

AI-900: Microsoft Azure AI Fundamentals Cheat Sheet

Quick Bytes for you before the exam!

The information provided in Cheat Sheets is for educational purposes only; created in our efforts to help aspirants prepare for the Exam AI-900 certification. Though references have been taken from Microsoft documentation, it's not intended as a substitute for the official docs. The document can be reused, reproduced, and printed in any form; ensure that appropriate sources are credited and required permissions are received.

Are you Ready for MS Azure AI Fundamentals? “AI-900” Certification?



Self-assess yourself with

[Whizlabs FREE TEST](#)



800+ Hands-on-Labs and Cloud Sandbox

[Hands-on Labs Cloud Sandbox environment](#)



Index		
Domain Name	Topics	Page no
Describe Artificial Intelligence workloads and considerations	AI workloads: Common Features	
	Azure AI Services – Using AI for Intelligent Apps	5
	What is Azure Content Moderator? - AI Services	7
	What Is Computer Vision?	8
	Natural language processing technology	10
	Knowledge Mining	12
	What is Azure AI Document Intelligence?	15
	Responsible AI: guiding principles	
	What is generative AI?	16
	What is Responsible AI - Azure Machine Learning	17
	Machine learning fairness	18
	AI shared responsibility model - Azure	20
	Transparency Note for Azure AI Studio safety evaluations	22
Describe fundamental principles of machine learning on Azure	Identify common machine learning techniques	
	Azure Machine Learning - ML as a Service	23
	Machine Learning Algorithms	25
	Linear Regression: Component Reference	27
	How to select algorithms for Azure Machine Learning	29
	Create an Azure Machine Learning compute cluster	30
	Core machine learning concepts	
	What Is Deep Learning? - Microsoft Azure	31
	Create and explore datasets with labels	33
	Create a labeled dataset using Azure Machine Learning	34
	Data splits and cross-validation in automated machine	35

	Describe Azure Machine Learning capabilities	
	What is automated ML? AutoML - Azure Machine Learning	36
	Azure's automated machine learning capabilities	37
	Machine Learning Solutions for Data Scientists	39
	MLOps: Machine learning model management	41
Describe features of computer vision workloads on Azure	Identify Azure tools and services for computer vision tasks	
	What is Azure AI Vision?	42
	Azure AI Vision with OCR and AI	43
	What is the Azure AI Face service?	45
	Identify common types of computer vision solution	
	Image classification on Azure	47
	Object detection - Azure AI Vision	49
	OCR - Optical Character Recognition - Azure AI services	50
	Face detection, attributes, and input data	51
Describe features of generative AI workloads on Azure	Identify features of generative AI solutions	
	Fundamentals of Generative AI - Training	53
	Generative AI in Azure Machine Learning	56
	10 ways generative AI and Azure OpenAI Service	58
	Responsible AI Solutions	59
	Identify capabilities of Azure OpenAI Service	
	Overview of Responsible AI practices for Azure OpenAI	62
	What is Azure OpenAI Service?	64
	Azure openAI for code generate and development task	66
	Quickstart: Generate images with Azure OpenAI Service	67

Describe features of Natural Language Processing (NLP) workloads on Azure	NLP Workload Scenarios and Common Features	
	What is key phrase extraction in Azure AI Language?	69
	What is Named Entity Recognition (NER) in Azure AI Language?	71
	What is sentiment analysis and opinion mining?	72
	What is Azure AI Language - Azure AI services	73
	What is the Speech service? - Azure AI services	75
	What is Document Translation? - Azure AI services	79
	Azure tools and services for NLP workloads	
	Azure AI Language - Capabilities	82
	Azure AI Speech Capabilities	83

AI workloads: Common Features

Azure AI Services

Azure AI Services Platform

- **Integrated Platform:** Azure AI Services provides a comprehensive suite of tools and APIs for constructing, deploying, and overseeing AI models.
- **Audience:** Designed to cater to developers, data scientists, and business analysts alike, Azure AI delivers solutions tailored to individual requirements.

Azure AI Vision: Empowering Visual Intelligence

- **Cutting-Edge Capabilities:** Under the umbrella of Azure AI Vision, users can leverage advanced computer vision capabilities to dissect images and videos seamlessly.
- **Image Tagging:** Identify and categorize objects, scenes, and concepts within visual content effortlessly.
- **OCR Functionality:** Extract printed and handwritten text from images across diverse languages and writing styles.
- **Facial Recognition:** Develop intelligent applications capable of discerning and authenticating human identity seamlessly.

Enhancing Content Discoverability

- **Automated Image Analysis:** Azure AI Vision's image analysis component offers automated generation of natural language captions for images.
- **Smart Cropping:** Employ smart cropping techniques to optimize visual aesthetics.
- **Content Classification:** Classify images based on content to enrich user experiences and render content more accessible and engaging.

Real-Time Spatial Analysis

- **Movement Monitoring:** Empowers users to monitor movements and scrutinize environments in real-time.
- **Applications:** Invaluable tool for applications spanning surveillance, retail analytics, and smart spaces.

Facilitating Text Extraction with OCR

- **Digitization:** Instrumental in extracting text from images for the digitization of documents, forms, and handwritten notes.

Secure User Experiences with Facial Recognition

- **Authentication:** Ensure secure user experiences by enabling seamless authentication, access control, and personalized services.

Customization Options

- **Tailored Models:** Offer the flexibility to tailor computer vision models to specific requirements through custom image classification and object detection.
- **Adaptability:** Train models with minimal image inputs to adapt classification and detection mechanisms to unique needs.

Responsible AI Practices

- **Guidance:** Microsoft emphasizes responsible AI practices, providing comprehensive guidance on ethical and accurate utilization of AI capabilities.
- **Adherence:** Adhering to responsible AI standards ensures that users achieve optimal outcomes while upholding ethical principles.

What is an Azure Content Moderator?

Azure Content Moderator: Empowering Content Safety

Use Cases:

- Online Platforms: Deployed for content moderation in online marketplaces, ensuring the integrity of product catalogs and user-generated content.
- Gaming Industry: Integral for gaming companies to moderate user-generated game artifacts and maintain safe chat room environments.
- Social Messaging: Essential for social messaging platforms to uphold standards of safety and appropriateness in content shared among users.
- Enterprise Media: Employed by media companies to effectively manage content across diverse channels, maintaining brand reputation and audience engagement.
- Education Sector: Critical for K-12 education solution providers to foster a safe online environment for students, safeguarding against exposure to inappropriate content.

Key Features:

- Text Analysis: Scans text for various forms of harmful content including sexual content, violence, hate speech, and self-harm, with the option to apply severity levels to flagged content.
- Image Moderation: Detects and classifies potentially harmful images, automatically applying content flags for further action.
- Video Content Management: Extends content moderation capabilities to video streams, analyzing video frames for inappropriate content.
- Content Safety Studio: An online tool equipped with the latest content moderation machine learning models, providing templates and customizable workflows for handling offensive or risky content effectively.

Language Support:

- Wide Language Coverage: Supports over 100 languages, ensuring comprehensive moderation capabilities for global audiences.
- Specific Language Expertise: Specifically trained in major languages including English, German, Japanese, Spanish, French, Italian, Portuguese, and Chinese, ensuring accurate detection and moderation across diverse linguistic contexts.

Transition to Azure AI Content Safety:

- Advanced Features: Azure AI Content Safety represents the evolution of Azure Content Moderator, offering advanced AI capabilities and enhanced performance for a wide range of scenarios.
- Versatility: Suitable for various applications, including those mentioned above, Azure AI Content Safety provides comprehensive solutions for managing and ensuring the safety of digital content.

What is computer vision?

- Computer vision, a captivating domain within AI, enables computers to recognize and comprehend objects and individuals in visual media.
- It seeks to emulate human capabilities by automating tasks like object recognition, pattern detection, and visual interpretation.
- Operating at the intersection of various fields, it aims to automate processes akin to human perception.
- Tasks include recognizing objects, identifying patterns, and interpreting visual data.

How It Works:

Data-Driven Learning:

- Computer vision relies on extensive datasets to learn and generalize effectively.
- It iteratively analyzes data to identify differences and patterns.

Deep Learning and Convolutional Neural Networks (CNN):

- Deep learning algorithms empower computers to self-educate on visual data context.
- CNNs decompose images into pixels, assigning labels and predicting visual content.
- The neural network refines its predictions over time to match real-world observations, mimicking human visual perception.

Applications:

Object Identification:

- Entails the recognition of distinct objects depicted in images or video content.

Image Categorization:

- Involves sorting images into predefined categories, like discerning between various animal species.

Facial Detection and Identification:

- Involves pinpointing and recognizing human faces captured in photos or video feeds.

Medical Imaging Analysis:

- Comprises the examination of medical scans, such as X-rays and MRIs, to aid in medical diagnoses.

Self-Driving Vehicle Technology:

- Empowers autonomous vehicles to interpret and navigate through their surroundings.

Manufacturing Quality Assurance:

- Involves the identification of defects and irregularities in manufacturing processes.

Public Space Monitoring:

- Encompasses the surveillance of public areas to ensure safety and security.

Challenges:

Image Variability:

- Images exhibit considerable diversity stemming from factors like lighting disparities, varying angles, and occlusions.

Scale and Perspective Recognition:

- Involves the ability to identify objects across a spectrum of sizes and viewing angles.

Noise and Uncertainty Management:

- Entails effectively handling noisy or incomplete data sets that may contain ambiguities.

Real-Time Analysis:

- Requires the efficient processing of video streams without delay or lag.

Natural language processing technology

Natural Language Processing (NLP)

- NLP represents a distinct domain within the realm of AI, aiming to streamline communication between machines and humans through the use of natural language.
- Utilizing computational techniques, NLP delves into the analysis and understanding of both textual and verbal communication to derive significant interpretations.

Applications of Natural Language Processing (NLP)

- **Voice-Activated Assistants:** NLP enables virtual helpers such as Siri and Alexa to comprehend and execute verbal instructions effectively.
- **Language Translation Tools:** NLP technology forms the backbone of translation services, allowing seamless communication across different languages and promoting global connectivity.
- **Emotional Analysis:** NLP plays a pivotal role in dissecting the emotional context embedded in textual content, particularly in the realm of social media analysis.
- **Interactive Chatbots:** NLP-driven conversational agents engage users using natural language, offering assistance and resolving queries across various customer service platforms.
- **Summarization Algorithms:** NLP algorithms condense lengthy textual content into succinct summaries, facilitating rapid information retrieval and comprehension.
- **Entity Recognition:** NLP methodologies identify and classify entities within the text, enriching the understanding and structuring of information.

Challenges:

- **Interpretation Complexity:** Human language poses challenges due to its inherent ambiguities, including homonyms, idiomatic expressions, and sarcasm, making accurate comprehension difficult.
- **Expression Diversity:** Individuals communicate differently, leading to variations in accents, grammar, and sentence structures, which can complicate NLP tasks.
- **Contextual Understanding:** Accurate NLP relies on understanding the context of language, as words may carry different meanings depending on their context.

NLP Tasks:

- **Speech Recognition:** Converting spoken language into text, a fundamental function for voice-controlled systems.
- **Part-of-Speech Tagging:** Identifying the grammatical function of each word (e.g., verb, noun, adjective) within a sentence.
- **Word Sense Disambiguation:** Resolving the correct meaning of a word with multiple interpretations based on its context.
- **Language Modeling:** Predicting the subsequent word in a sentence based on contextual cues to enhance coherence and fluency.

Knowledge Mining

Overview:

Knowledge mining is an emerging field within artificial intelligence (AI) focused on leveraging intelligent services to rapidly learn from vast data volumes. Its primary aim is to empower organizations with profound insights, enabling easy exploration of information, uncovering hidden patterns, and identifying relationships on a large scale.

Components and Capabilities:

Semantic Search:

Azure provides semantic search functionalities, discerning user intent and contextually ranking search results for relevance. This allows users to navigate content more efficiently and access contextually pertinent information.

Content Ingestion:

Knowledge mining enables content ingestion from diverse sources such as Azure Blob storage, Azure Table storage, Azure SQL Database, and Azure Cosmos DB. Dedicated connectors facilitate data extraction from hundreds of third-party sources, streamlining the process.

Content Enrichment:

Utilizing AI capabilities, knowledge mining extracts insights, discovers patterns, and enhances understanding of content. Techniques include text extraction from file formats like PDF, Microsoft Word, PowerPoint, and CSV, augmenting data comprehensiveness.

Exploration and Discovery:

Once content is indexed, users can explore it through various interfaces including search interfaces, bots, existing business applications, and data visualizations. This facilitates seamless exploration and discovery of relevant information.

Use Cases:**Content Research:**

Knowledge mining aids employees in swiftly reviewing dense technical materials, enhancing research efficiency.

Auditing, Risk, and Compliance Management:

Attorneys utilize knowledge mining to identify critical entities from discovery documents and highlight important ideas across various documents, streamlining compliance processes.

Business Process Management:

Companies leverage knowledge mining to prevent costly errors in competitive bidding or real-time problem diagnosis, optimizing business processes.

Customer Support and Feedback Analysis:

Efficient analysis of customer support interactions and feedback is facilitated through knowledge mining, enabling prompt and effective responses to customer queries and concerns.

Azure Services for Knowledge Mining:

Azure AI Search:

The sole cloud search service embedded with AI capabilities, facilitating the identification and exploration of relevant content.

Azure AI Services:

Utilizes cognitive capabilities to comprehend diverse content types, aiding in information extraction and analysis.

Azure Machine Learning:

Empowers the application of custom machine learning models as skills tailored to specific requirements, enhancing data processing and insights generation.

Azure AI Bot Service:

Enables the design of interactive experiences that allow users to extract information seamlessly through bot interfaces, enhancing user engagement and information retrieval efficiency.

What is Azure AI Document Intelligence?

Azure AI Document Intelligence Overview:

- **Automated Extraction:** Leveraging AI and OCR technology, Azure AI Document Intelligence automates text and structure extraction from documents, ensuring swift data processing.
- **Streamlined Information Extraction:** The primary objective is to streamline information extraction, effectively converting documents into actionable data for improved decision-making.

Key Features:

- **Automated Extraction:** The system automatically extracts text, key-value pairs, and structured data from documents, facilitating efficient data processing.
- **Focus Optimization:** By automating data extraction, users can redirect their focus towards utilizing extracted information instead of spending time on manual compilation.

- **Prebuilt and Custom Models:** Azure AI Document Intelligence offers both prebuilt and custom models tailored to various document types, enhancing flexibility and accuracy.

Deployment Flexibility:

- **Cloud or Edge Deployment:** The system can be deployed in the cloud or at the edge, providing flexibility in data ingestion from diverse sources.
- **Enhanced Security:** Azure AI Document Intelligence ensures enterprise-grade security measures for both data and trained models, adhering to privacy regulations and safeguarding data integrity and confidentiality.

Responsible AI: guiding principles

What is generative AI?

- Artificial Intelligence (AI) replicates human behavior via machine learning algorithms, enabling autonomous task execution without explicit instructions.
- Generative AI, a subset of AI, specializes in generating original content based on encountered context, exemplified by Microsoft Copilot, enhancing web browsing experiences.
- Generative AI applications accept natural language inputs and produce relevant responses across various formats (text, images, code), employing sophisticated algorithms for contextual understanding.

Features:

- Natural Language Processing (NLP): Enables accurate comprehension and interpretation of human language inputs.
- Content Generation: Capable of producing diverse content types like text, images, and code snippets.
- Contextual Understanding: Understands conversational context to generate coherent and relevant responses.
- Adaptability: Continuously learns from interactions and feedback, improving over time.

Use Cases:

- Conversational Assistants: Powers virtual assistants and chatbots for natural language interactions, facilitating various tasks.
- Creative Content Generation: Used in creative industries to generate artwork, music, and literature.
- Code Generation: Assists developers with code snippets, debugging, and optimization suggestions.
- Personalized Recommendations: Analyzes user preferences to generate tailored recommendations across platforms.
- Content Creation and Automation: Simplifies content creation by generating text, images, and multimedia content for various platforms and campaigns.

What is Responsible AI - Azure Machine Learning

Responsible AI, also known as Ethical AI, places a strong emphasis on ensuring that AI development and deployment aligns with human values and respects individual rights. As AI becomes more prevalent, it becomes increasingly important to address unintended biases and potential harm that may arise from its use.

Features

- **Equity and Inclusion:** Azure Machine Learning prioritizes fair treatment across diverse demographics, facilitated by its fairness assessment tool.
- **Reliability and Safety:** Azure Machine Learning ensures consistent, secure, and error-free AI system operation, supported by an error analysis module.
- **Privacy and Security:** Azure Machine Learning implements robust security measures to protect user data and privacy.
- **Transparency:** Azure Machine Learning offers tools to elucidate model predictions and behavior, promoting transparent AI systems.
- **Accessibility:** Azure Machine Learning promotes the development of inclusive models, ensuring accessibility for all individuals.
- **Accountability:** Azure Machine Learning fosters accountability and trustworthiness by adhering to Responsible AI principles, and holding stakeholders accountable for AI system outcomes.

Machine learning fairness - Azure

What is Machine Learning Fairness?

- Definition: Machine learning fairness pertains to ensuring AI models avoid discriminatory behavior or bias against particular demographic groups.
- Objective: The goal is to develop models that treat all individuals equitably, irrespective of their demographic attributes like gender, race, or age.

Types of AI-Caused Harms

1. Harm of Allocation

- Definition: AI systems allocate or withhold opportunities, resources, or information for specific groups.
- Examples: Examples include biased hiring decisions, unfair school admissions, and discriminatory lending practices.

2. Harm of Quality-of-Service

- Definition: AI systems exhibit performance disparities across different groups.
- Examples: For instance, a voice recognition system may perform better for men than for women.

Group Fairness Approach

Identifying At-Risk Groups

- **Focus:** Group fairness examines which demographic groups are susceptible to harm.
- **Role of Sensitive Features:** Sensitive features such as gender or ethnicity are crucial in assessing fairness.

Quantifying Fairness

- **Disparity Metrics:** These metrics evaluate and compare model behavior across subgroups.
- **Examples:** Metrics include accuracy rate, error rate, precision, recall, and mean absolute error.

Azure Machine Learning and Fairness

Responsible AI Dashboard

- **Purpose:** The dashboard evaluates model fairness during both training and deployment phases.
- **Disparity Assessment:** It quantifies disparities across subgroups based on sensitive features.

Fairlearn Open-Source Package

- **Functionality:** Fairlearn aids in assessing impact and devising mitigation strategies.
- **Flexibility:** It allows for trade-offs tailored to specific scenarios.

Trade-Offs and Human Decision-Making

- **Complexity:** Fairness is a socio-technical challenge, with quantitative metrics unable to capture all fairness aspects.
- **Human Judgment:** Humans must navigate trade-offs based on context and values in addressing fairness concerns.

AI shared responsibility model - Azure

Division of Responsibility

- **Shared Responsibility Model:** When implementing AI capabilities, it's crucial to understand the shared responsibility model, which outlines the tasks handled by the AI platform or application provider versus those managed by the user.
- **Variation by Service Type:** Responsibilities differ based on whether the AI integration uses Software as a Service (SaaS), Platform as a Service (PaaS), or Infrastructure as a Service (IaaS).

AI Layer Overview

AI Platform

- **Functionality:** Supplies AI capabilities to applications.

- **Responsibilities:** Includes constructing and securing the infrastructure for the AI model, managing training data, and configuring model parameters (e.g., weights and biases).
- **Access:** Typically accessed via APIs.

AI Application

- **Functionality:** Utilizes AI capabilities to deliver services or interfaces to users.
- **Complexity:** Components can range from simple to highly sophisticated, depending on the application.

AI Security Considerations

- **Platform Protection:** Safeguard the AI platform from malicious inputs.
- **Safety Systems:** Implement measures to filter harmful instructions sent to the AI model (inputs) and prevent the generation and return of harmful content (outputs).

Microsoft's Commitment

- **Security Focus:** Microsoft prioritizes security and provides built-in safety systems for both PaaS and SaaS AI offerings:
 - PaaS (Azure OpenAI Service): Protects against harmful inputs and outputs.
 - SaaS (Microsoft Security Copilot): Ensures safety during AI interactions.

Transparency Note for Azure AI Studio safety evaluations

What is a Transparency Note?

- Purpose: Microsoft's Transparency Notes aim to help users understand AI technology, including how it operates, decisions that system owners can make to influence performance and behavior, and the importance of considering the entire system (technology, people, and environment).
- Content: These notes offer insights into AI models, including their capabilities, limitations, and potential risks.

Azure AI Studio Safety Evaluations

Purpose:

- **Objective:** These evaluations enable users to assess the outputs of their generative AI applications for potential textual content risks.

Content Risks Evaluated:

- **Hateful and Unfair Content:** Evaluates language that promotes hate or unfair representations of individuals or social groups.
- **Sexual Content:** Assesses language related to anatomical organs, relationships, erotic acts, and similar topics.
- **Violent Content:** Reviews language associated with physical harm or violence.
- **Self-Harm-Related Content:** Monitors language pertaining to self-harm or suicide.
- **Jailbreak Vulnerability:** Assesses the AI's susceptibility to security breaches.

How It Works:

- **Evaluation Method:** Users evaluate their generative AI application's outputs using specific prompts.
- **AI-Assisted Simulator:** The simulator stimulates adversarial interactions to generate a test dataset with the application's responses.
- **Groundedness Evaluation:** Provides insights into ungrounded parts of generated responses, helping to identify and correct issues.

Responsible Innovation:

- **Microsoft's Commitment:** Azure AI Studio safety evaluations demonstrate Microsoft's dedication to responsible AI.
- **Risk Mitigation:** By assessing and mitigating risks, developers can create AI systems that adhere to ethical and safety standards.

Identify common machine-learning techniques

What is Azure Machine Learning?

- Azure Machine Learning is a cloud-based service designed to support the large-scale development, deployment, and management of machine learning models.
- It provides a wide range of tools and services tailored for data scientists, developers, and business analysts.
- Facilitates the creation and management of complete, end-to-end machine learning workflows.

Key Features and Capabilities:

Model Training and Experimentation:

- Construct and train machine learning models using a diverse selection of algorithms and frameworks.

- Investigate different hyperparameters to assess and enhance model effectiveness.

Automated ML:

- Streamline the process of model selection and hyperparameter tuning with automated machine learning functionalities.
- Azure ML independently evaluates numerous algorithms and configurations to pinpoint the most effective model.

Model Deployment and Management:

- Deploy machine learning models as either web services or within containers.
- Monitor model performance, implement necessary modifications, and efficiently manage versioning.

Integration with Azure Services:

- Seamlessly merge with various Azure services such as Azure Databricks, Azure Data Factory, and Azure Synapse Analytics.
- Utilize data pipelines and coordination tools to enable thorough machine learning workflows.

Use Cases and Scenarios:

Predictive Analytics:

- Employ Azure ML to develop predictive models for tasks such as projecting sales, predicting customer churn, and forecasting demand.

Computer Vision and Natural Language Processing:

- Integrate pre-existing AI functionalities into your systems, including capabilities like recognizing images, performing optical character recognition (OCR), and analyzing sentiment.

Anomaly Detection:

- Identify irregularities in diverse areas such as manufacturing processes, network traffic patterns, or financial transactions.

Recommendation Systems:

- Generate individualized recommendations for users by analyzing their behavior and preferences.

Security and Compliance:

Data and Model Security:

- Azure ML implements strong security protocols to safeguard both data and trained models.
- Guarantee adherence to privacy regulations and industry standards to maintain compliance.

Machine Learning Algorithms

Machine learning algorithms are computational models that enable computers to identify patterns and forecast outcomes from data without explicit programming. These algorithms serve as the foundation of contemporary artificial intelligence and are applied across various domains such as image recognition, natural language processing, recommendation systems, fraud detection, and autonomous vehicles.

Supervised Learning Algorithms:

In supervised learning, models undergo training using labelled datasets to understand the relationship between input data and output labels. This enables the model to predict or classify new, unseen data.

Supervised Learning Algorithms:

- **Linear Regression:** A straightforward method commonly used for tasks involving regression.

- **Logistic Regression:** Employed for tasks requiring classification with associated probabilities.
- **Support Vector Machines (SVM):** Highly adaptable and robust for both classification and regression tasks.
- **Decision Trees:** Models structured in a hierarchical manner for interpretability.
- **Random Forest:** An ensemble technique consisting of multiple decision trees.
- **Naive Bayes Classifier:** A probabilistic algorithm known for its computational efficiency.
- **K-Nearest Neighbors (KNN):** An unsupervised and iterative approach utilized for both classification and regression problems.

Unsupervised Learning Algorithms:

In unsupervised learning, algorithms operate on data without predefined labels. The objective is to unveil inherent patterns or cluster similar data points.

Examples

- Clustering (K-Means): Organizes data into groups based on similarity.
- Dimensionality Reduction (PCA): Condenses data while retaining crucial information.
- Association Rule Mining (Apriori): Uncovers interesting connections among variables.

Ensemble Learning Algorithms:

Ensemble methods amalgamate multiple models to enhance overall performance.

Examples:

- Random Forest: Merges several decision trees.
- Boosting (AdaBoost, XGBoost): Sequentially improves weak models and aggregates their predictions.

Generative Models:

These models discern the underlying probability distribution of the data.

Examples:

- Naive Bayes: Applied in text classification and spam detection.
- Gaussian Processes: Utilized for regression and estimating uncertainty.

Time Series Forecasting Algorithms:

These algorithms prognosticate future values based on historical data.

Examples:

- ARIMA (AutoRegressive Integrated Moving Average): Handles time series data effectively.
- Exponential Smoothing: Captures trends and seasonal variations.

Deep Learning Algorithms:

These neural network architectures acquire intricate representations from data.

Examples:

- Convolutional Neural Networks (CNN): Specialized in image recognition tasks.
- Recurrent Neural Networks (RNN): Tailored for sequential data processing, like natural language tasks.

Linear Regression: Component Reference

What is Linear Regression?

Linear regression is a statistical method used to establish the connection between independent variables, often referred to as features, and a numeric outcome, known as

the dependent variable. It's widely used in both statistics and machine learning for basic predictive tasks. Despite its simplicity, linear regression is effective even on datasets with few data points and straightforward relationships.

Fundamental Concepts:

- **Simple Regression:**
 - In simple regression, there's one independent variable that influences a dependent variable.
 - Azure Machine Learning's Linear Regression component accommodates this scenario.
- **Multiple Linear Regression:**
 - Multiple linear regression entails two or more independent variables affecting a single dependent variable.
 - Situations where numerous inputs forecast a numerical outcome are termed as multivariate linear regression.
- **Error and Fit Measurement Methods:**
 - The Linear Regression component supports two methods:
 - **Ordinary Least Squares (OLS):** Aim to minimize the sum of squared differences between observed and predicted targets.
 - **Gradient Descent:** An iterative strategy for reducing error during model training.

How to use Linear Regression Component?

Model Initialization:

- Access the Initialize Model section within Azure Machine Learning Designer.
- Navigate to Regression by expanding the options.

- Drag the Linear Regression Model component into your pipeline.

Configuring Properties:

- Within the Properties pane, opt for the Ordinary Least Squares method as the solution approach.

How to select algorithms for Azure Machine Learning

When choosing algorithms for Azure Machine Learning, it's crucial to take into account your specific data science situation and project needs. Here are some important aspects to consider:

- **Business Inquiry and Objective:**
 - Clearly articulate the business inquiry targeted with your dataset.
 - Determine if your task falls into classification, regression, or clustering categories.
- **Characteristics of Data:**
 - Gain insight into your data's characteristics:
 - Ascertain whether it's structured or unstructured.
 - Evaluate the number of features (variables) available.
 - Address any missing data appropriately.
 - Consider the scale of your dataset, whether it's small, medium, or large.
- **Model Complexity and Interpretation:**
 - Acknowledge the varying complexities of algorithms, some of which may require more computational resources.
 - Strike a balance between model intricacy and interpretability.
 - Understand that simpler models like linear regression offer transparency, while deep learning models may lack interpretability.
- **Accuracy and Efficiency:**

- Assess the equilibrium between model accuracy and training time.
- Recognize that some algorithms prioritize speed over accuracy, while others prioritize accuracy over speed.
- **Assumptions and Constraints:**
 - Different algorithms rely on distinct assumptions about data distribution.
 - Account for constraints related to interpretability, scalability, and real-time prediction needs.
- **Experimentation and Comparative Analysis:**
 - Employ Azure Machine Learning's automated tools, such as AutoML, to explore various algorithms.
 - Compare their performance metrics such as accuracy, precision, recall, and F1-score to determine the most fitting algorithm for your use case.

Create an Azure Machine Learning Compute Cluster

A Compute Cluster in Azure Machine Learning is a centrally managed setup that simplifies the creation of either a single or multiple node computing environment. It facilitates the distribution of tasks like training models or performing batch inference across a collection of CPU or GPU-based compute nodes hosted in the cloud. Here's a guide on how to set up and oversee one:

Prerequisites:

- Access to an Azure Machine Learning workspace.
- Create one via Azure portal or Azure CLI if not available.
- Installation of Azure CLI extension for Machine Learning service (v2).
- Execute command: `az extension add -n azure-cli-ml`.

Steps to Create a Compute Cluster:

1. Initialize Your Workspace in Python:

- Connect to your Azure ML workspace using the Azure Machine Learning Python SDK in your Python script.
- Replace <SUBSCRIPTION_ID>, <RESOURCE_GROUP>, and <AML_WORKSPACE_NAME> with your actual values.
- Example code snippet:

```
from azure.ai.ml import MLClient
from azure.identity import DefaultAzureCredential
subscription_id = "<SUBSCRIPTION_ID>"
resource_group = "<RESOURCE_GROUP>"
workspace = "<AML_WORKSPACE_NAME>"
ml_client = MLClient(DefaultAzureCredential(), subscription_id, resource_group,
workspace)
```

2. Create the Compute Cluster:

- Utilize the ml_client handler to create a compute cluster.
- Specify the desired configuration such as VM size and number of nodes.
- **Example code snippet:**

```
from azure.ai.ml import ComputeConfiguration
compute_config = ComputeConfiguration(target_nodes=2,
vm_size="Standard_DS3_v2")
compute_cluster = ml_client.compute.create("my-cluster", compute_config)
```

3. Use the Compute Cluster:

- Once created, utilize the compute cluster as a target for training or inference jobs in your Azure Machine Learning workspace.

Core machine learning concepts

What Is Deep Learning? - Microsoft Azure

Deep learning, as employed in Microsoft Azure, refers to a form of machine learning characterized by the utilization of artificial neural networks. These networks empower digital systems to glean insights and make determinations using unstructured and unlabeled data. Unlike traditional machine learning, which relies on structured data,

deep learning enables AI systems to learn from experiences, identify patterns, offer recommendations, and evolve autonomously.

Key Points on Deep Learning:

How Deep Learning Functions:

- Deep learning operates through intricate neural network architectures comprising multiple layers.
- The cloud or cluster-based deployment of high-performance GPUs is pivotal for achieving remarkable accuracy in tasks such as text, speech, and image recognition.
- The effective training of deep learning models necessitates substantial volumes of labeled data.

Significance of Deep Learning:

- Data scientists and developers harness deep learning for:
 - Analyzing vast and intricate datasets.
 - Executing intricate and nonlinear operations.
 - Delivering prompt and precise responses to text, voice, or image inputs surpassing human capabilities.
- Deep learning finds practical utility across diverse sectors such as autonomous vehicles, smart home technology, and healthcare.

Frameworks and Enhancements:

- Developers employ deep learning frameworks like TensorFlow or PyTorch to construct sophisticated machine learning models.
- Accelerators like ONNX Runtime enhance the performance of these frameworks, expediting both training and inference processes.

Create and explore datasets with labels

What are Datasets with Labels?

In Azure Machine Learning, datasets that contain labels are known as labeled datasets. These datasets are a type of TabularDataset with a specific label column and are exclusively generated as a result of Azure Machine Learning data labeling projects.

Prerequisites:

1. Azure Subscription:

- If you don't have an Azure subscription, create a free account before you begin.

2. Azure Machine Learning SDK for Python:

- Install the Azure Machine Learning SDK for Python or access Azure Machine Learning Studio.

3. Machine Learning Workspace:

- Create a Machine Learning workspace. If you haven't already, see how to create workspace resources.

4. Access to an Azure Machine Learning Data Labeling Project:

- If you don't have a labeling project, first create one for image labeling or text labeling.

5. Export Data Labels:

- After finishing a data labeling project, export the label data.
- This captures both data references and labels.
- Export formats:
 - i. COCO Format: Saved in the default blob store of the Azure ML workspace under export/coco.
 - ii. Azure ML Dataset: Available in the Datasets section of Azure ML studio, with sample Python code provided on the Details page.

6. Explore Labeled Datasets with Pandas DataFrame:

- Load labeled datasets into a pandas DataFrame.

- Use the `to_pandas_dataframe()` method from the `azureml-dataprep` class.

Create a Labeled dataset using Azure Machine Learning

What is a Labeled Dataset?

A labeled dataset consists of data samples (such as images, text, or tabular data) along with corresponding labels or annotations. These labels provide ground truth information about the data, allowing machine learning models to learn patterns and make predictions.

Prerequisites:

1. Azure Subscription:
 - If you don't have an Azure subscription, create a free account before you begin.
2. Azure Machine Learning Workspace:
 - Create an Azure Machine Learning workspace. If you haven't already, see how to create workspace resources.

Steps to Create a Labeled Dataset:

1. **Data Labeling Project:**
 - Create a data labeling project in Azure Machine Learning.
 - Choose the type of labeling task (e.g., image classification, object detection, text classification).
2. **Upload Data:**
 - Upload your data samples to the labeling project.
 - Assign labels or annotations to each sample based on the task requirements.
3. **Review and Export Labels:**
 - Review the labeled data to ensure accuracy.
 - Export the labeled data as an Azure Machine Learning dataset.
4. **Explore the Labeled Dataset:**
 - Load the labeled dataset into your machine-learning pipeline.
 - Use it for training, validation, or testing your machine-learning models.
 -

Data splits and cross-validation in automated machine

Data Splits:

1. Training Data:

- The training data is used to train the machine learning model.
- It contains labeled examples (input features and corresponding target labels).

2. Validation Data:

- The validation data is used to tune hyperparameters and evaluate model performance during training.
- It helps prevent overfitting by assessing how well the model generalizes to unseen data.

3. Test Data:

- The test data is used to evaluate the final model's performance after training.
- It provides an unbiased estimate of the model's accuracy on new, unseen data.

Cross-Validation:

K-Fold Cross-Validation:

- Partition the dataset into K equally sized groups.
- Train the model K times, using K-1 groups for training each time and the remaining group for validation.
- Calculate the average performance over all K groups.

Stratified K-Fold Cross-Validation:

- Ensure each fold reflects the overall distribution of target labels.
- Especially useful for datasets with imbalanced classes.

Leave-One-Out Cross-Validation (LOOCV):

- Use each individual data point once as the validation set, with the rest used for training.

- Best suited for smaller datasets.

Shuffle Split Cross-Validation:

- Randomly split the data into training and validation sets multiple times.
- Ideal for large datasets.

Describe Azure Machine Learning capabilities

What is automated ML?

AutoML, known as Automated Machine Learning, simplifies the process of selecting the optimal machine learning algorithms and parameters for a given dataset.

- It enhances productivity and efficiency in creating models while maintaining quality standards.

AutoML in Azure Machine Learning:

- Azure Machine Learning's AutoML creates multiple pipelines simultaneously during training, experimenting with diverse algorithms and parameters.
- Through iterative testing, it explores various algorithm-feature combinations to optimise for specific metrics until predefined criteria are met.

Key Steps in the AutoML Process:

1. **Identifying the Problem:** Recognize the machine learning problem type, such as classification, regression, or NLP.
2. **Experience Selection:** Choose between a code-first approach with Azure Machine Learning SDKv2 or a no-code studio web interface.
3. **Specifying Data Source:** Define the labeled training data source.
4. **Configuring AutoML:** Set parameters including model iterations, hyperparameters, preprocessing, and evaluation metrics.
5. **Submitting Training Job:** Submit the training job to Azure Machine Learning.

6. **Reviewing Results:** Analyze the results, inspect job details, and understand feature importance in the generated models.

When to Use AutoML:

- AutoML is suitable for various tasks including classification, regression, forecasting, computer vision, and NLP.
- It's particularly valuable for rapid prototyping, exploring different algorithms and feature engineering techniques, and minimizing manual effort in model selection and tuning.

Benefits of AutoML:

- **Efficiency:** Accelerates the model development process.
- **Accessibility:** Makes machine learning accessible to non-experts.
- **Scalability:** Handles large datasets and complex problems effectively.
- **Quality:** Ensures high-quality models through comprehensive algorithm and parameter testing.

Azure's automated machine learning capabilities

Azure's Automated Machine Learning (AutoML) simplifies model development by intelligently selecting and training models without manual intervention. It enhances scalability, efficiency, and productivity while maintaining quality standards.

Key Features:

- **Automated Model Selection:** Recommends the best models tailored to the dataset and scenario.

- **Simplified Workflow:** Focuses on business objectives by automating part of the workflow.
- **Increased Accessibility:** Makes AI accessible to a broader audience.
- **Enhanced Model Quality:** Improves accuracy and performance with advanced features.

Enhanced Feature

- **No-Code Web Interface:** User-friendly interface for training models without code.
- **Time Series Forecasting:** Offers forecasting capabilities with advanced features.
- **Model Transparency:** Provides insights into the machine learning pipeline for better interpretability.
- **ONNX Models:** Supports seamless integration with C# applications.
- **.NET Integration:** Enables .NET developers to leverage AutoML within Visual Studio.
- **PowerBI Integration:** Allows building, deploying, and using ML models within PowerBI.
- **SQL Server and Spark Integration:** Integrates with familiar environments like SQL Server and Spark.

Benefits:

- **Intelligent Stopping:** Halts experiments based on progress, optimizing resources.
- **Sub-Sampling Strategy:** Explores more model pipelines within a shorter time.
- **Preprocessing Optimization:** Searches for the best approach for data preprocessing.
- **XGBoost Support:** Expands learners to include high-performing XGBoost models.
- **Enhanced Dataset Support:** Supports datasets up to 10GB, catering to diverse use cases.

Machine Learning Solutions for Data Scientists

Machine Learning on Azure:

Leverage tools for data scientists and engineers to build and deploy scalable machine learning solutions.

Features:

- **Develop Models on Your Terms:** Create models with your favorite languages, environments, and frameworks. Deploy to the cloud, on-premises, or at the edge with Azure AI.
- **Build Responsibly:** Understand your models, secure data with differential privacy and confidential computing, and manage the lifecycle with audit trails and datasheets.
- **Confident Deployment:** Deploy and manage scalable, fault-tolerant, and reproducible machine learning solutions with confidence.

Utilising Machine Learning for Mission-Critical Applications:

Scaling Training Jobs:

- Utilize Azure's vast compute resources, such as VM Scale Sets and Azure Machine Learning clusters, to handle extensive training datasets.
- Implement distributed training techniques to enhance efficiency and reduce the time required for model training.

Model Deployment and Inferencing:

- Explore deployment options such as Azure Kubernetes Service (AKS) and Azure Functions to handle varying inference loads.
- Optimize models for deployment using techniques like model quantization and pruning to improve performance and reduce latency.

Understanding MLOps:

- Learn how MLOps facilitates continuous integration and delivery (CI/CD) of machine learning models, ensuring consistent and reliable updates.
- Understand the importance of monitoring and managing model performance post-deployment to maintain accuracy and relevance.

Securing Machine Learning Environments:

- Implement Azure's identity and access management (IAM) to control and monitor access to your machine learning resources.
- Use Azure Security Center to detect and respond to threats, ensuring the protection of sensitive data and machine learning workflows.

Hybrid and Multi-Cloud Machine Learning:

- Set up Azure Arc to extend Azure services and management to any infrastructure, enabling hybrid and multi-cloud machine learning deployments.
- Leverage Azure Stack to run machine learning workloads on-premises while maintaining cloud compatibility and integration.

Open and Interoperable Machine Learning:

- Integrate Azure Machine Learning with open-source frameworks like TensorFlow, PyTorch, and Scikit-Learn for flexible model development.
- Utilize Azure Data Factory and Azure Synapse Analytics for seamless data integration and orchestration across various Azure services and open-source tools.

MLOps: Machine learning model management

Understanding MLOps:

- MLOps comprises best practices for deploying and managing machine learning models in production efficiently, drawing inspiration from DevOps principles.
- It ensures collaboration and communication between software developers and IT professionals.

The Need for MLOps:

- With machine learning models increasingly integrated into business operations, a systematic approach to their management becomes essential.
- MLOps automates the machine learning lifecycle, fosters team collaboration, and ensures reliable model deployment.

Key Components of MLOps:

1. **Version Control:** Manages changes to models, data, and code for reproducibility and rollback capabilities.
2. **Continuous Integration and Delivery (CI/CD):** Automates model testing and deployment similar to software pipelines.
3. **Monitoring and Operations:** Continuously tracks model performance to address issues promptly.
4. **Governance and Compliance:** Ensures models comply with regulatory requirements and ethical standards.

MLOps Lifecycle Stages:

- **Development:** Data preparation, feature engineering, and model training.
- **Testing:** Model validation and evaluation against business requirements.

- **Deployment:** Serving models for batch or real-time predictions.
- **Monitoring:** Tracking model performance and data drift.
- **Maintenance:** Updating models with new data, retraining, and redeployment.

Benefits of MLOps:

- **Efficiency:** Reduces time to deploy models to production.
- **Scalability:** Manages multiple models across various environments.
- **Reproducibility:** Enables model recreation and auditing.
- **Quality:** Maintains high standards for model accuracy and performance.

MLOps Best Practices:

- **Collaboration:** Promote open communication among data scientists, engineers, and operations teams.
- **Automation:** Automate lifecycle processes to minimize manual errors and accelerate workflows.
- **Monitoring:** Implement robust monitoring systems to detect issues early.

Identify Azure tools and services for computer vision tasks

What is Azure AI Vision?

Azure AI Vision integrates functionalities that were formerly part of Cognitive Services and Azure Applied AI Services. It offers a range of tools and services designed to help applications interpret and analyze visual data, effectively allowing them to "see" and comprehend visual information.

Services Offered by Azure AI Vision:

- **Text Extraction (OCR):** Captures and extracts text from images, including both printed and handwritten content from various surfaces and backgrounds.

- **Image Feature Analysis:** Identifies and detects visual elements in images, such as objects, faces, and explicit content, and generates automatic text descriptions.
- **Face Recognition and Analysis:** Detects, identifies, and analyzes human faces in images, suitable for scenarios like identification and contactless access control.
- **Video Content Analysis:** Features include Spatial Analysis, which monitors the presence and movement of people in video streams, and Video Retrieval, allowing for the indexing and searching of videos using natural language queries.

Azure AI Vision for Digital Asset Management (DAM):

Azure AI Vision optimizes Digital Asset Management (DAM) by effectively organizing, storing, and retrieving rich media assets while handling digital rights and permissions. It can classify and identify images based on visible features such as logos, faces, objects, and colors. Additionally, it can automatically generate captions for images, enhancing their searchability.

Getting Started with Azure AI Vision:

Developers can easily explore the features of Azure AI Vision using Vision Studio in a web browser. This interactive approach is perfect for understanding and experimenting with the service's capabilities.

Azure AI Vision with OCR and AI

Azure AI Vision is a comprehensive service offering advanced computer vision capabilities, enabling applications to analyze images, read text, and detect faces using pre-built models. The service integrates OCR (Optical Character Recognition) and AI,

providing developers with a broad range of vision features, even if they lack machine learning expertise.

Key Features of Azure AI Vision

- **Image Analysis:** Utilizes over 10,000 concepts and objects to detect, classify, caption, and generate insights from images.
- **Spatial Analysis:** Monitors the presence and movements of people within physical spaces in real-time.
- **OCR:** Extracts text, both printed and handwritten, from images, supporting a variety of languages and writing styles.
- **Facial Recognition:** Facilitates the creation of applications that can recognize and verify human identities.

Enhancing Content Discoverability with Image Analysis

Azure AI Vision's image analysis functionality automatically captions images using natural language, performs smart cropping, and classifies images, thereby enhancing the discoverability and accessibility of content.

Real-Time Video Streaming with Spatial Analysis

The service supports real-time tracking of movement and environmental analysis using computer vision, image analysis, and object detection.

Reading Text from Images with OCR

The OCR technology in Azure AI Vision excels at extracting both printed and handwritten text from images, accommodating various languages and writing styles, which is essential for applications that need to interpret textual information within visual data.

Verifying Identities with Facial Recognition

Facial recognition within Azure AI Vision provides a highly secure and seamless user experience, enabling efficient identity verification in applications.

Customizing Computer Vision Models

Developers can train custom computer vision models tailored to their specific needs with only a few images, maintaining high accuracy. This feature is currently available in preview.

Applying AI Responsibly

Microsoft offers comprehensive guidelines for the responsible use of AI Vision to ensure goals are met and results are accurate, adhering to responsible AI principles.

What is the Azure AI Face service?

Azure AI Face Service Overview:

- Provides AI algorithms for detecting, recognizing, and analyzing human faces in images.
- Useful for scenarios like identification, touchless access control, and privacy-related face blurring.

Accessing the Face Service:

- Use through a client library SDK.
- Directly call the REST API.

Getting Started:

- Follow the quickstart guide.
- Try out the Face service capabilities in your browser using Vision Studio.

Use Cases

1. **Identity Verification:**

- Verify user identity by comparing their face against a trusted image, like a government-issued ID or an enrollment photo.
- Liveness detection ensures the authenticity of the user's presence, preventing spoofing attacks with printed photos or masks.

2. **Touchless Access Control:**

- Implement opt-in face identification for enhanced access control in various settings, such as airports, stadiums, offices, and schools.
- Facial recognition streamlines the check-in process while minimising hygiene and security risks associated with traditional methods like cards or tickets.
- **Face Detection and Analysis:** The Detect API identifies human faces in images, providing their location and extracting attributes for accurate detection, facilitating various adjustments in applications.
- **Liveness Detection:** Essential for distinguishing real faces from counterfeits in video streams, crucial for verifying authentication and combating spoofing attempts, continually evolving to enhance security against emerging threats.
- **Identification:** Face identification allows matching one face to multiple faces in a secure repository, typically used for granting access or verifying users.
- **Verification:** This involves matching one face to a single face from a secure repository, verifying if they belong to the same person, commonly used in access control or identity verification processes.
- **Find Similar:** This operation matches a target face with a set of candidate faces, identifying similar faces, useful for face searches by image.
- **Group:** Divides unknown faces into smaller groups based on similarity, providing a "messyGroup" array for faces with no similarities found, useful for organizing faces into likely groups belonging to the same person.

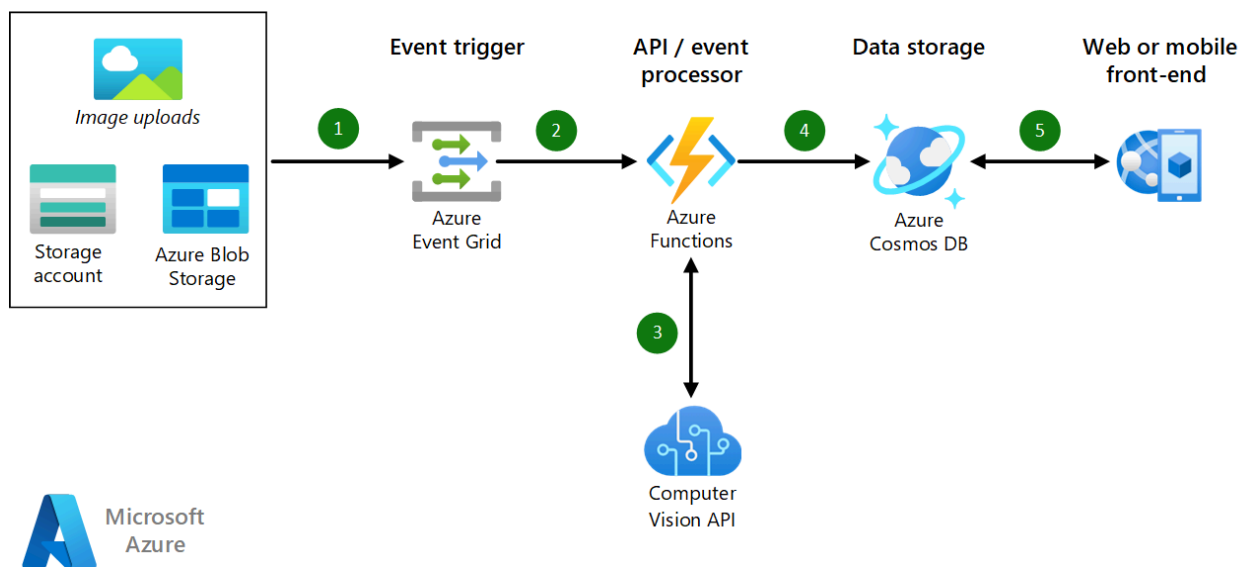
- Input Requirements:** Supported image formats include JPEG, PNG, GIF (first frame), BMP, with a maximum file size of 6 MB. For face detection, the minimum detectable face size is 36 x 36 pixels, with a maximum of 4096 x 4096 pixels, while for recognition, factors like extreme lighting or facial obstructions may affect recognition accuracy.

Identify common types of computer vision solution

Image classification on Azure

By leveraging Azure services like the Computer Vision API and Azure Functions, companies can avoid managing individual servers, cut costs, and benefit from Microsoft's expertise in image processing through Azure AI services.

This example scenario focuses on an image-processing application. For other AI requirements, explore the comprehensive range of Azure AI services.



[Source: Microsoft Documentation]

Workflow

The scenario outlines the backend processes for a web or mobile application. The data flow follows these steps:

Step	Description
1	Uploading images to Blob storage triggers an event in Azure Event Grid.
2	Azure Event Grid notifies and activates Azure Functions.
3	Azure Functions invoke the Azure Computer Vision API to analyze the image, accessing it via the blob URL.
4	Azure Functions store the Computer Vision API response, including analysis results and image metadata, in Azure Cosmos DB.
5	Analysis results are displayed on the web or mobile front end, but not the uploaded image itself.

1. Retail Industry:

- **Fashion Website:**
 - Automate image classification for product uploads.
 - Enhance product search via visual impression.

2. Gaming Industry:

- **Game Screenshot Classification:**
 - Automate classification of video games from screenshots.
 - Improve the accuracy of stream information on platforms like Twitch.
 - Aid in detecting novel games with custom models.

3. Finance and Insurance Industry:

- **Insurance Claims Processing:**
 - Speed up claims processing and underwriting.
 - Analyze damage from natural disasters and vehicle accidents.
 - Identify residential and commercial properties.

Object detection - Azure AI Vision

Azure's object detection uses **deep learning to identify objects in images**, marking them with bounding boxes. It then classifies these objects and integrates seamlessly into applications **through APIs**. This functionality enables tasks like image tagging, content moderation, and inventory management, driving innovation and enhancing user experiences.

Object Detection Functionality	Tagging Functionality
Identifies objects and living things in an image.	Applies tags based on objects or living things found in the image.
Returns bounding box coordinates (in pixels) for each detected object.	No formal relationship exists between the tagging and object detection taxonomies.
Useful for processing relationships between objects and determining multiple instances of the same object.	Object detection focuses solely on finding objects, while tagging can include additional contextual terms.

Detection Constraints

- Small objects, typically less than 5% of the image, **may not be accurately detected**.
- Objects **arranged closely together**, such as in stacks or clusters, might pose challenges for detection.
- **Brand or product name differentiation is not supported**, such as distinguishing between various soda types on a store shelf.

Utilising the API

Access the API through a **native SDK or via REST calls**.

- Specify the inclusion of Objects in the **visualFeatures query parameter** to enable object detection.
- Upon receiving the full JSON response, **extract and interpret the "objects"** section to access relevant information.

OCR - Optical Character Recognition - Azure AI services

OCR, or Optical Character Recognition, uses machine learning to extract text from images like posters, street signs, product labels, and various documents. It organizes the extracted text into words, lines, and paragraphs, providing digital access and reducing the need for manual data entry.

Microsoft's Read OCR Engine:

Supporting Multiple Languages

- Utilizes advanced machine-learning models for multilingual support.
- Extracts printed and handwritten text, accommodating diverse writing styles and mixed languages.

Deployment Flexibility

- Available as both a cloud service and an on-premises container, ensuring deployment flexibility.

Latest Update: Introducing Read as a Synchronous API

- Enhances OCR-assisted user experiences with performance improvements.
- Catering specifically to single, image-only scenarios.

Feature	Azure AI Vision	Document Intelligence
Integration into Platforms	Integrated into Azure AI Vision platform.	Integrated into the Document Intelligence platform.
Key Capabilities	- Extracts printed and handwritten text.	- Extracts printed and handwritten text.
	- Provides comprehensive information including page layouts, text lines, individual words, and their locations with confidence scores.	- Provides comprehensive information including page layouts, text lines, individual words, and their locations with confidence scores.
Support for Mixed Languages	Can effectively handle content in mixed languages and recognize both print and handwritten formats.	Can effectively handle content in mixed languages and recognize both print and handwritten formats.
On-Premises Deployment	Available as a Distroless Docker container for on-premises deployment.	Available as a Distroless Docker container for on-premises deployment.
Compatibility and Integration	Ensures compatibility and facilitates integration with existing infrastructure.	Ensures compatibility and facilitates integration with existing infrastructure.

Face detection, attributes, and input data

Face Detection:

- Face detection refers to the identification of human faces within an image.
- This process involves locating the positions of faces and can optionally include retrieving various face-related data.
- The Detect API is utilized to perform face-detection tasks on an image.

Face Attribute Data:

- Face attribute data encompasses different characteristics or features associated with detected faces.
- These attributes may include details such as age, gender, emotion, facial landmarks, and facial hair.

- By using the Detect API, users can access and retrieve face attribute data along with the detected faces in an image.

Face Attributes in Face Detection

Attributes	Description
Face Rectangle	<ul style="list-style-type: none"> - Represents each detected face. - Provides pixel coordinates for the left, top, width, and height of the face. - Faces listed in size order from largest to smallest.
Face ID	<ul style="list-style-type: none"> - Unique identifier for each detected face. - Requires limited access approval, obtainable via an intake form. - Can be requested in the Detect API call.
Face Landmarks	<ul style="list-style-type: none"> - Identifiable points on a face, such as pupils or nose tip. -The default set includes 27 predefined landmarks. - Coordinates provided in pixel units. - Detection_03 model offers precise landmark detection.

Identify features of generative AI solutions

Fundamentals of Generative AI - Training

Artificial Intelligence and Generative AI

- **Artificial Intelligence (AI):**
 - Mimics human behaviour using machine learning.
 - Interacts with the environment to perform tasks without detailed instructions.
- **Generative AI:**
 - A specialised branch of AI that creates new content.
 - It is commonly found in chat applications, such as Microsoft Copilot, which aids in more effective web browsing.
 - Processes natural language inputs and produces suitable responses in various forms, including text, images, and code.

Language Models:

Generative AI applications leverage language models, which are specialized machine learning models designed for natural language processing (NLP) tasks, such as:

- Assessing sentiment or classifying text.
- Summarizing content.
- Evaluating semantic similarity across different text sources.
- Creating new text.

Although the mathematical concepts behind these language models are intricate, having a basic grasp of their architecture can provide a conceptual understanding of their functionality.

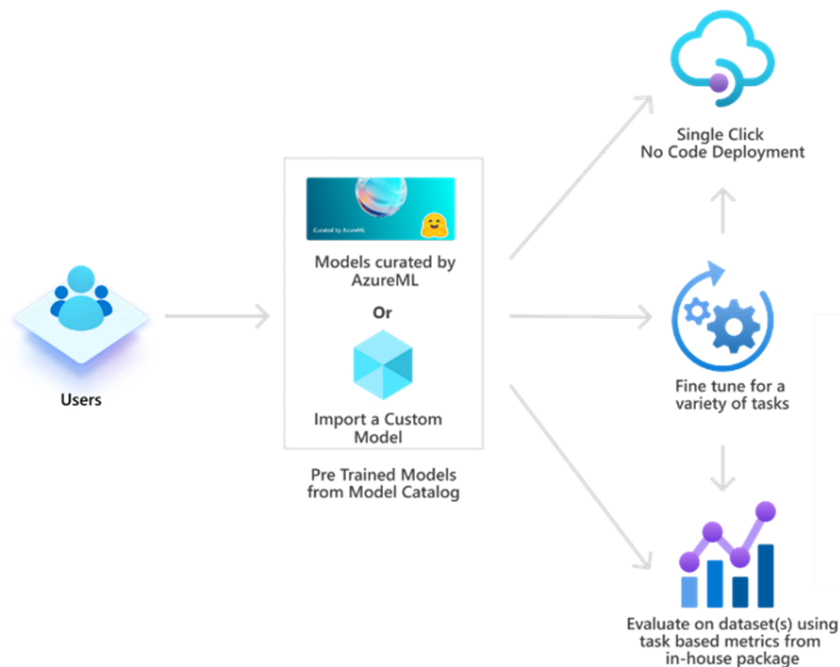
How transformer model work?

- **Transformer Architecture:** Modern state-of-the-art language models are based on transformer architecture, enhancing previously successful techniques in vocabulary modeling for NLP tasks, particularly language generation.
- **Training with Large Text Corpora:** Transformer models are trained on extensive amounts of text, allowing them to understand and represent semantic relationships between words.
- **Predictive Text Sequences:** These models use the learned relationships to predict probable and coherent text sequences.
- **Human-like Language Generation:** Transformer models with sufficiently large vocabularies can generate responses that are nearly indistinguishable from those produced by humans.

Transformer Models:

Large Language Models (LLMs)	Small Language Models (SLMs)
LLMs are trained on extensive and diverse datasets from various sources, covering a broad spectrum of topics.	SLMs focus on narrower datasets, honing in on specific subject areas.
Trained LLMs boast billions or trillions of parameters, facilitating the generation of complex language sequences.	SLMs typically have fewer parameters compared to LLMs.
LLMs exhibit versatile language generation capabilities across diverse conversational contexts.	SLMs excel in generating content tailored to specific conversational topics but may lack versatility in broader language tasks.
The significant size of LLMs can impact performance and present challenges for local deployment.	SLMs offer greater deployment flexibility, including local deployment and quicker fine-tuning, owing to their smaller size.
Fine-tuning LLMs with additional data can be resource-intensive and time-consuming.	Fine-tuning SLMs is generally less demanding in terms of time and resources compared to LLMs.

Azure AI Model Resources and Catalog



[Source: Microsoft Documentation]

Azure OpenAI Service:

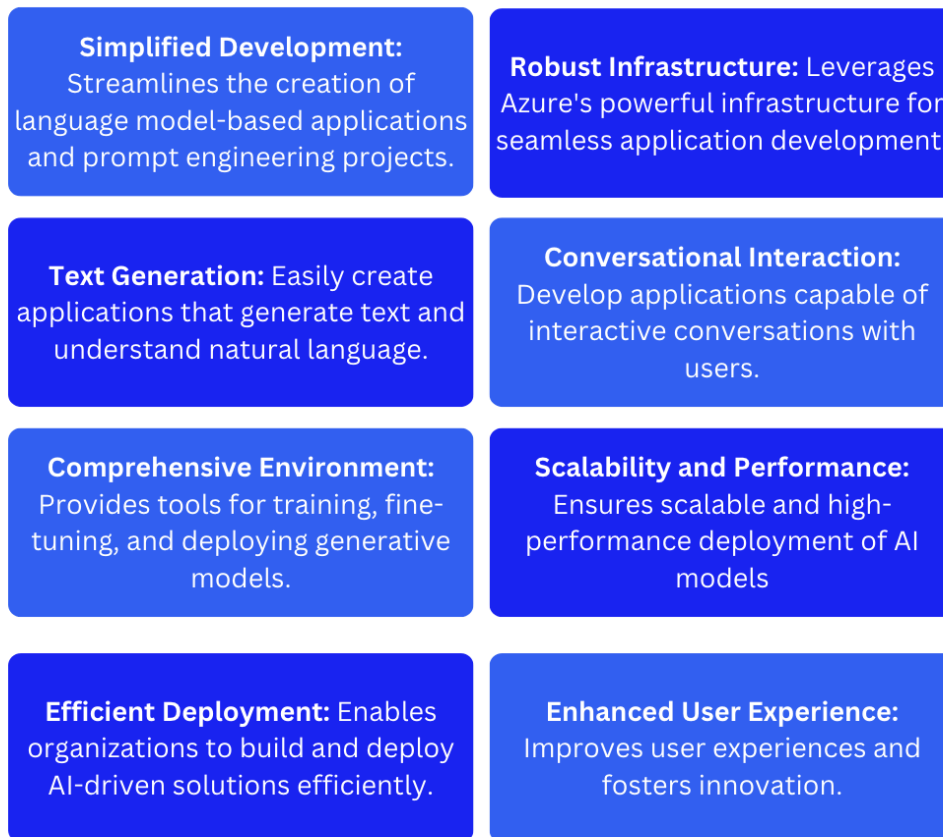
- Provides access to foundation models like GPT and DALL-E.
- Offers secure and scalable hosting on Azure.
- Enables developers to fine-tune models with custom training data.

Model Catalog in Azure:

- Curated repository of pre-trained AI models.
- Accessible to data scientists and developers.
- Includes latest open-source models from various partners.
- Enables easy integration of advanced AI capabilities into projects.

Generative AI in Azure Machine Learning

What is Generative AI and its features?



Integrate Generative AI with Applications

Foundation Model Catalog

- Model Collection: Access a variety of **open-source foundation models** and customize them with your own data.

Prompt Flow

- AI Development Streamlining: Simplify AI app creation and prompt management with interactive **visual flows, variations, and assessments**.

Open-Source Framework Integration

- Flow Integration: Incorporate your [LangChain or Semantic Kernel-developed flows](#) into the prompt flow for seamless scalability of experiments.

Azure AI Infrastructure

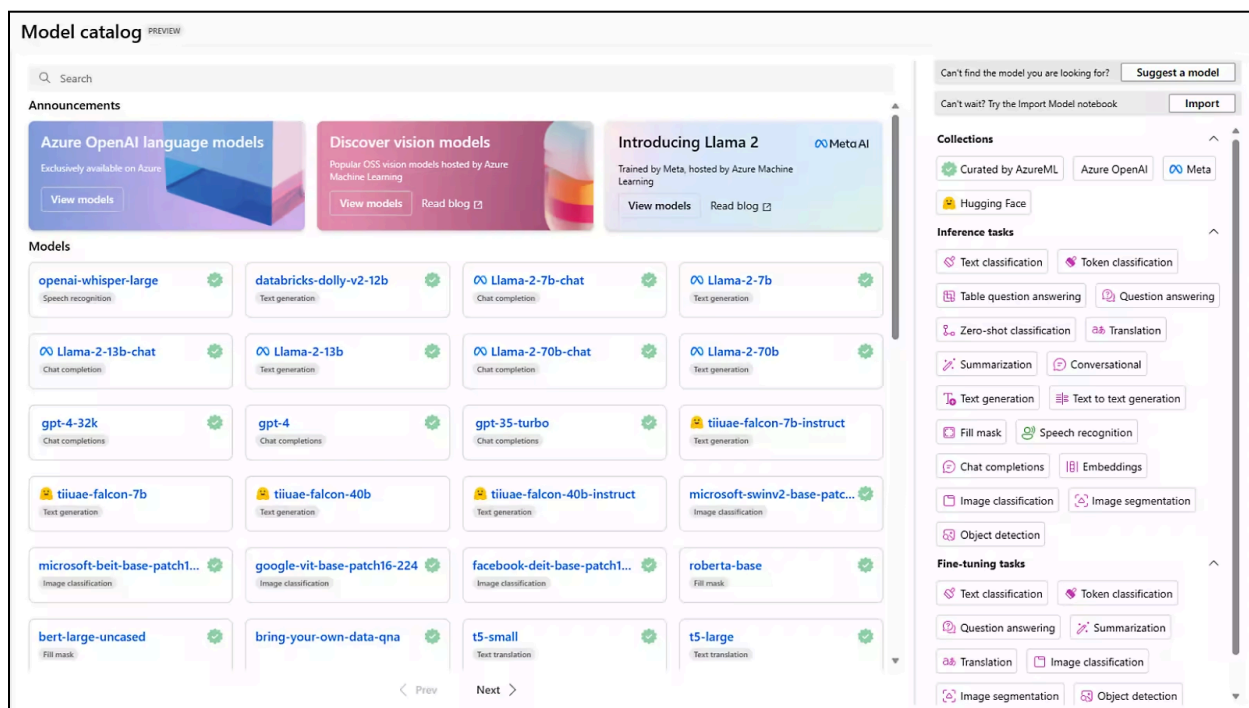
- Purpose-Built Optimization: Enhance model fine-tuning on specialized infrastructure using [DeepSpeed and ONNX Runtime](#) for improved performance.

Content Safety

- Harmful Content Filtering: Detect and filter harmful content using severity scores with an [integrated connection to Azure AI Content Safety](#).

Model Monitoring

- Performance Tracking: [Track safety, quality metrics](#), and [token consumption](#) for generative AI applications in production.



[Source: Microsoft Documentation]

10 ways generative AI and Azure OpenAI Service

1. Content Creation & Design:

- Generative AI revolutionizes content creation with **high-quality images, videos, and graphics**.
- **Example:** Companies like Typeface leverage AI to customize content, enhancing marketing efforts.

2. Accelerated Automation:

- Generative AI automates IT tasks, **improving efficiency and employee experiences**.
- **Example:** AT&T streamlines operations with Azure OpenAI Service, boosting productivity.

3. Personalized Marketing:

- Generative AI tailors **marketing campaigns** to individual preferences, boosting engagement.
- Take Blip and Azure OpenAI Service to optimize customer interactions, enhancing marketing strategies.

4. Chatbots & Virtual Assistants:

- **AI-powered chatbots** deliver instant, accurate customer support, enhancing satisfaction.
- **Example:** Johnson & Johnson deploys cost-effective chatbots using Azure Bot Services.

5. Product & Service Innovation:

- AI **optimizes content, transforming it** into engaging videos.
- Nota uses Microsoft Azure OpenAI Service to enhance content distribution and revenue generation.

6. Language Translation & NLP:

- Generative AI improves **real-time translation**, facilitating global communication.

- **Example:** Microsoft Azure AI enhances HelloTalk’s language learning tools, enabling cultural exchange.

7. Fraud Detection & Cybersecurity:

- AI detects fraud by **analyzing patterns in data**, ensuring secure transactions.
- **Example:** Swift and Microsoft use Azure Machine Learning for anomaly detection, and safeguarding financial transactions.

8. Predictive Analytics & Forecasting:

- AI aids decision-making with **predictive analytics, optimizing operational efficiency**.
- Azure IoT supports Husky in system optimization and infrastructure maintenance.

9. Creative Writing & Content Generation:

- AI automates content creation, **generating articles and marketing materials**.
- CarMax efficiently creates website content with Azure OpenAI Service.

10. Medical Research & Diagnosis:

- AI analyzes **medical data** for disease diagnosis, improving patient outcomes.
- Generative models assist researchers in inaccurate diagnoses and treatment development.

What is Responsible AI?

Responsible AI Implementation: Azure AI facilitates the development, utilization, and governance of AI solutions in a responsible manner, ensuring compliance with ethical standards and regulations.

Features

Feature	Description
<p>Scalability and Confidence</p>	<p>Organizations can confidently scale their AI initiatives with Azure AI, utilizing tools and methodologies developed by Microsoft experts to ensure reliability and confidence in their projects.</p>
<p>Built-in Governance Tools</p>	<p>Azure AI provides configurable controls and governance tools, allowing organizations to proactively manage AI-related risks and transition from reactive to proactive risk management approaches.</p>
<p>Industry-Leading Technologies</p>	<p>Azure AI incorporates industry-leading technologies and best practices to address various challenges associated with AI deployment, including risk management, accuracy enhancement, privacy protection, transparency, and compliance.</p>
<p>Integration of Responsible AI</p>	<p>Azure AI enables cross-functional teams to seamlessly integrate responsible AI practices into their workflows, offering built-in tools and templates for open-source, MLOps, and generative AI workflows.</p>
<p>Trusted Foundation</p>	<p>Azure AI serves as a trusted foundation for building AI applications, offering enterprise-grade privacy, security, and compliance capabilities developed by experts across Microsoft's research, policy, and engineering divisions.</p>

Principles of Responsible AI

Microsoft's Responsible AI Commitment

- Microsoft prioritizes **transparency, equity**, and accountability in AI development, embedding responsible AI practices company-wide.
- Collaboration among AI experts, policymakers, and engineers drives the **establishment of governance frameworks for ethical AI deployment.**

MLOps: Streamlining ML Workflows

- MLOps automates ML processes, **reducing manual efforts** in model deployment and enhancing workflow management.
- It fosters cross-functional collaboration, improving **communication and teamwork** among data scientists and engineers for enhanced reproducibility.

Model Interpretability: Ensuring Transparency

- Model interpretability aids in understanding **algorithmic decisions**, facilitating the detection of biases and ethical concerns.
- Prioritizing model interpretability ensures transparency and confidence in AI systems, **enabling informed decision-making and ethical deployment.**

Identify capabilities of Azure OpenAI Service

Overview of Responsible AI practices for Azure OpenAI

Responsible AI Implementation Framework:

Identify Potential Harms:

- Conduct impact assessments and iterative testing to identify potential harms.
- Prioritize identified harms based on risk factors.
- Collaborate with stakeholders to share findings and insights.

Measure Harm Occurrences:

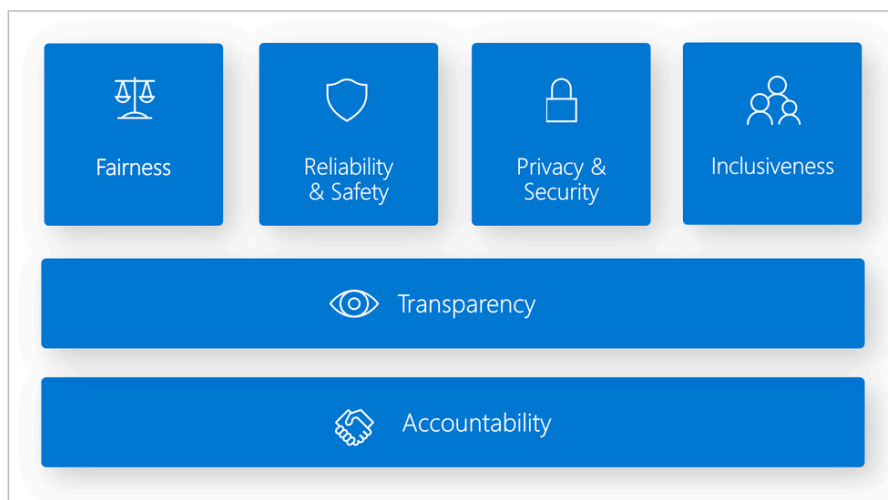
- Develop clear metrics to measure harm frequency and severity.
- Implement manual and automated measurement approaches.
- Continuously monitor and update measurement results for ongoing refinement.

Mitigate Identified Harms:

- Implement layered mitigations at model, safety system, application, and positioning levels.
- Educate users on AI limitations and encourage responsible usage.
- Establish mechanisms for incident response, rollback, and user feedback channels.

Operational Readiness:

- Deploy systems gradually using phased delivery approaches.
- Develop incident response and rollback plans for quick action.
- Collect telemetry data to monitor system performance and user satisfaction.
- Maintain compliance with legal and regulatory requirements through continuous evaluation and adaptation.



[Source: Microsoft Documentation]

Responsible AI practices for Azure OpenAI:

Identifying AI Harms:

- Assess **potential harms** based on model, application, and context.
- **Prioritize risks** by frequency, severity involving experts.
- Conduct red team testing, stress testing for **anticipatory insights**.
- **Share findings** to refine mitigation strategies and ensure compliance

Systematic Measurement of AI Harms:

- Use **manual methods** initially for small, prioritized issues.
- **Develop measurement sets**, clear metrics for harmful outputs.
- **Share results, refine metrics** iteratively for ongoing tracking.

Mitigating AI Harms:

- **Model Level:** Understand model capabilities, experiment with **versions to assess harms**.
- **Safety System Level:** **Evaluate platform solutions** like content filters for risk

mitigation.

- **Application Level: Implement prompt engineering**, user-centered design to prevent misuse.
- **Positioning Level: Provide transparent documentation**, guidelines for user education.

What is Azure OpenAI Service?

Contrast Between Azure OpenAI Service and OpenAI

Feature	Azure OpenAI Service	OpenAI
Model Access	GPT-4, GPT-3, Codex, DALL-E, Whisper, text-to-speech models	GPT-4 and GPT-3 models
Development	Co-develops APIs with OpenAI for compatibility	Independently develops APIs
Security	Integrates Microsoft's Azure security features	Uses proprietary security measures
Enterprise	Provides enterprise-grade solutions with Microsoft's reliability	Focuses on AI research and model development
Private Networking	Supports secure private networking	Not inherently included
Regional Availability	Models available regionally through Azure's global infrastructure	Availability varies based on policies
Responsible AI	Implements content filtering for responsible AI	Focuses on AI ethics and principles
Compatibility	Ensures compatibility through collaborative API development	Models directly managed by OpenAI
Content Filtering	Provides content filtering for responsible AI usage	Uses own content moderation practices

Azure OpenAI Service for Advanced Language and Image Processing Tasks

Resources

Azure OpenAI Setup:

- Create a resource/instance in your Azure Subscription.
- Refer to Azure's resource management design for details.

Deployments

Model Deployment:

- Create an Azure OpenAI Resource.
- Deploy a model using Deployment APIs before making API calls.

Prompt Engineering

Prompt-Based Models:

- Models respond to text prompts with text completions.
- Effective prompt crafting (prompt engineering) is crucial.
- Example models: GPT-3, GPT-3.5, GPT-4.

Models

Model Varieties:

- DALL-E: Generates images from text prompts.
- Whisper: Transcribes and translates speech to text.
- Text to Speech: Synthesizes speech from text (in preview).

Azure OpenAI for Code Generation and Development Tasks

Overview

- **Azure OpenAI Service:** Leverages large language models (LLMs) for content generation, including programming code.
- **Functionality:** Allows developers to generate and improve existing code across various programming languages, enhancing efficiency and understanding.

Utilizing Azure OpenAI to Generate Code and Assist with Development Tasks

Steps	Description
Construct Code from Natural Language	<ul style="list-style-type: none"> - Generate code from natural language prompts. - Handle tasks ranging from simple commands to full applications. - Edit and update provided code or previous responses.
AI Models for Code Generation	<ul style="list-style-type: none"> - Previous Generations: Some models specifically trained for code (Codex models). - Recent Generations: Models like GPT-3.5-Turbo and GPT-4 handle both natural language and code without specialized models. - Example Used: GPT-3.5-Turbo.
Write Functions	<ul style="list-style-type: none"> - Create functions and applications by describing desired functionality.
Change Coding Language	<ul style="list-style-type: none"> - Translate code from one programming language to another.
Understand Unknown Code	<ul style="list-style-type: none"> - Interpret and understand unclear or unfamiliar code.
Complete Code and Assist Development	<ul style="list-style-type: none"> - Assist in writing unit tests, completing partial code, commenting, and generating documentation.
	<ul style="list-style-type: none"> - Enable developers to focus on complex programming and problem-solving tasks.
Complete Partial Code	<ul style="list-style-type: none"> - Generate code based on comments, function names, and partially written code.
	<ul style="list-style-type: none"> - More complex functions require additional context for useful responses.
Fix Bugs and Improve Code	<ul style="list-style-type: none"> - Fix Bugs: Analyze and suggest changes to fix issues in the code.
	<ul style="list-style-type: none"> - Improve Performance: Identify more efficient ways to accomplish tasks.
	<ul style="list-style-type: none"> - Refactor Inefficient Code: Suggest refactoring for better maintenance and reduced bugs.

Quickstart: Generate images with Azure OpenAI Service

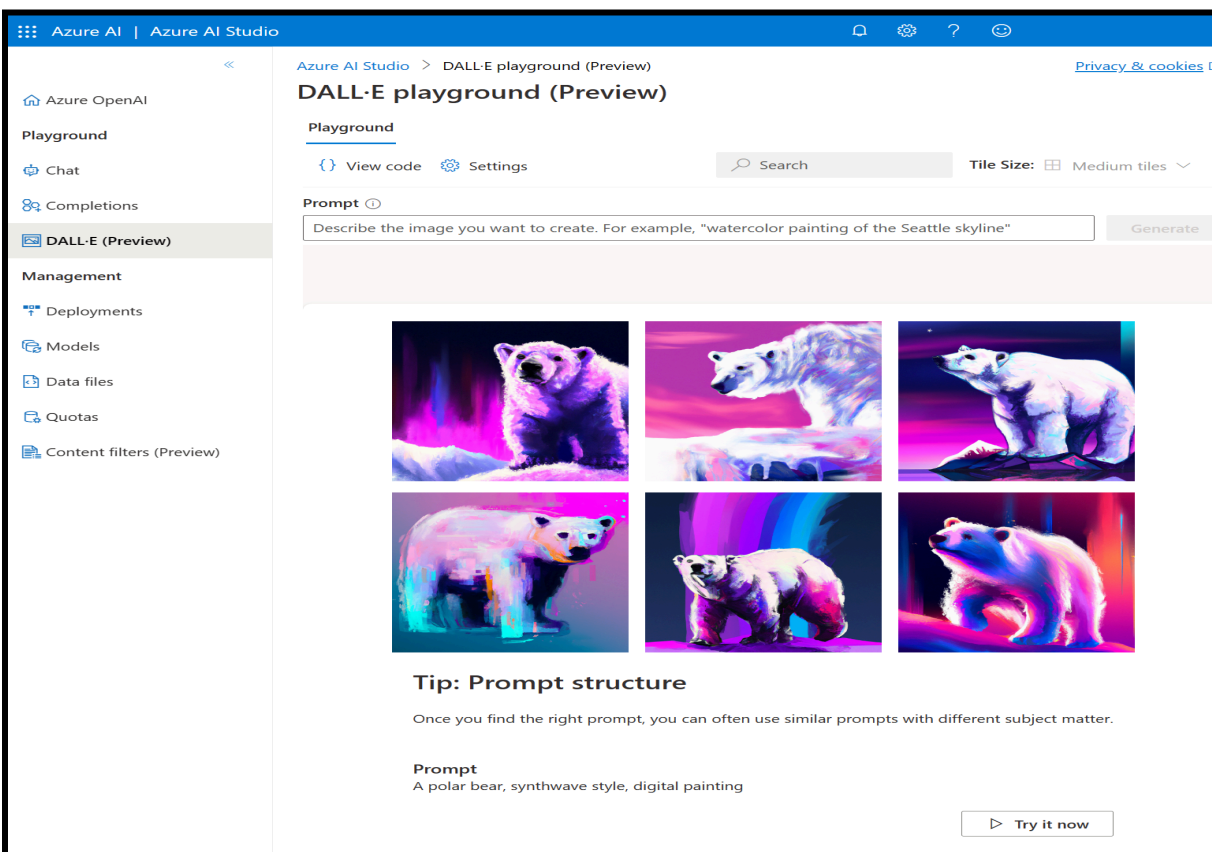
Navigate to Azure OpenAI Studio and login using the credentials linked to your Azure OpenAI resource. During or after the login process, choose the correct directory, Azure subscription, and Azure OpenAI resource.

Begin exploring Azure OpenAI capabilities through a no-code approach with the DALL·E playground (Preview).

- Type your image prompt into the text box. Then click generate
- The AI-generated image will display on the page when it's ready.

Additionally, in the DALL·E playground (Preview), you can:

- Access pre-configured Python and cURL code samples based on your settings.
- Click View code at the top of the page to see the samples.
- Utilize this code to build an application that accomplishes the same task.



[Source: Microsoft Documentation]

NLP Workload Scenarios and Common Features

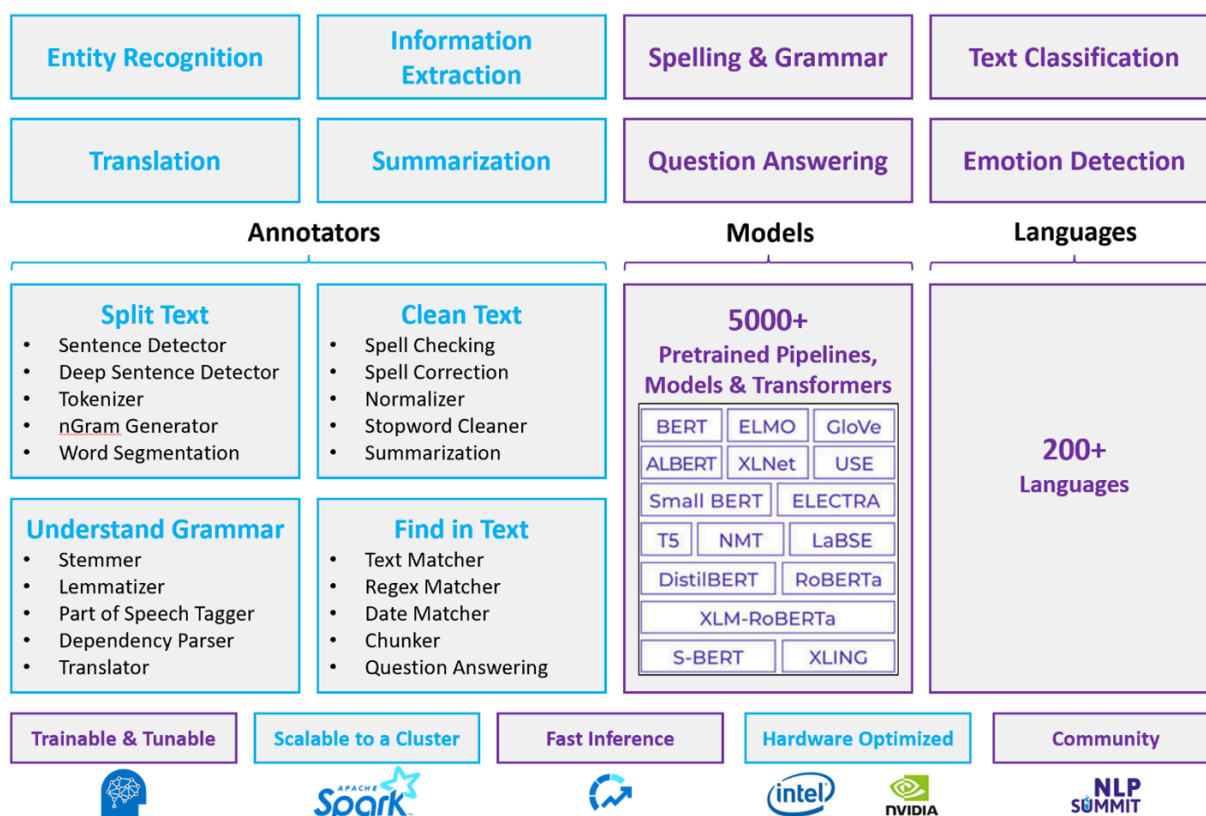
What is NLP?

- Natural language processing (NLP) enables applications to interact with users by seeing, hearing, speaking, and understanding them.
- Microsoft Azure simplifies the creation of NLP-supported applications through text analytics, translation, and language understanding services.
- NLP has various applications, including:
 - Sentiment analysis
 - Topic detection
 - Language detection
 - Key phrase extraction
 - Document categorization

Customized NLP Frameworks with Apache Spark

- Apache Spark boosts big-data analytics with efficient in-memory processing.
- Platforms like Azure Synapse Analytics, Azure HDInsight, and Azure Databricks utilize Spark's capabilities.
- Spark NLP is ideal for customized NLP tasks, processing large text datasets.
- This open-source library supports Python, Java, and Scala, offering full NLP functionalities similar to spaCy, NLTK, Stanford CoreNLP, and OpenNLP.
- Key features of Spark NLP include:
 - Spell checking
 - Sentiment analysis
 - Document classification

Spark NLP excels with superior accuracy, speed, and scalability compared to earlier solutions.



[Source: Microsoft Documentation]

What is key phrase extraction in Azure AI Language?

Key phrase extraction is a feature provided by Azure AI Language, which encompasses a suite of machine learning and AI algorithms available in the cloud for creating intelligent applications that process written language. This feature allows you to rapidly identify the central ideas and concepts within the text.

Methods to use key phrases extraction:

Here are three ways you can utilize Azure's Language Studio for keyphrase extraction:

1. Language Studio Platform:

- Accessible online without needing an Azure account.
- Allows experimentation with NLP features using provided text examples or your own uploaded data.
- Users can visit the Language Studio website and follow a quickstart guide for detailed instructions.

2. REST API or Client Library (Azure SDK):

- Enables developers to integrate keyphrase extraction directly into their applications.
- Offers a REST API for making HTTP requests and a Client Library in various programming languages.
- Developers can utilize the REST API by sending HTTP requests to specific endpoints or use the Client Library's pre-built methods in their preferred programming language.

3. Docker Container Deployment:

- Allows for running the key phrase extraction service locally on-premises infrastructure via Docker containers.
- Provides greater control over data and operations, suitable for environments with strict compliance and security requirements.
- Users can download and run the Docker container image, beneficial for organizations with specific regulatory or operational constraints.

Each of these options caters to different use cases and technical requirements, providing flexibility in how you deploy and use the key phrase extraction feature.

Summary of the typical workflow for using this feature in Azure AI

Language:

- ➔ Begin by establishing an **Azure AI Language resource**, granting access to Azure's language processing functionalities. This resource provides a unique key and endpoint URL for authentication.
- ➔ Proceed by creating a request using either the **REST API or supported client libraries (C#, Java, JavaScript, Python)**. For efficiency, consider combining multiple feature requests into a single batch request.
- ➔ Subsequently, **submit your request** containing the text data to be analyzed. Utilize the generated key and endpoint for authentication purposes.
- ➔ Lastly, handle the **API response locally**, either through direct streaming or storage for further examination or utilization.

What is Named Entity Recognition (NER) in Azure AI

Language?

- Named Entity Recognition (NER) is a feature available in Azure AI Language, which comprises machine learning and AI algorithms hosted in the cloud for building intelligent applications involving written language.
- NER can identify and categorize entities within unstructured text, such as people, places, organizations, and quantities.
- The prebuilt NER feature comes with a predefined list of recognized entities, allowing for immediate use.
- Additionally, the custom NER feature enables users to train the model to identify specialized entities tailored to their specific use cases.

Methods to use Name Entity Recognition (NER):

Two methods are available for utilizing NER:

Language Studio Platform: An online tool for NLP experimentation, accessible without an Azure account, offering text examples and a quickstart guide.

REST API or Client Library (Azure SDK): Enables NER integration via REST API or Client Library in multiple languages for developers, streamlining application development.

Use Cases:

Enhance search functionalities and indexing: Users can leverage detected entities within documents to construct knowledge graphs, enriching document search by employing them as tags.

Customer insights and analysis: Analyze customer feedback across various channels like reviews, emails, and calls to identify prevalent topics and trends over time, aiding in understanding customer preferences and behaviors.

Automate operational workflows: For instance, in insurance claim assessment, identified entities such as names and locations could be highlighted, expediting the review process. Alternatively, email content could trigger the automatic generation of support tickets with customer details.

What is sentiment analysis and opinion mining?

Sentiment analysis and opinion mining are features within the Language service, a cloud-based suite of AI algorithms for analyzing written language. They help discern public sentiments about brands or topics by detecting positive or negative cues in text and associating them with specific elements. Compatible with multiple languages.

Sentiment analysis:

This feature assigns sentiment labels like "negative," "neutral," and "positive" based on the highest confidence score identified by the service at both sentence and document

levels. Additionally, it provides confidence scores ranging from 0 to 1 for each document and its sentences for positive, neutral, and negative sentiments.

Opinion mining:

A subset of sentiment analysis, opinion mining, also known as aspect-based sentiment analysis in NLP, offers detailed insights into opinions related to specific words, such as attributes of products or services, within the text.

Methods to use sentiment analysis:

Language Studio Platform: Online tool for NLP experimentation with provided or uploaded data, accessible without an Azure account, with a quickstart guide available.

REST API or Client Library (Azure SDK): Integrates sentiment analysis into applications via HTTP requests or client libraries in multiple languages.

Docker Container Deployment: Runs sentiment analysis service locally via Docker containers, ideal for strict compliance and security needs.

What is Azure AI Language - Azure AI services?

Azure AI Language is a cloud service offering Natural Language Processing (NLP) capabilities for text analysis and comprehension. It aids in developing intelligent applications through the web-based Language Studio, REST APIs, and client libraries. This service consolidates several former Azure AI services, including Text Analytics, QnA Maker, and LUIS.

The Language service introduces new features that can be:

- **Preconfigured:** These features use fixed AI models, requiring only your data input to generate outputs for your applications.
- **Customizable:** These features allow you to train AI models using provided tools to tailor them specifically to your data.

Features:

The Language service also offers several new features, which can be either:

- **Preconfigured:** These features use AI models that are not customizable. You simply send your data and use the feature's output in your applications.
- **Customizable:** These features allow you to train an AI model using our tools to specifically fit your data.

Preconfigured Features:

- Named Entity Recognition (NER)
- Personally-Identifying Information (PII) and health (PHI) information detection
- Language detection
- Sentiment Analysis and Opinion Mining
- Summarization
- Key phrase extraction
- Entity linking
- Text analytics for health
- Question answering

Customizable Features:

- Custom text classification
- Custom Named Entity Recognition (Custom NER)
- Custom text analytics for health
- Conversational language understanding
- Orchestration workflow

Use Cases:

Objective	Document Format	Recommended Solution
Extract categories of information without a custom model.	Unstructured text	Preconfigured NER feature
Extract information categories using a custom model.	Unstructured text	Custom NER
Identify main topics and key phrases.	Unstructured text	Key phrase extraction
Analyze sentiment and opinions in text.	Unstructured text	Sentiment analysis and opinion mining
Clarify entities and provide Wikipedia links.	Unstructured text	Entity linking
Classify documents into categories.	Unstructured text	Custom text classification
Extract medical data from clinical documents without a custom model.	Unstructured text	Text analytics for health
Develop a conversational app that responds to user inputs.	Unstructured user inputs	Question answering
Identify user intent and extract information from user inputs.	Unstructured user inputs	Conversational language understanding

What is the Speech service? - Azure AI services

Speech-to-Text and Text-to-Speech Functionalities

The Speech service offers both speech-to-text and text-to-speech functionalities via a Speech resource. You can accurately convert spoken words into text, generate natural-sounding voices from text, translate spoken audio, and recognize individual speakers in conversations.

Customization and Model Development

You can create custom voices, expand your base vocabulary with specific words, or develop personalized models. The service runs seamlessly in the cloud or on the edge in containers.

Easy Integration

Enabling speech capabilities in your applications, tools, and devices is straightforward with the Speech CLI, Speech SDK, Speech Studio, or REST APIs.

Deployment Options for Azure AI Speech:

Cloud Deployment:

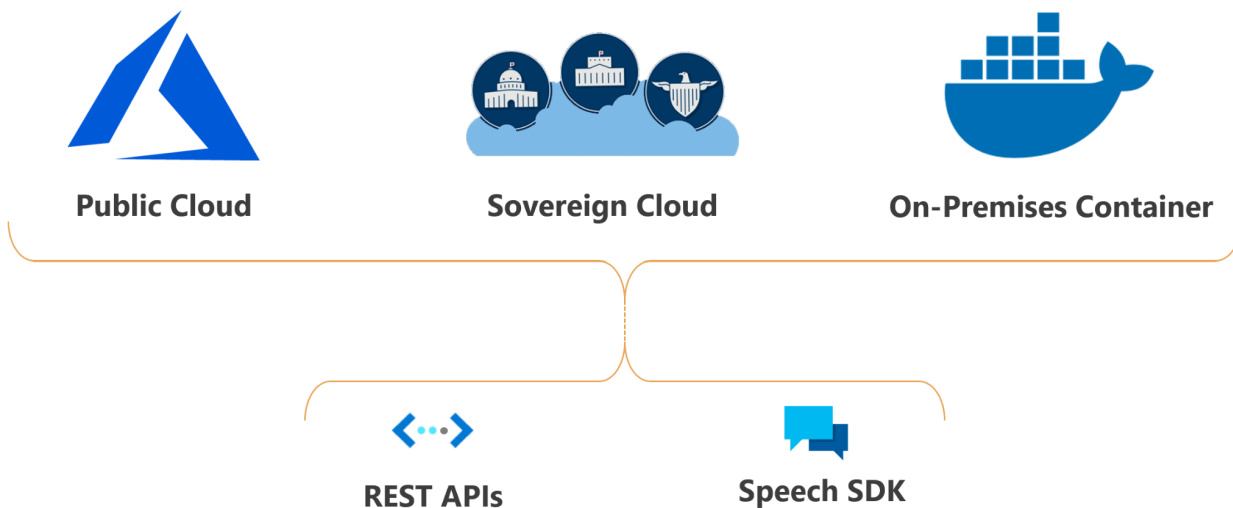
Implement Azure AI Speech functionalities in the cloud, leveraging Microsoft's scalable and dependable infrastructure. This method ensures flexibility and accessibility, enabling users to utilize the service from anywhere with an internet connection.

On-Premises Deployment Using Containers:

Deploy Azure AI Speech on-premises by utilizing containers, thereby bringing the service closer to your data to meet compliance, security, or other operational needs. This approach grants enhanced control over data governance and privacy.

Sovereign Cloud Deployment:

Opt for Azure AI Speech deployment in sovereign clouds, catering specifically to certain government entities and their affiliates. This alternative guarantees adherence to regulatory stipulations and addresses data sovereignty concerns pertinent to specific geographic regions or jurisdictions.



[Source: Microsoft Documentation]

Common Use Cases for Speech

Captioning

Learn how to sync captions with your audio input, apply profanity filters, obtain partial results, make customizations, and identify languages spoken in multilingual contexts.

Audio Content Creation

Use neural voices to enhance interactions with chatbots and voice assistants, convert e-books into audiobooks, and improve in-car navigation systems.

Call Centers

Transcribe calls in real-time or in batches, redact personal information, and extract insights such as sentiment to optimize call center operations.

Language Learning

Give pronunciation feedback to language learners, support live transcription for remote learning, and read aloud educational materials using neural voices.

Voice Assistants

Create natural, human-like conversational interfaces for applications and experiences, providing quick and reliable interactions between devices and assistant implementations.

Microsoft Applications

Microsoft uses Speech for various purposes, including captioning in Teams, dictation in Office 365, and the Read Aloud feature in the Microsoft Edge browser.



[Source: Microsoft Documentation]

What is Document Translation? - Azure AI services

Document Translation is a feature offered by the Azure AI Translator service, allowing cloud-based machine translation of various documents. It supports translation of multiple and intricate documents across all available languages and dialects while maintaining the original document structure and data format intact.

The Document Translation API facilitates two primary translation operations:

1. Asynchronous Batch Document Translation
2. Synchronous Document Translation

Objective	Asynchronous Batch Document Translation	Synchronous Document Translation
Processing	Supports asynchronous processing of numerous documents and large files.	Enables real-time processing of individual file translations
Storage Requirement	Requires an Azure Blob storage account with dedicated containers for both the with dedicated containers for both the	Doesn't require an Azure Blob storage account
Return of Translation [Results]	Translated documents are stored in the Azure Blob storage and can be accessed after processing.	Translated document is directly returned to the calling client as part of the final response.

Asynchronous Batch Document Translation

Utilize asynchronous document processing to translate multiple documents and large files seamlessly.

For batch development options:

1. **REST API:** This language-agnostic interface allows you to craft HTTP requests and authorization headers to initiate document translation.
2. **Client-library SDKs:** These are language-specific tools comprising classes, objects, methods, and code snippets that you can easily integrate into your project. Currently, Document Translation supports programming languages like C#/.NET and Python.

Supported Document Formats:

To obtain a list of supported document formats, utilize the "Get supported document formats" method, which includes common file extensions and their corresponding content types when using the upload API.

Example:

Document Format	File Extension	Description
Adobe PDF	pdf	Adobe's Portable Document Format, widely utilized for various document types. Document Translation utilizes OCR technology to extract and translate text from scanned PDFs while maintaining the original layout.
Comma-Separated Values	csv	A file format commonly employed in spreadsheet applications, featuring data separated by commas.
Microsoft Excel	xls, xlsx	A spreadsheet format extensively used for data analysis and documentation purposes.
Microsoft PowerPoint	ppt, pptx	A file format designed for creating presentations, often utilized for professional or educational presentations.
Microsoft Word	doc, docx	A standard file format for textual documents, frequently used for word processing tasks.

Synchronous Document Translation

Utilize synchronous translation processing by including the document within the HTTP request body and retrieving the translated document directly in the HTTP response.

Features:

Translate individual-page documents

Description: The synchronous translation request accommodates only one document per request.

Maintain original document formatting

Description: Translate files while conserving their initial layout and structure.

Personalized translation

Description: Translate documents using both standard and customized translation models.

Customized glossaries

Description: Enhance translations by leveraging custom glossaries tailored to specific terminology.

Single-language translation

Description: Support translation between two languages without multilingual capabilities.

Automatic language detection

Description: Allow the Document Translation service to identify the language of the provided document automatically.

Leverage customized glossaries

Description: Integrate custom glossaries into the translation process for improved accuracy and consistency.

Azure tools and services for NLP workloads

Azure AI Language - Capabilities

What is AI Language?

Azure AI-Language is a managed service tailored for crafting applications that process natural language. It empowers users to recognize key terms, assess sentiment, condense text, and construct conversational interfaces. With Azure AI-Language, you can annotate, train, assess, and implement adaptable AI models without extensive knowledge of machine learning.

- Utilize state-of-the-art natural language processing driven by advanced AI exploration.
- Employ multilingual models trainable in one language and applicable across numerous others.
- Access a visual interface facilitating annotation, training, assessment, and deployment of models.
- Ensure data security and confidentiality, with text inputs exclusively utilized for model training objectives.

Use Cases:

Recognize core concepts:

- Identify fundamental ideas and principles within the text.
- Extract, categorize, and mask essential details from various types of content.

Comprehend information at a glance:

- Quickly grasp the essence of documents or discussions.
- Summarize content from customer interactions like calls and reviews for swift comprehension.

Analyze customer sentiment:

- Gain insights into customer perceptions and emotions.
- Understand customer opinions and attitudes through sentiment analysis, aiding decision-making and improving customer experiences.

Azure AI Speech - Capabilities

What is AI Speech?

A managed service provides cutting-edge speech capabilities, including speech-to-text, text-to-speech, speech translation, and speaker recognition.

Confidently and swiftly develop voice-enabled applications using the Speech SDK.

Achieve highly accurate speech-to-text transcription, generate natural-sounding text-to-speech output, translate spoken audio, and implement speaker recognition in conversations. Explore a no-code environment and tailor custom models to your application's needs using Speech Studio.

Features:

- **Exceptional Quality:** Access top-tier speech-to-text, lifelike text-to-speech, and renowned speaker recognition.
- **Compliance and Security:** Maintain data ownership and confidentiality—your speech input isn't stored during processing.
- **Customization Options:** Personalize voices, expand vocabulary with specific words, or create bespoke models.
- **Deployment Flexibility:** Deploy Speech solutions anywhere—whether in the cloud or at the edge using containers.

Use Cases:

Transcribe Speech:

- Swiftly and accurately convert audio into text across more than 100 languages and variations.
- Utilize call center transcription for customer insights, enhance user experiences with voice-enabled assistants, and capture important discussions during meetings and other interactions.

Text-to-Speech:

- Enable your applications and services to speak naturally using text-to-speech functionality.

- Create lifelike audio content, enhance accessibility through read-aloud features, and develop customized voice assistants.

Speech Translation:

- Translate spoken audio in real-time from over 30 languages.
- Tailor translations to include your organization's specific terms, all within your preferred programming language.

Speaker Recognition:

- Verify individuals' identities or identify speakers within meetings and other settings.
- Enhance your application with speaker verification and identification capabilities.

Custom Keyword Activation:

- Personalize your brand's presence by establishing custom keywords for IoT devices and voice-enabled assistants.
- Elevate brand recognition, enhance user interaction, and bolster security with unique keywords.

Voice Command Integration:

- Implement hands-free voice commands to facilitate touchless interactions and improve safety.
- Create a voice-first experience to support various scenarios, including back-to-work environments.