Global's efforts to publish key architectural guidance, ran Microsoft customer feedback programs for Azure Development and Data Services, and was a technical writer at Microsoft for six years, helping lead TechNet Wiki. Ed now leads Microsoft's efforts to publish reference architectures on the Azure Architecture Center (including a strong focus on AI architectures). He is an instructor at Bellevue College, where he teaches design and computer science. At Microsoft, he also helps lead volunteer

efforts to teach thousands of students how to code each year, focusing on girls and minorities. Ed is a co-author of six books, including *Azure Cloud Native Architecture Mapbook*, *Cloud Debugging and Profiling in Microsoft Azure* (Apress), and *Learn to Program with Small Basic*.



Adnan Masood, PhD, is an artificial intelligence and machine learning researcher, software engineer, Microsoft Regional Director, and Microsoft MVP for Artificial Intelligence. An international speaker and thought leader, Adnan currently works at UST as Chief AI Architect and collaborates with Stanford Artificial Intelligence Lab and MIT AI Lab on building enterprise solutions. Adnan has authored four books, including *Automated Machine Learning* and *Cognitive Computing Recipes* (Apress).

ABOUT THE AUTHORS



Gaurav Arora is Chief Technology Officer at SCL, with a doctorate in Computer Science. Gaurav is a Microsoft MVP award recipient. He is a lifetime member of the Computer Society of India (CSI), an advisory member, and senior mentor at IndiaMentor, certified as a Scrum trainer and coach, ITIL-F certified, and PRINCE-F and PRINCE-P certified. Gaurav is an open source developer and a contributor to the Microsoft TechNet community. He has authored ten books, including *Cloud Debugging and Profiling in Microsoft Azure* (Apress).

About the Technical Reviewer



Rohit Mungi is an Azure Service Engineer with Microsoft from Hyderabad, India, with more than 13 years of experience in various roles in development, support, and DevOps. His passion for cloud technologies helped him understand and bridge the gap between on-prem and cloud migrations for various customers. His regular day work involves helping customers onboard new cloud technologies in the Data and AI (artificial intelligence) space, particularly AI and machine learning offerings of Azure. He is very active on Microsoft Q&A forums and ready to learn and share his knowledge on these topics. He has also been recognized as a community ninja on Microsoft's tech community. Follow him on LinkedIn at [@rohit-mungi-0953678b](#).

Acknowledgments

To Rohit, our technical editor, for being so thorough, meticulous, detailed, careful, conscientious, diligent, scrupulous – we are running out of synonyms, but seriously, you rock.

To Shrikant for his extreme patience and regular follow-ups; thank you for putting up with us. Ghosting you is virtually impossible; we tried!

To Smriti for believing in us and this project!

Introduction

In about 400 BC, the Greek tale of Talos was told (and vividly illustrated). It was later chronicled in the epic poem, *Argonautica*, and its namesake is a giant automaton made of bronze, by Hephaestus, as requested by his father Zeus. This first robot (mythically speaking) sported wings, protected Europa (a woman) in Crete (an island), threw rocks at ships, and met his demise when Medea's spirits drove him crazy enough to pull out his own nail (thus, disassemble himself). In the 300s (BC), Aristotle brought us his syllogistic logic, a deductive reasoning system. Hero of Alexandria applied mechanics, basic robotics, and wrote *Automata* in the first century AD. AI in robots continued in myths with Pandora, who was also created by Hephaestus (you know, it's the story about the box), Pygmalion's Galatea (thanks to Venus, this is a statue of a woman who came alive to bear a child; don't ask), and then Jewish lore brought us Golems (in the early 1200s, AD), who are intelligent clay automatons that cannot speak.

Likely, in 1495, Leonardo da Vinci expanded on the many clockmakers' mechanical novelties to construct an automaton knight, and in 1515, we believe he continued his efforts to create a walking lion (presented to Francis I, the new King of France, and it's said to have been able to open its chest to present the gift of lilies). In 1642, Pascal invented a mechanical calculator. Swift's *Gulliver's Travels* (1726) seemed to have described an engine that's an AI computer and that forces knowledge on the poor Lilliputians. Punched cards came out in 1801, Mary Shelley scared us all with her intelligent *Frankenstein* in 1818, and in 1884, IBM got its start when Herman Hollerith invented a punched card tabulating machine to conduct the 1890 US census (that's the simplified explanation). At the 1939 World's Fair (in New York), Westinghouse unveiled Elektro, the mechanical man (he smoked), and Sparko, his mechanical dog.

Isaac Asimov wrote about robotics in 1950. Walt Disney (after having his Imagineers hack European wind-up novelties) developed a Dancing Man in 1961, he unveiled his tiki birds (which are in the tiki room) in 1963, he also recreated Abe Lincoln in 1963 for the 1964–1965 World's Fair, and he put the first animatronics in film with a bird in 1964's *Mary Poppins*. Ever since then, as computers evolved, we've been focused on building increasingly improved artificial intelligence in software and robotics (and we've

INTRODUCTION

been really hoping that our intelligent creations never take over and destroy us, like in *Terminator*, *Matrix*, or even *WALL-E*).

We intended this book to act as an introduction to Azure Cognitive Services, but also for it to become your guide to implementing Cognitive Services in real-world application scenarios. That means that we take you on a tour of everything that Cognitive Services has to offer, and, as you can imagine, that's quite a journey.

Azure Cognitive Services are broken down into services and APIs that explore decision, language, speech, and vision. With these APIs and technologies, you can create the future software that powers our computers and robots for centuries to come! (Or you can just make a silly website.) In this book, we take you on a journey to build an application, one step at a time. Please join us. After all, you're now walking in the footsteps of Aristotle, Hero, Hephaestus, Rabbi Loew (chief Golem mischief maker), Leonardo (not the turtle), Pascal, Gulliver, Mary Shelley, Hollerith, Elektro, Asimov, Disney, Bill Gates, Steve Jobs, and countless others. It's time to make a creation that's all your own.

CHAPTER 1

The Power of Cognitive Services

The terms **artificial intelligence (AI)** and **machine learning (ML)** are becoming more popular every day. Microsoft Azure Cognitive Services provides an opportunity to work with the top cutting-edge AI and ML technologies. To work with these technologies, we require some framework.

The aim of this first chapter is to set up the values, reasons, and impacts that you can achieve through Azure Cognitive Services. The chapter provides an overview of the features and capabilities. In the upcoming sections, you will understand how Azure Cognitive Services is helpful and how it makes it easy for you to work with AI and ML.

We also introduce you to our case study and the structures that we'll use throughout the rest of the book.

In this chapter, we cover the following topics:

- Overview of Azure Cognitive Services
- Exploring the Cognitive Services APIs: Vision, Speech, Language, Web Search, and Decision
- Overview of machine learning
- Understanding the use cases
- The COVID-19 SmartApp scenario

Overview of Azure Cognitive Services

Microsoft Azure Cognitive Services provides you with the ability to develop smart applications. You can build these smart applications with the help of APIs, SDKs (software development kits), services, and so on.

Microsoft Azure Cognitive Services is a set of APIs, SDKs, and services that facilitate developers to create smart applications (without the prior knowledge of AI or ML).

Azure Cognitive Services provides everything that developers need in order to work on AI solutions, without the knowledge of data science. A developer can create a smart application that can converse, understand, or train itself.

Why Azure Cognitive Services

Azure Cognitive Services is backed by world-class model deployment technologies, and it is built by top experts in the area. There are a lot of plans and offers that use the pay-as-you-go model. You no longer have to invest in the development and infrastructure that you may need in order to build and host your models. Cognitive Services provides all this for you.

The following list shows the advantages you gain when you use Azure Cognitive Services:

- You don't need to build your own custom machine learning model.
- You gain a required AI service for your app. Azure Cognitive Services, as a **Platform as a Service (PaaS)**, can offer these required features.

*You can build upon a **Platform as a Service (PaaS)** without being concerned about the infrastructure used to support the service.*

- You can invest your development time in the core app and release a stronger product.

Note You should not use Azure Cognitive Services if it doesn't meet your requirements. For example, your data might have regulatory requirements that stop you from using an external service, like Azure. Or your organization might have a long-term commitment toward developing its own data science practices and product.

In the next section, we will discuss the Cognitive Services APIs in more detail.

Exploring the Cognitive Services APIs: Vision, Speech, Language, Web Search, and Decision

In the preceding section, we discussed Cognitive Services and the advantages that it provides. In this section, we will explore the APIs that are available to help developers.

Figure 1-1 provides a pictorial overview of these APIs.



Figure 1-1. Pictorial overview of Cognitive Services APIs

Figure 1-1 displays the following elements of the Azure Cognitive Services APIs:

- I – Represents all the Azure Cognitive Services APIs
- II – Represents the developer who consumes these APIs to build smart apps
- 1 – Vision APIs
- 2 – Speech APIs
- 3 – Language APIs
- 4 – Web Search APIs
- 5 – Decision APIs

The Vision APIs provide insights on images, handwriting, and videos. The Speech APIs analyze and convert audio voices. The Language APIs offer you text analysis, they can make text easier to read, they help you create intelligent chat features, and they can translate text. The Bing Web Search APIs allow you to search and pull content from the entire Internet, leveraging pages, text, images, videos, news, and more. Finally, the Decision APIs help your app make intelligent decisions to moderate and personalize content for your users, and they help you detect anomalies in your data.

In the upcoming sections, we will briefly introduce you to each of these sets of APIs.

Vision APIs

First, let's explore the Vision APIs from Azure Cognitive Services. Use these APIs whenever you need to work with images or videos, to understand or analyze their contents. These APIs help you to get information like a facial analysis (determining age, gender, and more), feelings (e.g., through facial expressions), and more visual contents. Furthermore, with the help of these APIs, you can read the text from images, and thumbnails can be easily generated from images and/or videos. The Cognitive Services Vision APIs are divided into the following APIs, as detailed in the following.

Computer Vision

The Computer Vision API allows the developer to analyze an image and its contents. In the previous section, we discussed that with the help of Vision APIs, you can understand and collect the image contents. You can decide what content and information to retrieve

from an image, based on your requirements. For example, a business might need to access the images in order to help make sure kids using their web app will avoid viewing adult content.

This API can also read printed text, hand-written text via optical character recognition (OCR).

Note The current version of the Computer Vision API is v3.0, at the time of writing this book.

From a development perspective, you can either use RESTful (representational state transfer) APIs or you can build applications using an SDK. We will cover the development instructions and details in [Chapter 3](#).

Custom Vision

The Cognitive Services APIs for Custom Vision provide a way to customize images with various customizations. You can customize the images with labels, and you can assess and improve these images based on the customization classifiers. The Custom Vision APIs use machine learning algorithms and apply labels to assess and improve the images. Furthermore, it is divided into two parts:

1. **Image classification** – Applies the label to an image.
2. **Object detection** – Applies the label, and it returns the coordinates from the image where the label is located.

Face

This API helps you detect and analyze human faces from an image. The algorithm can detect and analyze the data.

This service provides the following features:

- **Face detection** – Detects a human face and provides the coordinates of where the face is located in the image. Based on the algorithm, you can also get various properties of face detection, such as gender, the head pose, emotions, age, and so on. [Figure 1-2](#) shows the face detection of a human (the author, Gaurav Arora).



Figure 1-2. *Human face detection using the Face API*

- **Face verification** – Verifies two similar faces from images of one human face to compare them, in order to find out whether it belongs to the same human. Figure 1-3 shows two faces of the same human.




Figure 1-3. *Face verification*

- **Face grouping** – Groups the similar faces from an available database or set of faces.

Note During your development cycle, the Face API and its data must meet the requirements of the privacy policies. You can refer to the Microsoft policies on customer data here: <https://azure.microsoft.com/en-us/support/legal/cognitive-services-compliance-and-privacy/>.

Form Recognizer

Form Recognizer extracts the data in key/value pairs and extracts the table data from a form-type document.

This is made with the following components:

- **Custom models** – Enables you to train your own data by providing five-form samples.
- **Prebuilt receipt model** – You can also use the prebuilt receipt model. Currently, only English sales receipts from the United States are available.
- **Layout API** – It enables Form Recognizer to extract the text and table structure data, by using optical character recognition (OCR).

Video Indexer

Video Indexer provides a way to analyze a video's contents, by using three channels: voice, vocal, and visual. In this way, you will get insights about the video, even if you don't have any expertise on video analysis. It also minimizes your efforts, as there is no need to write any additional or custom code.

Video Indexer provides us a way to easily analyze our videos, and it covers the following categories:

- Content creation
- Content moderation
- Deep search
- Accessibility
- Recommendations
- Monetization

We will cover video analysis more thoroughly in Chapter [3](#).

Speech APIs

Speech APIs provide you a way to make your application smarter. Thus, your application can now listen and speak. These APIs filter out the noise (words and sounds that you don't want to analyze), detect speakers, and then perform your assigned actions.

Speech Service

Microsoft introduced the Speech service to replace the Bing Speech API and Translator Speech. These are the services that provide an extraordinary effect to your application, in such a way that your application can hear users and speak/interact with your users.

Note You can also customize Speech services by using frameworks. For speech to text, refer to <https://aka.ms/CustomSpeech>. For text to speech, refer to <https://aka.ms/CustomVoice>.

The Speech service enables the following scenarios:

- Speech to text
- Text to speech
- Speech translation
- Voice assistants

With the help of different frameworks, you can also customize your Speech experience.

Speaker Recognition (Preview)

Speaker Recognition is in Preview, at the time of writing this book. This service enables you to recognize the speakers; you can determine who is talking. With the help of this service, your application can also verify that the person that is speaking is who they claim to be. So, it is now much easier for your application to identify unknown speakers from a group of potential speakers.

It can be divided into these two parts:

- Speaker verification
- Speaker identification

We will cover voice recognition in detail in Chapter 5.

Language APIs

With the help of prebuilt scripts, the Language APIs enable your application to process the natural language. Also, they provide you the ability to learn how to recognize what users want. This would add more capabilities to your application, like textual and linguistic analysis.

Immersive Reader

Immersive Reader is a very intelligent service that builds a tool to help every reader, especially people affected with dyslexia.

Note Dyslexia affects that part of the brain that processes language. People with dyslexia have difficulty reading, and they can find it very challenging to identify the sound in written speech.

Immersive Reader is designed to make it easier for everyone to read.

It provides the following features:

- Reads textual content out loud
- Highlights the adjectives, verbs, nouns, and adverbs
- Graphically represents commonly used words
- Helps you understand the content in your own translated language

Language Understanding (LUIS)

Think of a scenario where you need to make your application smart enough, so that it can understand user input (such as speech, text, and so on). The Speech service makes your application smart enough to listen and speak with the user. But your application might need to be smart enough to answer a question that your user asks it, such as, *“What is my health status?”* Even after implementing the Speech APIs, your application will not be ready to understand commands like that. To achieve such a complicated requirement, we have the Language Understanding (LUIS) services. (LUIS stands for Language Understanding Intelligent Service.) With the help of LUIS, you can build an application that interacts with users and pulls the relevant information out of the

conversations. For a question like, “*What is my health status?*”, your application can assess the stored data, and it then provides the status of the user’s health. Or it asks a few questions, and based on the user’s answer, it would then provide the user’s health status.

You can work with the following two types of models:

- Prebuilt model
- Custom model

Learn more about LUIS in Chapter 4.

QnA Maker

QnA is very relevant, when you have an FAQ and want to make it interactive. This means that you have a predefined set of QnA (questions and answers). QnA is mostly used in chat-based applications, where the user enters queries, and then your application answers the question. You can try using Microsoft’s www.qnamaker.ai/ to enable your experience with QnA Maker.

Text Analytics

With the help of the Text Analytics service, you can build an application that analyzes the raw text and then gives you the result. It includes the following functions:

- Sentiment analysis
- Key phrase extraction
- Language identification
- Name identifications

Translator

Translator enables text-to-text translation, and it provides a way to build translation into your application. With the help of Translator, you can add multilingual capabilities to your application. Currently, more than 60 languages are supported. If you want to translate a spoken speech, you will need to use the Speech service.

Web Search APIs

The Web Search APIs enable you to build more intelligent applications, and they give you the power of Bing Search. They allow you to access data from billions of web pages, images, and news articles (and more), in order to build your search results.

Bing Search APIs

Bing Search facilitates your application by providing the ability to do a web search. You can imagine that with the implementation of Bing Search APIs, you now have a wide range of web pages with which to build out your search results. The code implementation is very easy as well (see Listing 1-1).

Listing 1-1. The sample code to implement Bing Web Search

```
//Sample code
public static async void WebResults(WebSearchClient client)
{
    try
    {
        var fetchedData = await client.Web.SearchAsync(query: "Tom
        Campbell's Hill Natural Park");
        Console.WriteLine("Looking for \"Tom Campbell's Hill Natural
        Park\"");

        // ...
    }
    catch (Exception ex)
    {
        Console.WriteLine("Exception during search. " + ex.Message);
    }
}
```

Bing Web Search

With the Bing Web Search API, you can suggest search terms while a user is typing, filter and restrict search results, remove unwanted characters from search results, localize search results by country, and analyze search data.

Bing Custom Search

The Bing Custom Search API allows you to customize the search suggestions, the image search experience, and the video search experience. You can share and collaborate on your custom search, and you can configure a unique UI for your app to display your search results.

Bing Image Search

The Bing Image Search API enables you to leverage Bing's image searching capabilities. You can suggest image search terms, filter and restrict image results, crop and resize images, display thumbnails of the images, and showcase trending images.

Bing Entity Search

The Bing Entity Search API gets you search results of entities and locations, such as restaurants, hotels, or stores. You can provide real-time search suggestions, entity disambiguation (providing multiple search results), and return information on businesses and other entities.

Bing News Search

The Bing News Search API lets you search for relevant news articles. You can suggest search terms, return news articles, showcase the day's top news articles and headlines, and filter news results by category.

Bing Video Search

The Bing Video Search API returns high-quality videos. You can suggest search terms, filter and restrict the results of the video query, create thumbnail previews of the videos, showcase trending videos, and access video insights.

Bing Visual Search

The Bing Visual Search API provides metrics and insights for an individual image, whether one that is uploaded or shared via a URL. You can identify similar images and products, identify shopping sources for the image, access related searches made by others (based on the content of your image), find the web pages that display the image, find recipes for a dish featured in your image, and use the image to automatically gather information about an entity (such as information about an actor in an image or directions to a location in the image).

Bing Autosuggest

The Bing Autosuggest API can improve your search experience by returning a list of suggested queries, based on a partial query.

Bing Spell Check

The Bing Spell Check API performs grammar and spell checking on any text. You can check for slang or informal language, differentiate between similar words, and track new brands, titles, and popular expressions.

Bing Local Business Search

The Bing Local Business Search API helps you find information about a business. You can find the nearest restaurant to you (such as a specific chain or a type of food), locate and map out specific businesses and places, limit the distance parameters of your search, and filter the business results by category.

Decision APIs

Decision APIs help you create a rich user experience that is personalized, with content moderation. They help you perform efficient decision making.

Anomaly Detector (Preview)

Anomaly Detector is in Preview, at the time of writing this book. It helps you detect abnormalities in the data, in order to remove them. When you use this service, you don't need to think about what model is fit for your business. The Anomaly Detector service automatically determines the best-fit model of your data.

Content Moderator

Content Moderator helps you identify the content that your business does not allow or content that does not fit your business. It checks text, image, and video contents as well. For example, you can remove profanity and undesirable text, you can identify videos that contain profanity, you can moderate adult and racy video and image content, you can check text for personally identifiable information (PII), and you can detect other offensive or unwanted images.

Personalizer

Personalizer helps you give a personal experience to your user individually. It helps you show different and personalized content, to individual users. The content can be within text, images, URLs, or emails.

Overview of Machine Learning

There are a lot of possible definitions of what machine learning is. According to all the available definitions, we can conclude the following.

Machine learning is a subset AI that provides you with a way to study computer algorithms. In this process, after studying the data, computer programs can forecast future events.

Our learning process is simple and continuous. The process happens anytime and anywhere, like while you're walking, making a coffee, or purchasing items. Every learning has its own scope and boundaries. For example, while you are walking, your smartwatch (a machine fitted with a program) collects your data (such as step counts, miles, heartbeats, and your blood pressure). For another example, while you visit the department store for purchases, technology systems (like loyalty cards and supporting software) learn about your purchase behavior, based on the items that you've purchased in the past.

With this information, machines can predict the future. For example, the machine that learns purchase behavior can predict which items would be ordered by their customers. For another example, it helps the store owners understand how much a customer consumes a particular item.

Note In summary, machine learning is a study that helps us understand future behavior. We achieve this by studying data in various aspects, such as analyzing text for emotional sentiment or analyzing images to recognize objects or faces.

There are several phases in the machine learning process. These phases are iterative and require repetition, based on the requirement and/or training. Here are the most common phases:

- **Data collection phase** – In this phase, the data is being prepared or collected. You determine what data to collect, based on the need of training or the component that is required for training.
- **Test and develop training models** – This is the phase where the learning is being tested and where the training models are developed. For example, to build a new training model for a department store, we need to analyze the grocery items. You would gather and access the data by using algorithms.
- **Deploy and manage trained models** – This is the phase where you can see the actual output. You can also call this phase *machine learning in action*.

Microsoft made this easy for developers, with the help of various plans. You can follow the preceding phases and manage your machine learning models, using the available Microsoft technologies.

The following various Microsoft technologies are available, to work with ML:

- **Cloud based** – Microsoft Azure provides a lot of options, where you can use cloud services and work with machine learning.
- Cognitive Services APIs (such as Vision, Speech, Language, and Web Search) provide sophisticated pretrained ML models.