



AI on Azure

AI 900



Presenter

Mohammed Arif, PhD GenAI Architect & Data Scientist



Mohammed Arif has more than eighteen (18+) years of working experience in Information Communication and Technology (ICT) industry. The highlights of his career are more than nine (9) years of holding various senior management and/or C-Level and had six (6) years of international ICT consultancy exposure in various countries (APAC and Australia), specially on Big Data, Data Engineering, Machine Learning and AI arena.

He is also Certified Trainer for Microsoft & Cloudera.





01

AI Foundations

02

Machine Learning Core

03

Generative AI & Agents

04

Natural Language

05

Computer Vision



CONTENTS

<https://arif.works/mbb/>



Introduction to AI concepts

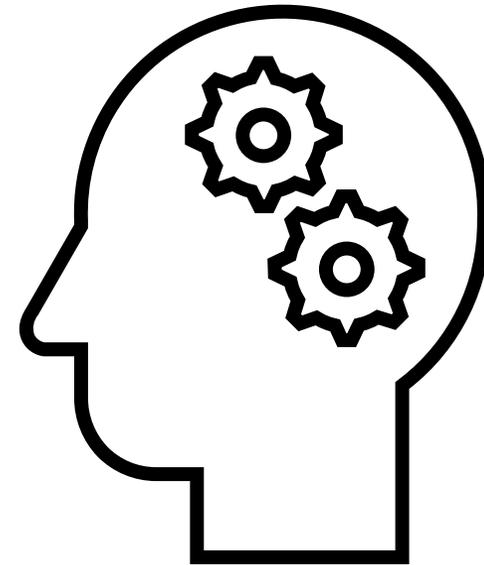
<https://aka.ms/mslearn-ai-concepts>



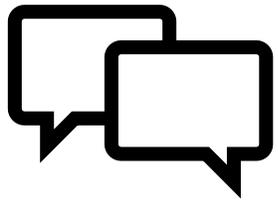
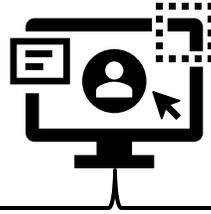
What is Artificial Intelligence?

Software that imitates human capabilities

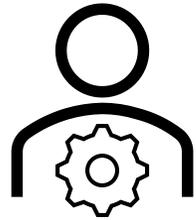
- Predicting outcomes and recognizing patterns based on historic data.
- Recognizing abnormal events, making decisions, and taking appropriate action.
- Interpreting visual input.
- Understanding language and engaging in conversations.
- Extracting information from sources to gain knowledge.



What is Artificial Intelligence?



Generative
AI



Agents and
automation



Natural
Language



Computer
Vision



Information
Extraction



$$f(x) = y$$

Machine Learning

Principles of responsible AI

	Challenge or Risk	Example
 Fairness	Bias can affect results.	A loan-approval model discriminates by gender due to bias in the data with which it was trained.
 Reliability & safety	Errors may cause harm.	An autonomous vehicle experiences a system failure and causes a collision.
 Privacy & security	Private data could be exposed.	A medical diagnostic agent is trained using sensitive patient data, which is stored insecurely.
 Inclusiveness	Solutions may not work for everyone.	A predictive app provides no audio output for visually impaired users.
 Transparency	Users must trust a complex system.	An AI-based financial tool makes investment recommendations – what are they based on?
 Accountability	Who's liable for AI-driven decisions?	An innocent person is convicted of a crime based on evidence from facial recognition – who's responsible?

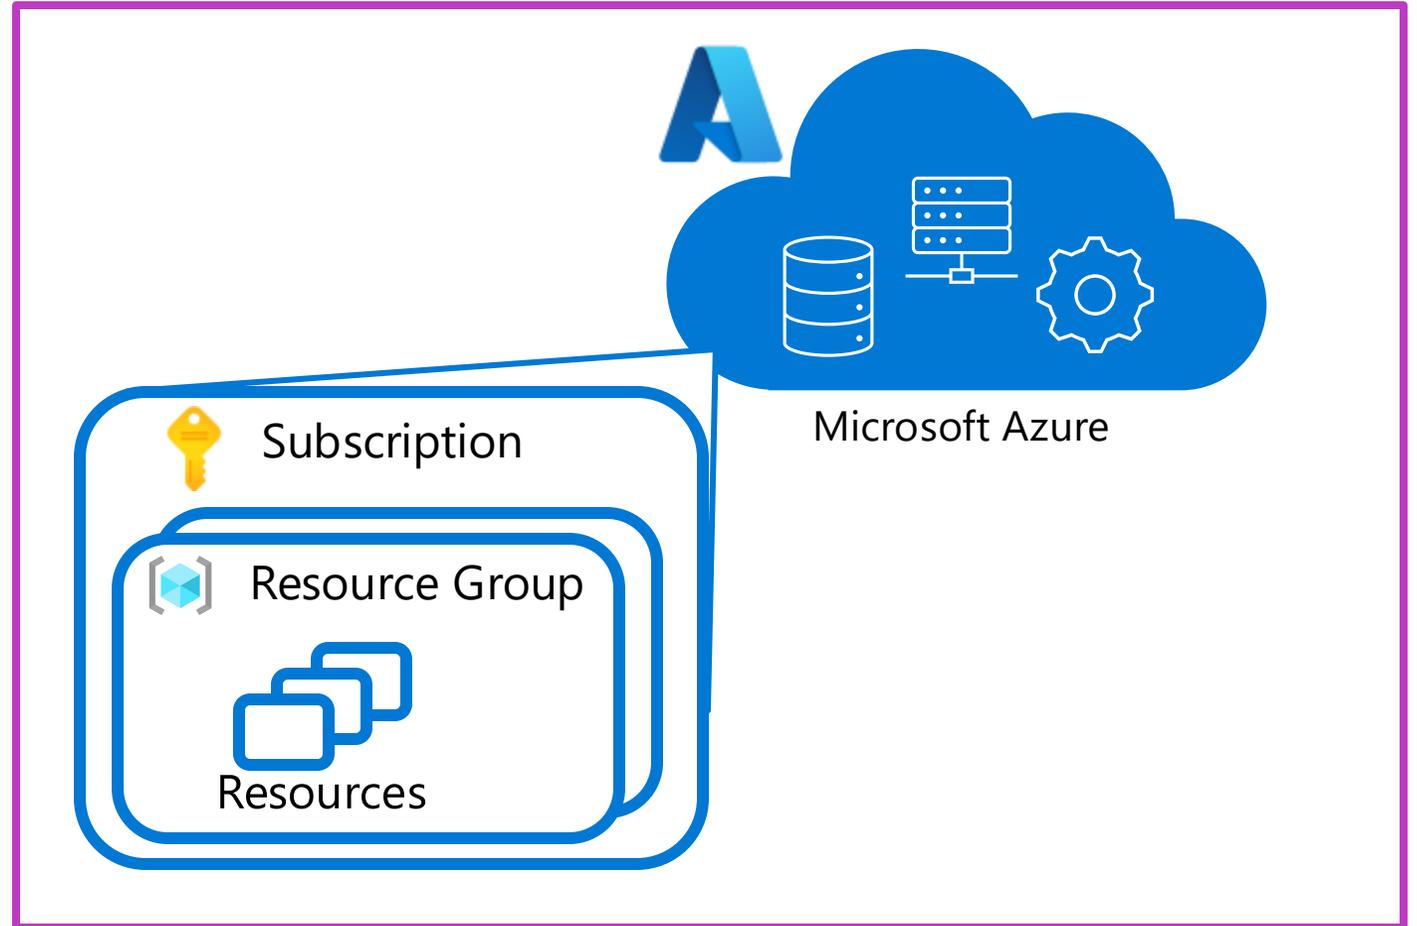
Get started with Azure AI services



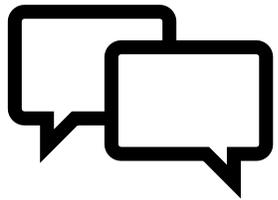
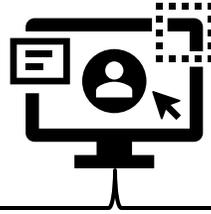
Microsoft Azure basics

Microsoft's Azure cloud platform provides scalable and reliable:

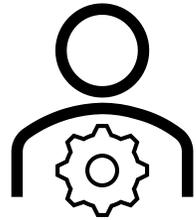
- Data storage
- Compute
- Services



Common Workloads or "Areas" of AI



Generative
AI



Agents and
automation



Natural
Language



Computer
Vision



Information
Extraction



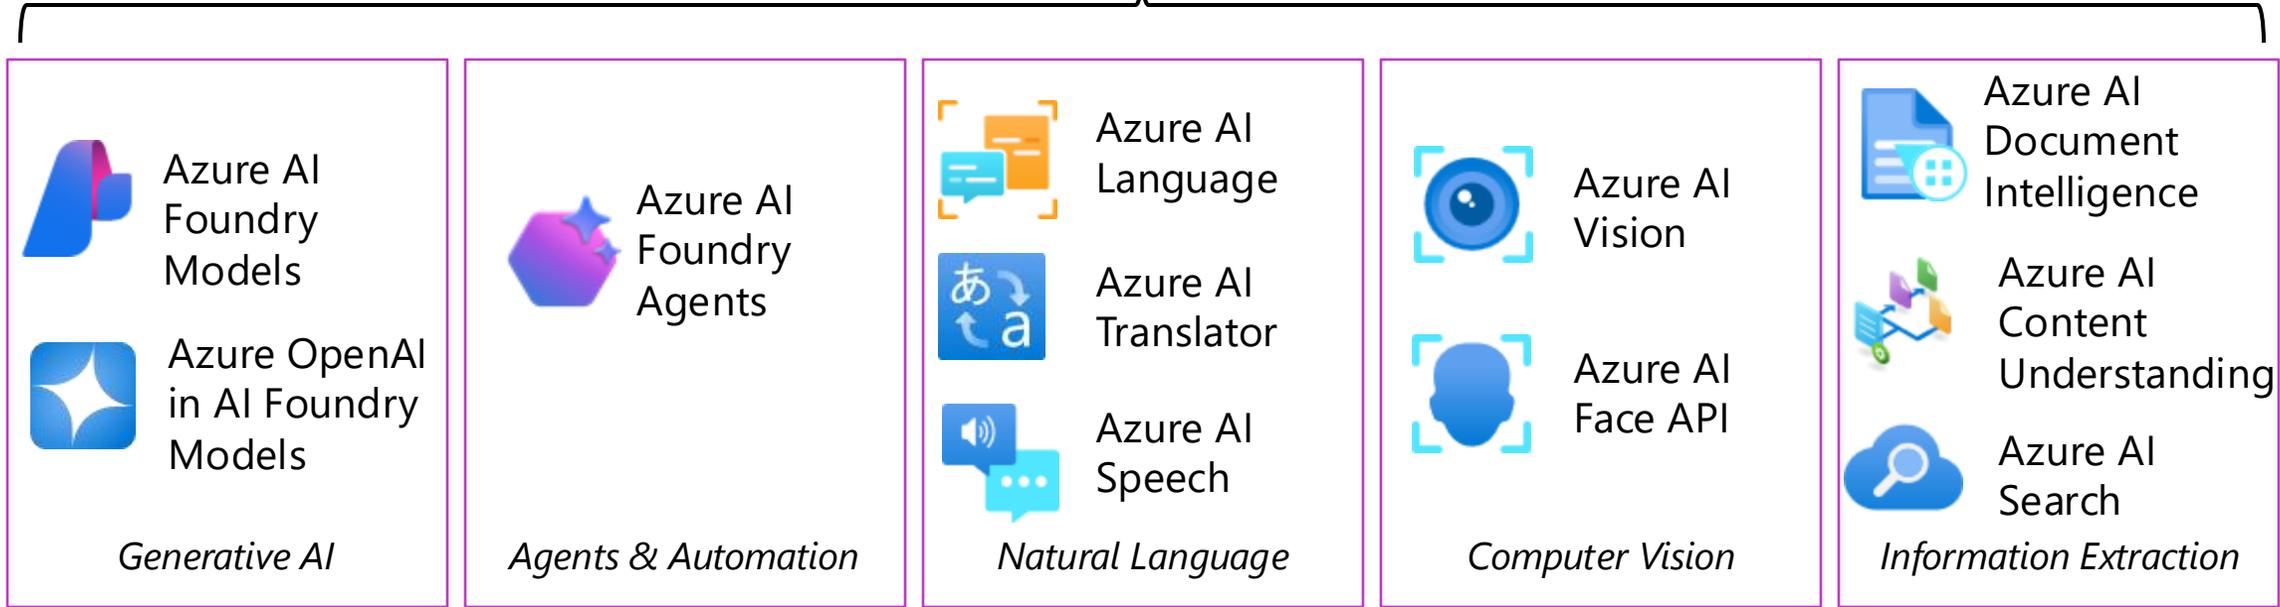
$$f(x) = y$$

Machine Learning

AI services in Azure



Azure AI Foundry

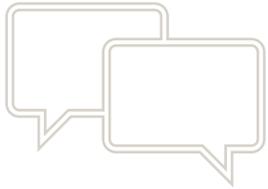
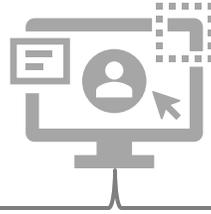


Azure Machine Learning

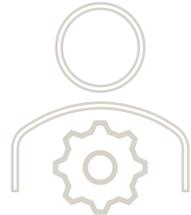
Introduction to AI in Azure: Machine Learning



Machine learning in context



Generative
AI



Agents and
automation



Natural
Language



Computer
Vision



Information
Extraction



$$f(x) = y$$

Machine Learning

Agenda



- Introduction to machine learning concepts
- Get started with machine learning in Azure

Introduction to machine learning concepts

<https://aka.ms/mslearn-ml-concepts>

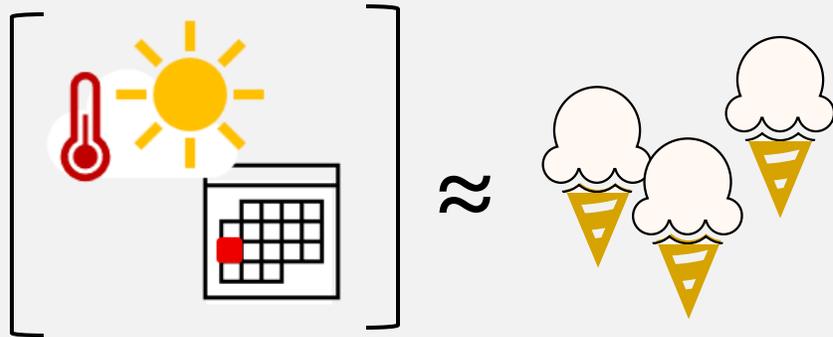


What is machine learning?

Predictive models that encapsulate relationships between:

- **Features** (known characteristics of something)
- **Label** (the thing we want to predict)

Example: Predicting ice cream sales for a given day



Features

- Day of week
- Month
- Temperature
- ...

Label

- Number of ice creams sold

Generalized function

$$f(\mathbf{x}) = \hat{y}$$

$\underbrace{[\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3]}$

(typically, an array of multiple features)

Types of machine learning

Machine Learning

Supervised machine learning

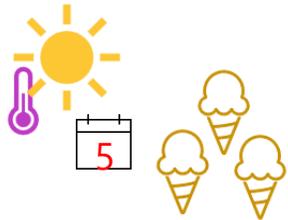
Training data includes known labels

Unsupervised machine learning

Training data is unlabeled

Regression

Label is a numeric value



Predict the number of ice creams sold based on day, season, and weather

Classification

Label is a categorization (or *class*)

Binary classification

Label *is* or *is not* a class



Predict whether a patient is at-risk for diabetes based on clinical data

Multiclass classification

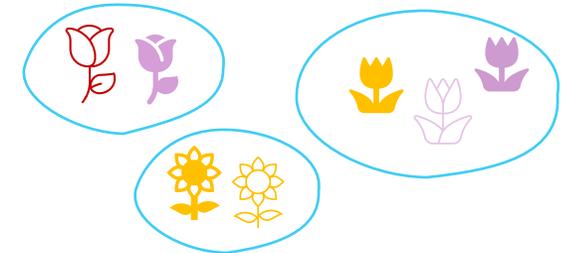
Label is one of multiple classes



Predict the species of a penguin based on its measurements

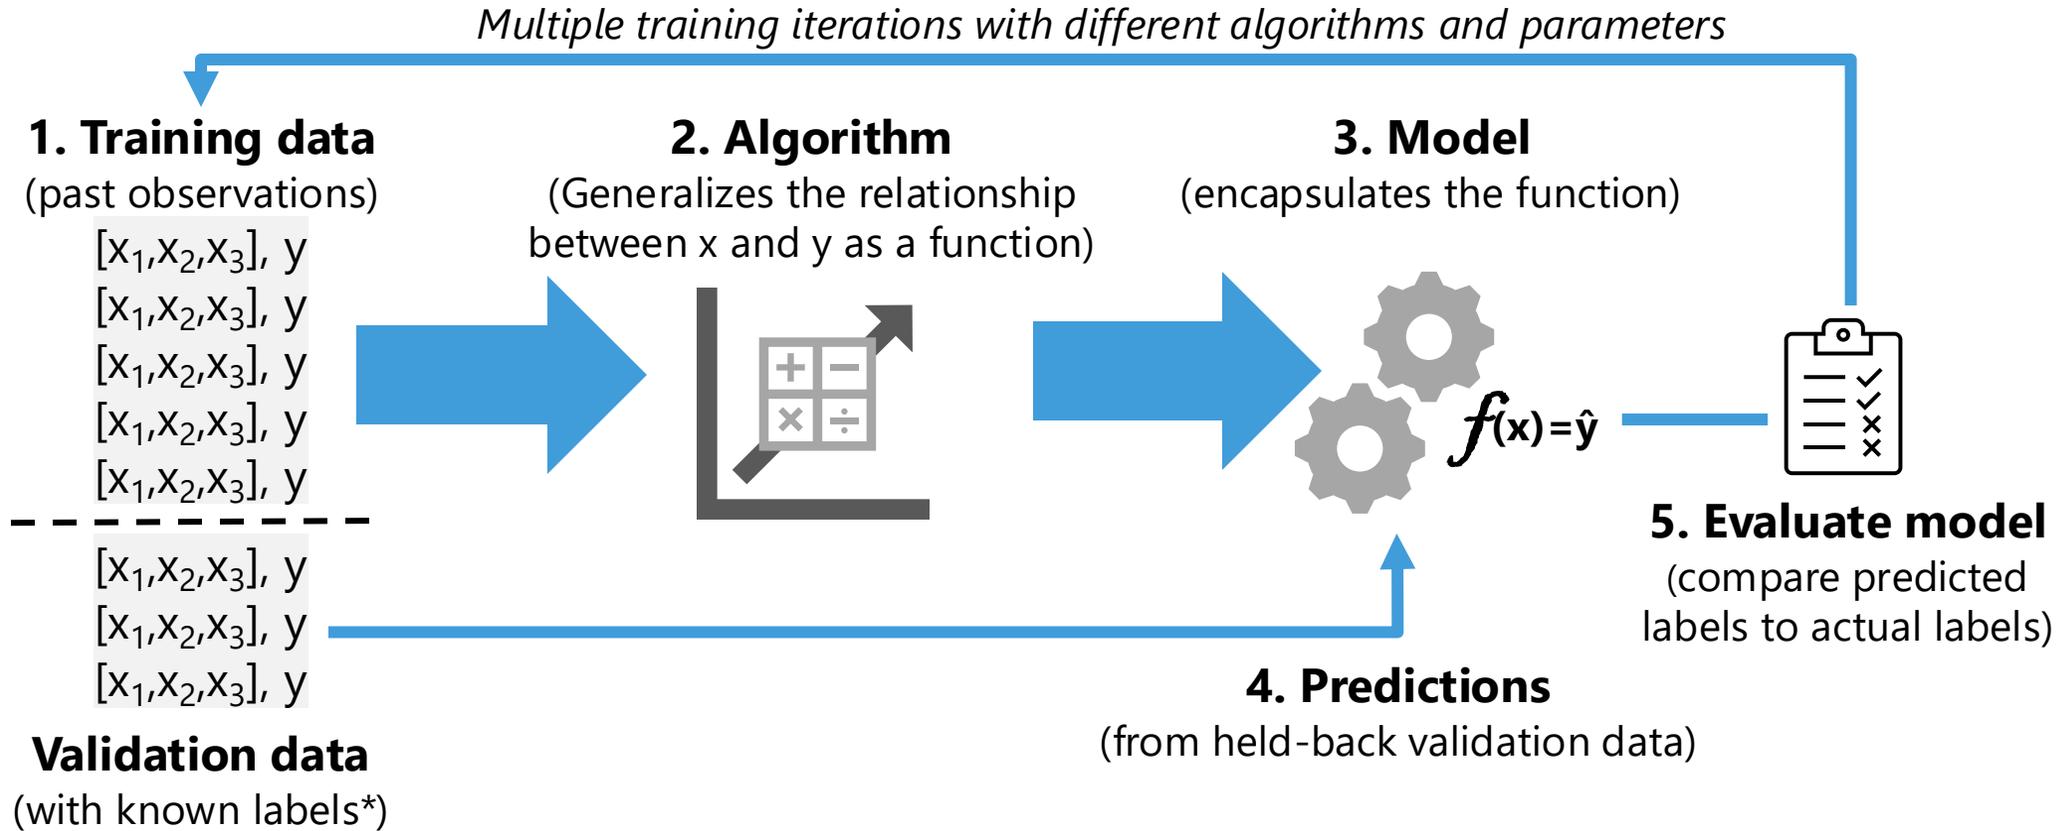
Clustering

Similar items are grouped together



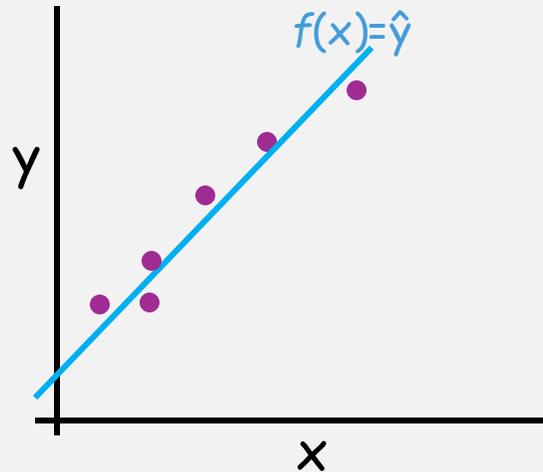
Separate plants into groups based on common characteristics

Model training and evaluation



Model algorithms and evaluation metrics

Regression

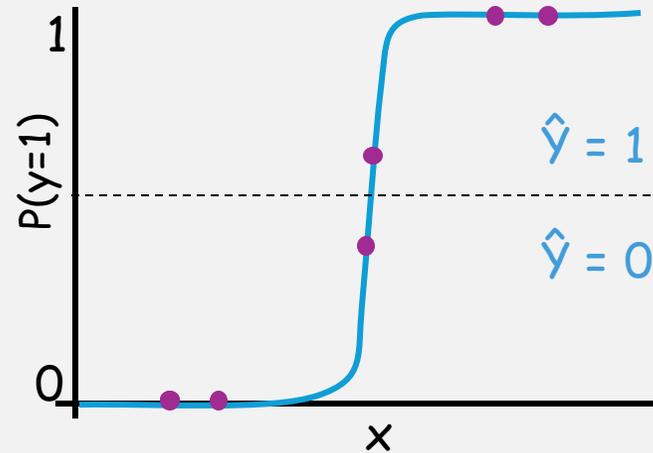


Example, *Linear Regression*

Evaluation metrics:

- *Mean absolute error (MAE)*
- *Mean squared error (MSE)*
- *Root mean squared error (RMSE)*
- *Coefficient of determination (R^2)*

Classification

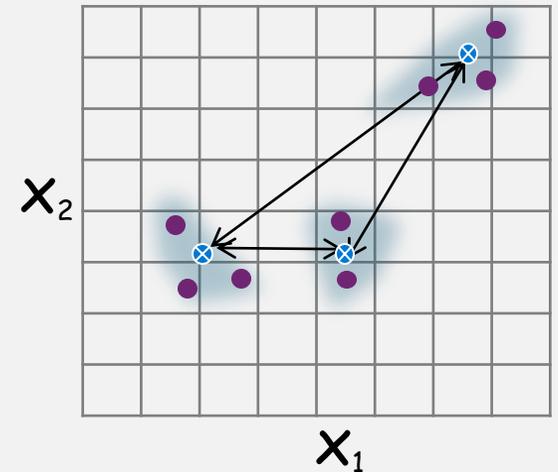


Example, *Logistic Regression*

Evaluation metrics (per class and overall):

- *Accuracy*
- *Recall*
- *Precision*
- *F1 Score*

Clustering



Example, *K-Means*

Evaluation metrics:

- *Cluster centroid distance*
- *Mean distance to centroid within cluster*
- *Silhouette score*

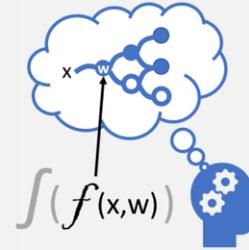
Deep learning

Human neural network

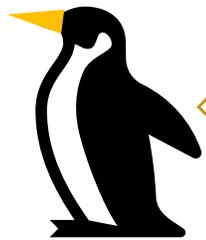


- Neurons fire in response to electrochemical stimuli
- When fired, the signal is passed to connected neurons

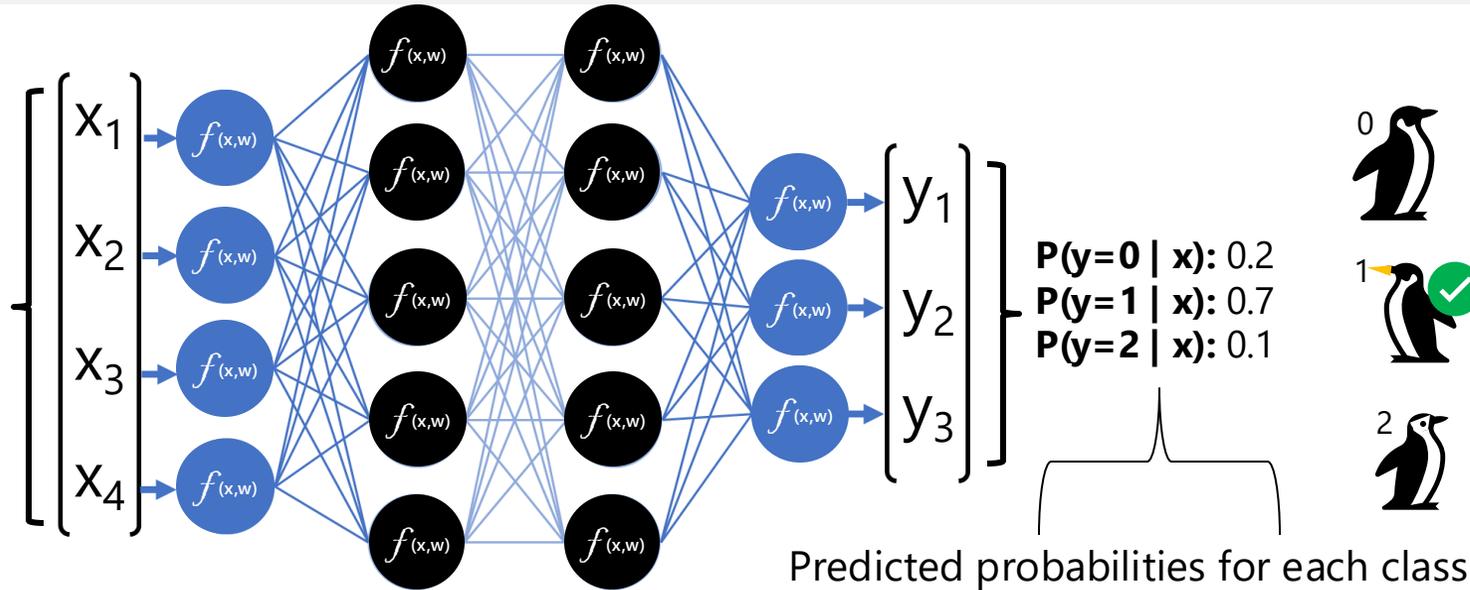
Artificial neural network



- Each neuron is a function that operates on an *input* value (x) and a *weight* (w)
- The function is wrapped in an *activation function* that determines whether to pass the output on to the next layer



$x_1: 37.3$
 $x_2: 16.8$
 $x_3: 19.2$
 $x_4: 30.0$



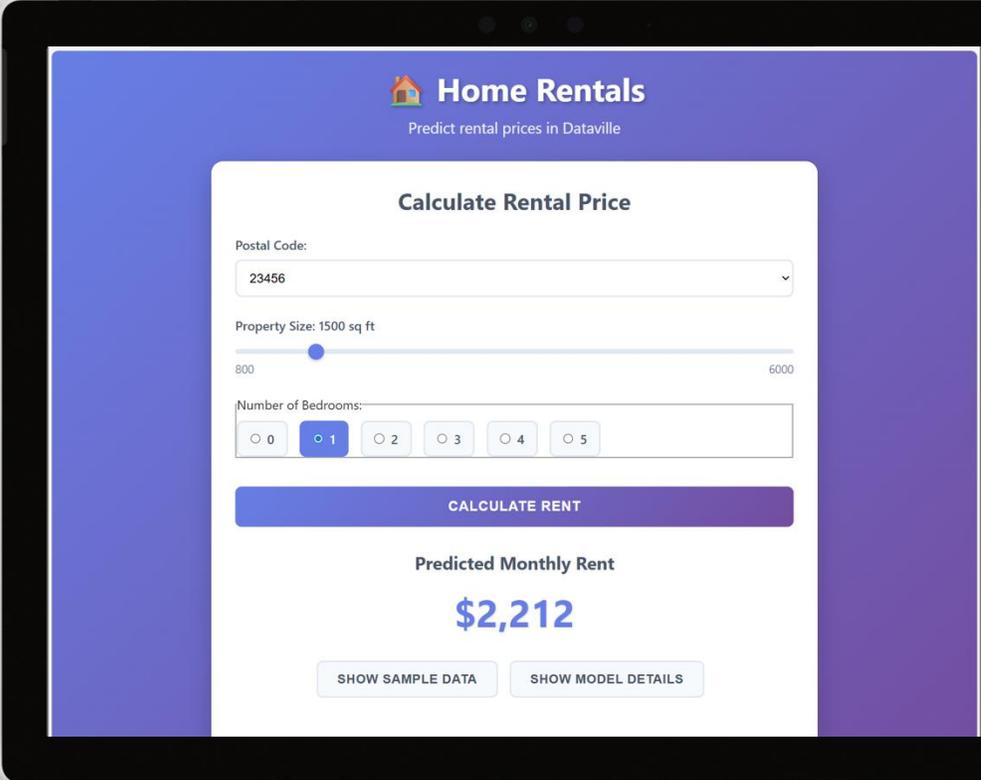
- The difference between predicted probabilities and actual label (0 or 1) is the **loss**.
- During training, weights are iteratively adjusted to minimize the loss

Explore machine learning scenarios

In this exercise, you'll explore applications that predict unknown values.

Start the exercise at:

<https://go.microsoft.com/fwlink/?linkid=2334118>



The screenshot shows a web application titled "Home Rentals" with the subtitle "Predict rental prices in Dataville". The main heading is "Calculate Rental Price". The form includes a "Postal Code" dropdown menu with "23456" selected, a "Property Size" slider set to "1500 sq ft" (ranging from 800 to 6000), and a "Number of Bedrooms" selection with radio buttons for 0, 1, 2, 3, 4, and 5. A prominent purple button labeled "CALCULATE RENT" is positioned below the input fields. The result, "Predicted Monthly Rent", is displayed as "\$2,212" in large blue text. At the bottom, there are two buttons: "SHOW SAMPLE DATA" and "SHOW MODEL DETAILS".

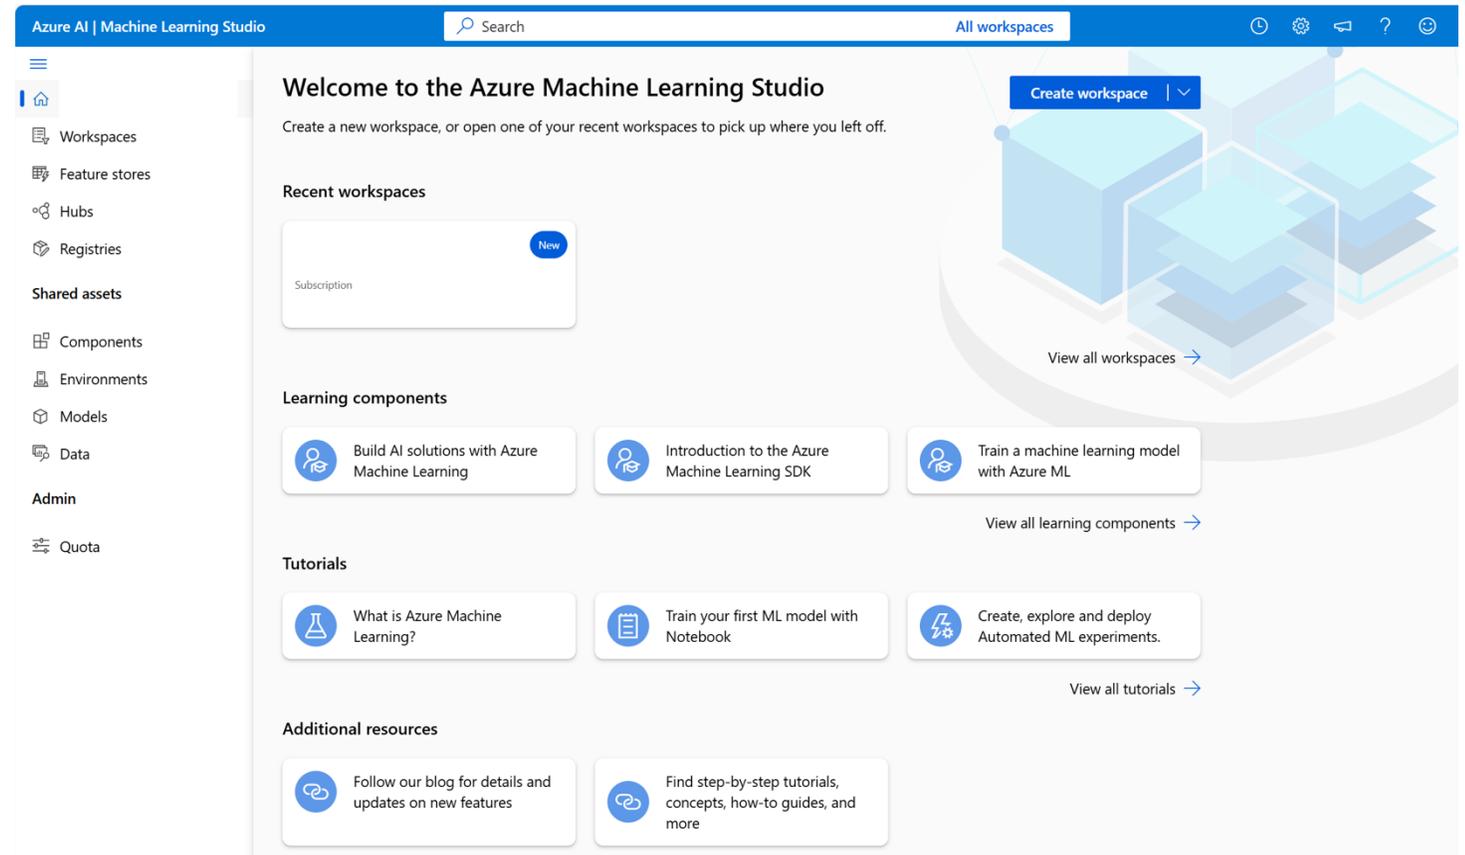
Get started with machine learning in Azure

<https://aka.ms/mslearn-azure-ml-intro>



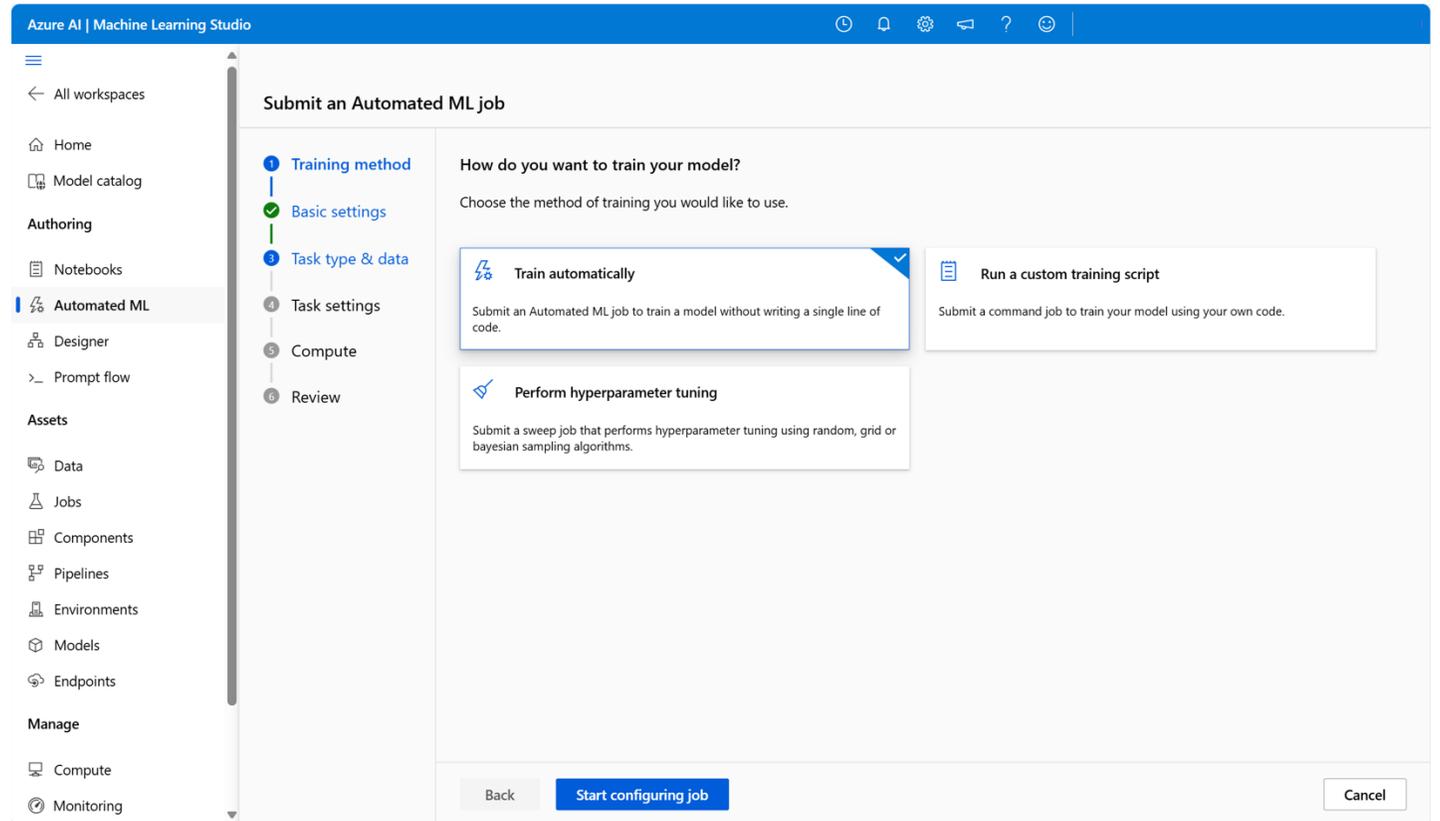
What is Azure Machine Learning?

- Azure Machine Learning is a cloud-based platform for machine learning
- Azure Machine Learning Studio is a user interface for accessing Azure Machine Learning capabilities
- Machine learning models trained with Azure Machine Learning can be published as services and consumed by applications



What is Automated Machine Learning?

- A step-by-step wizard that helps you run machine learning training jobs
- Supports multiple machine learning types, including regression, time-series forecasting, classification, computer vision, and natural language processing tasks
- Connect to your data, define the training job and target metrics, and deploy the best resulting model

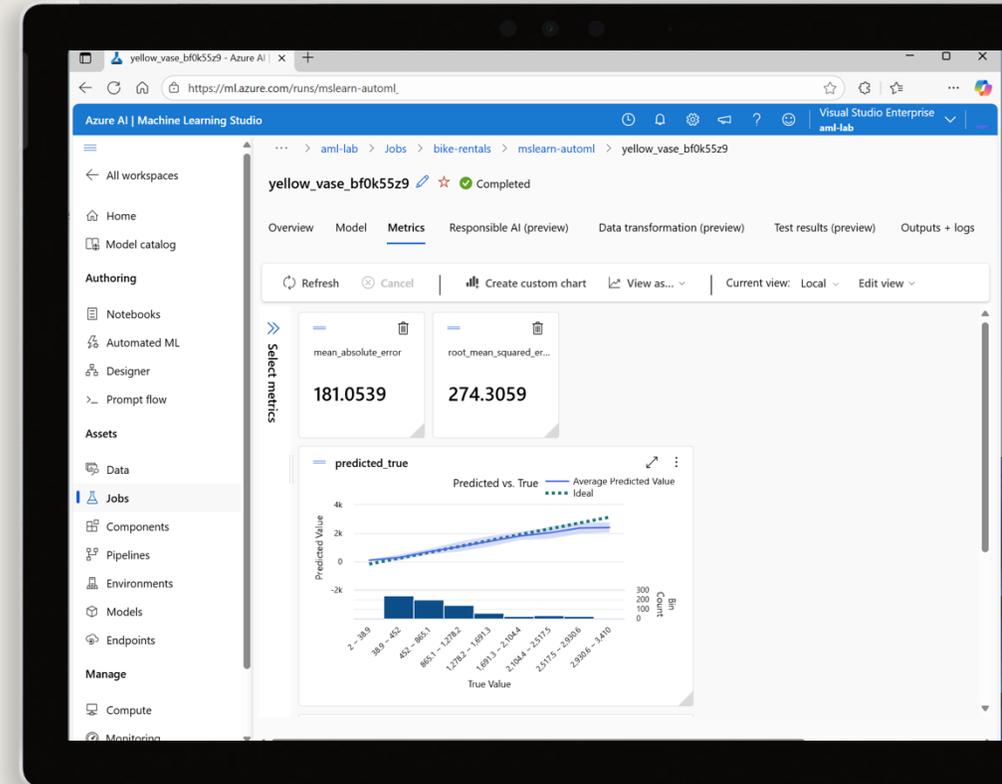


Explore Automated Machine Learning in Azure Machine Learning

In this exercise, you'll use the Automated Machine Learning capability of Azure Machine Learning to train a model.

Start the exercise at:

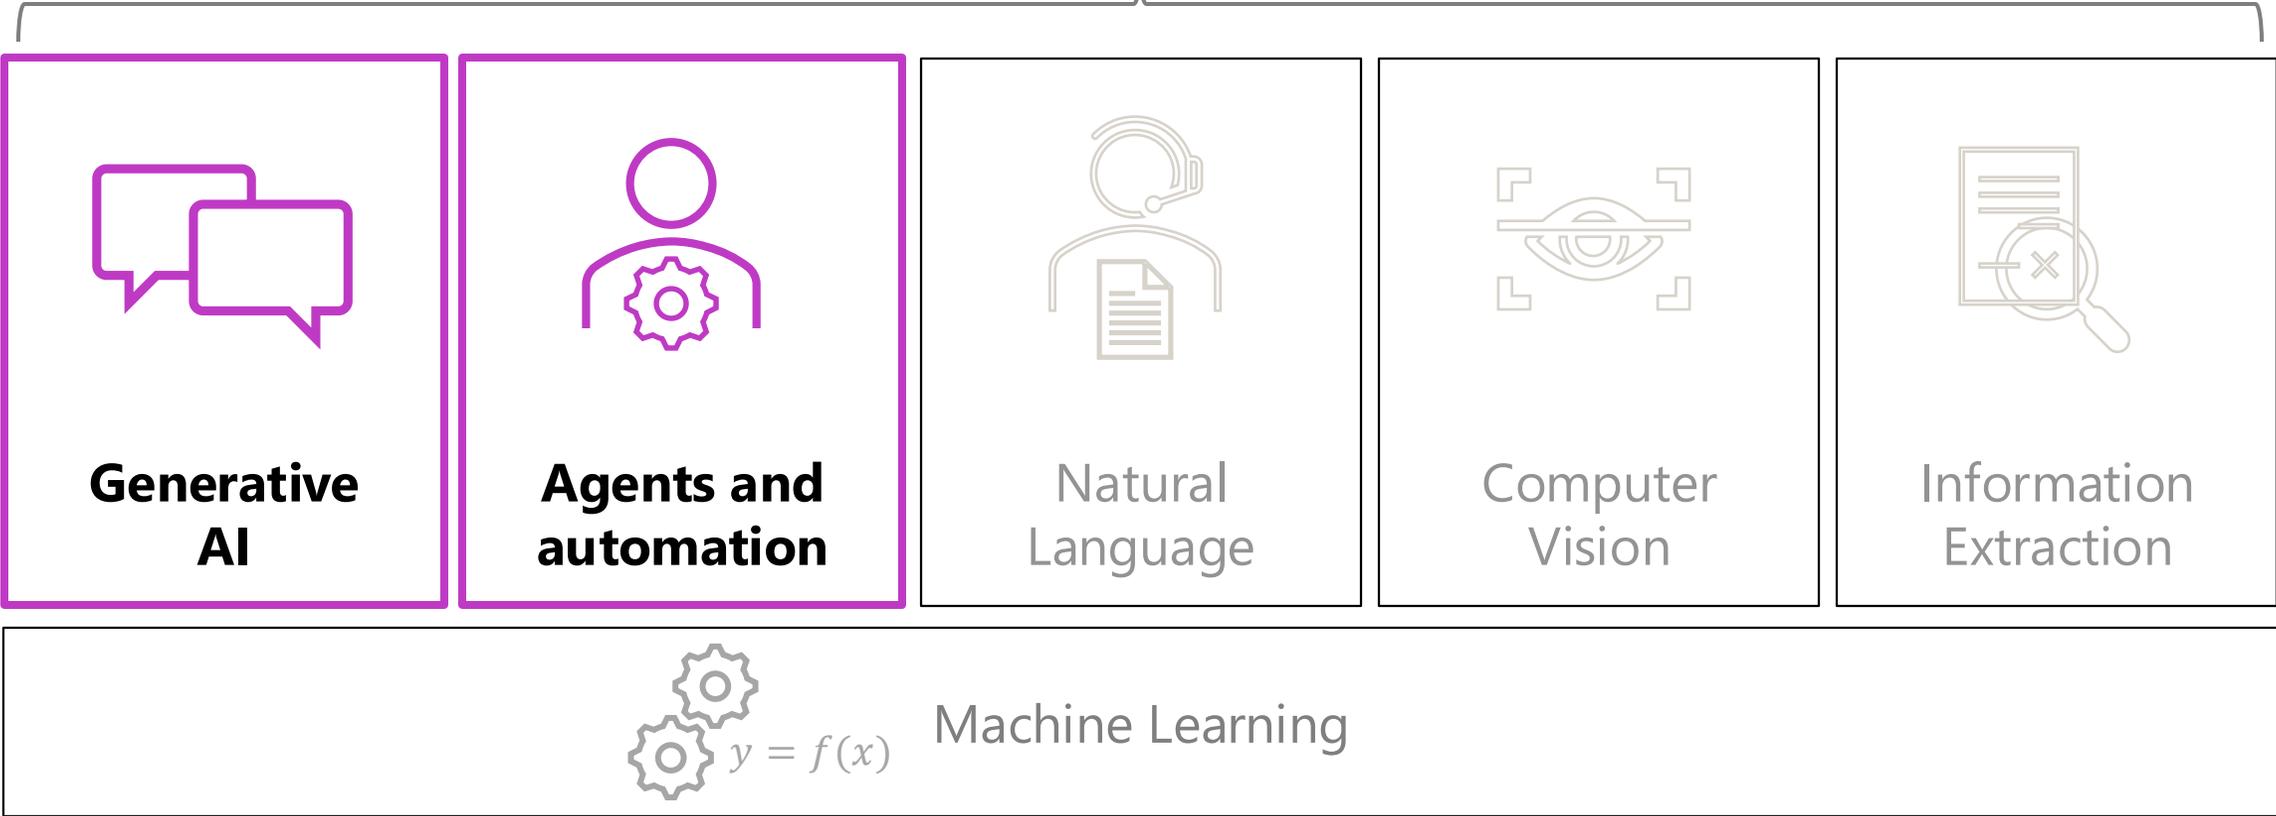
<https://go.microsoft.com/fwlink/?linkid=2250144>





Introduction to AI in Azure: Generative AI

Generative AI and agents in context



Agenda



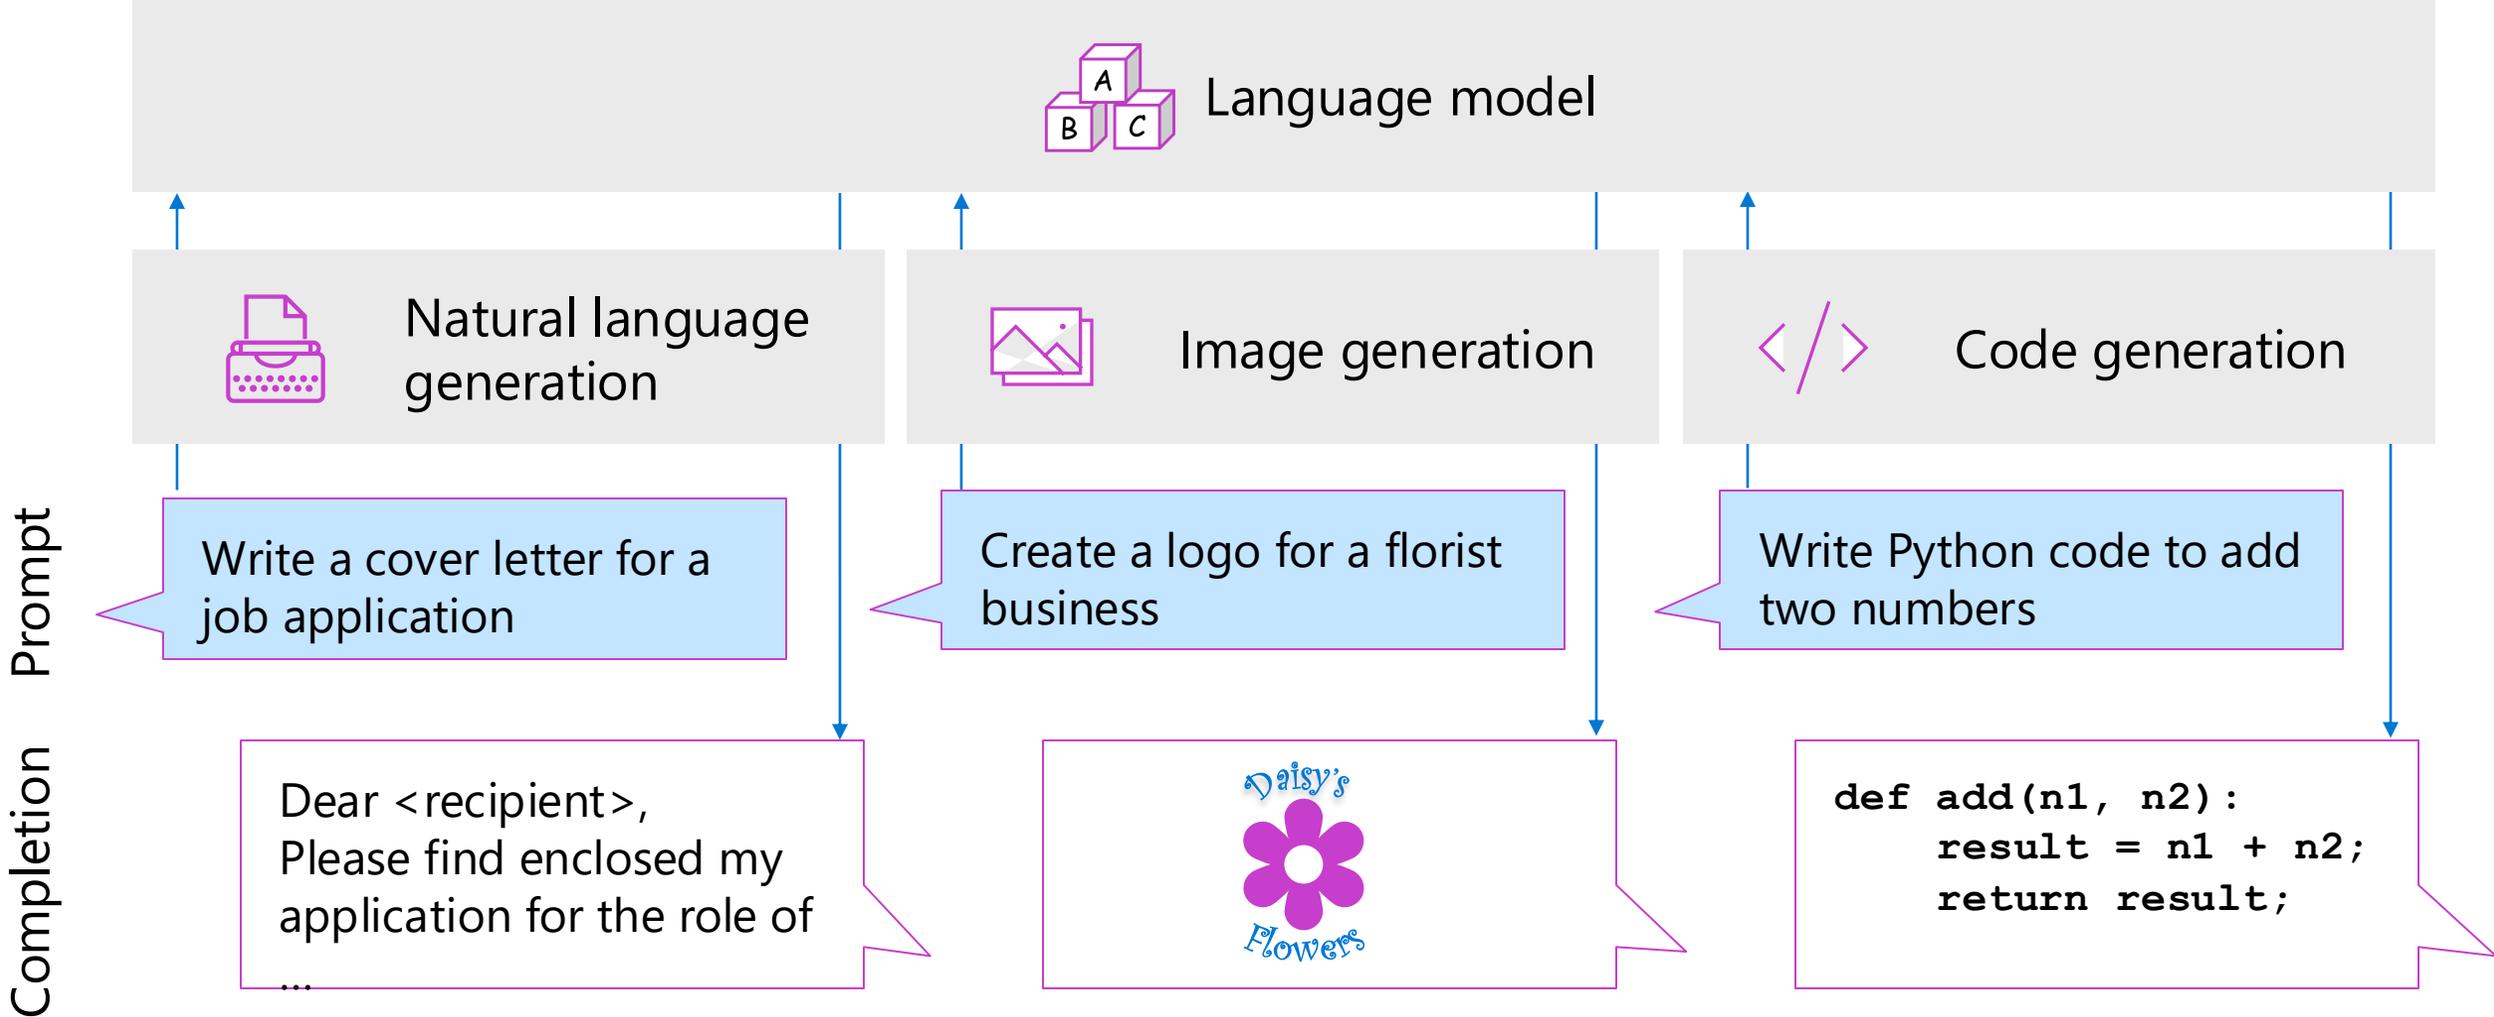
- Introduction to generative AI concepts
- Get started with generative AI in Azure

Introduction to generative AI concepts

<https://aka.ms/mslearn-intro-gen-ai>

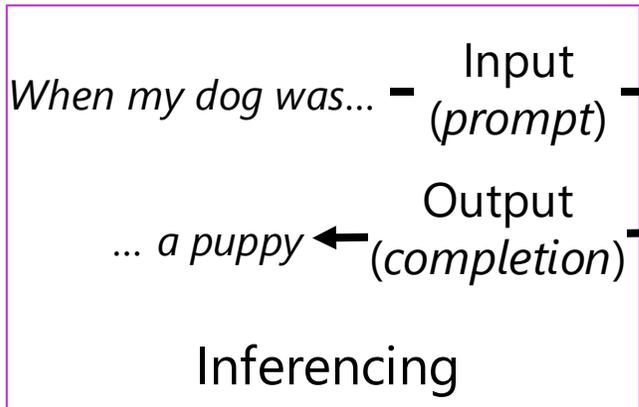
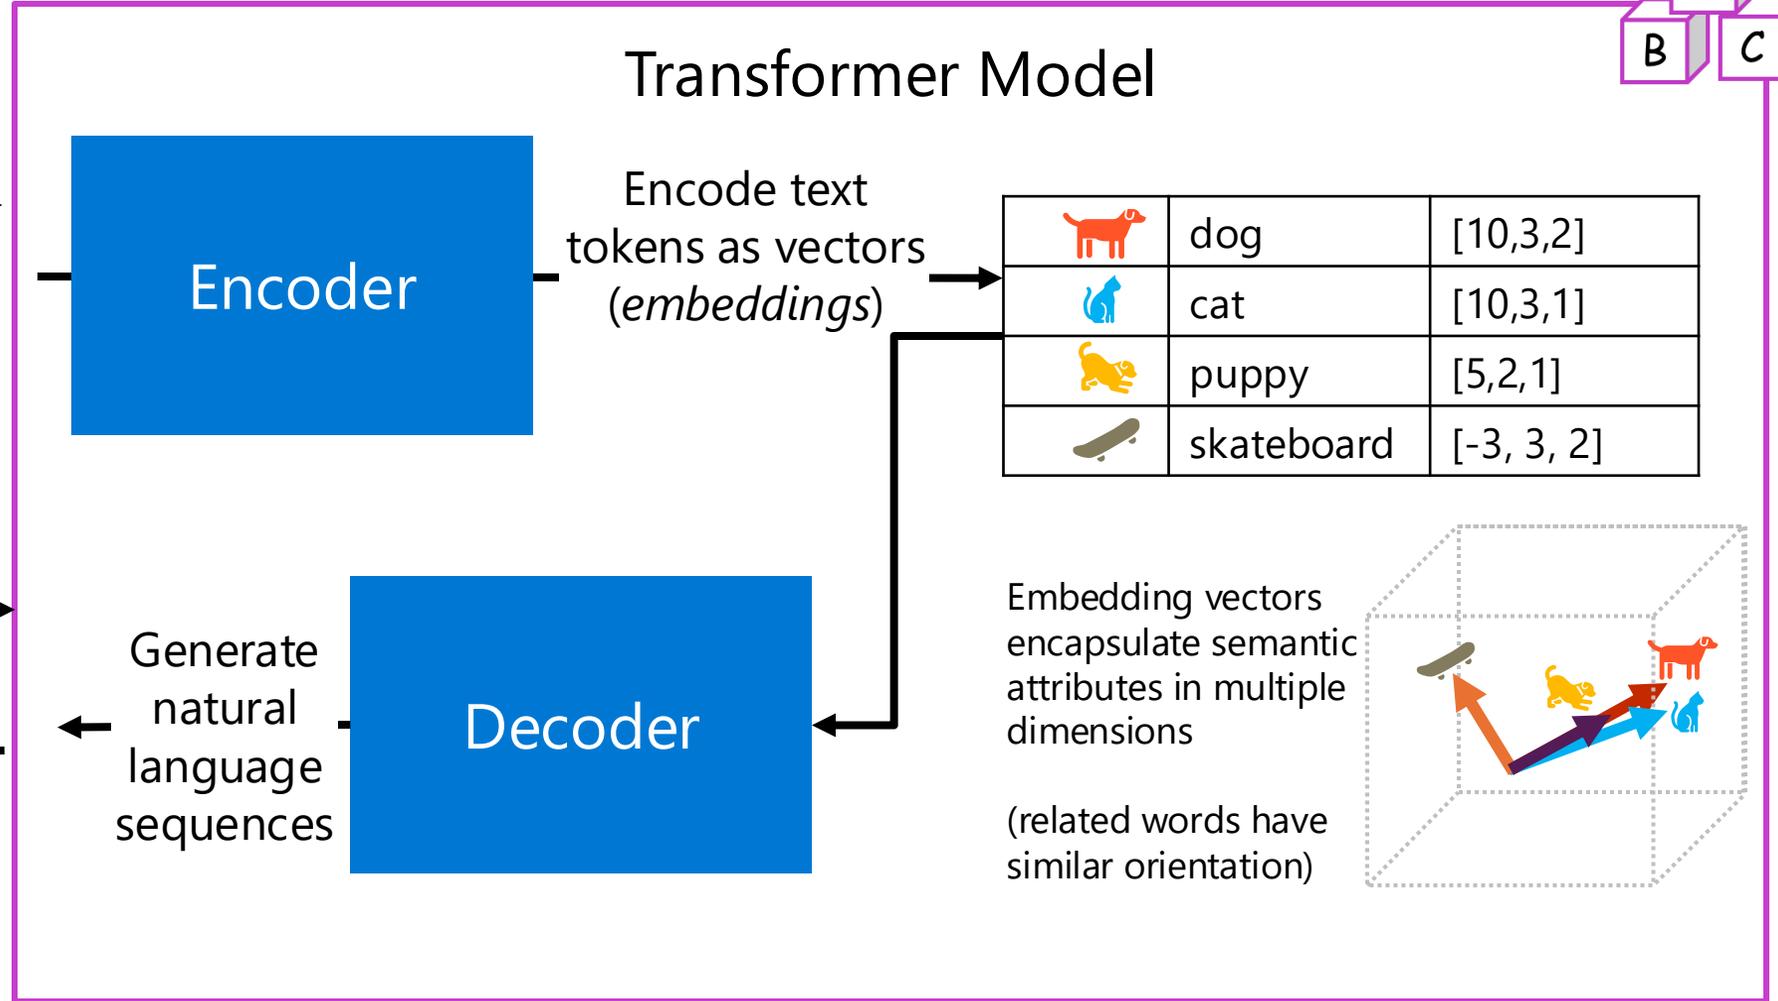
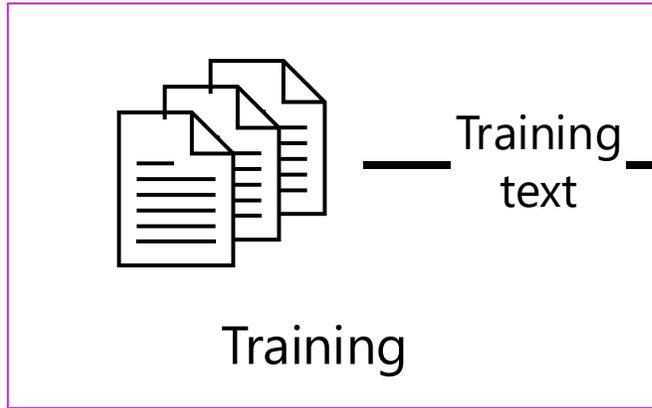
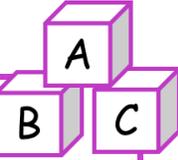


What is Generative AI?



Language models

A (very) high-level overview



Language models – tokenization

Step one: tokenization

The first step in training a transformer model is to decompose the training text into *tokens*.

Example sentence: *I heard a dog bark loudly at a cat.*

"I"=1

"heard"=2

"a"=3

"dog"=4

"bark"=5

"loudly"=6

"at"=7

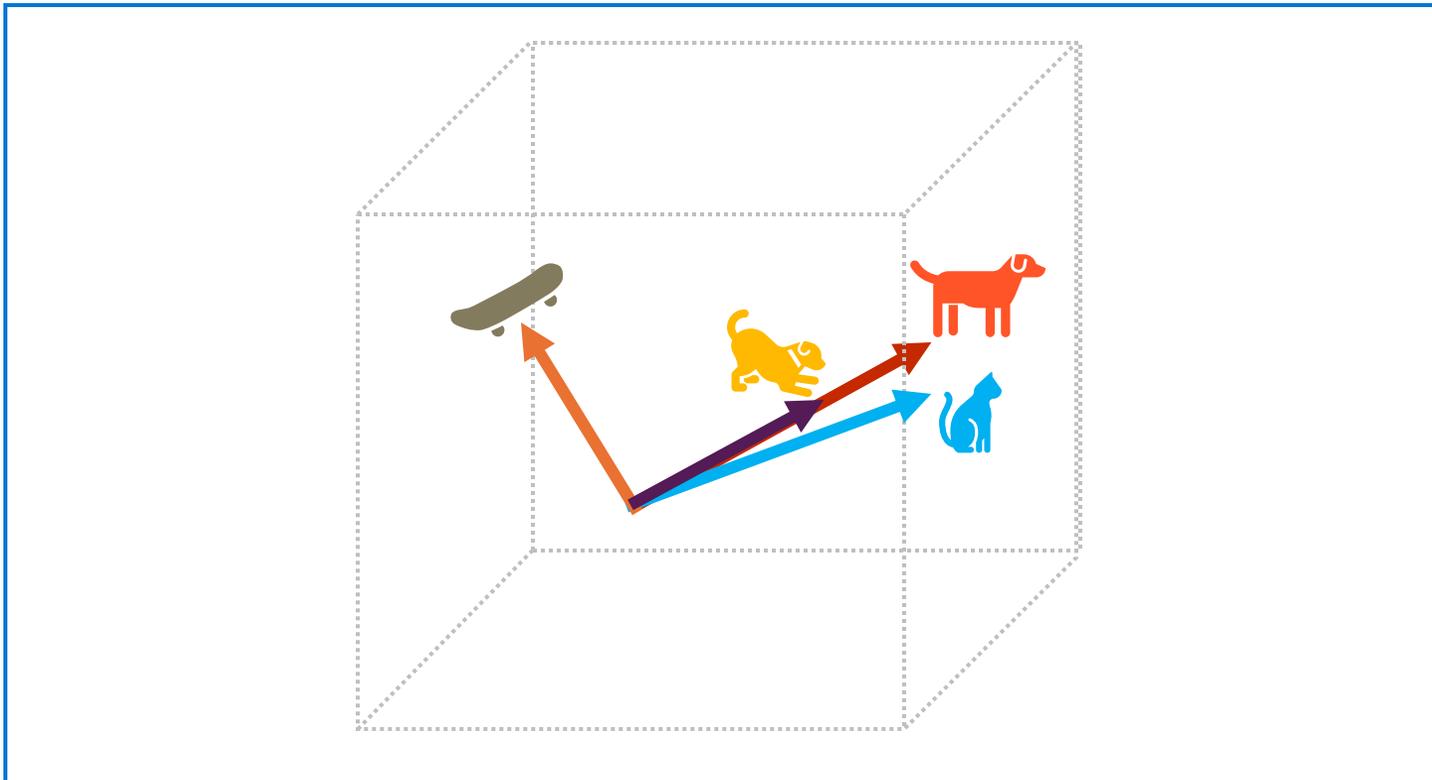
"cat"=8

- The sentence is now represented with the tokens: $[1\ 2\ 3\ 4\ 5\ 6\ 7\ 3\ 8]$.
- Note "a" is tokenized as 3 only once.
- Similarly, the sentence "I heard a cat" could be represented with the tokens $[1\ 2\ 3\ 8]$.

Language models – embeddings

Step two: embeddings

The semantic relationship between tokens is encoded in vectors, known as embeddings.



	Token	Word	Embedding
	4	dog	[10,3,2]
	8	cat	[10,3,1]
	9	puppy	[5,2,1]
	10	skateboard	[-3, 3, 2]

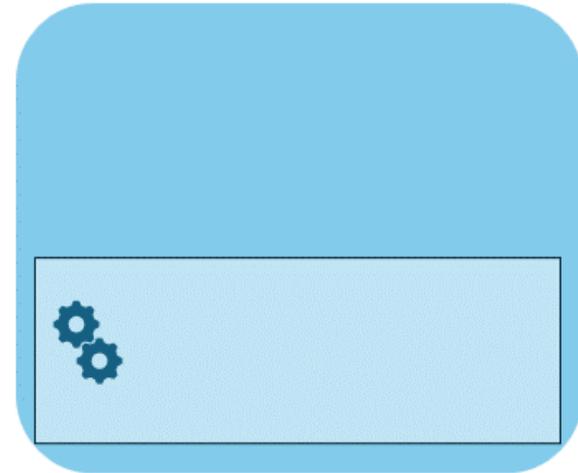
Language models – attention

Step three: attention

Capture the strength of relationships between tokens using the attention technique.

Example:

- Goal: predict token after "**dog.**"
- Represent "**I heard a dog**" as vectors.
- Assign "**heard**" and "**dog**" more weight.
- Several possible tokens can come after dog.
- The most probable token is added to the sequence, in this case "**bark.**"

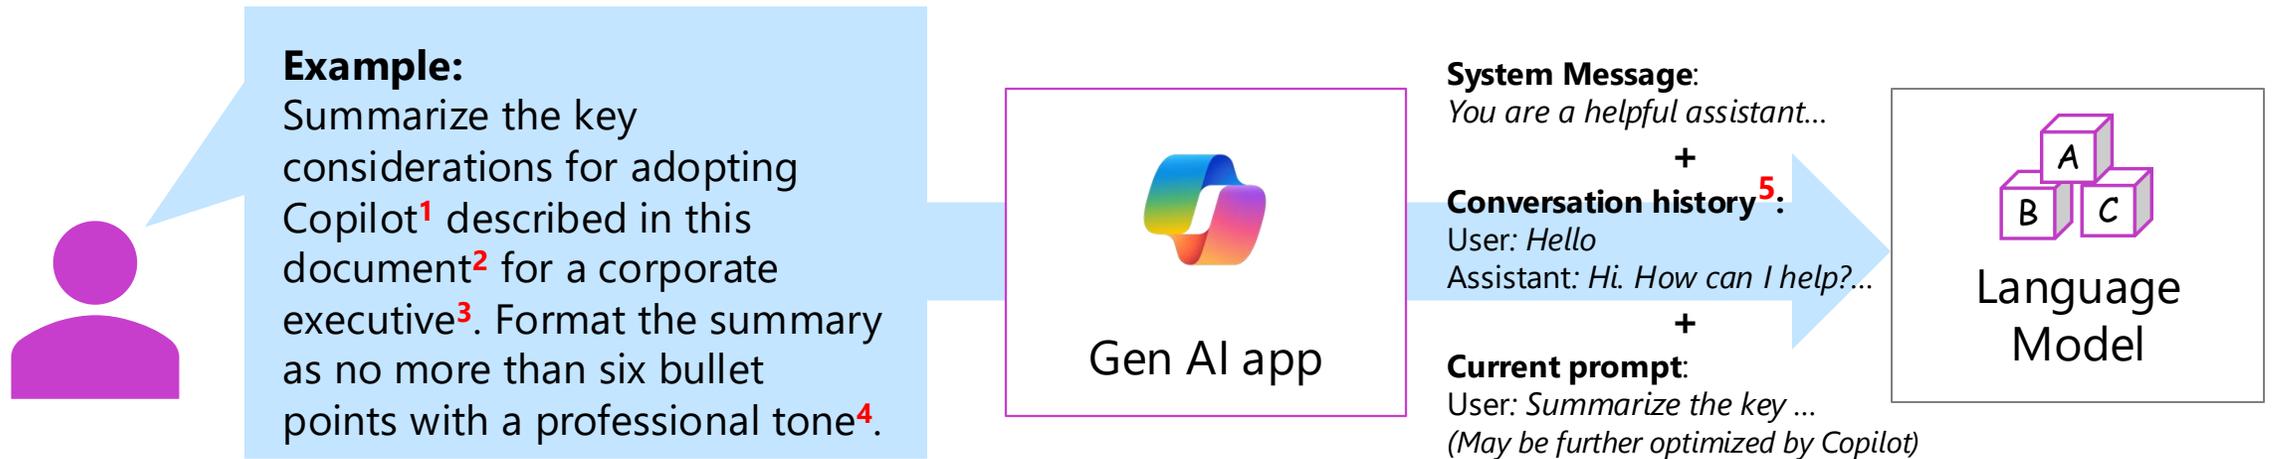


Language models

Large and small language models

Large language models (LLMs)	Small language models (SLMs)
Trained with large volumes of general text data	Trained with focused text data
Many billions of parameters	Fewer parameters
Comprehensive language generation capabilities in multiple contexts	Focused language generation capabilities in specialized contexts
Large size can impact performance and portability	Fast and portable
Time-consuming (and expensive) to fine-tune with your own training data	Faster (and less expensive) to fine-tune with your own training data
Examples include: Azure OpenAI in Azure AI Foundry GPT 5 Mistral 7B Meta Llama 3	Examples include: Microsoft Phi 4 Microsoft Orca 2 Hugging Face GPT Neo

Considerations for prompts



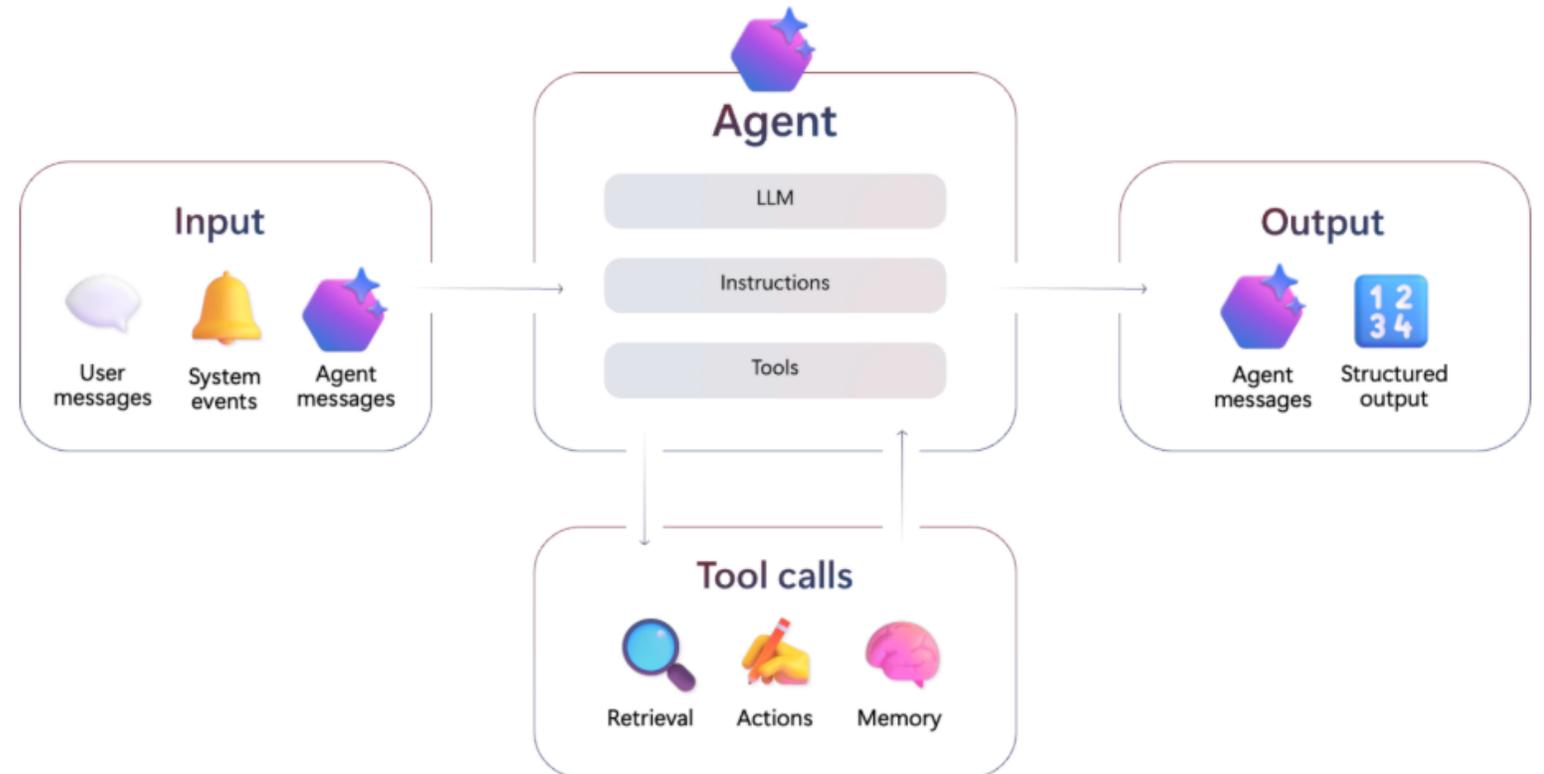
1. Start with a specific goal for what you want the output to be
2. Provide a source to *ground* the response in a specific scope of information
3. Add context to maximize response appropriateness and relevance
4. Set clear expectations for the response
5. Iterate based on previous prompts and responses to refine the result

What are agents?

Agents are generative AI applications that can respond to user input or assess situations *autonomously* and take appropriate *actions*.

Agents have three core components:

- **Model (LLM)**: Powers reasoning and language understanding
- **Instructions**: Define the agent's goals, behavior, and constraints
- **Tools**: Let the agent retrieve *knowledge* or take *action*

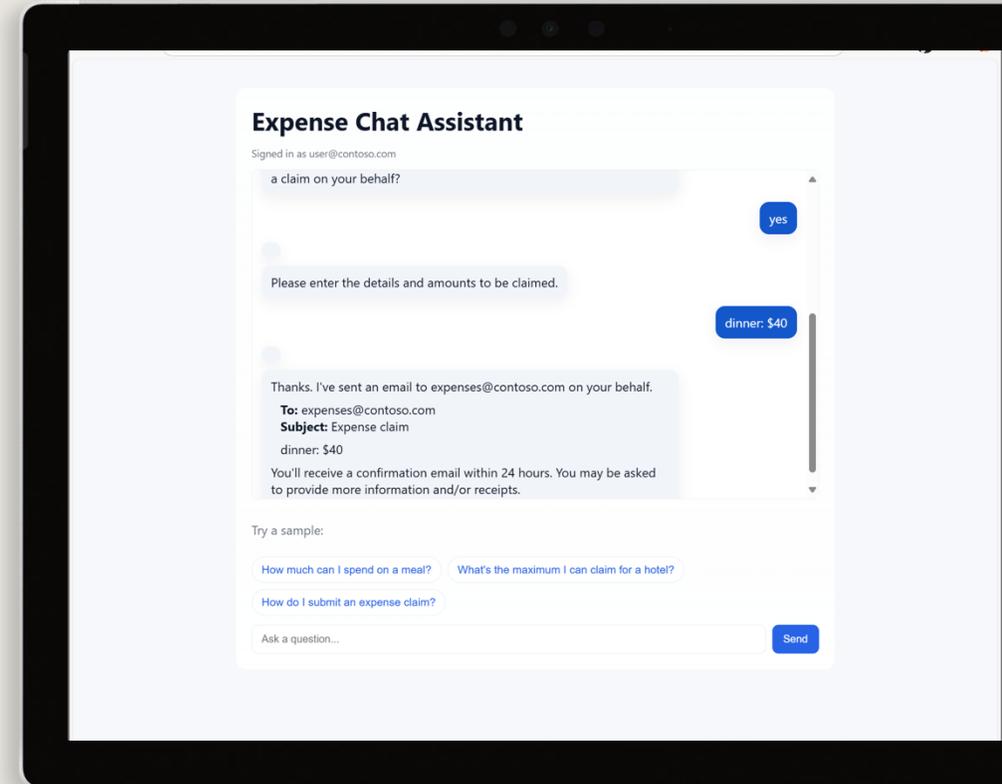


Explore generative AI agent scenarios

In this exercise, you'll explore generative AI apps and agents.

Start the exercise at:

<https://go.microsoft.com/fwlink/?linkid=2334224>



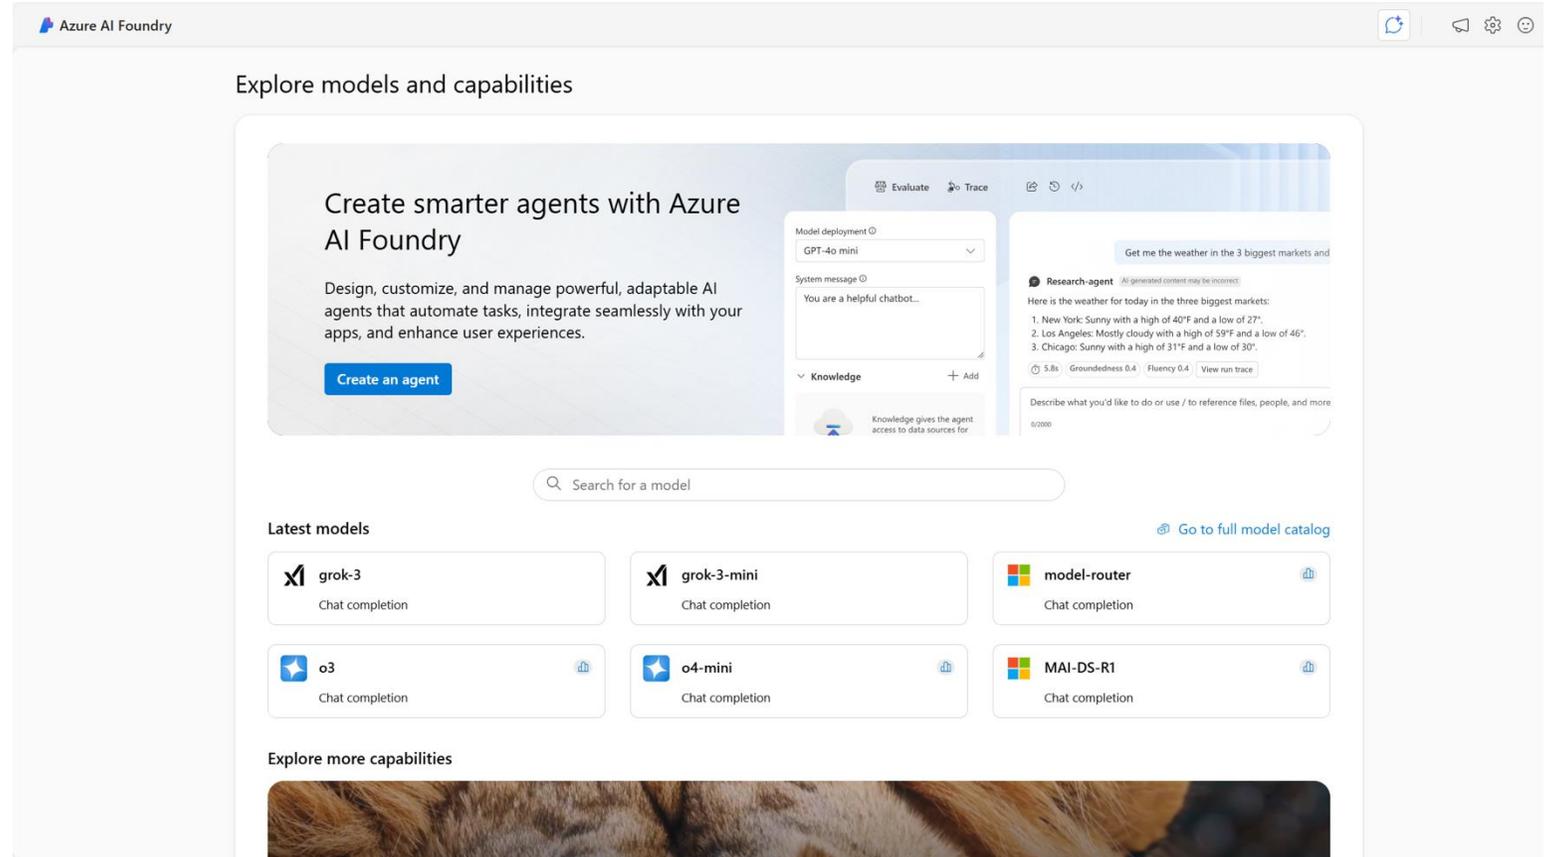
Get started with generative AI in Azure

<https://aka.ms/mslearn-get-started-gen-ai>



What is Azure AI Foundry?

- A cloud-based platform for AI development
- Provides a collaborative pro-code environment with enterprise-grade security for AI innovation
- Use to explore, build, test, and deploy AI solutions that are reliable, scalable, and secure



Azure AI Foundry projects



Azure AI Foundry resource



Azure AI Foundry Models



Azure AI Foundry Agents

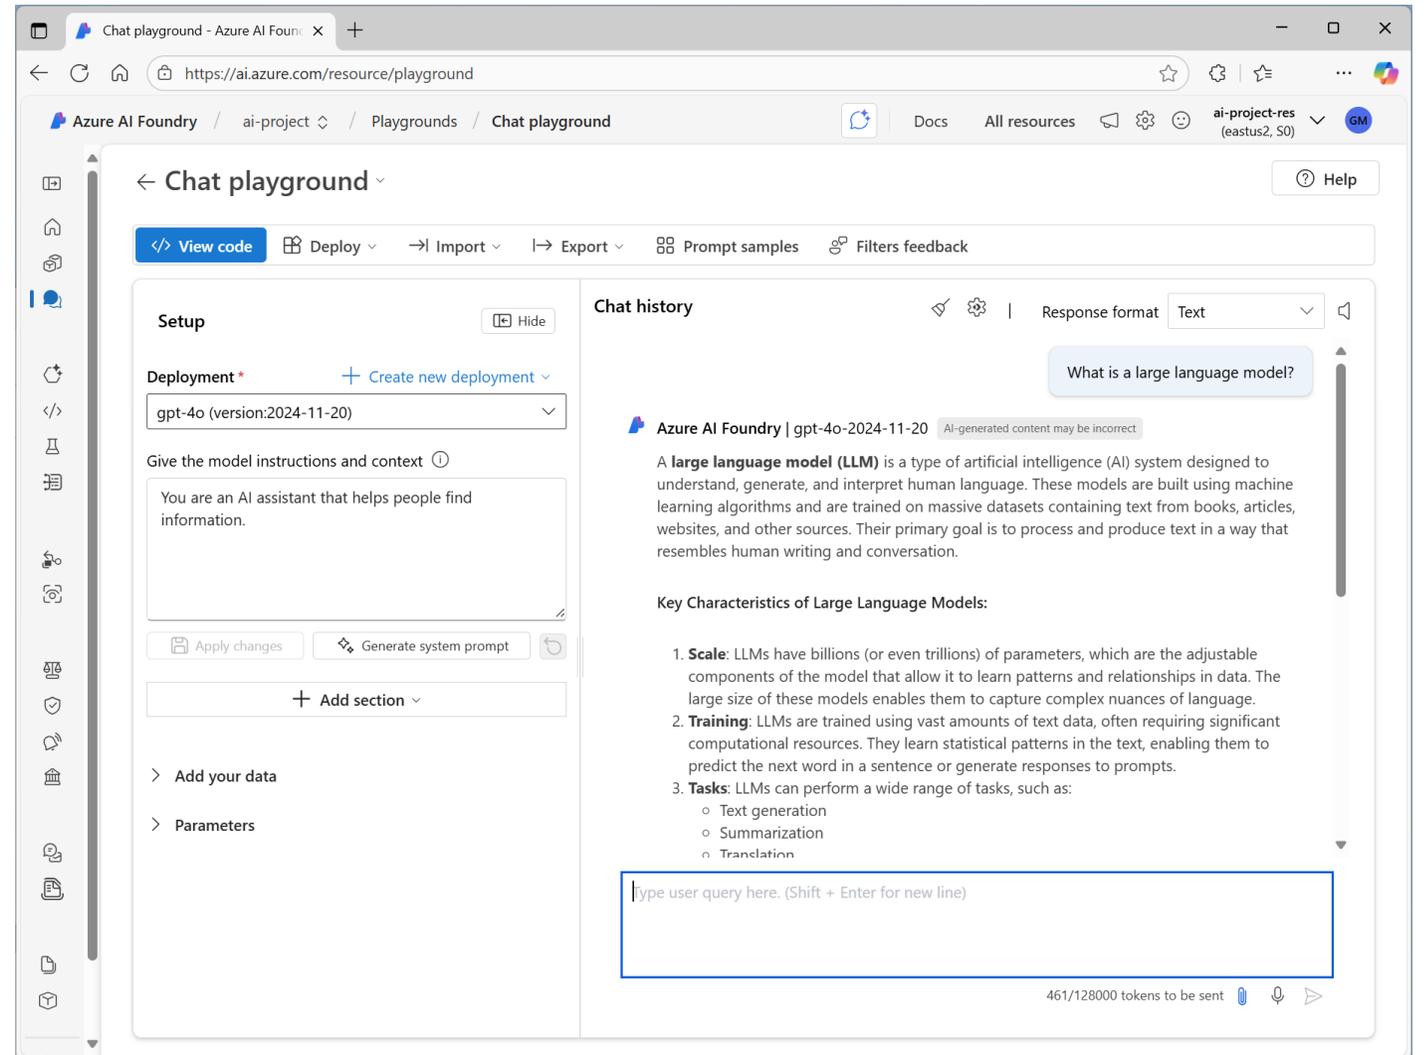


Azure AI services



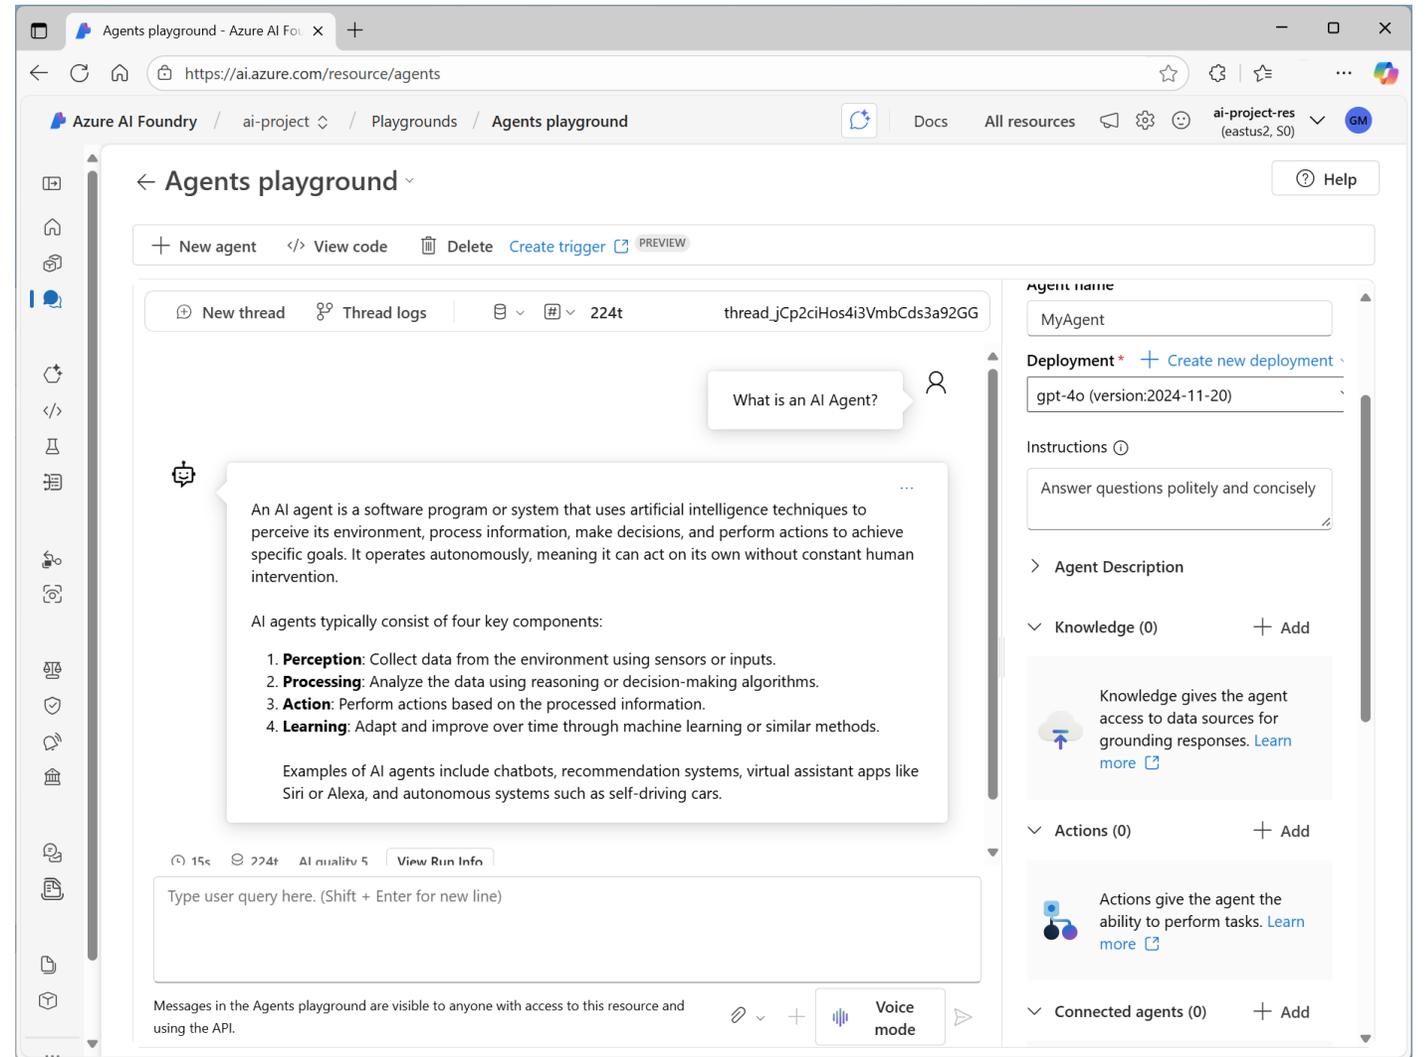
Azure AI Foundry Models

- Find the model that best suits your needs in the Azure AI Foundry model catalog:
 - OpenAI models (e.g. GPT 5)
 - Microsoft models (e.g. Phi 4)
 - Popular third-party models
- Deploy the model(s) you want to use in your apps
- Test model deployments in the *Chat playground*



Azure AI Foundry Agents

- Create agents in Azure AI Foundry portal, specifying:
 - Agent name
 - Model deployment
 - Knowledge tools
 - Action tools
 - Connected agents
- Test agents in the *Agents playground*

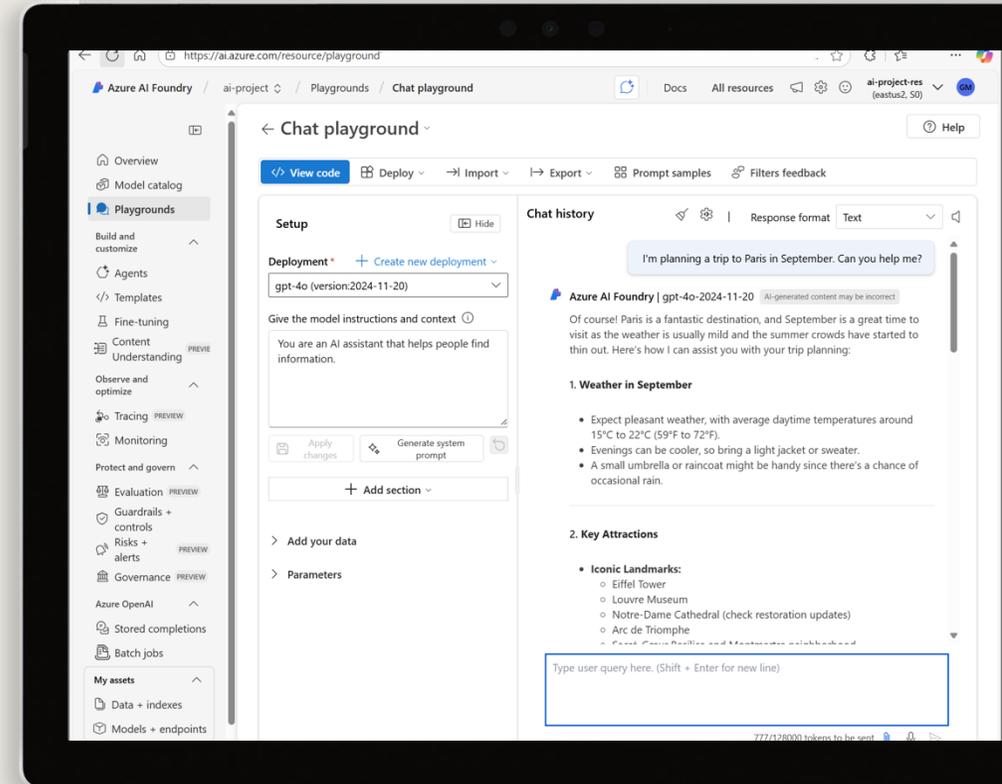


Explore generative AI in Azure AI Foundry Portal

In this exercise, you'll use Azure AI Foundry portal to explore generative AI models.

Start the exercise at:

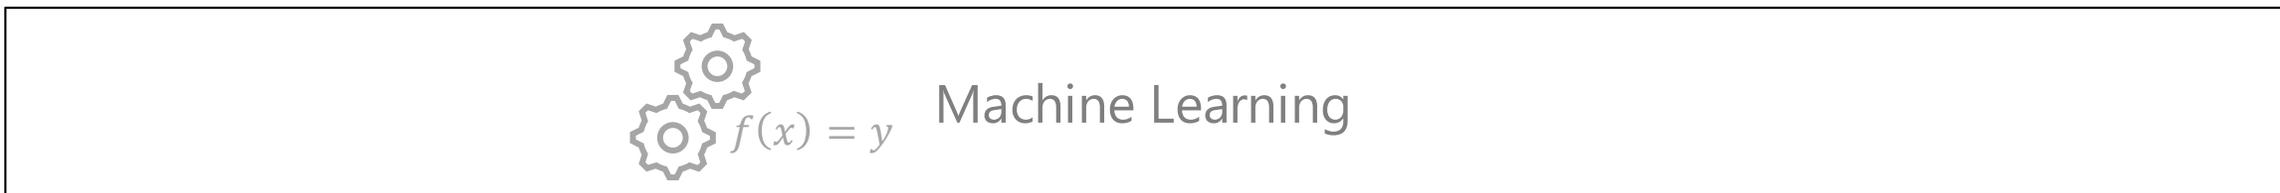
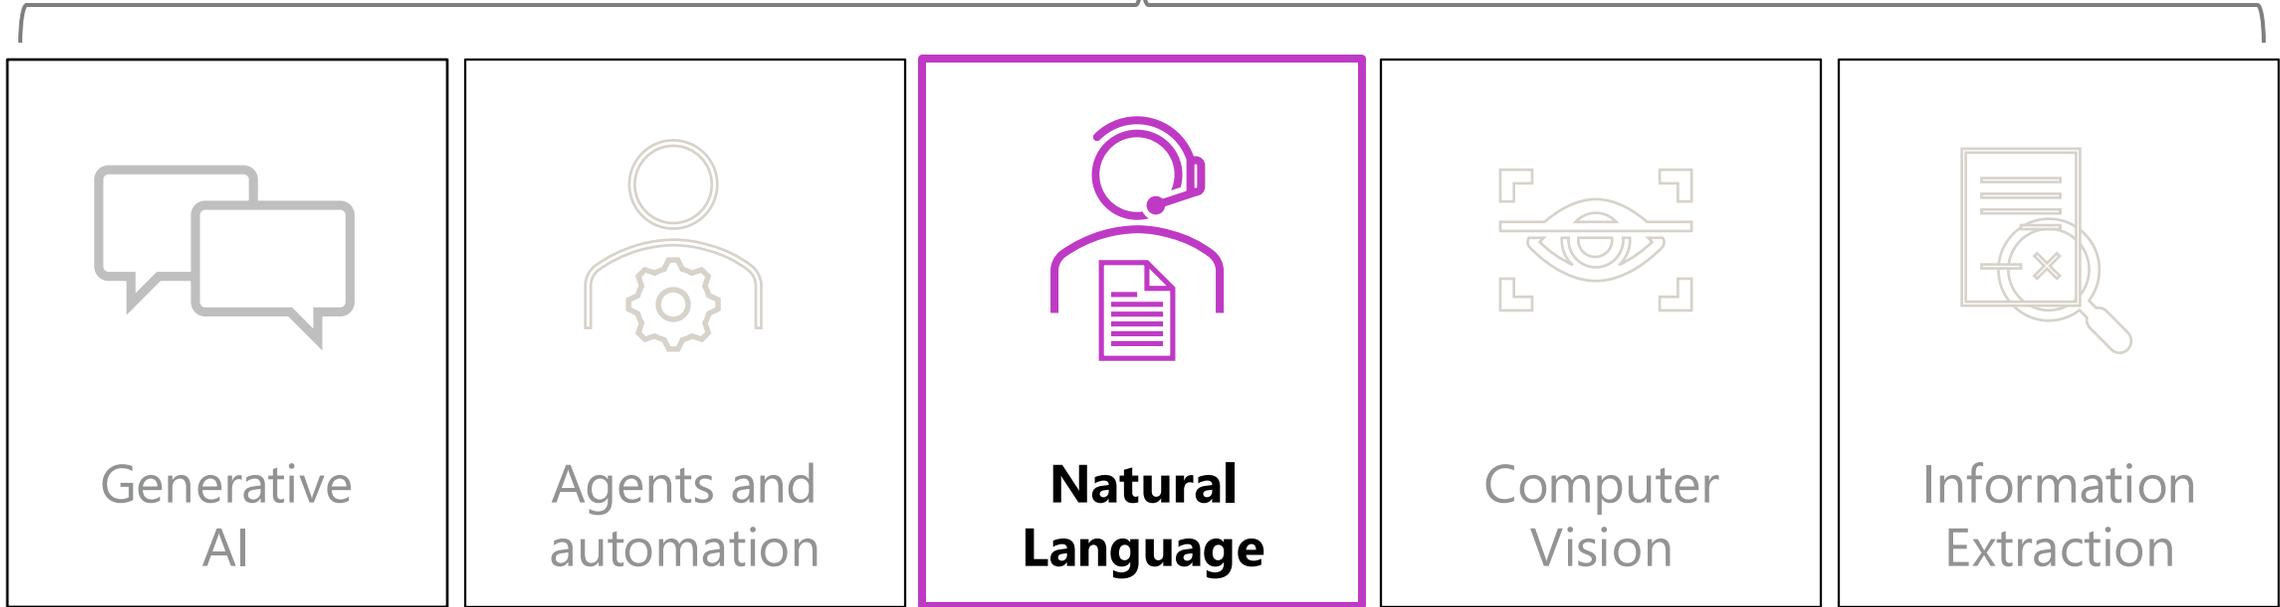
<https://go.microsoft.com/fwlink/?linkid=2249955>





Introduction to AI in Azure: Natural Language Processing

Natural Language Processing in context



Agenda



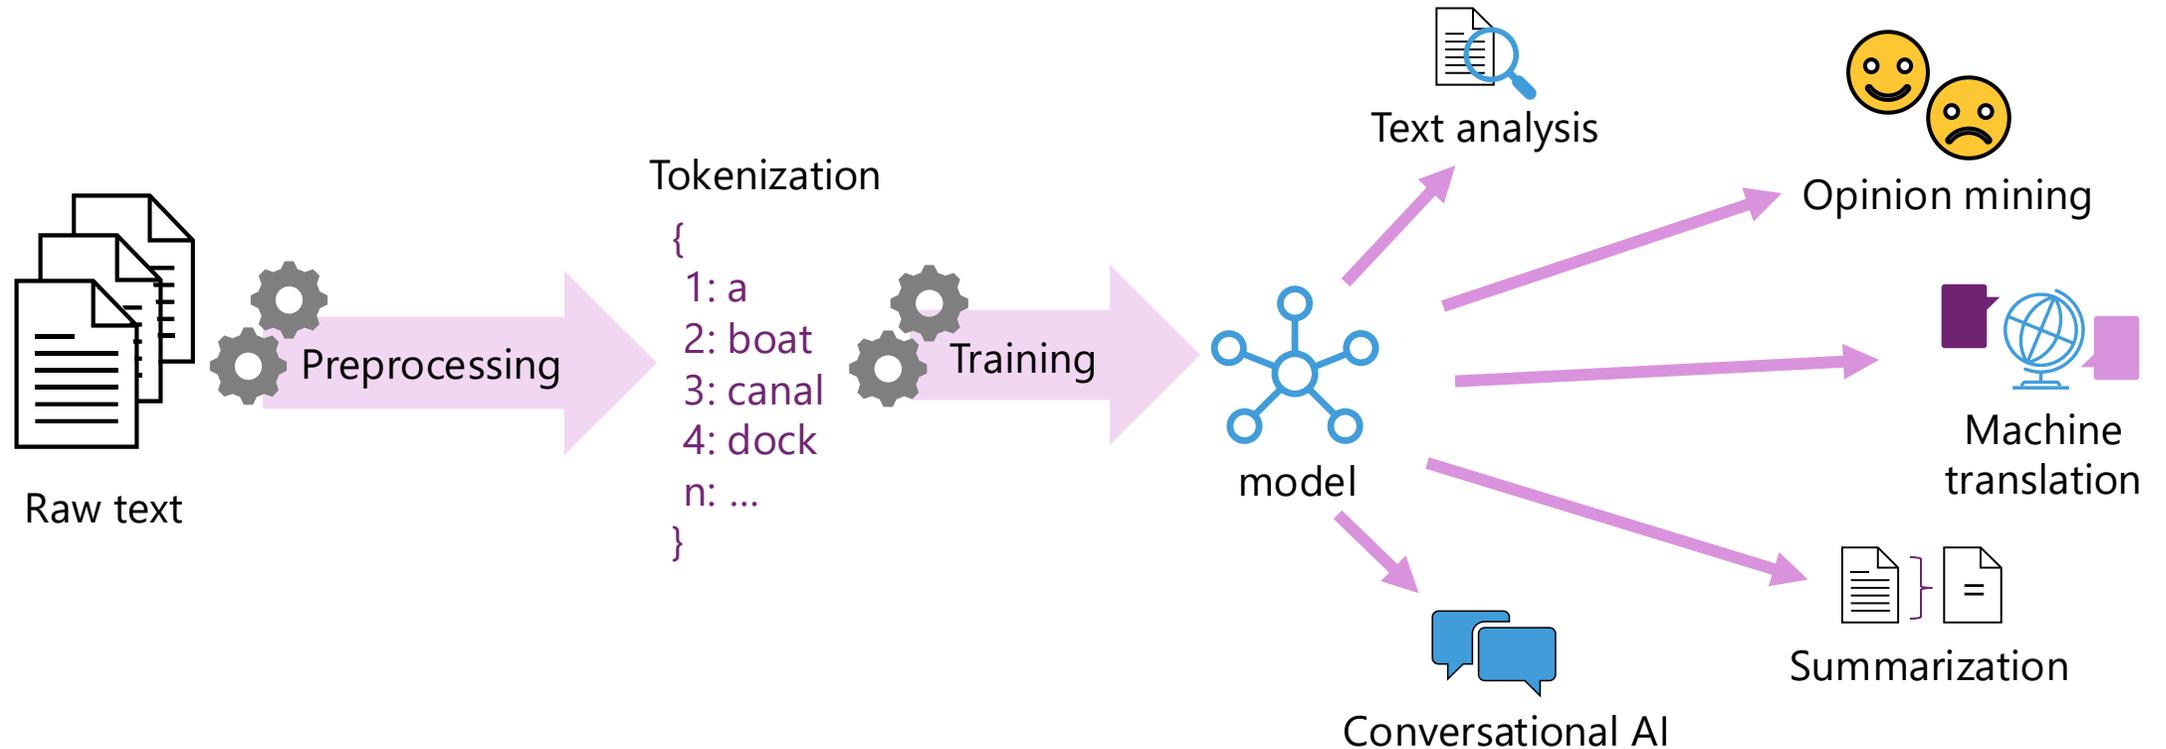
- Introduction to natural language processing concepts
- Get started with natural language processing in Azure
- Get started with speech in Azure

Introduction to natural language processing concepts

<https://aka.ms/mslearn-nlp>



What is natural language processing?

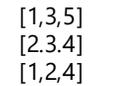


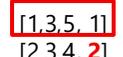
Techniques for language modeling and text analysis

Statistical techniques

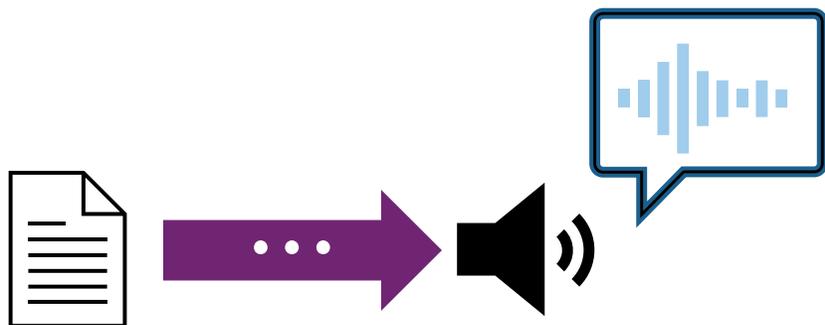
-  *Classification*: Use algorithms like Naïve Bayes and Logistic Regression to classify text based on the presence of specific tokens (words). For example, implementing an email spam filter based on terms that often appear in junk mail but not in “regular” email.
-  *Term Frequency / Inverse Document Frequency*: A statistical technique that compares the frequency of a specific term in individual documents with its frequency across a full *corpus* of documents. Used to determine the core *subject* of an individual document

Semantic modeling techniques

-  *Transformer models*: Represent language tokens as vector-based *embeddings* that encapsulate semantic relationships. Related terms have similar vector directions. Used for tasks like *translation* and *next-word-completion*.


- Attention*: A technique used in transformer models to improve next-word prediction based on token context. Weights are applied to tokens to reflect their influence on the next word. Using attention in the transformer architecture enables the latest advances in generative AI models

Speech processing



Speech Synthesis

(Text to speech)

1. *Tokenize* text
2. Map tokens to *phonemes*
3. Generate audio signal



Speech Recognition

(Speech to text)

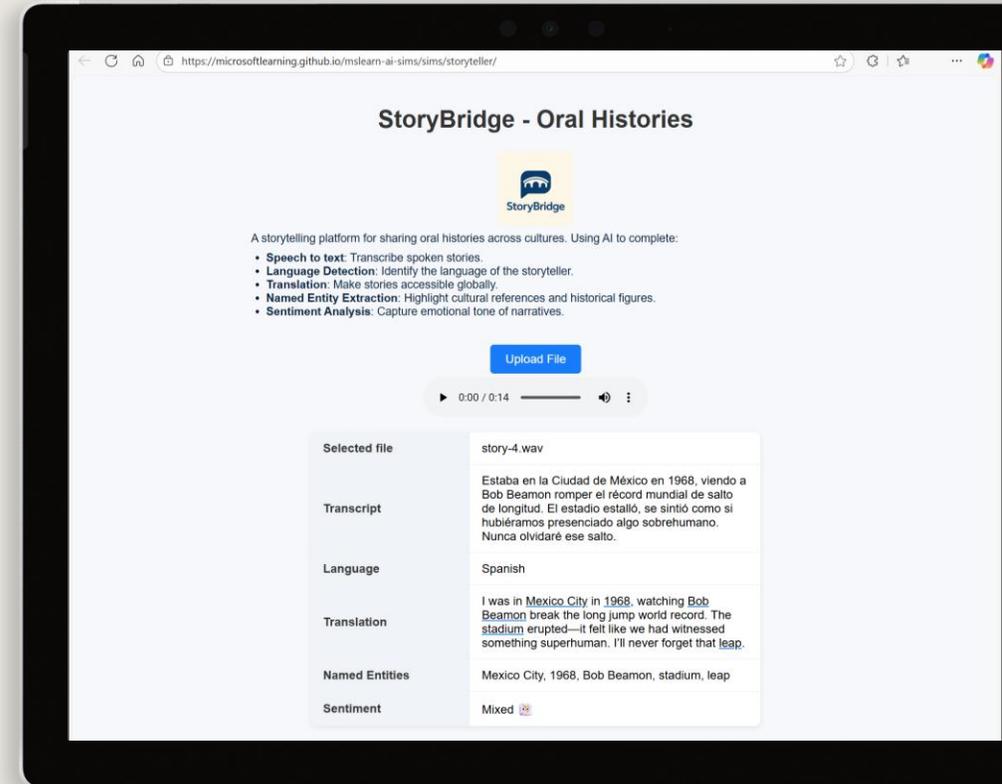
1. Capture audio signal
2. Break into *phonemes*
3. Map *phonemes* to text *tokens*

Explore natural language processing scenarios

In this exercise, you'll explore solutions that include natural language processing functionality.

Start the exercise at:

<https://go.microsoft.com/fwlink/?linkid=2334225>



Get started with natural language processing in Azure

<https://aka.ms/mslearn-azure-language>



AI Language services in Azure



Azure AI Language

- Language detection
- Key phrase extraction
- Named entity detection
- Sentiment analysis and opinion mining
- Personal information detection
- Summarization
- ...



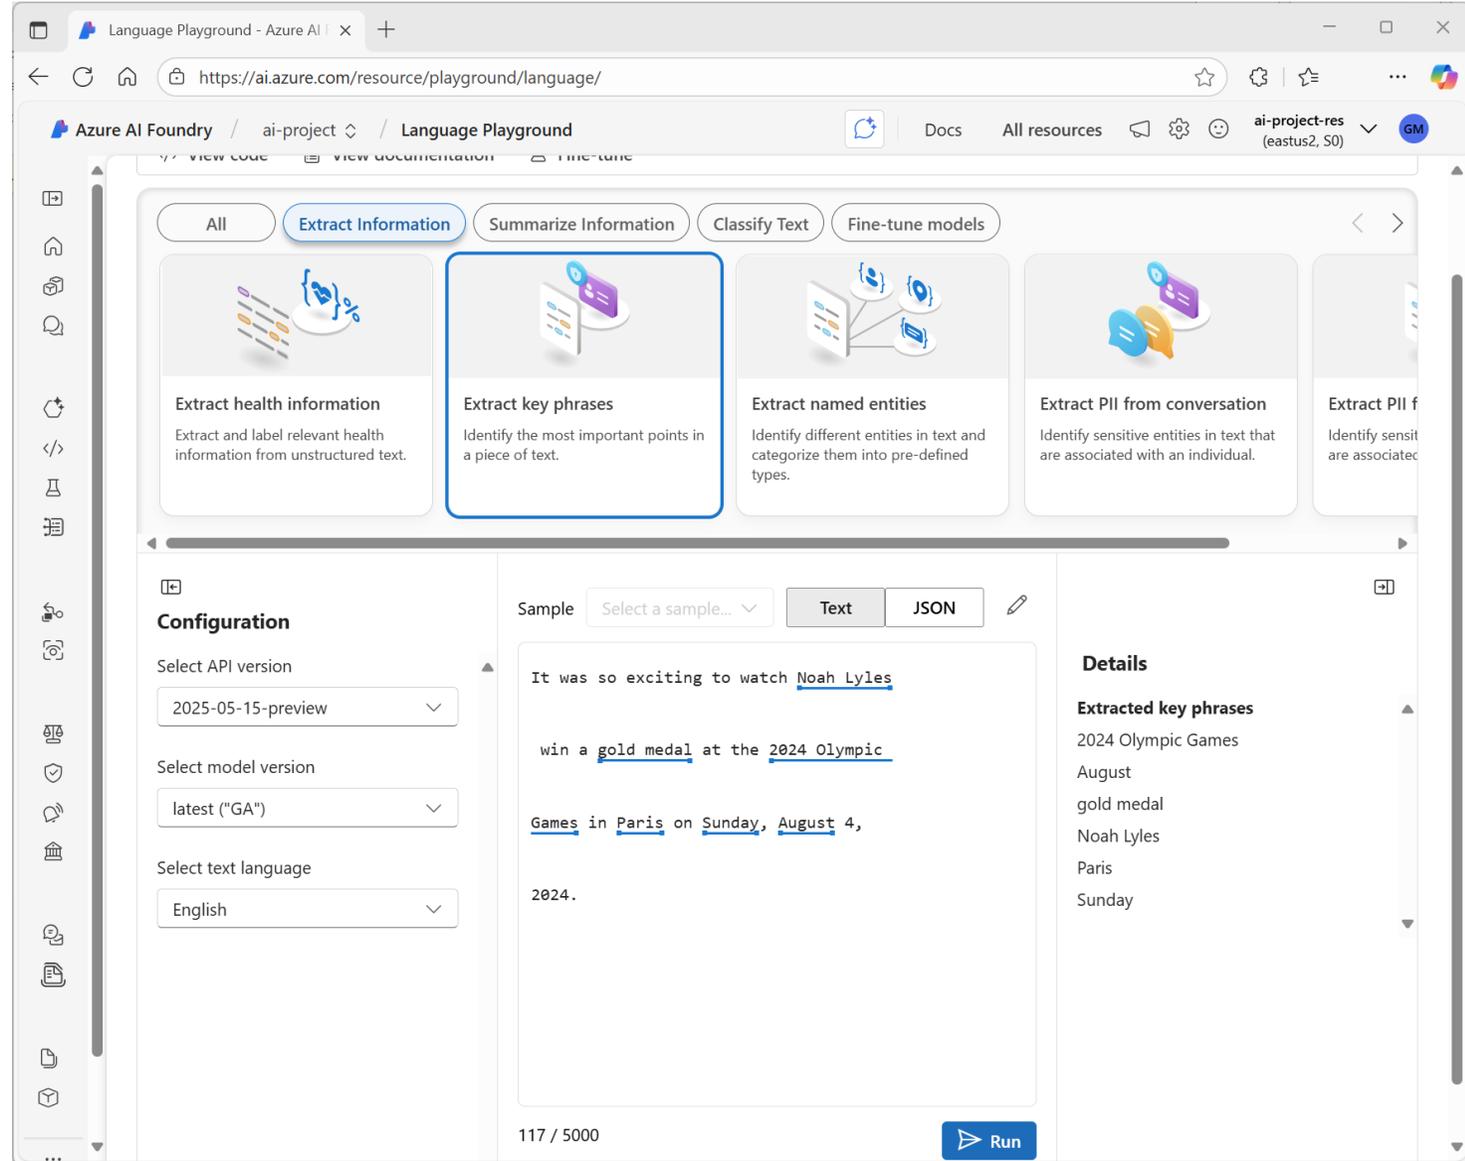
Azure AI Translator

- Text translation
- Document translation
- Custom translation
- ...

Analyzing text

"It was so exciting to watch Noah Lyles win a gold medal at the 2024 Olympic Games in Paris on Sunday, August 4, 2024."

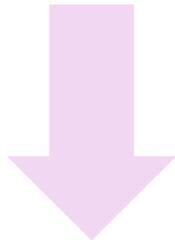
- **Language:** English
- **Sentiment:** Positive
- **Key Phrases:**
 - 2024 Olympic Games
 - gold medal
 - Noah Lyles
 - Paris
 - ...
- **Entities:**
 - *Person:* Noah Lyles
 - *Event:* Olympic Games
 - *City:* Paris
 - *Date:* Sunday, August 4, 2024
 - ...



The screenshot displays the Azure AI Foundry Language Playground interface. The browser address bar shows the URL <https://ai.azure.com/resource/playground/language/>. The page title is "Language Playground". The interface includes a navigation bar with "Docs", "All resources", and "ai-project-res (eastus2, S0)". The main content area features a "Language Playground" section with tabs for "All", "Extract Information", "Summarize Information", "Classify Text", and "Fine-tune models". The "Extract Information" tab is active, showing five analysis options: "Extract health information", "Extract key phrases", "Extract named entities", "Extract PII from conversation", and "Extract PII from text". The "Extract key phrases" option is selected, and its results are displayed in the "Details" panel on the right. The "Configuration" panel on the left shows the following settings: "Select API version" (2025-05-15-preview), "Select model version" (latest ("GA")), and "Select text language" (English). The "Sample" panel shows the input text: "It was so exciting to watch Noah Lyles win a gold medal at the 2024 Olympic Games in Paris on Sunday, August 4, 2024." The "Details" panel shows the "Extracted key phrases": "2024 Olympic Games", "August", "gold medal", "Noah Lyles", "Paris", and "Sunday". A "Run" button is visible at the bottom right of the interface.

Translating text

"It was so exciting to watch Noah Lyles win a gold medal at the 2024 Olympic Games in Paris on Sunday, August 4, 2024."



"C'était tellement excitant de voir Noah Lyles remporter une médaille d'or aux Jeux olympiques de 2024 à Paris le dimanche 4 août 2024."

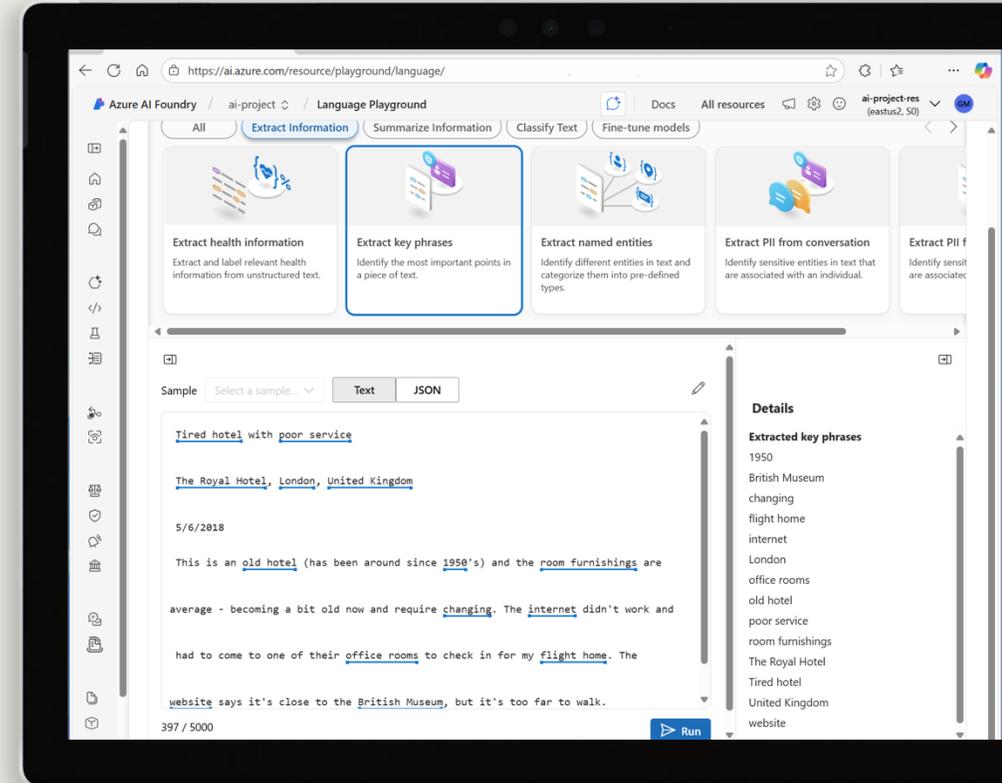
The screenshot displays the Azure AI Translator Playground interface. The browser address bar shows the URL <https://ai.azure.com/resource/playground/translator>. The page title is "Translator Playground". The interface includes a navigation bar with "Azure AI Foundry", "ai-project", "Playgrounds", and "Translator Playground". The main content area is titled "Translator Playground" and contains two tabs: "Text translation" and "Document translation (sync)". The "Text translation" tab is active, showing a configuration panel on the left with "Translate from" set to "Auto detect" and "Translate to" set to "French". The "Original" text is "It was so exciting to watch Noah Lyles win a gold medal at the 2024 Olympic Games in Paris on Sunday, August 4, 2024." and the "Translation" is "C'était tellement excitant de voir Noah Lyles remporter une médaille d'or aux Jeux olympiques de 2024 à Paris le dimanche 4 août 2024." The character count is 117 / 1000. A "Translate" button is visible at the bottom right.

Analyze text in Azure AI Foundry portal

In this exercise, you'll use Azure AI Foundry portal to explore Azure AI Language capabilities.

Start the exercise at:

<https://go.microsoft.com/fwlink/?linkid=2250314>

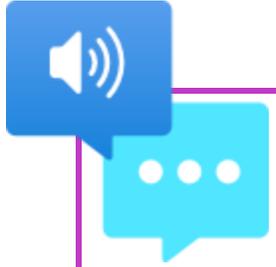


Get started with speech in Azure

<https://aka.ms/mslearn-azure-speech>



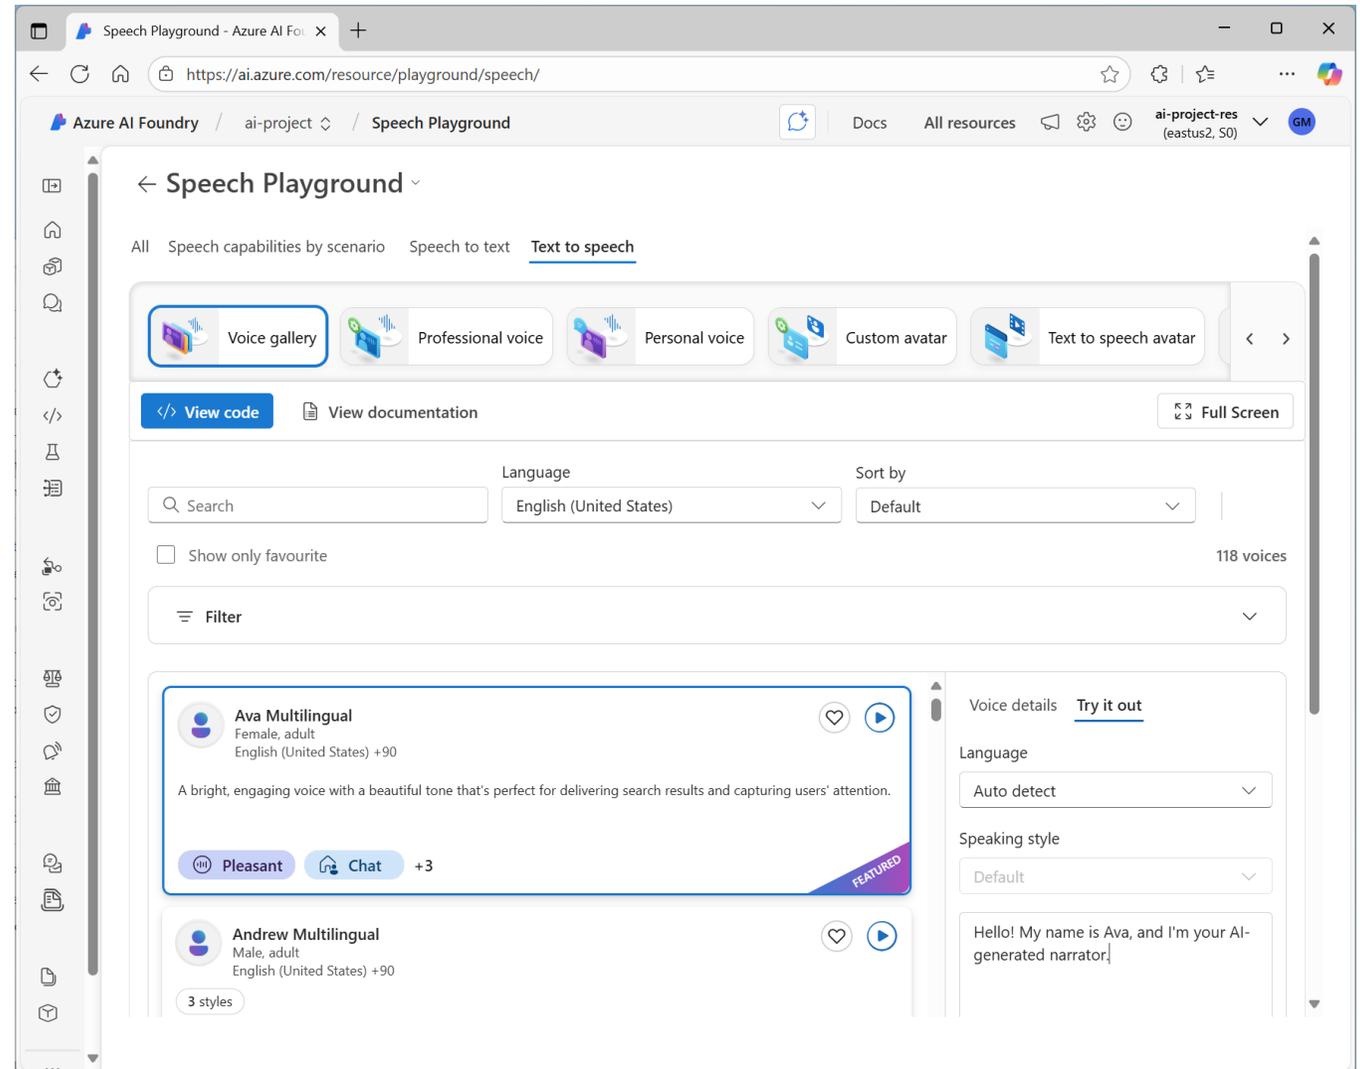
Azure AI Speech capabilities



- Speech to text transcription (real-time and batch)
- Text to speech synthesis
- Speech translation
- Custom speech models
- Built-in and custom voices
- Audio content creation
- Video translation
- Voice Live: real-time voice agents
- Built-in and custom avatars
- ...

Speech in Azure AI Foundry Portal

- Use the *Speech playground* to:
 - Transcribe recorded or live speech
 - Synthesize speech with built-in voices
 - Create custom voices
 - Create custom speech models
 - Add live voice capabilities to an agent
 - Translate audio and video
 - Create avatar-presented videos
 - Create custom avatars
 - ...

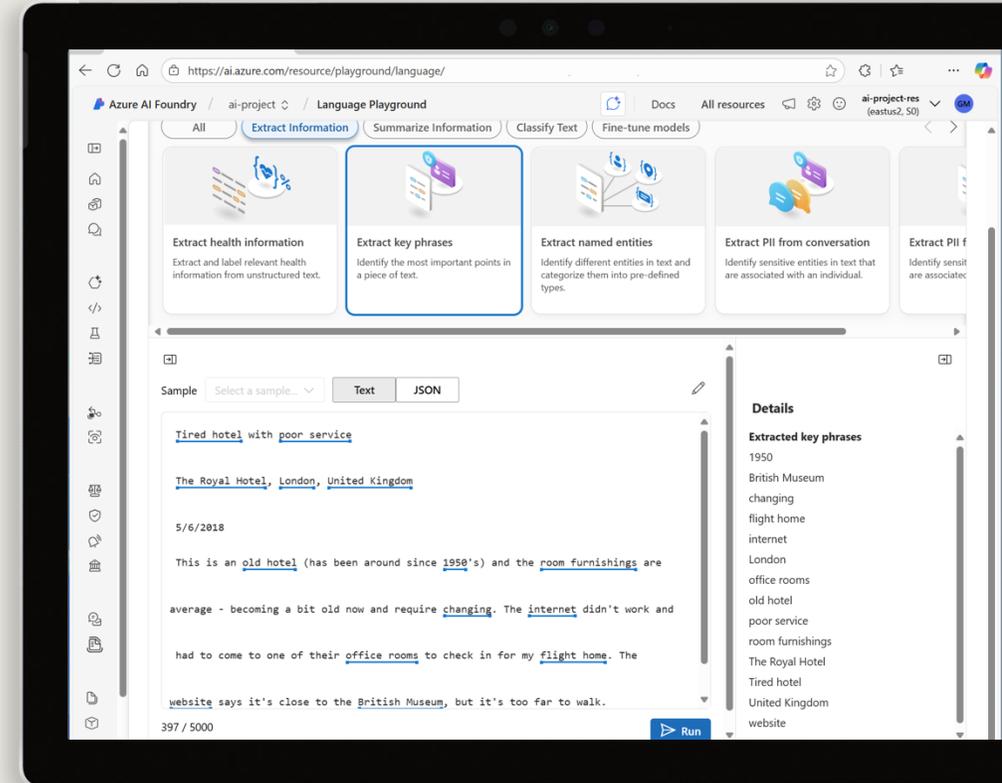


Explore Speech in Azure AI Foundry

In this exercise, you'll use Azure AI Foundry portal's speech playground to see speech capabilities in action.

Start the exercise at:

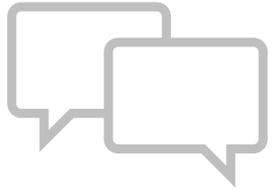
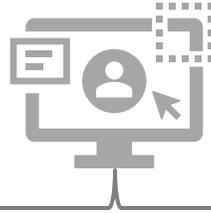
<https://go.microsoft.com/fwlink/?linkid=2250148>



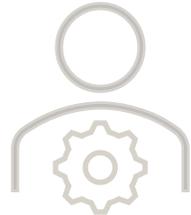


Introduction to AI in Azure: Computer Vision

Our focus



Generative
AI



Agents and
automation



Natural
Language



**Computer
Vision**



Information
Extraction



$$f(x) = y$$

Machine Learning

Agenda



- Introduction to computer vision concepts
- Get started with computer vision in Azure

Introduction to computer vision concepts

<https://aka.ms/mslearn-vision>



Images and image processing

An image is an array of pixel values

Single-channel (monochrome) or multi-channel (color)

0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	255	255	255	0	0
0	0	255	255	255	0	0
0	0	255	255	255	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0

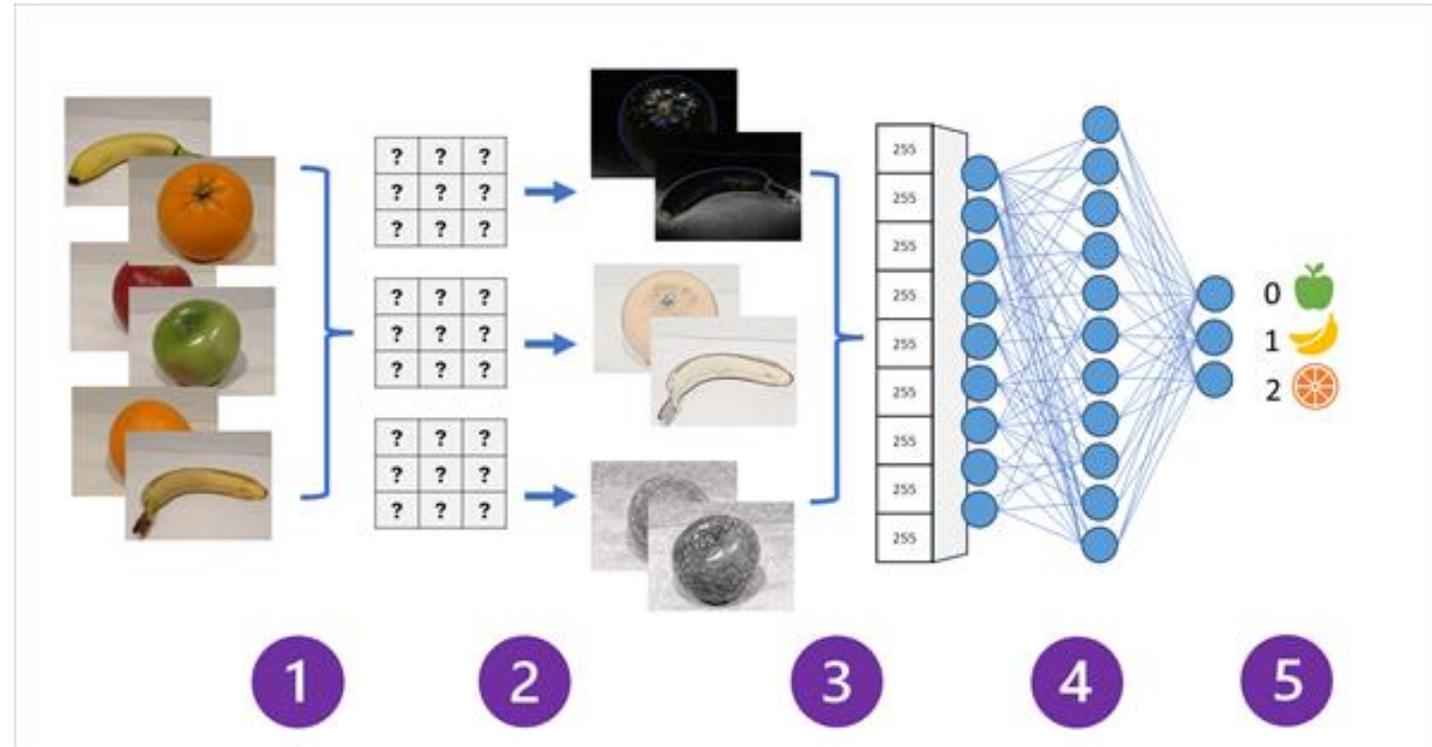
Filters are applied based on a *kernel*

Creating structure and extracting features

1	1	1	0	0	0	0
1	8	1	0	0	0	0
1	1	255	255	255	0	0
0	0	255	0	255	0	0
0	0	255	255	255	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0

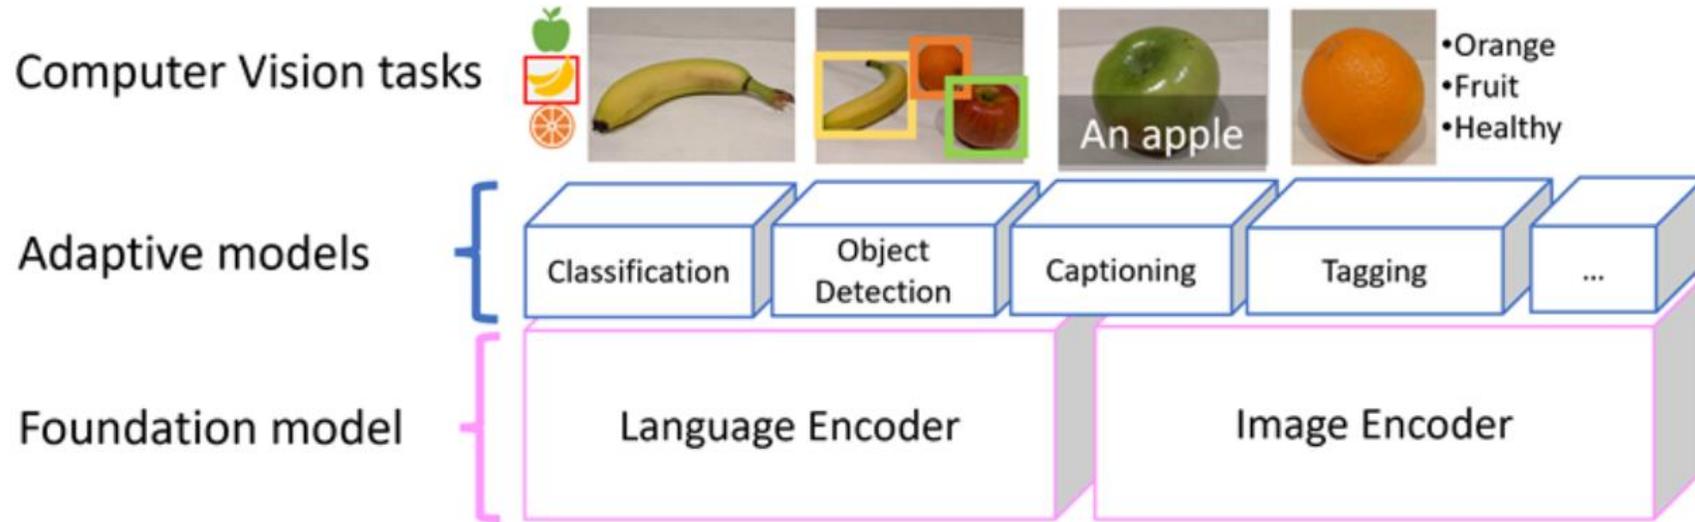
Convolutional Neural Networks

1. Labeled images are used to train the model
2. Filter layers extract *feature maps* from each image
3. The feature maps are flattened
4. The feature values are fed into a fully connected neural network
5. The output layer produces a probability value for each possible class label



- During training, the filter kernels start with random weights. These weights are iteratively adjusted to improve the accuracy of the predictions based on the known labels.
- The trained model uses learned weights to extract features from new images and predict their class.

Multi-modal models



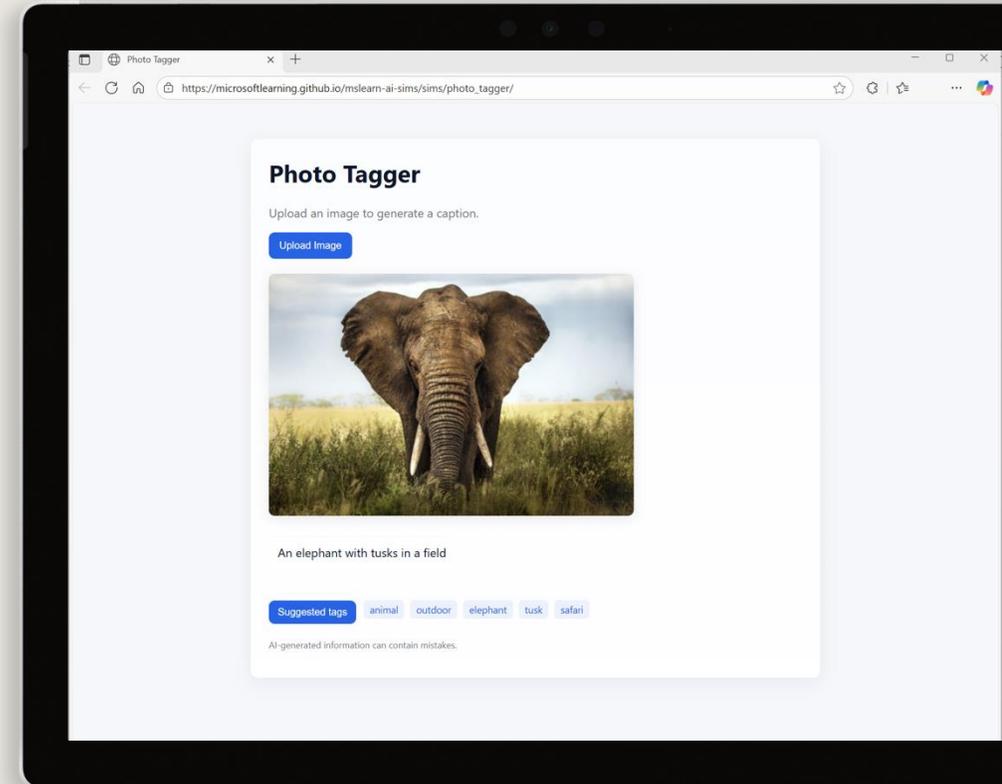
- A newer approach to modeling involves combining language and vision models that encode image and text data
- The model encapsulates semantic relationships between features extracted from the images and text extracted from related captions.
- A multi-modal model can be used as a *foundation* model for more specialized *adaptive* models.

Explore computer vision scenarios

In this exercise, you'll explore computer vision capabilities in an application.

Start the exercise at:

<https://go.microsoft.com/fwlink/?linkid=2334316>



Get started with computer vision in Azure

<https://aka.ms/mslearn-azure-vision>



Computer vision services in Azure



Vision

- Image Analysis:
 - Image tagging, captions, model customization, and more.
- Optical Character Recognition (OCR)
- Video analysis



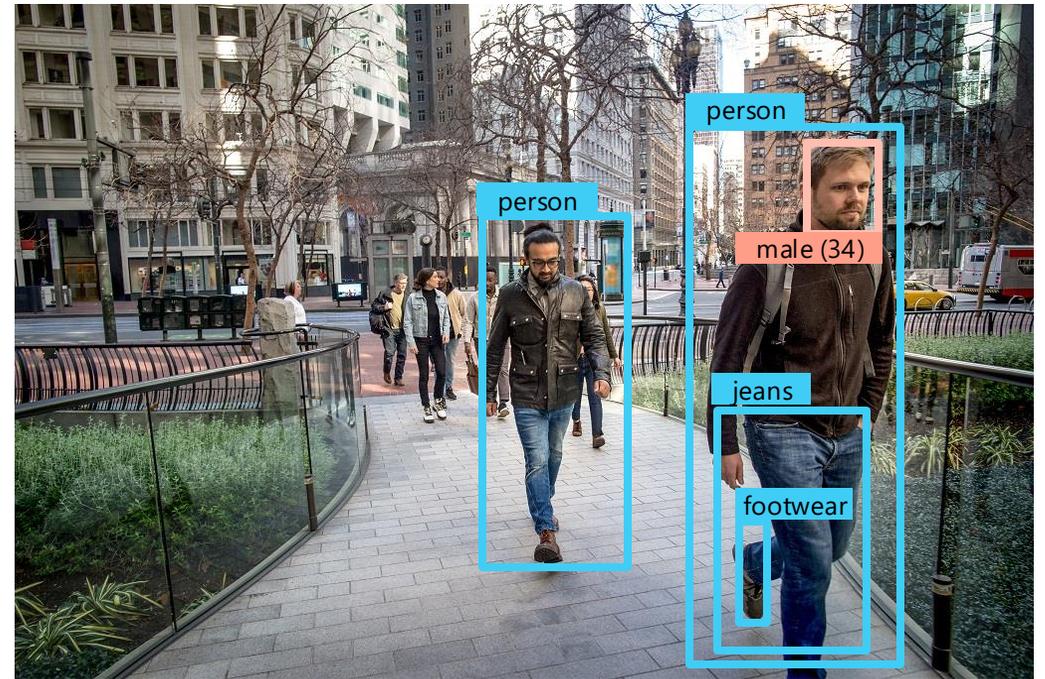
Face

- Face detection
- Face recognition

Image analysis 4.0 with the *AI Vision Service*

Capabilities include:

- Model customization
- Read text from images
- Detect people in images
- Generate image captions
- Detect objects
- Tag visual features
- Smart crop



Caption: A group of people walking on a sidewalk

Tags: Building, jeans, street, outdoor, jacket, city, person

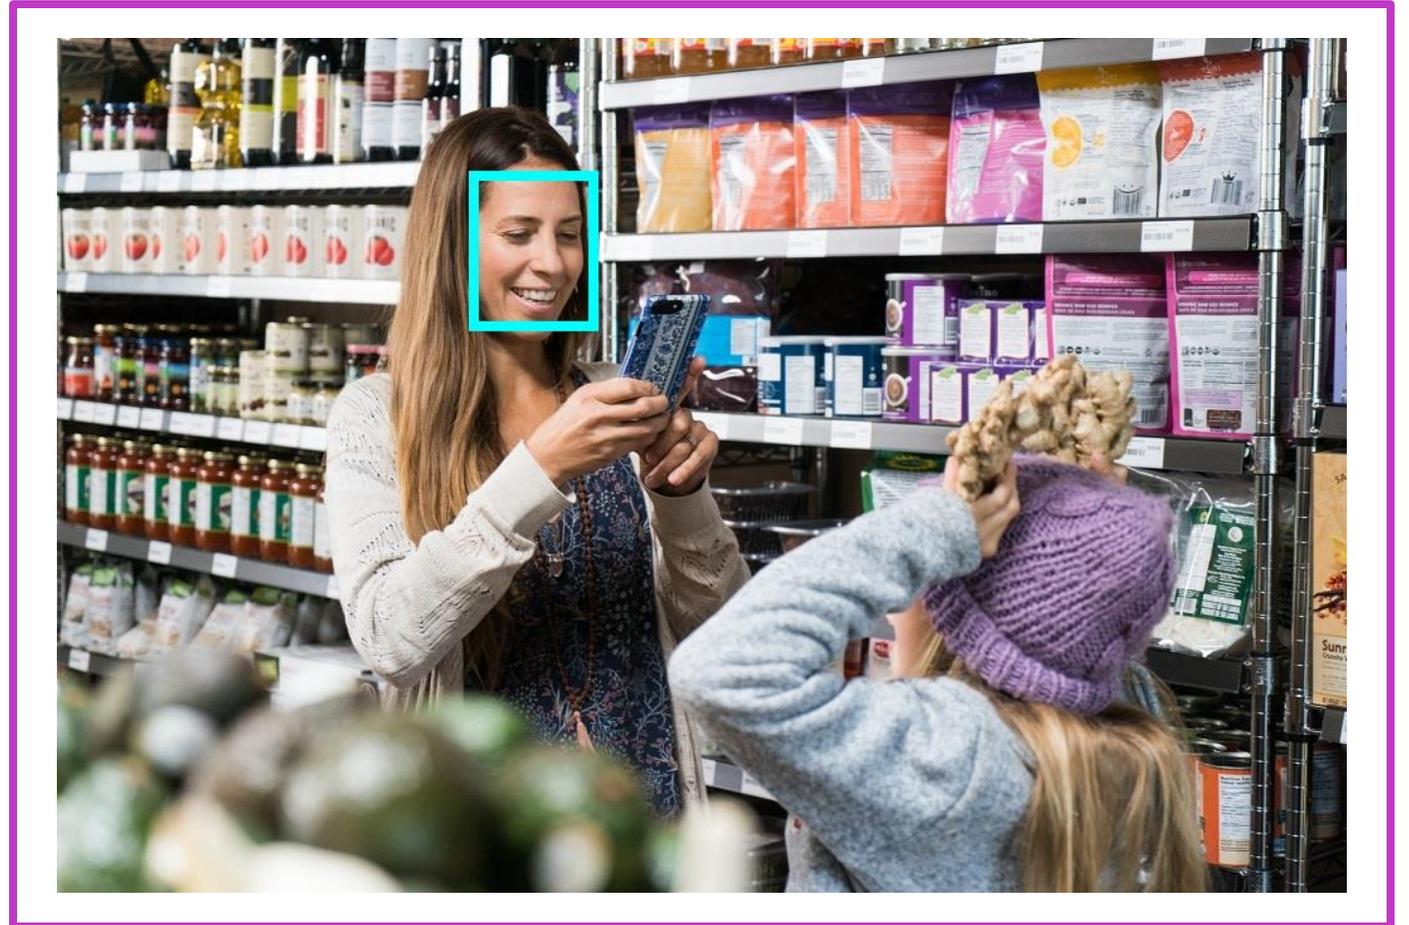
Detecting faces with the *Face* Service

Everyone can use the Face service to detect:

- Blur
- Exposure
- Glasses
- Head pose
- Noise
- Occlusion

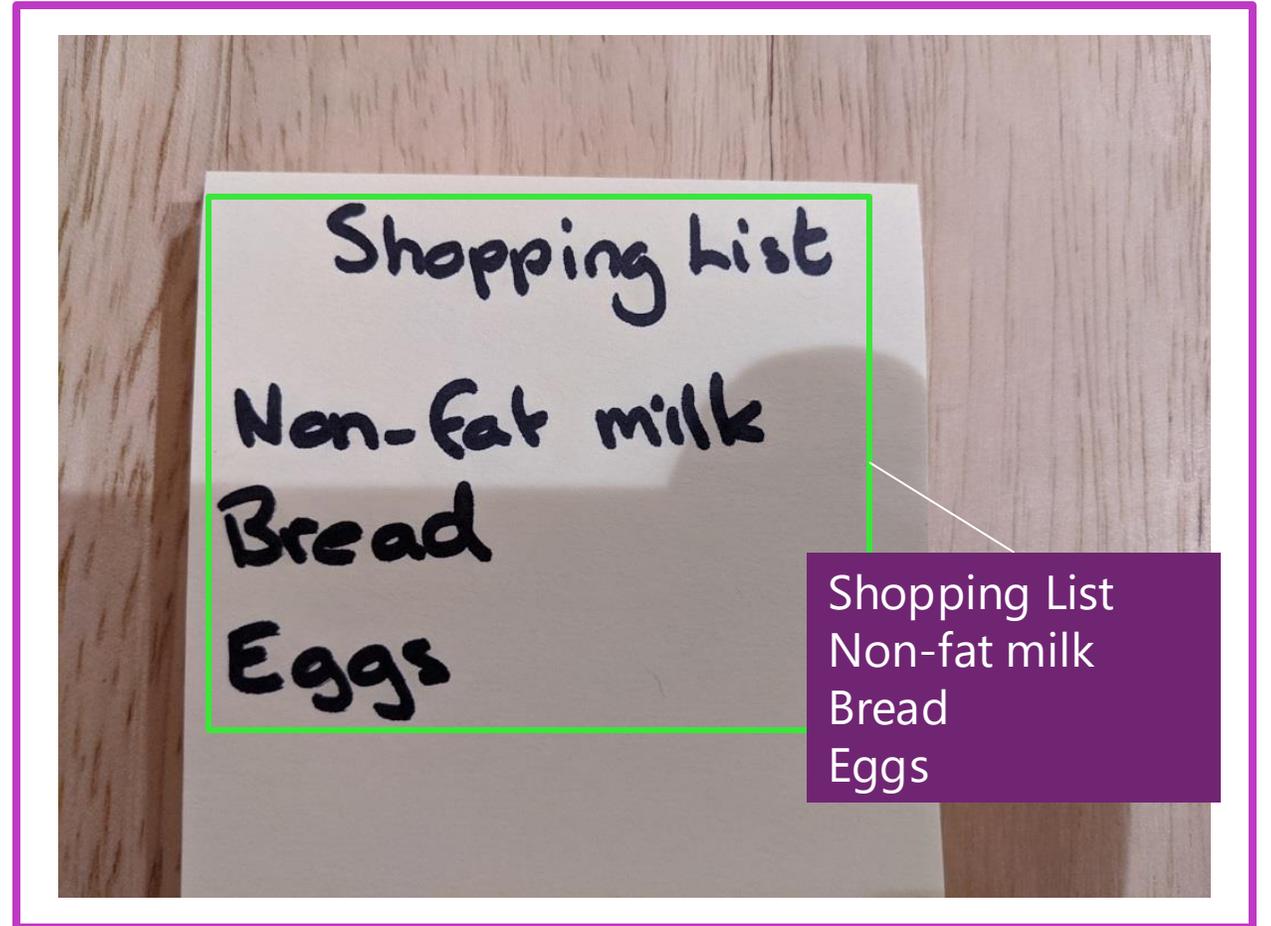
Only Managed Microsoft customers can access facial recognition capabilities:

- Similarity matching
- Identity verification



Reading text with Optical Character Recognition (OCR)

- Detect the location and characters of **printed** and **handwritten** text
- Options for quick text extraction from images, or asynchronous analysis of larger scanned documents

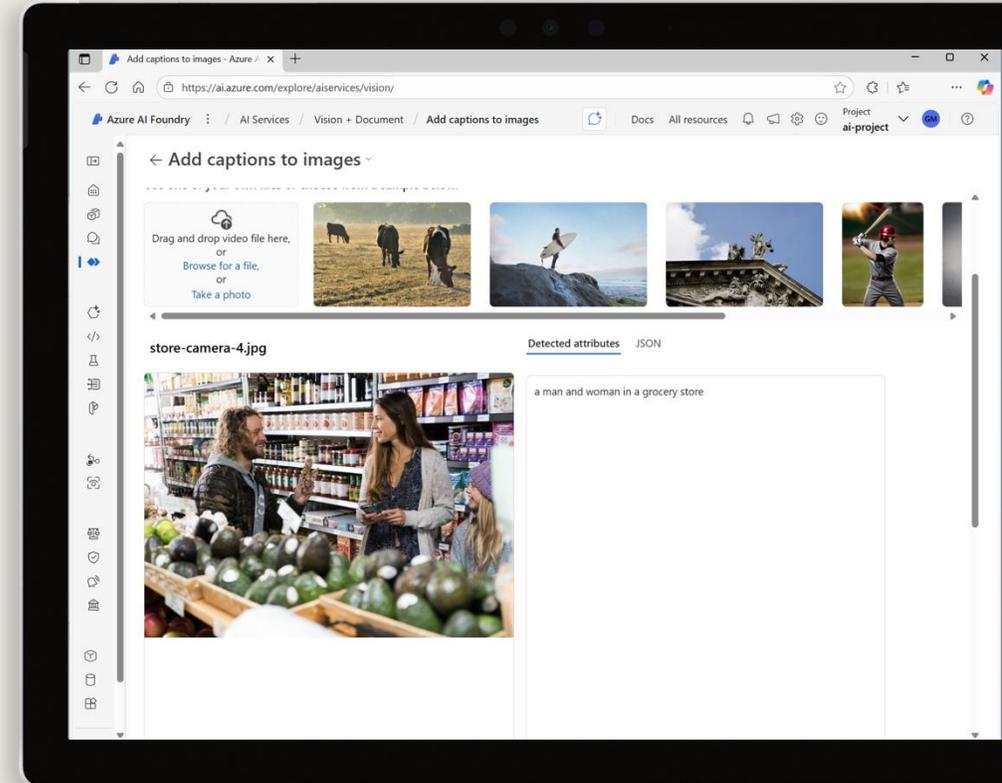


Analyze images in Azure AI Foundry portal

In this exercise, you'll use Azure AI Foundry portal to explore Azure AI Vision capabilities.

Start the exercise at:

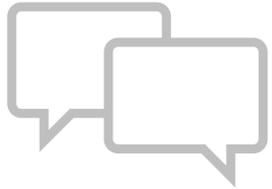
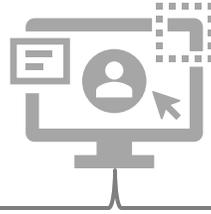
<https://go.microsoft.com/fwlink/?linkid=2250145>



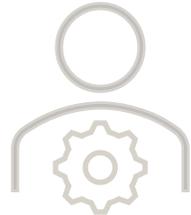


Introduction to AI in Azure: Information Extraction

Our focus



Generative
AI



Agents and
automation



Natural
Language



Computer
Vision



**Information
Extraction**



$$f(x) = y$$

Machine Learning

Agenda



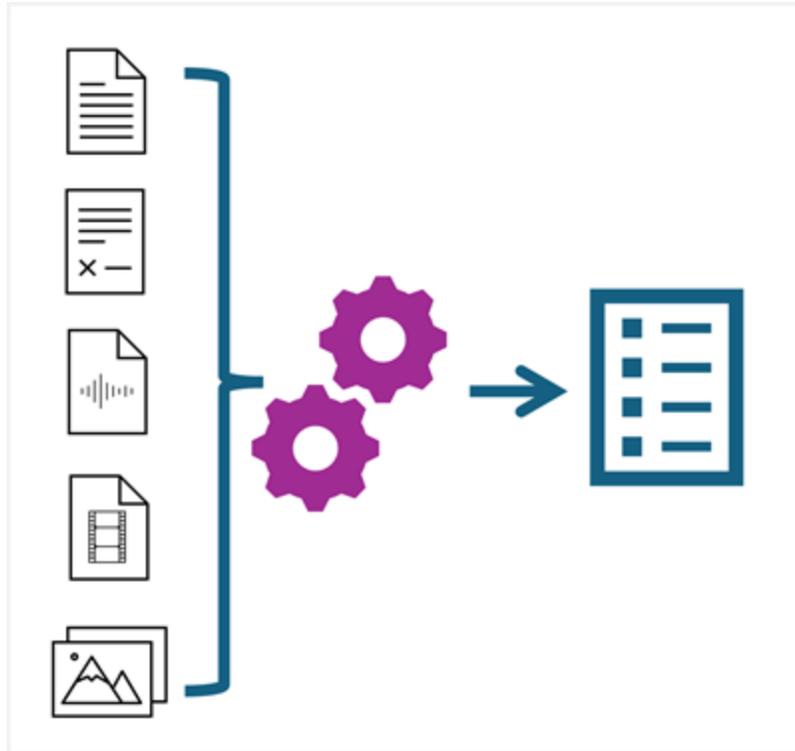
- Introduction to AI-powered information extraction concepts
- Get started with AI-powered information extraction in Azure

Introduction to AI-powered information extraction concepts

<https://aka.ms/mslearn-ai-info>



Information extraction processes

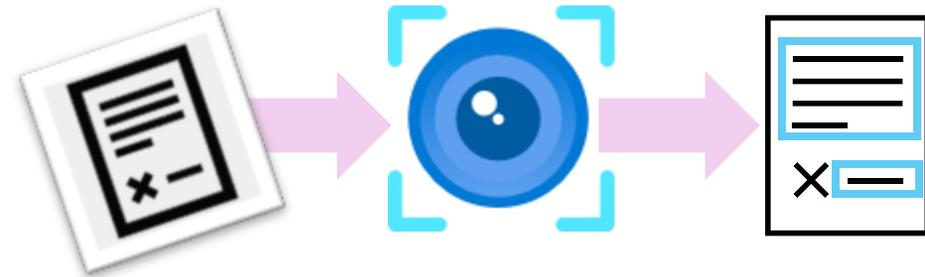


Step	Description
Source Identification	Determine where the information resides and if it needs to be digitized.
Extraction	Leverages many techniques based on machine learning to understand and extract data from digitized content.
Transformation & Structuring	Extracted data is transformed into structured formats like JSON or tables.
Storage & Integration	The processed data is then stored in databases, data lakes, or analytics platforms for further use.

Understand the extraction of data from images

Use optical character recognition (OCR) to:

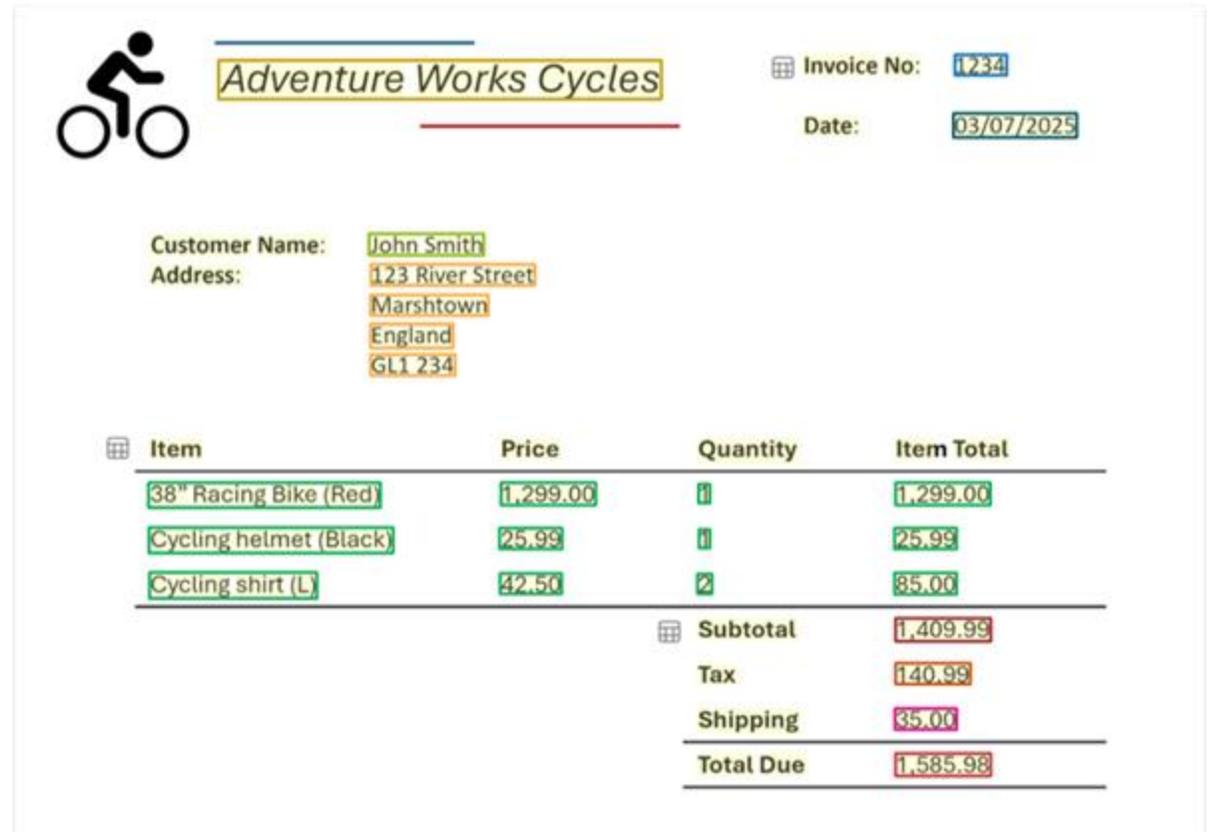
- Extract contact information from scanned business cards or conference badges
- Capture information from IDs to complete electronic forms
- Scan and store recipes or other text
- Photograph a sign or storefront so you can submit it to a translation service
- Digitize handwritten notes



Understand the extraction of data from forms

- The **field name** is the key or type of data entry.
- The **field description** is the definition of what the field name represents.
- The **value** corresponds with the field name and is the data specific to the content.

The data in forms is recognized with *bounding boxes*.

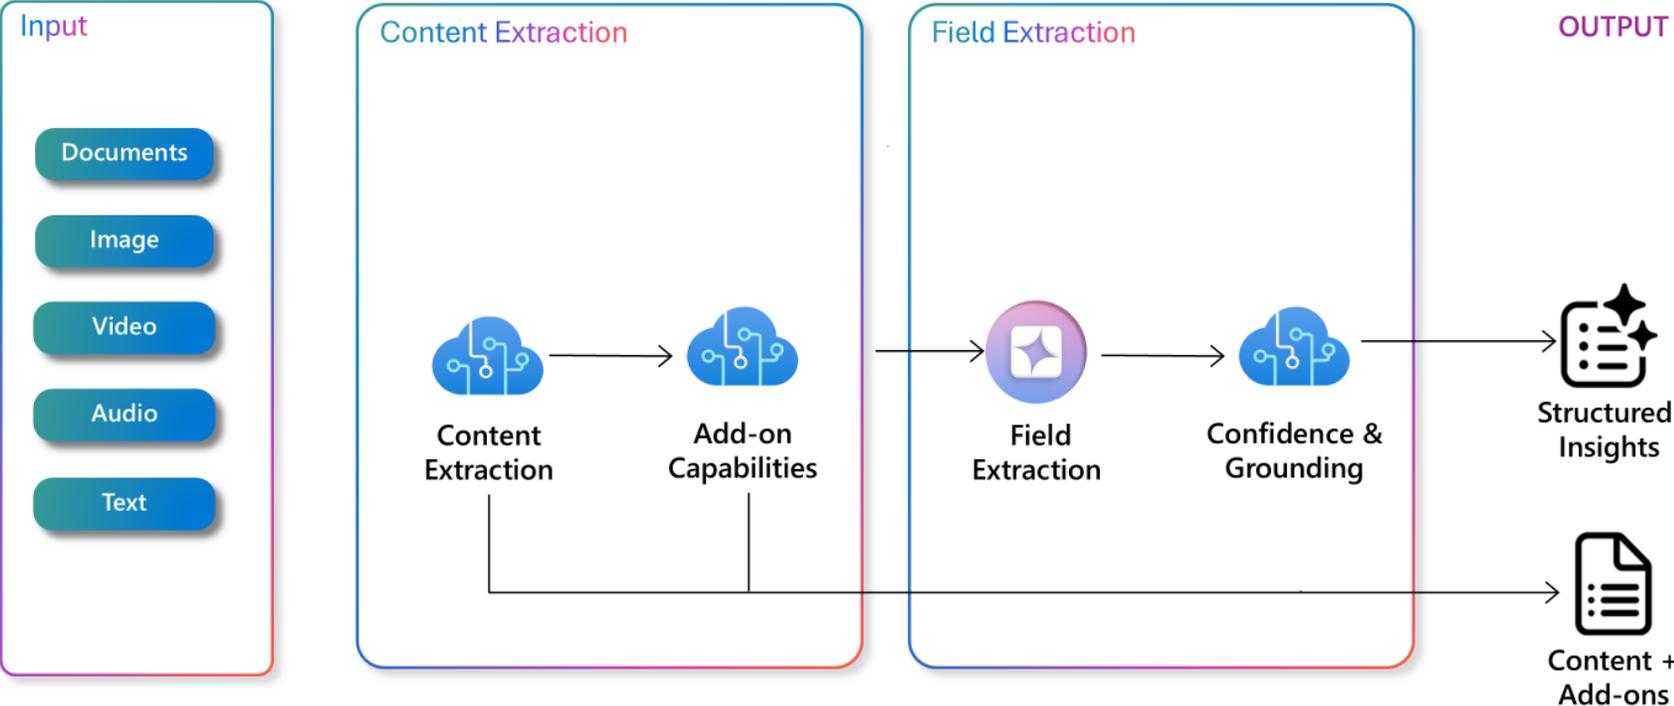


The image shows a screenshot of an invoice form for "Adventure Works Cycles". The form includes a logo of a cyclist, the company name, invoice number (1234), and date (03/07/2025). Customer information includes the name "John Smith" and address "123 River Street, Marshtown, England, GL1 234". A table lists items: "38\" Racing Bike (Red)" (1 unit, 1,299.00), "Cycling helmet (Black)" (1 unit, 25.99), and "Cycling shirt (L)" (2 units, 85.00). The total due is 1,585.98.

Item	Price	Quantity	Item Total
38" Racing Bike (Red)	1,299.00	1	1,299.00
Cycling helmet (Black)	25.99	1	25.99
Cycling shirt (L)	42.50	2	85.00
Subtotal			1,409.99
Tax			140.99
Shipping			35.00
Total Due			1,585.98

Understand multimodal data extraction

AI-powered information extraction techniques can be combined to perform data extraction on multiple modalities of content, from documents to video and audio.

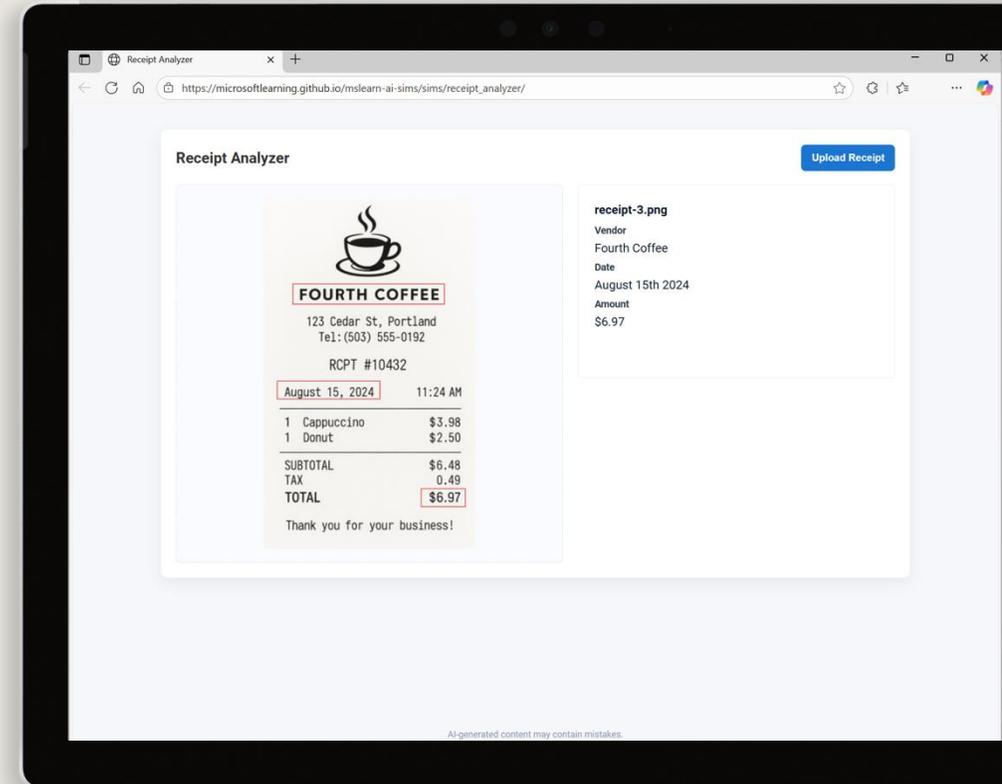


Explore AI information extraction scenarios

In this exercise, you'll explore applications that extract information

Start the exercise at:

<https://go.microsoft.com/fwlink/?linkid=2334119>



Get started with AI-powered information extraction in Azure

<https://aka.ms/mslearn-azure-info-extract>



Azure AI options for reading text



Azure AI Vision Image Analysis

- Identify text and its location in scanned documents
- Find and read text in photographs
- Combine with other image analysis features to implement a digital asset management (DAM) solution



Azure AI Document Intelligence

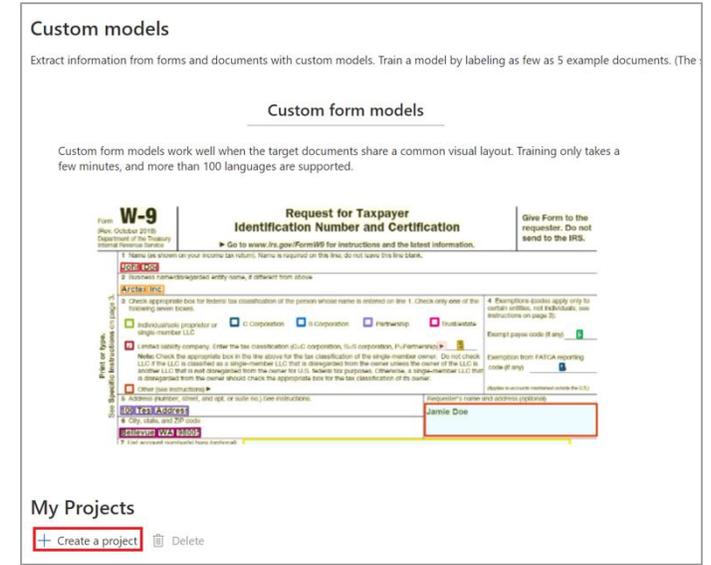
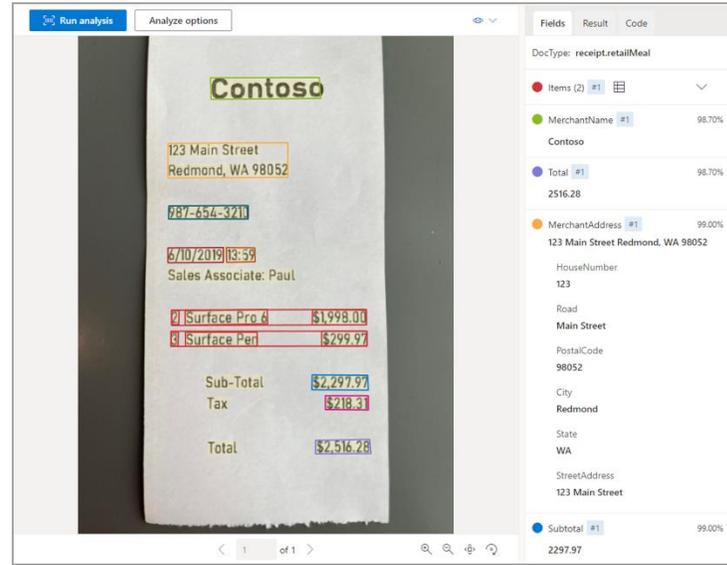
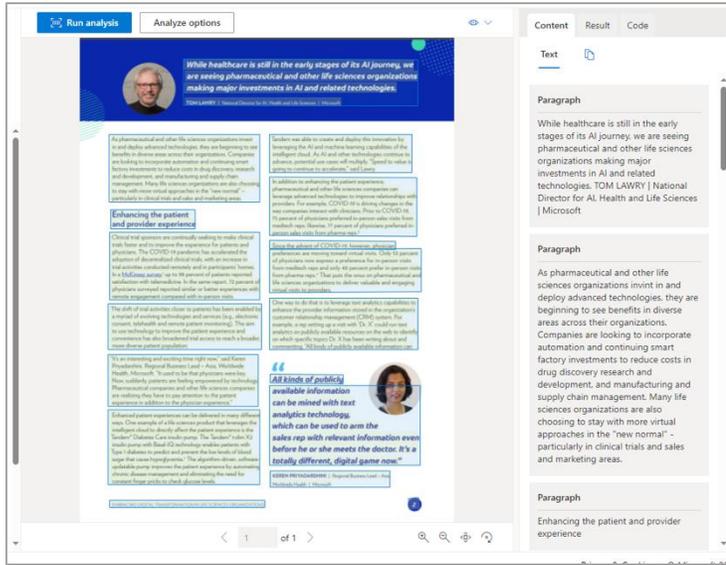
- Design to support form processing by extracting fields and associated values from documents
- Use prebuilt models for common document types
- Create custom models for your specific requirements



Azure AI Content Understanding

- Use multimodal content extraction capabilities to analyze documents, forms audio, video, and images.
- Create custom analyzers to extract specific content or fields tailored to business needs

Azure AI Document Intelligence service



- Document analysis
- Returns structured data representations.
- Regions of interest and relationships.
- Configure **Analyze options** for free and chargeable analysis

- Prebuilt models
- Invoices
- Receipts
- ID
- Recognizes and extracts key-value pairs.

- Custom models
- Train models with at least five sample data.
- Identify fields of interest to your organization.

Azure AI Content Understanding service

The screenshot shows the Azure AI Foundry interface for the Content Understanding service. The main view is titled "Test analyzer" and displays a video player for the file "wind-energy.mp4". The video shows a sunset scene with several wind turbines on a hill. The interface includes a navigation sidebar on the left with options like Overview, Model catalog, Playgrounds, AI Services, Build and customize, Agents, Templates, Fine-tuning, Content Understanding (selected), Prompt flow, Assess and improve, Tracing, Evaluation, and Safety + security. The main area has tabs for "Define schema", "Test analyzer", and "Build analyzer". Below the video player, there is a "Transcript" section with the text "Shot 00:00.000 => 00:09.976" and "Transcript WEBVTT". A "Key Frames" section is also visible. On the right, a "Fields" sidebar shows the analysis results for "description", "background", "shotType", and "videoCategories".

Fields Result

- description
A serene sunset scene featuring multiple wind turbines on a hill, silhouetted against a purple sky. The sun is setting on the horizon, creating a picturesque and tranquil atmosphere.
- background
The video is set outdoors during sunset, showcasing a landscape with wind turbines on a hill.
- shotType
WideAngle
- videoCategories (1)

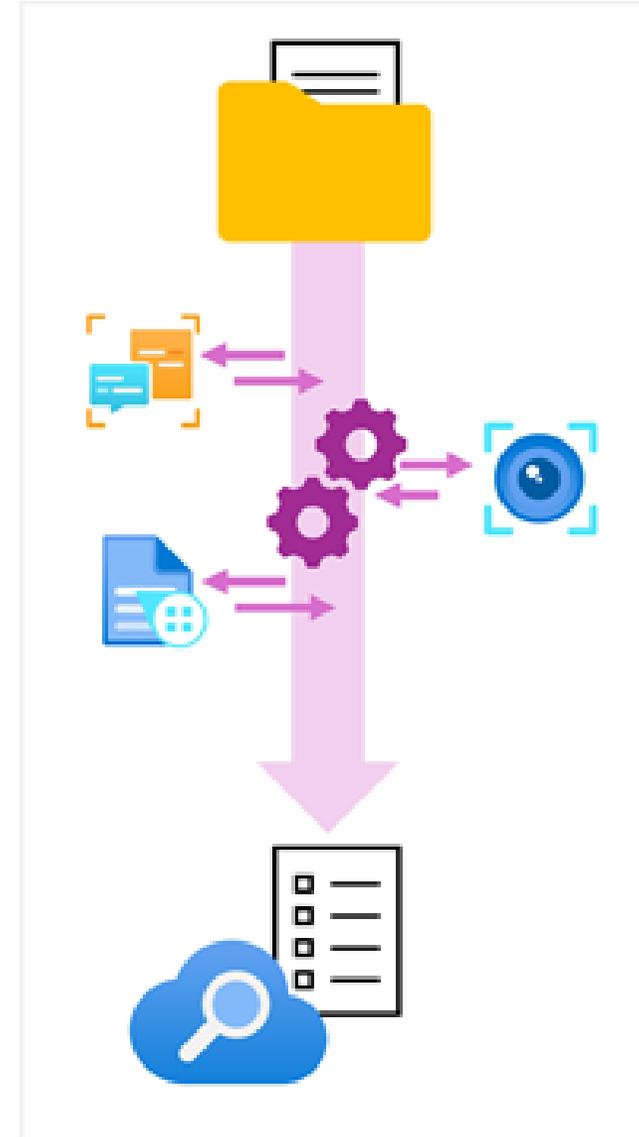
What is knowledge mining?

Organizations have a lot of content:

- Data is *locked away* in documents, PDFs, hand-written notes, etc.
- Time consuming and labour intensive to find data.
- Knowledge mining finds insights – *at scale*.

Azure AI Search is Azure's AI-powered knowledge mining platform.

- Azure AI Search results contain only your data, and can include new insights powered by AI
- Offered as Platform as a Service (PaaS) solution, meaning Microsoft manages the infrastructure and availability

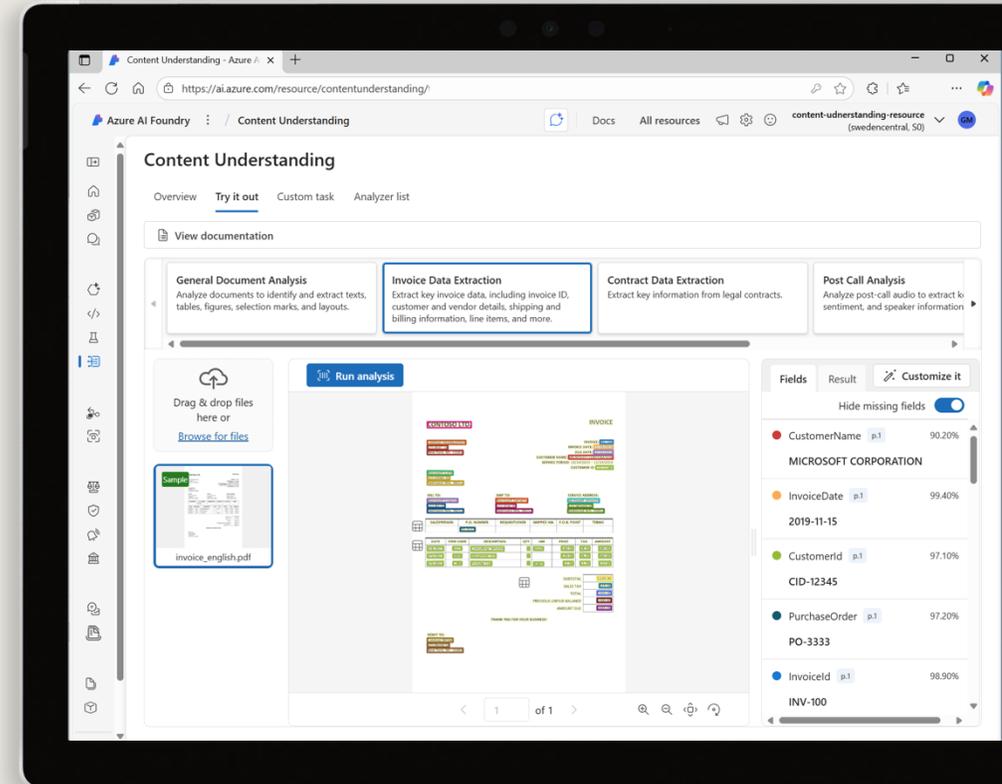


Extract data with Content Understanding in Azure AI Foundry portal

In this exercise, you'll use Azure AI Content Understanding to extract information from an invoice.

Start the exercise at:

<https://go.microsoft.com/fwlink/?linkid=2320420>



Q&A Session

