



Building Intelligent Apps with .NET and Azure AI Services

Start Your Journey in Building
Intelligent Solutions

—
Ashirwad Satapathi

Apress®

Building Intelligent Apps with .NET and Azure AI Services

**Start Your Journey in Building
Intelligent Solutions**

Ashirwad Satapathi

Apress®

Building Intelligent Apps with .NET and Azure AI Services: Start Your Journey in Building Intelligent Solutions

Ashirwad Satapathi
Gajapati, Odisha, India

ISBN-13 (pbk): 979-8-8688-0434-2

ISBN-13 (electronic): 979-8-8688-0435-9

<https://doi.org/10.1007/979-8-8688-0435-9>

Copyright © 2024 by Ashirwad Satapathi

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
Editorial Assistant: Kripa Joseph

Cover designed by eStudioCalamar

Cover image designed by Freepik (www.freepik.com)

Distributed to the book trade worldwide by Springer Science+Business Media New York, 1 New York Plaza, Suite 4600, New York, NY 10004-1562, USA. 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.

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. For more detailed information, please visit <https://www.apress.com/gp/services/source-code>.

If disposing of this product, please recycle the paper

*This book is dedicated to my father, Mr. Upendra Satapathi,
and mother, Mrs. Sabita Panigrahi, for supporting me through
each and every phase of my life.*

Table of Contents

About the Author	ix
About the Technical Reviewer	xi
Chapter 1: Introduction.....	1
Structure	1
Objectives	2
What Is Artificial Intelligence?.....	2
Introduction to Azure AI Services	3
A Quick Tour of Azure AI Services	4
Azure AI Search	4
Azure OpenAI	4
Language.....	5
Speech.....	5
Translator.....	6
Bot Service	6
Content Safety	6
Document Intelligence.....	7
Immersive Reader	7
Video Indexer.....	8
Vision	8
Custom Vision.....	8
Face.....	9
Summary.....	9

TABLE OF CONTENTS

- Chapter 2: Build a Language-Based Document Classifier with Azure Functions 11**
 - Structure 12
 - Objectives 12
 - Introduction to Azure AI Language Service 12
 - Problem Statement 14
 - Proposed Solution 14
 - Create an Azure AI Language Service 15
 - Create an Azure Storage Account 22
 - Create a Language-Based Document Classifier with Azure Functions 30
 - Test the Language-Based Document Classifier Function 38
 - Summary..... 42

- Chapter 3: Build a Multi-language Text Translator App with Azure Functions 43**
 - Structure 44
 - Objectives 44
 - Introduction to Azure AI Translator Service 44
 - Problem Statement 46
 - Proposed Solution 47
 - Create an Azure AI Translator Service 47
 - Create a Multi-language Text Translator App 53
 - Test the Multi-language Text Translator Function with Postman 64
 - Summary..... 65

- Chapter 4: Build a Desktop App with .NET MAUI to Generate Texts from Audio Files 67**
 - Structure 67
 - Objectives 68
 - Introduction to Azure AI Speech Service 68
 - Problem Statement 70
 - Proposed Solution 70
 - Create an Azure AI Speech Service 71

Create a Desktop App with .NET MAUI to Generate Texts from Audio Files.....	77
Test the .NET MAUI App.....	85
Summary.....	87
Chapter 5: Build a Desktop App with .NET MAUI to Extract Text from Images	89
Structure	89
Objectives	90
Introduction to Azure AI Vision Service	90
Problem Statement	92
Proposed Solution	92
Create an Azure AI Computer Vision Service.....	93
Create a Desktop App with .NET MAUI to Extract Text from Images	99
Test the .NET MAUI App.....	107
Summary.....	110
Chapter 6: Build a Web App to Extract Data from Invoices Using Azure AI Document Intelligence	111
Structure	111
Objectives	112
Introduction to Azure AI Document Intelligence Service.....	112
Problem Statement	114
Proposed Solution.....	114
Create an Azure AI Document Intelligence Service	115
Build a Web App to Extract Data from Receipts Using Azure AI Document Intelligence.....	120
Test the Web App.....	141
Summary.....	143
Chapter 7: Build a Content-Flagging App with Azure AI Content Safety	145
Structure	145
Objectives	146
Introduction to Azure AI Content Safety Service	146
Problem Statement	147

TABLE OF CONTENTS

- Proposed Solution 148
- Create an Azure AI Content Safety Service 148
- Build a Content-Flagging App with Azure AI Content Safety Service 155
- Test the Content-Flagging App..... 166
- Summary..... 167
- Chapter 8: Build a Text Summarizer with Azure OpenAI 169**
- Structure 169
- Objectives 170
- Introduction to Azure OpenAI Service 170
- Problem Statement 172
- Proposed Solution 172
- Create an Azure OpenAI Service 173
- Build a Text Summarizer with Azure OpenAI 184
- Test the Text Summarizer App..... 195
- Summary..... 197
- Index..... 199**

About the Author



Ashirwad Satapathi works as a software engineer at Microsoft and has expertise in building scalable applications with .NET Core. He has a deep understanding of building full-stack applications using .NET and Azure PaaS and serverless offerings. He is also an active blogger in the C# developer community. He has been awarded the C# Corner Most Valuable Professional (MVP) in September 2020 and September 2021 for his remarkable contributions to the developer community.

Ashirwad is an active speaker and delivers sessions on Blazor and Microsoft Azure. He has spoken for multiple communities such as Microsoft Reactor Bangalore, UTF, KonfHub Tech Conferences, and ServerlessDays Amsterdam. He is also an active community organizer and member of Utkal Techies Forum (UTF), a developer community based out of Odisha, India, and helps organize events for the community. He is a member of the Outreach Committee of the .NET Foundation. In addition, he has started ServerlessDays Bhubaneswar and hosts virtual sessions to build awareness among the developer community in the region about serverless technologies and make them proficient enough to build highly scalable and efficient serverless applications.

About the Technical Reviewer



Viktoria is an experienced team leader of blockchain projects. An expert in .NET and NestJS, she designs architecture and develops projects from scratch to finish. She has authored several articles and books as well as is a frequent speaker at online conferences.

CHAPTER 1

Introduction

With the rapid development in the technology world, AI has become an integral part of our lives. Starting from unlocking your smartphones with facial recognition to getting personalized product recommendations while shopping on ecommerce websites, the impact of AI can be felt everywhere. Nowadays, most of the applications have become AI powered in one way or the other. Instead of just being a good feature, it has become one of the core functionalities of modern apps to provide better user experience.

Building AI-powered intelligent applications comes with its own fair set of challenges. To add the AI capabilities, you require a team of AI researchers and domain experts to build AI models which can be leveraged by the application. Training and developing such models require a large amount of data and computational power. This requires up-front investments in setting up infrastructure and collection and preparation of data, which are both capital and time intensive.

A potential solution to overcome these hurdles is to leverage the AI models offered as cloud services by different cloud vendors. Azure AI Service is one such cloud service offered by Microsoft Azure. It enables organizations and developers to leverage prebuilt AI models created by a team of researchers at Microsoft to build intelligent solutions without worrying about the overhead of managing the underlying infrastructure.

The focus of this chapter will be to get a brief understanding about the various services which are part of the Azure AI Services. In the subsequent chapters, we will learn about ways to integrate these services in our applications to build intelligent solutions.

Structure

In this chapter, we will explore the following aspects of Azure:

- What is artificial intelligence?
- Introduction to Azure AI Services
- A Quick tour of Azure AI Services

Objectives

After studying this chapter, you should be able to

- Grasp the essentials of Azure AI Services
- Identify applicable scenarios for Azure AI Services

What Is Artificial Intelligence?

Artificial intelligence is a field of computer science which works toward building computational systems which can perform tasks that require human intelligence. The term AI was first coined by John McCarthy in 1956 during the Dartmouth summer research project. As an outcome of this project, AI came out as a field of study to do collaborative research aimed at building computational machines that can simulate human intelligence to solve complex problems.

Early progress in the field led to the birth of expert systems which focused on leveraging rules and symbols for knowledge and reasoning processing. With dwindling research funding and technological limitations, the field saw a period of stagnation during the 1970s and 1980s. This period is referred to as the AI winter. With the affordability of computational power and the advent of big data, the field has seen substantial growth over a period of time.

The field of AI encompasses various subfields like machine learning, natural language processing, computer vision, speech analysis, expert systems, reinforcement learning, and robotics. Often, these subfields overlap, and advancement in one greatly contributes to the progress of the other.

Today, AI has transformed from a futuristic concept to being an integral part of our day-to-day lives. It helps in performing various tasks ranging from predicting the weather forecast for a location to analyzing videos and generating insights. Some popular use cases of AI that are part of our daily lives are as follows:

1. Voice assistants are a common use case of AI. For example, using Siri on your iPhone to set a timer involves simply saying, “Hey Siri, set a timer for ten minutes.” Siri understands your voice instructions and sets the timer accordingly.

2. Product recommendations in ecommerce are a prominent use case of AI. For example, after purchasing a book on Amazon, the platform suggests other books that you may like based on your purchase history and buying patterns of similar users.
3. Content moderation in social networking sites is a key use case of AI. Social media platforms like Facebook and Instagram utilize AI to identify and remove inappropriate content which in turn enhances user safety.

In this book, we will explore ways to add AI to our application by leveraging the power of Azure AI Services.

Introduction to Azure AI Services

The demand for building AI-powered solutions is at an all-time high. But creating AI models from scratch comes with its own set of challenges. Developing and optimizing AI models from scratch require deep technical and domain expertise in the field and a considerable amount of time and resources. This is a capital- and time-intensive process. For many developers, it was a big obstacle as learning and gaining AI had a steep learning curve which restricted their ability to integrate AI into their application. To address this issue, we can use services offered by third-party cloud vendors that enable developers to use their prebuilt AI models via REST API calls or through client libraries instead of building one from scratch. Azure AI Service is one such offering of Microsoft Azure which enables developers and organizations to add AI to their applications with ease.

Utilizing Azure AI Services enables the development of solutions characterized by reliability, strong security measures, high availability, and fault tolerance, enriched with intelligent functionalities spanning areas such as vision, speech, language understanding, and search. It comes with an easy-to-use UI interface, with readily accessible solutions combined with responsible AI principles, empowering organizations to overcome the challenges of integrating AI in their applications. Organizations, regardless of size, can utilize them to create intelligent, responsible, and market-ready applications, delivering not only business value but also addressing real-world challenges.

A Quick Tour of Azure AI Services

In this section, we are going to explore the suite of services available through Azure AI Services.

Azure AI Search

Azure AI Search, formerly known as Azure Cognitive Search, is a fully managed enterprise-grade, search-as-a-service offering by Microsoft Azure. With Azure AI Search, organizations and developers are empowered to build information retrieval systems which can find relevant data for queries from their data sources without going through the hassles of building a search engine of their own. Apart from supporting full text search, it is also able to find semantically similar information by performing vector search throughout the vector representations of your data.

Azure AI Search has various applications across industries like ecommerce and healthcare, to name a few. We leverage its power to build solutions for enterprise search, knowledge mining, or geospatial search. By integrating Azure AI Search, we reduce the time to market of our solutions and provide a rich search experience for our end users.

Azure OpenAI

Azure OpenAI is a managed service offered by Microsoft Azure which enables enterprises and developers to use the AI models offered by OpenAI. With Azure OpenAI, we can integrate various pretrained models like GPT 4, GPT 3.5, Embedding, and DALL-E in our applications with minimal effort. With GPT 4 and GPT 3.5 series of models, you can create a conversational experience like ChatGPT in your applications. Developers can also fine-tune the foundational models to train with their own data to retrieve answers based on their data. In the entire process, the data you use to train the models stays within the boundary of your tenant and is not used to train any of the foundational models.

Azure OpenAI has various applications and is one of the most popular cloud services in Microsoft Azure. With most of the models, we can integrate various capabilities into our application, such as text generation, summarization, and language translation, to name a few. With the DALL-E model, we can also generate visual content by providing

prompts. Prompts are textual queries that we pass to the models as input. With models like GPT 4 with Vision, we can pass visual content, and it can generate answers for our prompts by analyzing the visual content. At the time of writing this book, GPT 4 with Vision was in public preview. We will explore more about Azure OpenAI in Chapter 8.

Language

The Azure AI Language service is an enterprise-grade managed cloud service offered by Microsoft Azure to process and gather insights from unstructured textual data. With the help of the Azure AI Language service, organizations and developers can enable natural language processing capabilities such as sentiment analysis, named entity recognition, personally identifiable information extraction, language detection, and text summarization in their applications. It also enables them to customize AI models to generate insights specific for our use cases. A potential use case for this could be a scenario where we want to categorize documents into user-defined categories.

One of the most common use cases of Azure AI Language that can be applied across industries is analyzing customer sentiments from the feedback received for the service or product offered by the companies. We are going to explore more about the Azure AI Language service in Chapters 2 and 3 where we build solutions to perform sentiment analysis and language detection.

Speech

The Azure AI Speech service is a fully managed cloud service offered by Microsoft Azure which enables developers to add speech capabilities to their applications. Speech to text, text to speech, speech translation, and speaker recognition are some of the key functionalities of the Speech service. At the time of writing this book, it can accurately analyze audios from 100 languages and generate their transcripts. With functionality like custom neural voice, you can train a custom model on your voice and generate audio in multiple languages.

The Azure AI Speech service can be leveraged for industry use cases such as caption generation, audio content creation, and adding voice capabilities to bots. A key use case for the Azure AI Speech service would be to build accessible solutions for individuals with visual or hearing impairments. With the help of the speech-to-text and text-to-speech capabilities, we can make our apps more accessible. In Chapter 4, we are going to explore ways to integrate the speech-to-text feature of the Speech service in our application.

Translator

The Azure AI Translator service is a fully managed cloud-based machine translation service of Microsoft Azure. With the Translator service, developers can easily integrate machine translation capabilities like text translation, transliteration, and language detection to their applications with minimal development efforts. We can also create custom models with the custom translator feature to build systems which can handle domain- or industry-specific terminologies.

The Azure AI Translator service has various use cases across industries. A potential use case can be content localization in ecommerce websites. With the help of the Translator service, we can translate product descriptions, reviews, and other information in the local language of the customers. In Chapter 3, we are going to explore ways to integrate the language translation feature of Translator services in our application.

Bot Service

Azure Bot Service provides a development platform to build, test, manage, and deploy chatbots to enable conversational experience for business. The platform comes up with tools like Bot Framework SDK and Bot Framework Composer. We can build, test, and maintain bots by leveraging these tools. Azure Bot Services enable developers and organizations to deploy their bots to multiple channels such as web apps, Microsoft Teams, Skype, or Slack. It easily integrates with other Azure AI Services. This enables us to build intelligent bots which can understand natural language, interpret speech, and address queries. Azure Bot Service provides metrics like latency and traffic to monitor the health of the bots.

Azure Bot Service enables organizations and developers to build chatbots which can cater to various use cases, including but not limited to bots which can automate internal processes like IT support or handle common scenarios like answering commonly asked customer queries.

Content Safety

The Azure AI Content Safety service is a fully managed cloud-based solution for content moderation. It can moderate text and image content to detect the presence of any harmful content. With the help of the Content Safety service, developers can build systems which can analyze user-generated content to detect the presence of potentially

inappropriate or unsafe material. If detected, it provides the ability to filter them out. It classifies objectionable content into four main categories which are hate and fairness, sexual, violence, and self-harm with a severity score.

Content monitoring in social media platforms like Facebook and Twitter and educational platforms to prevent the spread of hate messages and inappropriate content is a potential use case of the Content Safety service. In Chapter 7, we are going to explore ways to integrate the Content Safety service in our solutions.

Document Intelligence

Azure AI Document Intelligence, formerly known as Form Recognizer, is a fully managed cloud service for text extraction from documents by using advanced machine learning algorithms. With the help of document intelligence, developers can build a document processing system which can process documents to extract data from documents like invoices, receipts, and forms irrespective of the layouts. It is capable of extracting printed as well as handwritten texts from documents. With the custom model features, we can train custom models on our documents to learn the structure of the document to extract data efficiently.

Invoice and claims processing systems are classic examples where we can leverage the power of document intelligence to process documents and extract key information. In Chapter 6, we are going to explore ways to integrate the Document Intelligence service in our solutions.

Immersive Reader

Azure AI Immersive Reader is an AI service of Microsoft Azure which is designed to enhance the reading experience for users with different abilities. With features like reading aloud, isolating content for improved readability, and translating content in real time, the Immersive Reader aims at making lives easier for individuals who are new to reading, people who are in the process of learning a new language, or people who are diagnosed with dyslexia. Developers can build accessible solutions by leveraging the Immersive Reader library.

Video Indexer

Azure AI Video Indexer is a cloud-based video analytics service of Microsoft Azure. Video Indexers leverage more than 30+ AI models to extract insights from videos. Developers can leverage the power of the Video Indexer service to build applications which can perform operations such as face detection and recognition, content creation, transcript generation, or textual logo detection over stored video files. One of the popular use cases of Video Indexers is to generate content like trailers and highlight reels from stored videos. It can also be used to build monitoring systems which can detect if a particular object or face was identified in a particular video.

Vision

The Azure AI Vision service is a cloud-based service provided by Microsoft Azure to perform image processing and analysis. It leverages advanced machine learning models to extract key insights from visual content. With the help of the Vision service, we can add a range of computer vision capabilities like image analysis, spatial analysis, Optical Character Recognition (OCR), and image classification in our application. The Vision service can understand and extract information from images in many different languages. It comes with a vision studio where you can explore the various features of the Vision service. Apart from that, we can leverage the client SDKs to interact with the AI models of the Vision service.

The Vision service has wide application across industries like healthcare, manufacturing, media, and retail. Facial detection and image captioning are some of the popular use cases of the Vision service. In Chapter 5, we are going to explore ways to integrate the text extraction feature of the Azure AI Vision service in our solutions.

Custom Vision

Azure Custom Vision is a cloud offering of Microsoft Azure which enables teams to build, train, test, and deploy their own custom image classification models. It provides an interactive UI interface to build and test our own custom vision models in the browser itself. We can leverage the prebuilt vision models to create our custom vision models. This feature is a lifesaver when we don't have a large corpus of labeled data to train our models. By leveraging the existing models, we can accelerate the training process of our custom models, resulting in shorter time to market.

Product recognition can be one of the many use cases of Azure Custom Vision. With the Custom Vision service, we can train prebuilt models with the labeled data of our products. This can enable teams in the retail industry to identify the product Stock-Keeping Units (SKUs) present on the shelf and get their counts. The possibilities are endless.

Face

The Azure AI Face service is a fully managed cloud service offered by Microsoft Azure which enables teams to build cloud-native solutions with facial recognition features. We can add capabilities like face verification for touchless access control and emotion detection in our application with the help of the Face service. It is also able to detect the gender and age of the person present in the image. A popular use case for the Face service is to perform liveness checks while performing authentication with facial recognition. With the liveness check features, we can validate if a user is actually present in front of the camera or someone is trying to impersonate the person's identity by using a picture or video of theirs. We use the REST APIs or client libraries of the Face service to integrate them in our application with minimal coding efforts.

Summary

Azure AI Service is one of the fastest-growing cloud services within the Microsoft Azure ecosystem. It enables enterprises and developers to build intelligent applications with a shorter time to market. With the help of Azure AI Services, we leverage prebuilt AI models offered by Microsoft Azure to solve complex problems which require applications to comprehend visual, audio, and textual context from structured or unstructured data. In this chapter, we gained a brief understanding about all the services falling under the umbrella of Azure AI Service. In the subsequent chapters, we are going to take a deep dive and learn ways to add cognitive capabilities to our applications by leveraging these services.

CHAPTER 2

Build a Language-Based Document Classifier with Azure Functions

Language serves as a crucial means of communication for individuals across the world to express their thoughts and enable information exchange. For applications to be truly intelligent, they must possess the ability to understand and comprehend human language. Much of the data generated over the past few decades have been in textual, audio, or video formats which are unstructured by nature. A lot of insights can be derived from these data. Natural language processing (NLP) is a subfield of AI which focuses on enabling computational systems to interpret data, extract insights, understand context, and generate content. With NLP, applications can perform various tasks ranging from sentiment analysis to machine translations.

Some of the popular industry use cases of NLP are Google Translate, ChatGPT, and Grammarly. Google Translate leverages the power of the NLP to translate texts from one language to another on a real-time basis. Grammarly uses NLP algorithms to analyze written text and provide recommendations to improve the textual content. ChatGPT leverages NLP-based large language models (LLMs) to understand, interpret, and generate human-like response for queries. The use cases are many but to build an NLP system from scratch require intensive investments on hiring domain experts and setting up the computational infrastructure to build, train, and deploy models. It can take years to design and build such systems. To overcome these challenges, we can leverage the powers of the Azure AI Language service. It is a specialized AI service offered by Microsoft Azure to build solutions which can understand and extract insights from textual data.

In this chapter, we are going to briefly discuss about the Azure AI Language service and its use cases and build a language-based document classifier by leveraging its client SDK.