

Azure AI | Hands-On

Version 25.04

Ahmed Abdelwahed

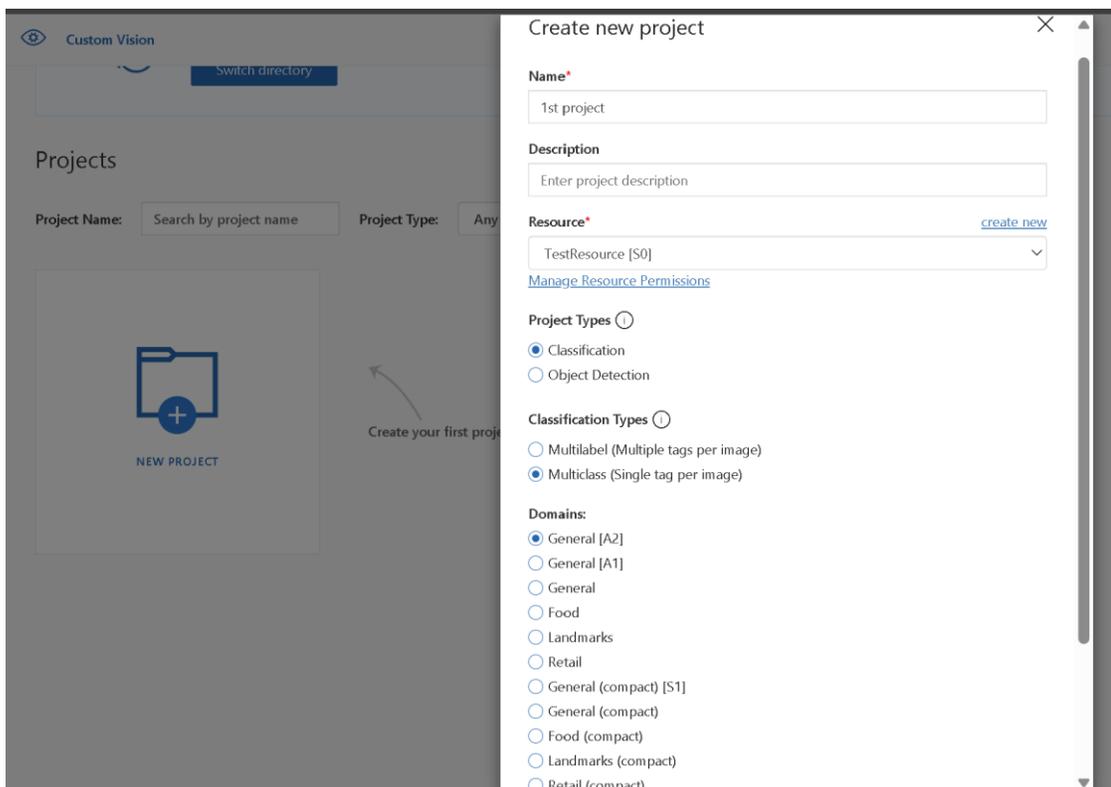
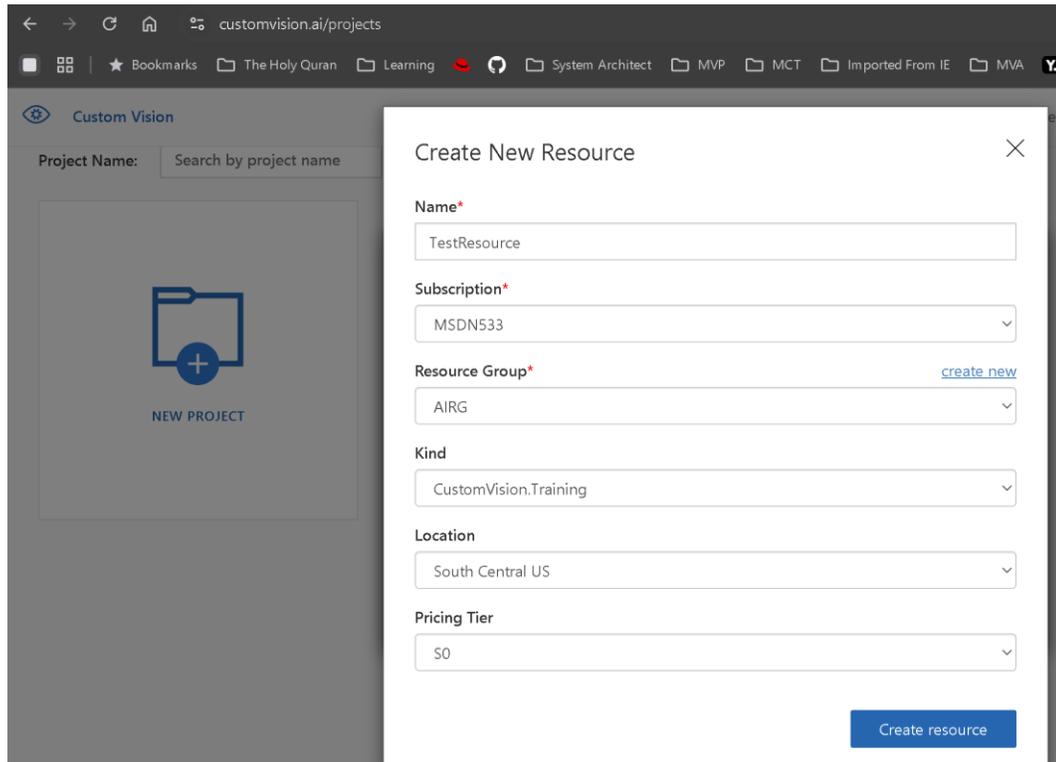
ahmed@abdelwahed.me

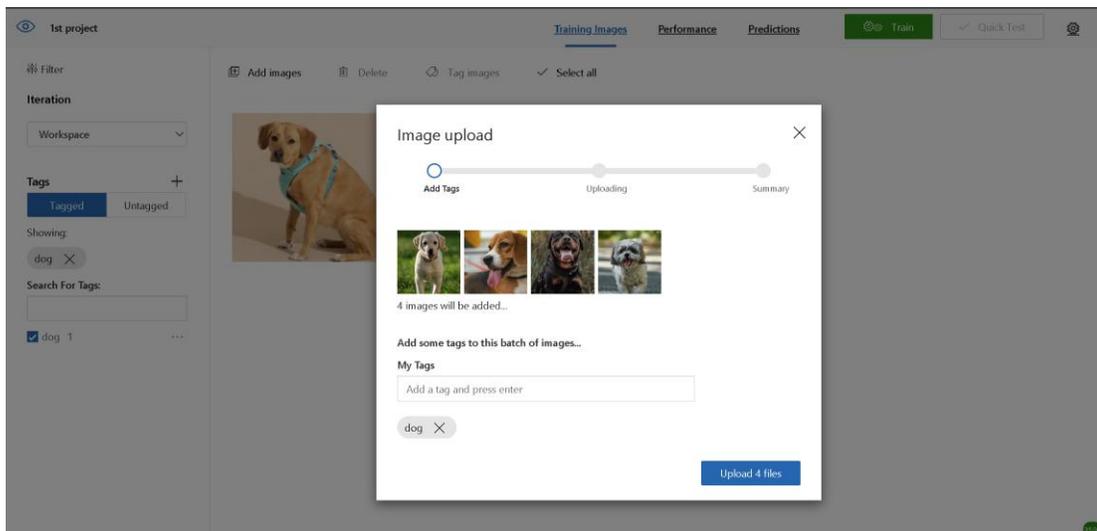
www.abdelwahed.me

[LinkedIn](#)

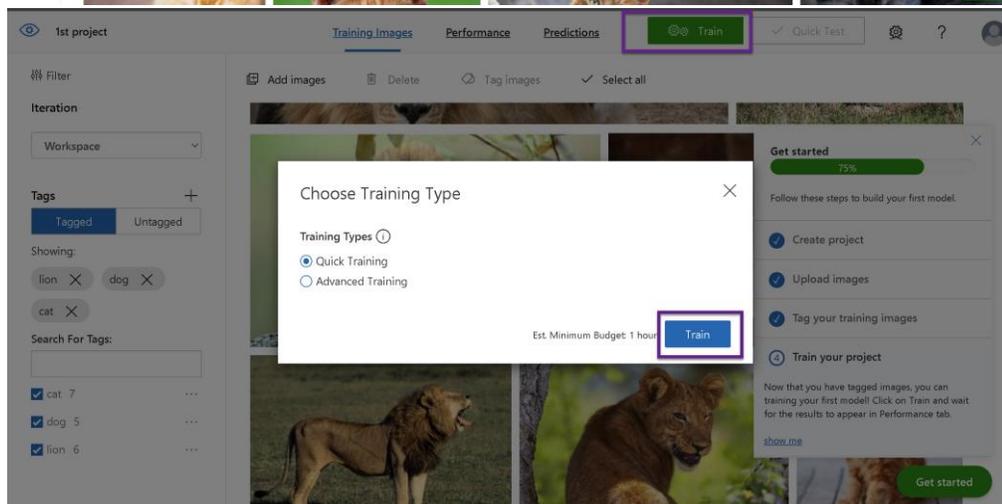
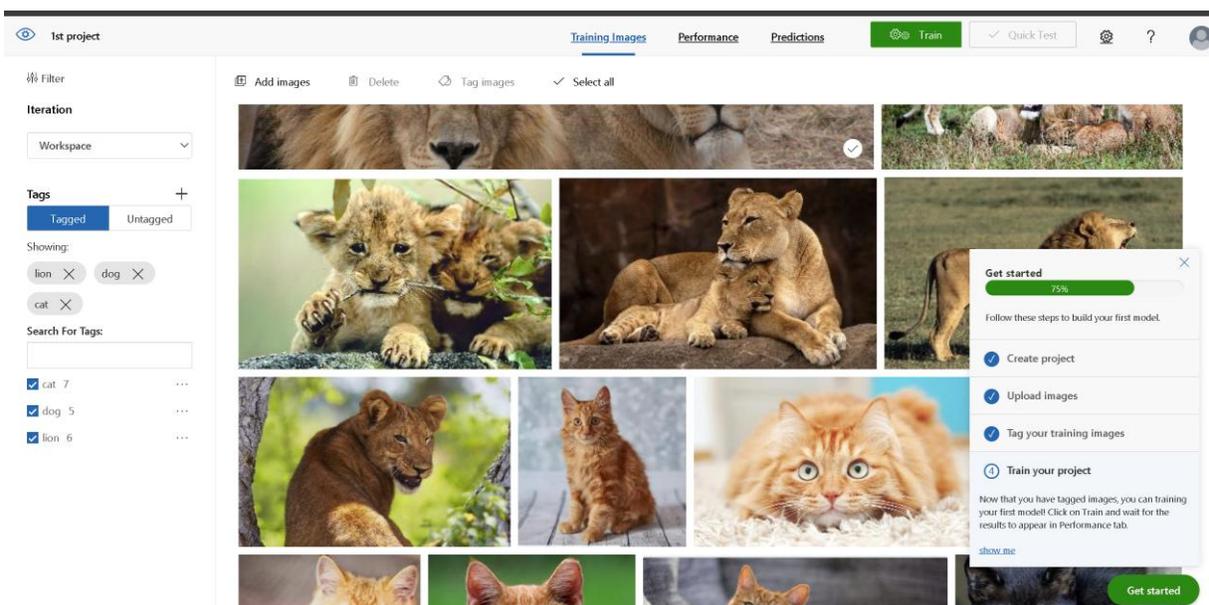
[GitHub](#)

Image Classification with Custom Vision





The more diverse and clean your data is, the better the model performs.



1st project

Training Images Performance Predictions Train Quick Test

Iterations

Probability Threshold: 50%

Iteration 1
Training...

Prediction URL Delete Export

Iteration 1
Training...

Training...

Last checked: 4/11/2025, 11:08:07 AM

Get started 100%

Follow these steps to build your first model.

- ✓ Create project
- ✓ Upload images
- ✓ Tag your training images
- ✓ Train your project

WHAT'S NEXT?

After training is completed, you can:

[Quick Test your model!](#)

Continue improving your model by providing images with different angles, backgrounds, object size, groups of photos, and other variants.

When you're ready to use your model:

Get started

1st project

Training Images Performance Predictions Train Quick Test

Iterations

Probability Threshold: 50%

Iteration 1
Trained : 2 minutes ago with General [A2] domain

Publish Prediction URL Delete Export

Precision 100.0%

Recall 100.0%

AP 100.0%

Performance Per Tag

Tag	Precision	Recall	A.P.	Image count
lion	100.0%	100.0%	100.0%	6
dog	100.0%	100.0%	100.0%	5
cat	100.0%	100.0%	100.0%	7

100%

1st project

Quick Test

Image URL

Enter image URL

or

Browse local files

File formats accepted: jpg, png, bmp
File size should not exceed: 4mb

Using model trained in

Iteration

Iteration 1

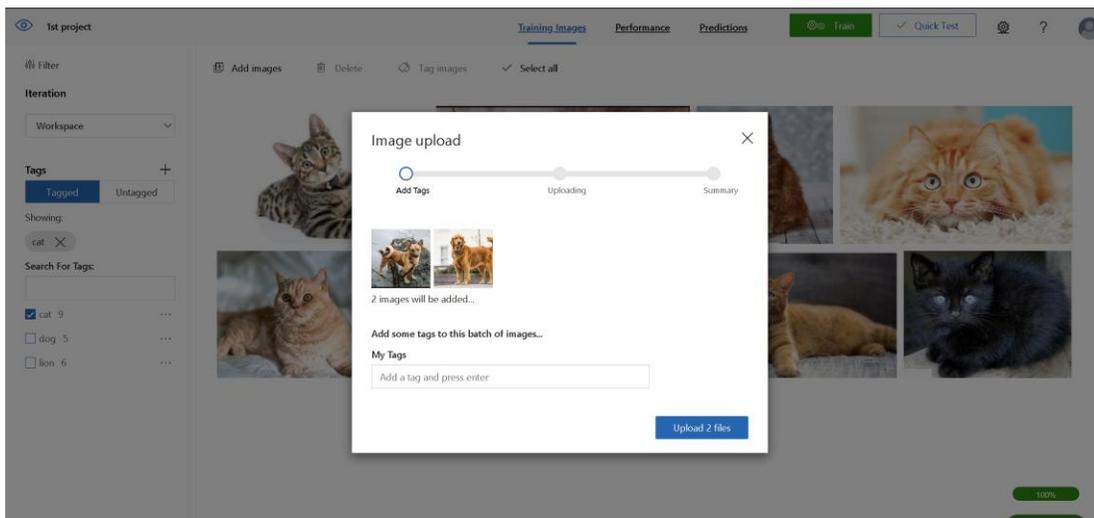
Predictions

Tag	Probability
cat	75.5%
dog	24%
lion	0.3%

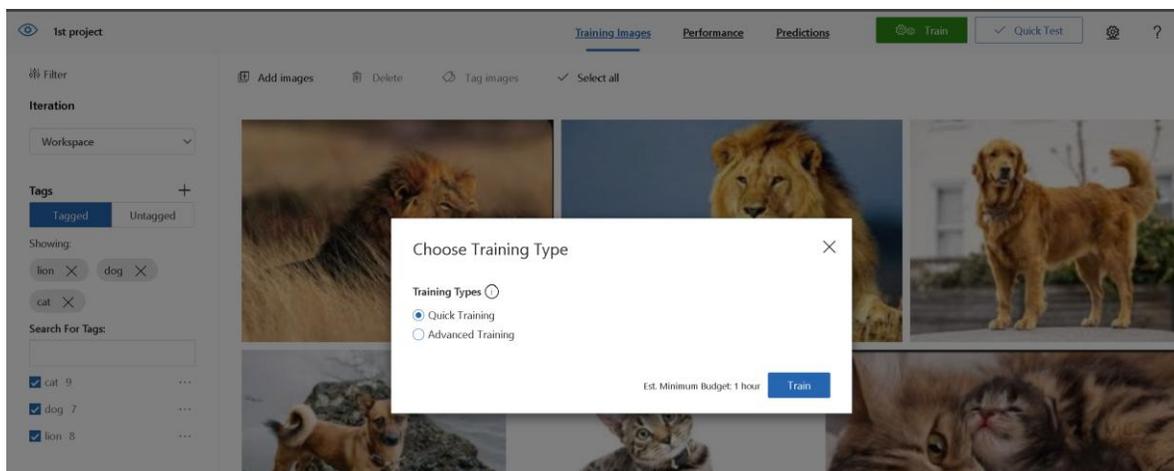
100%

Get started

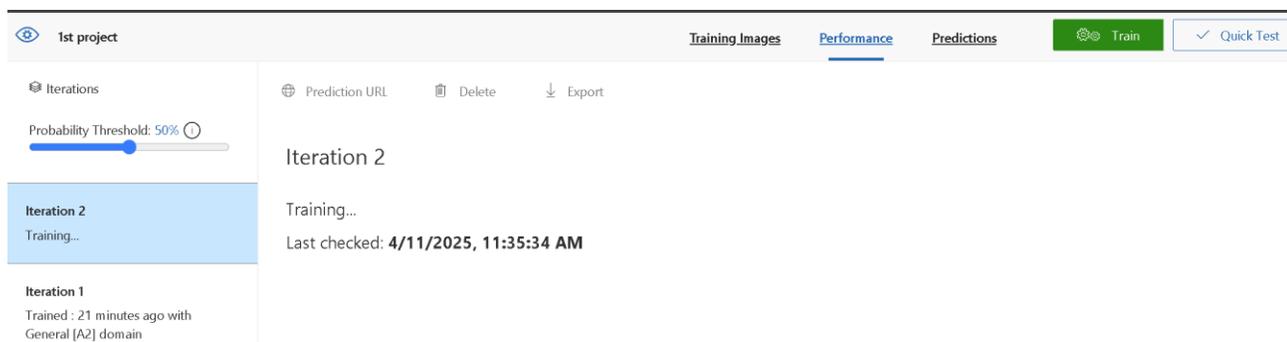
Add more images



And retrain your model



Creating new iteration



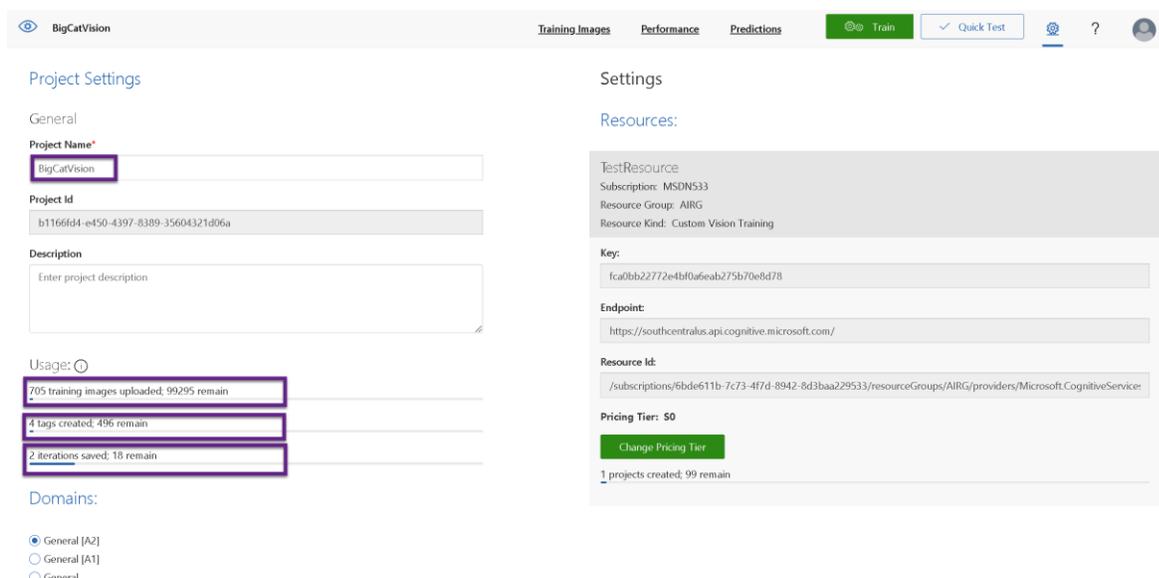
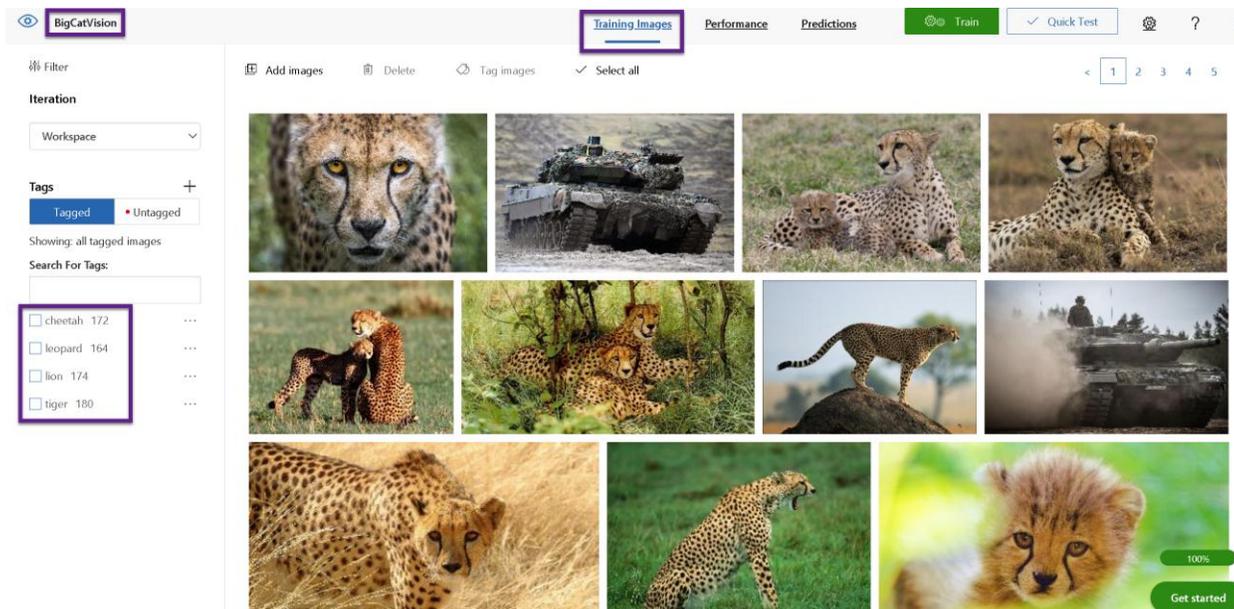
iteration is a version of your trained model.

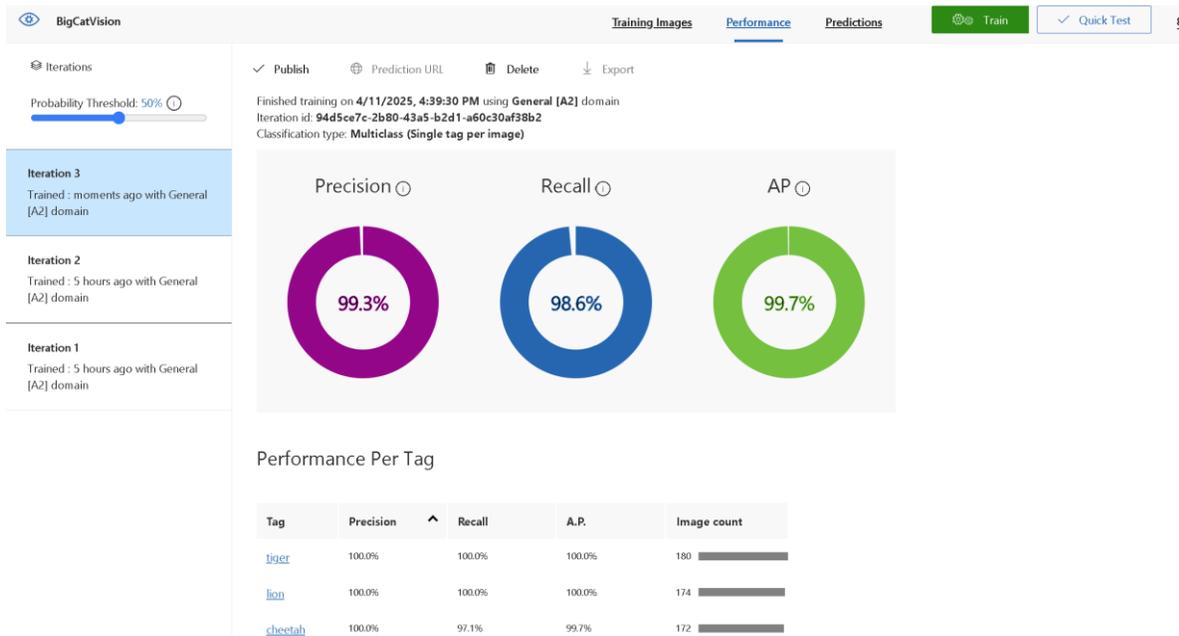
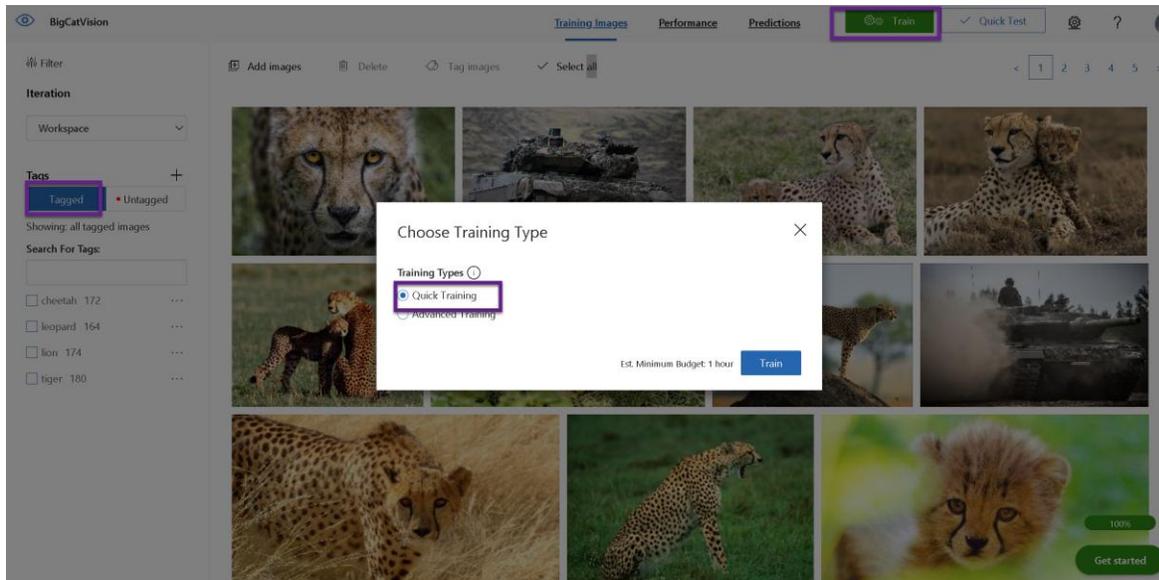
Classify Big Cats

Create new project and download big cat dataset from Kaggle to use it in the classifications

Image Datasets for Training:

1. **Kaggle Datasets**
2. Google Dataset Search
3. Open Images Dataset by Google
4. ImageNet





Test Prediction

Quick Test

Image URL

Enter Image URL →

or

Browse local files

File formats accepted: jpg, png, bmp
File size should not exceed: 4mb

Using model trained in

Iteration
Iteration 3

Predictions

Tag	Probability
cheetah	37.1%
tiger	35.5%
leopard	21.6%
lion	5.6%

Quick Test

Image URL

Enter Image URL →

or

Browse local files

File formats accepted: jpg, png, bmp
File size should not exceed: 4mb

Using model trained in

Iteration
Iteration 3

Predictions

Tag	Probability
tiger	56%
cheetah	31.1%
lion	6.9%
leopard	5.8%

Quick Test

Image URL

Enter Image URL →

or

Browse local files

File formats accepted: jpg, png, bmp
File size should not exceed: 4mb

Using model trained in

Iteration
Iteration 3

Predictions

Tag	Probability
lion	99.9%
tiger	0%
cheetah	0%
leopard	0%

Use a Custom Vision model for prediction

First, create an Azure Custom Vision Prediction resource in the same region you selected earlier. Then, go to customvision.ai, and during the publishing step of your iteration, link it to this prediction resource. This will provide you with both the Prediction URL and Prediction Key required in the Python code that follows.

Home > Create a resource > Marketplace > Custom Vision >

Create Custom Vision

optimize manufacturing processes, accelerate digital marketing campaigns -- and more. No machine learning expertise is required.

[Learn more](#)

Create options * Both Prediction Training

Project Details

Subscription *

Resource group * [Create new](#)

Instance Details

Region

Name *

[Previous](#) [Next](#) [Review + create](#)

BigCatVision Training Images **Performance** Predictions [Train](#) [Quick Test](#)

Iterations [Unpublish](#) [Prediction URL](#) [Delete](#) [Export](#)

Probability Threshold: 50%

Iteration 3 PUBLISHED
Trained : 2 hours ago with General [A2] domain

Finished training on 4/11/2025, 4:39:30 PM using General [A2] domain
Iteration id: 94d5ce7c-2b80-43a5-b2d1-a60c30af38b2
Classification type: Multiclass (Single tag per image)
Published as: Iteration3

How to use the Prediction API

If you have an image URL:

```
https://[resource-id].e.com/customvision/v3.0/Prediction/b
Set Prediction-Key Header to: [key]
Set Content-Type Header to: application/json
Set Body to: [{"url": "https://example.com/image.png"}]
```

If you have an image file:

```
https://[resource-id].azure.com/customvision/v3.0/Prediction/b
Set Prediction-Key Header to: [key]
Set Content-Type Header to: application/octet-stream
Set Body to: <image file>
```

[Got it!](#)

9.7%

100%

[Get started](#)

Change the prediction URL ending from /url to /image in Python code.

Create web app for bigcat prediction using streamlit on Linux

Step 1: Install Python

Rocky Linux comes with Python, but let's make sure you have a good version (3.8+ is great).

```
sudo dnf install python3 python3-pip -y
```

Step 2: Create Virtual Environment (Optional but Recommended)

```
python3 -m venv streamlit-env
source streamlit-env/bin/activate
```

Step 3: Install Streamlit

```
pip install streamlit
```

🔧 Step 4: Create Your Streamlit App File

```
import streamlit as st
import requests

st.title("🐾 BigCatVision - Identify Your Big Cat")
uploaded_file = st.file_uploader("Upload an image", type=["jpg", "jpeg", "png"])
PREDICTION_URL =
"https://bigcatprediction.cognitiveservices.azure.com/customvision/v3.0/Prediction/bxxxxxxxxxxxxx
xxxxx6a/classify/iterations/Iteration3/image"
PREDICTION_KEY = "xxxxxxxxxxxxxxxxx6C2DKmb3XvcKlQFaEiexxxxxxxxxxxxxxxxxxxxxx"
headers = {
    "Prediction-Key": PREDICTION_KEY,
    "Content-Type": "application/octet-stream"
}

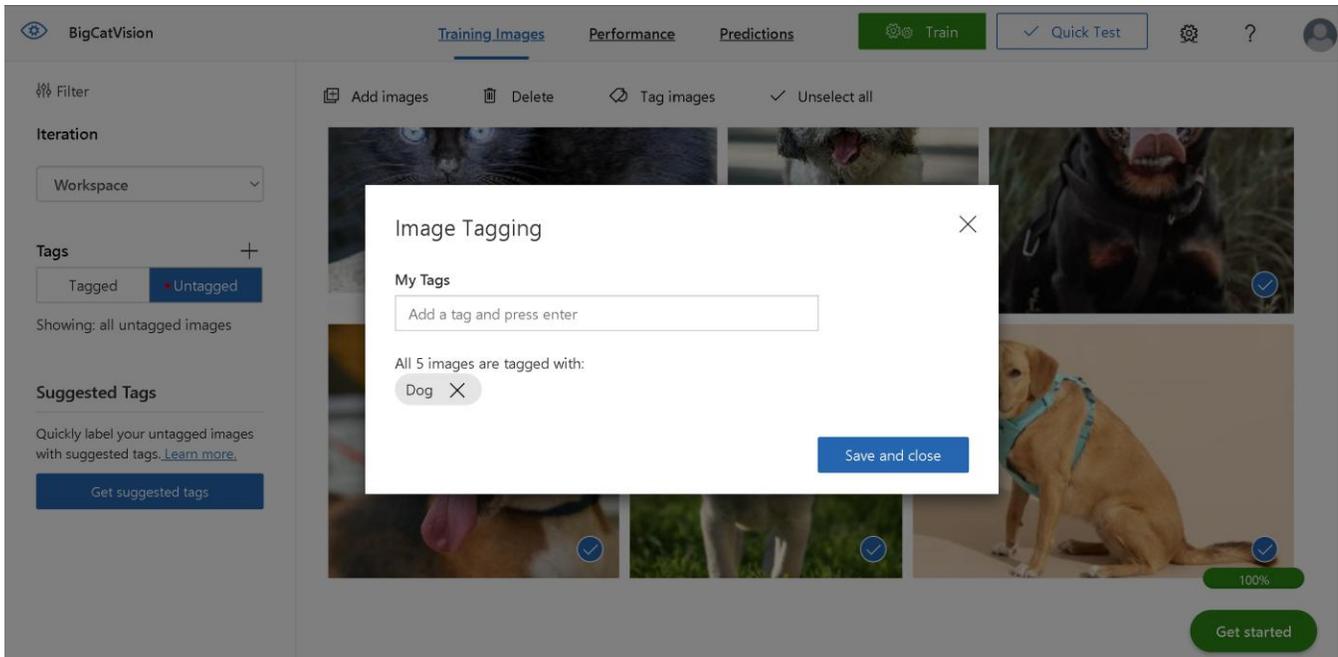
if uploaded_file:
    image_bytes = uploaded_file.read()
    st.image(image_bytes, caption="Uploaded Image", use_column_width=True)
    if st.button("🔍 Predict"):
        st.info("Classifying the image...")
        response = requests.post(PREDICTION_URL, headers=headers, data=image_bytes)
        try:
            result = response.json()
            predictions = result.get("predictions", [])
            if predictions:
                st.subheader("🔗 Top Predictions:")
                for pred in predictions[:3]: # Show top 3
                    tag = pred["tagName"]
                    confidence = pred["probability"] * 100
                    st.write(f"***{tag.title()}***: {confidence:.2f}%")
            else:
                st.warning("No predictions returned.")
        except Exception as e:
            st.error(f"❌ Prediction failed: {e}")
```

Step 5: Run Your App

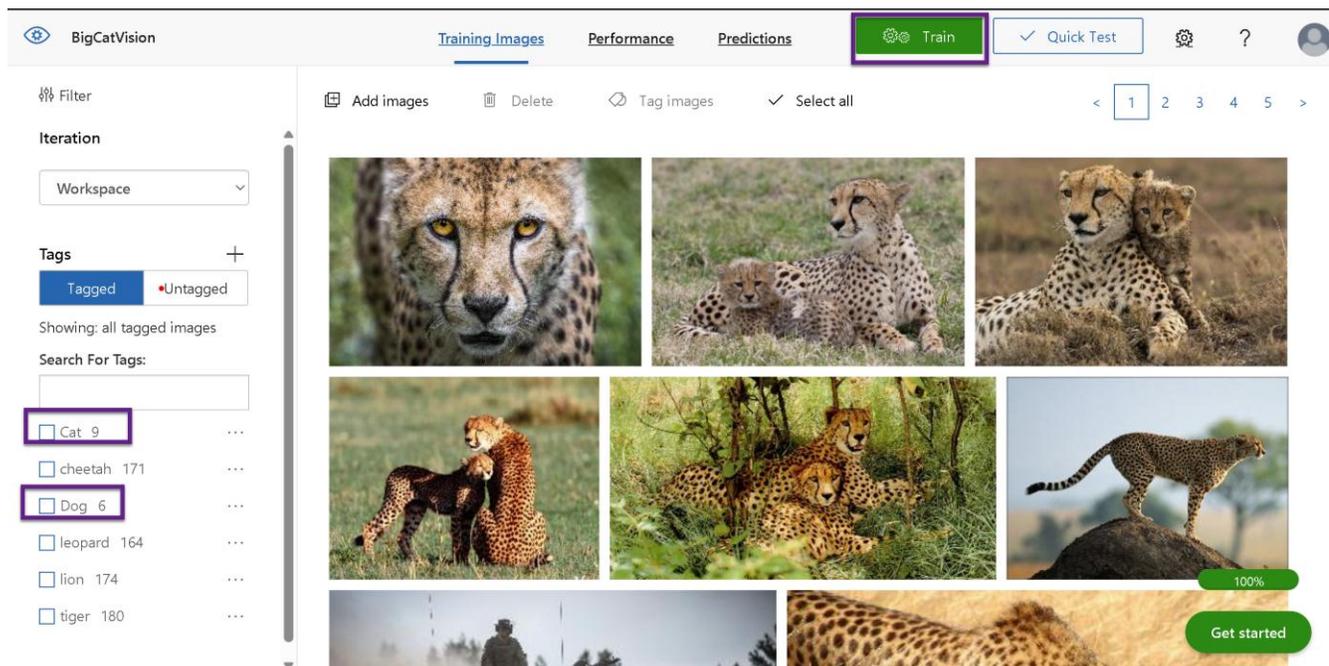
```
streamlit run app.py --server.address=0.0.0.0
```

Then access via http://<Rocky_Linux_IP>:8501

Update Images and link them to new tag



We have added two new tags and need to retrain the model so it can predict images with the new tags. This will produce a new iteration that includes the prediction URL and prediction key.



Training started, and a new iteration has been established.

BigCatVision

Training Images Performance Predictions Train Quick Test

Iterations

Probability Threshold: 50%

Iteration 4
Training...

Iteration 3
Trained : 3 hours ago with General [A2] domain

Iteration 2
Trained : 8 hours ago with General [A2] domain

Iteration 1
Trained : 8 hours ago with General [A2] domain

Prediction URL Delete Export

Iteration 4

Training...

Last checked: 4/11/2025, 7:25:24 PM

100%

Get started

It's time to publish and obtain the data required to use in python.

BigCatVision

Training Images Performance Predictions Train Quick Test

Iterations

Probability Threshold: 50%

Iteration 4
Trained : moments ago with General [A2] domain

Iteration 3
Trained : 3 hours ago with General [A2] domain

Iteration 2
Trained : 8 hours ago with General [A2] domain

Iteration 1
Trained : 8 hours ago with General [A2] domain

✓ Publish Prediction URL Delete Export

Iteration 4

Finished training on 4/11/2025, 7:32:53 PM using General [A2] domain
Iteration id: a2a26c0d-c1cc-40b8-9761-4bfb668ee500
Classification type: Multiclass (Single tag per image)

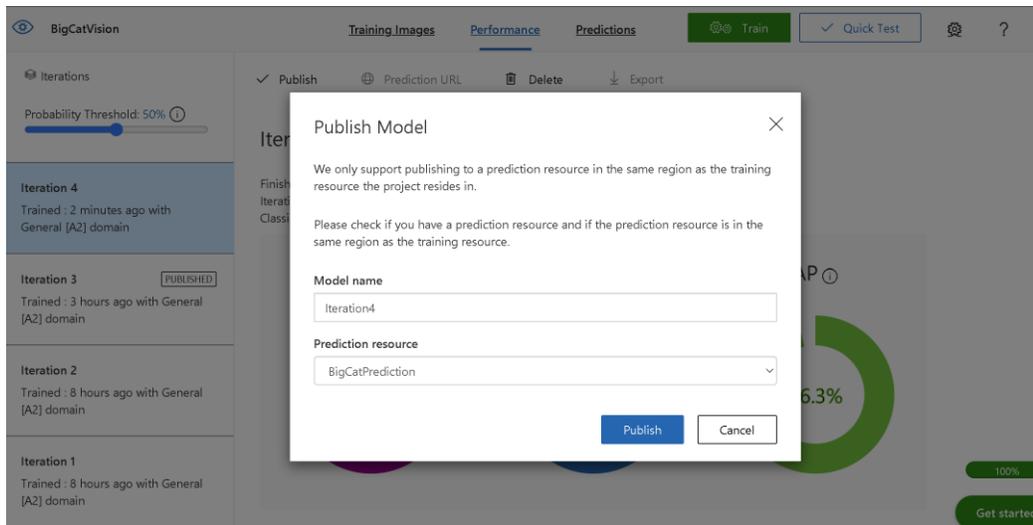
Precision 96.4%

Recall 95.7%

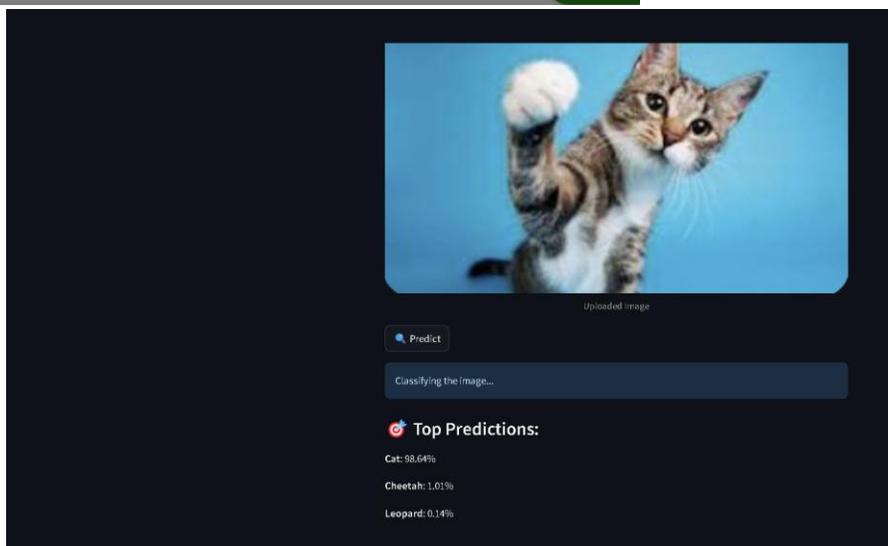
AP 96.3%

100%

Get started



It is now possible to predict the cats.

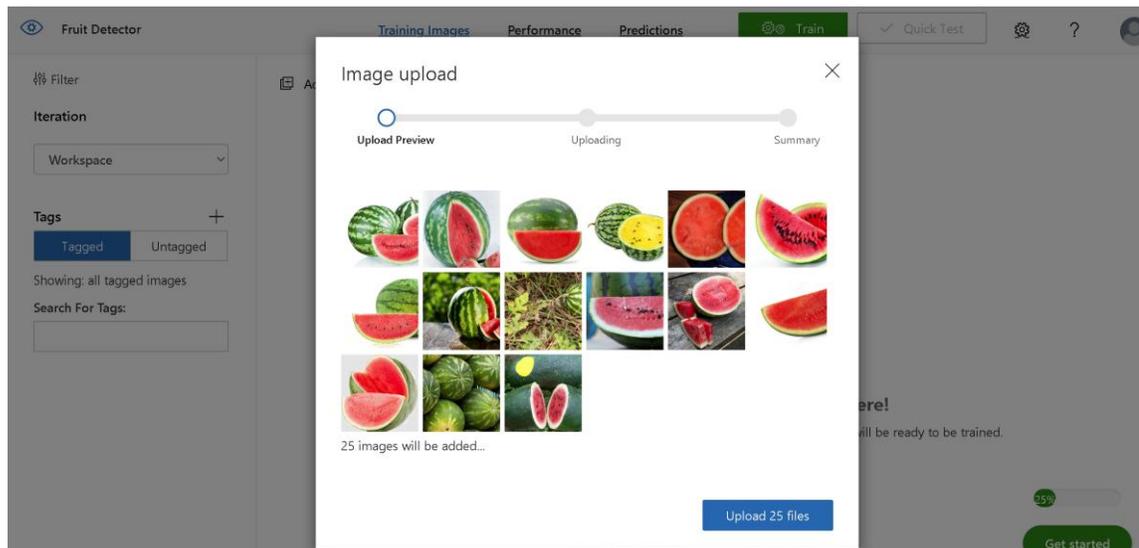
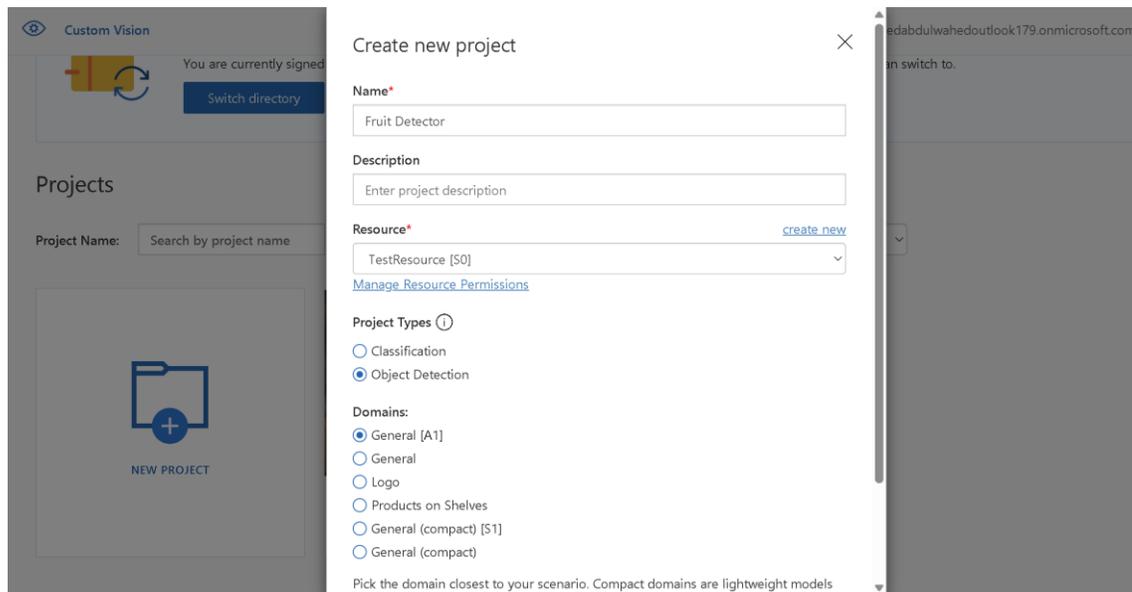


Object Detection with Custom Vision

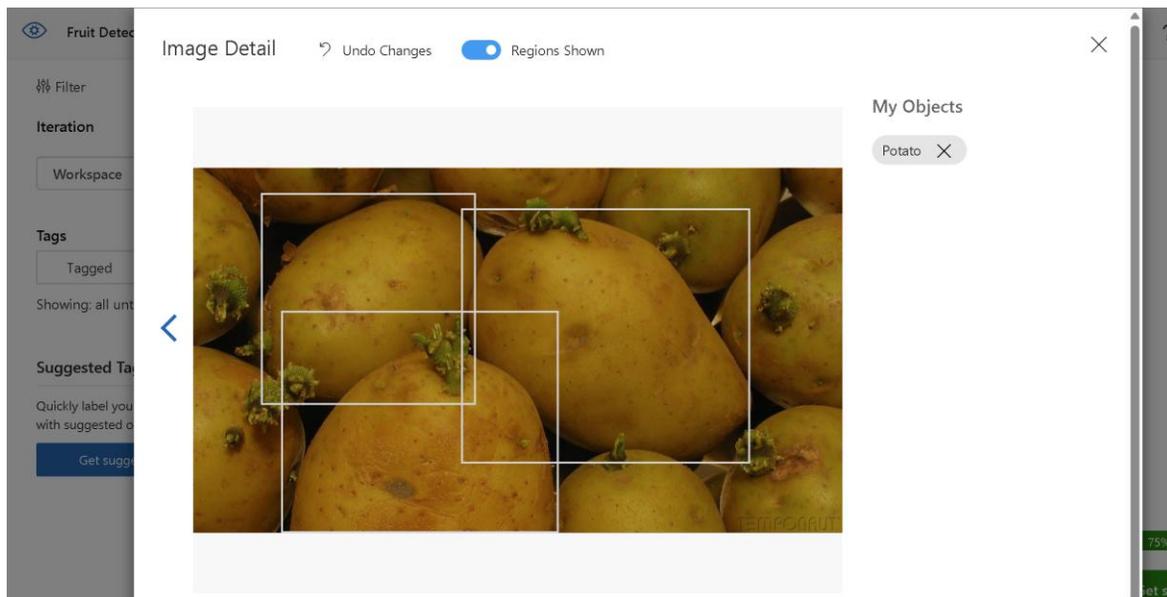
Unlike image classification (which labels the entire image), **object detection**:

- Detects **what** objects are in the image
- Locates **where** each object is (bounding boxes)

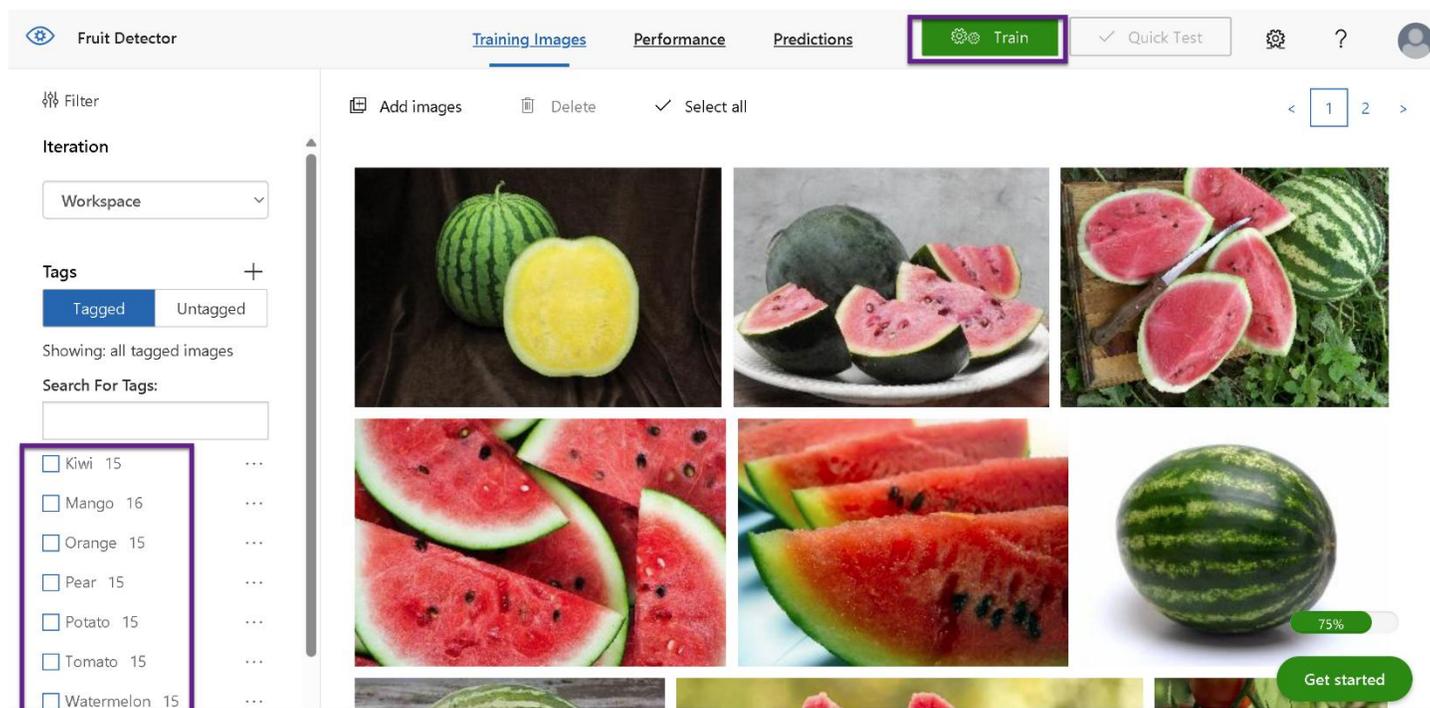
Instead of just saying “this is a lion,” object detection says: “There’s a lion at coordinates (x1, y1, x2, y2)” in the image.

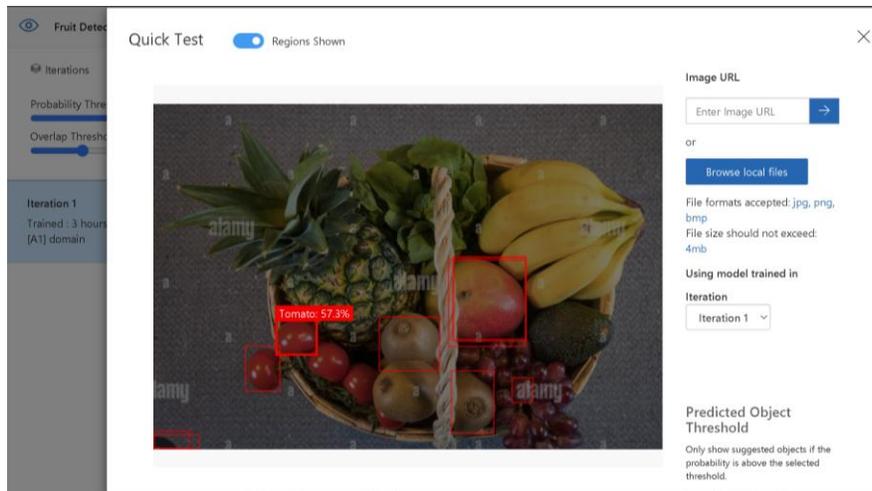
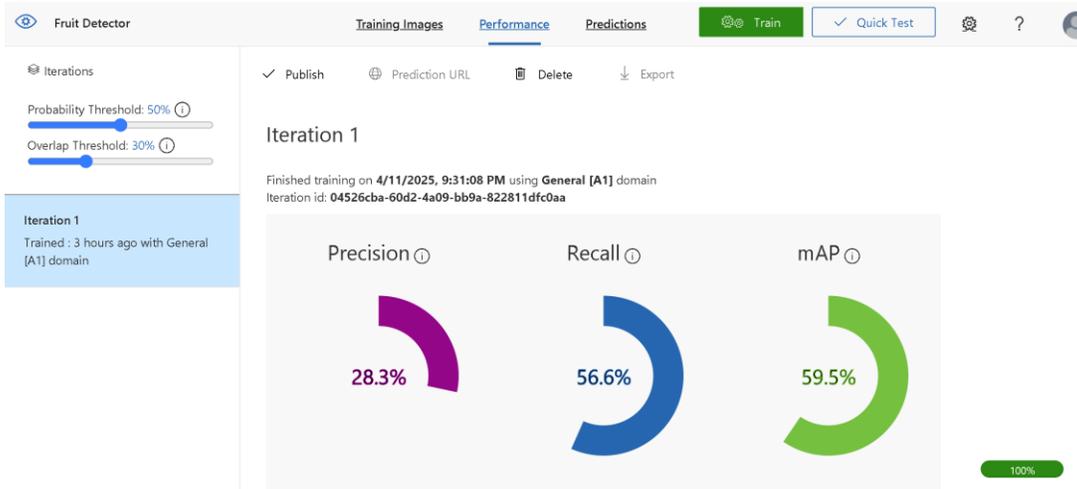


Tag each object in the image.



Added different fruits and tagged them all (to start train you have to add 15 image for each tag)





now we can publish with streamlit

How to use the Prediction API

If you have an image URL:

```
https://[subscription-id].initiveservices.azure.com/customvision/v3.0/Prediction/3  
Set Prediction-Key Header to: [key]  
Set Content-Type Header to: application/json  
Set Body to: {"Url": "https://example.com/image.png"}
```

If you have an image file:

```
https://[subscription-id].com/customvision/v3.0/Prediction/3  
Set Prediction-Key Header to: [key]  
Set Content-Type Header to: application/octet-stream  
Set Body to: <image file>
```

Got it!

Deploy Object Detection App with Streamlit on Linux

Follow the previous steps to build an app with Streamlit using this Python code.

```
import streamlit as st
import requests
from PIL import Image, ImageDraw
import io

st.title("🍌 Fruit Detector - Object Detection with Azure Custom Vision")
uploaded_file = st.file_uploader("Upload an image (JPG or PNG)", type=["jpg", "jpeg", "png"])
PREDICTION_URL =
"https://bigcatprediction.cognitiveservices.azure.com/customvision/v3.0/Prediction/39xxxx
xxxxxxxxxxxxxxxxxxxxxxxx/detect/iterations/Iteration1/image"
PREDICTION_KEY = "xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxvcklQFaEieGcbisaKsH1AW37UemBTzH5JQQJ99BDACLA"
headers = {
    "Prediction-Key": PREDICTION_KEY,
    "Content-Type": "application/octet-stream"
}
def draw_boxes(image, predictions):
    img = image.copy()
    draw = ImageDraw.Draw(img)
    w, h = img.size
    for pred in predictions:
        tag = pred['tagName']
        prob = pred['probability']
        bbox = pred['boundingBox']
        left = bbox['left'] * w
        top = bbox['top'] * h
        width = bbox['width'] * w
        height = bbox['height'] * h
        draw.rectangle([left, top, left+width, top+height], outline='red', width=3)
        draw.text((left, top), f"{tag} ({prob*100:.1f}%)", fill='white')
    return img
if uploaded_file:
    image_data = uploaded_file.read()
    image = Image.open(io.BytesIO(image_data))
    st.image(image, caption="Uploaded Image", use_column_width=True)

    if st.button("🔍 Detect"):
        st.info("Sending image to Azure Object Detection model...")
        try:
            response = requests.post(PREDICTION_URL, headers=headers, data=image_data)
            result = response.json()
            predictions = result.get("predictions", [])
```


1. Create the Azure Custom Vision Model

2. Create Your Streamlit App (locally)

1. On your local machine or VM:
`mkdir fruit-detector`
`cd fruit-detector`
2. Create a file app.py and paste the full app code (with camera, counter, snapshot features)
3. Create a requirements.txt:
`streamlit`
`streamlit-webrtc`
`pillow`
`requests`
`av`

3. Push to GitHub

1. Initialize Git:
`git init`
`git add .`
`git commit -m "initial commit with fruit detector app"`

2. Create a repo on GitHub:
👉 <https://github.com/new> → name: fruit-detector

3. Link to GitHub:

```
git remote add origin https://github.com/YOUR_USERNAME/fruit-detector.git
```

```
git push -u origin master
```

💡 If using HTTPS and GitHub blocks password login, use a **Personal Access Token** as password.

4. Deploy on Streamlit Cloud

1. Go to 👉 <https://share.streamlit.io>
2. Log in with GitHub and click "**New App**"
3. Select:
 - Repository: YOUR_USERNAME/fruit-detector
 - Branch: master
 - File: app.py
4. Click **Deploy**

5. App is Live!

You'll get a public URL like:

`https://your-username-fruit-detector.streamlit.app`

Open it from mobile → browser will ask to allow webcam
Snapshots will save locally in snapshots/ folder on server



← Back

What would you like to do?



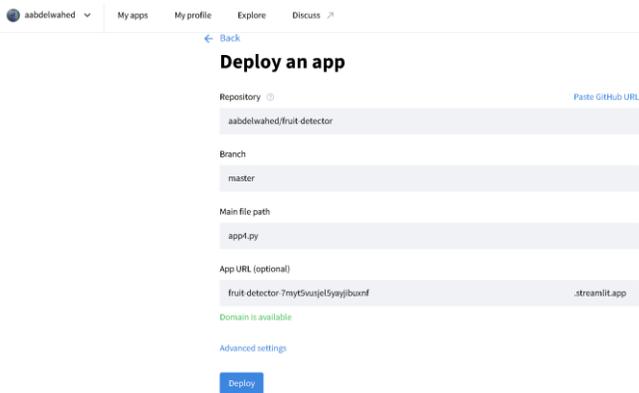
Deploy a public app from GitHub
My code is ready on a GitHub repo, and it is totally awesome.
[Deploy now](#)



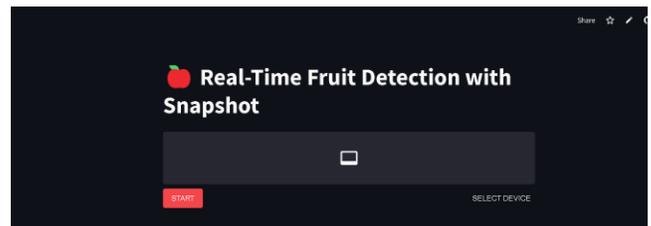
Deploy a public app from a template
I want to see what kind of amazing concoctions you have for me.
[Check out templates](#)



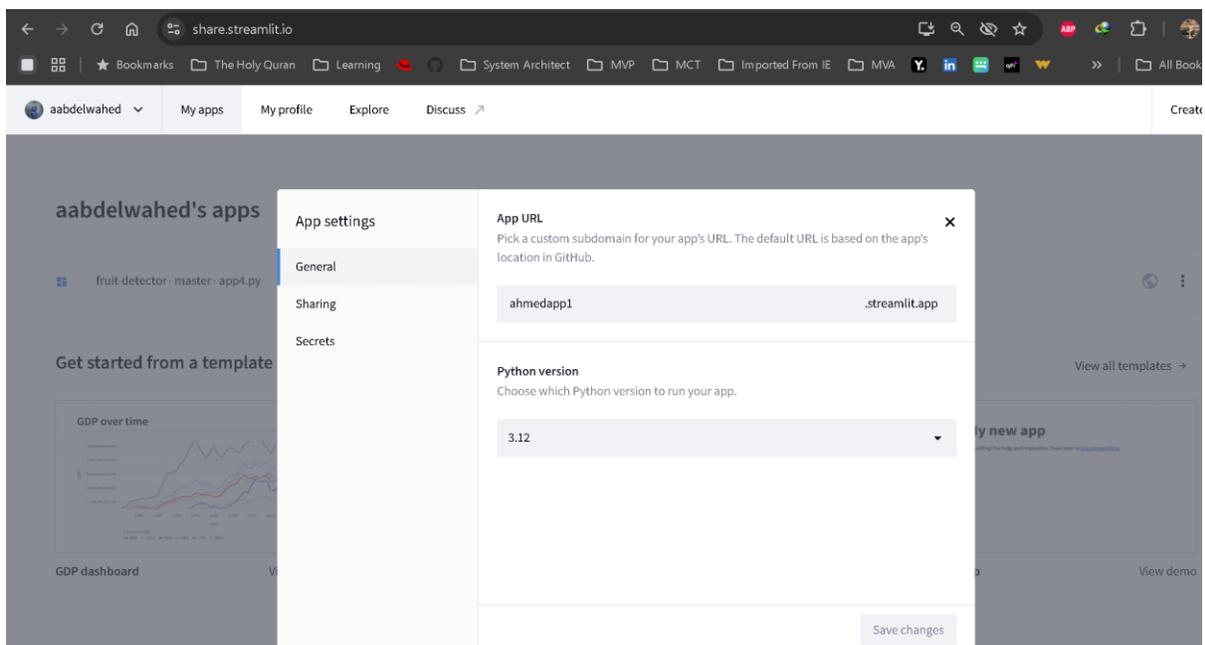
Deploy a private app in Snowflake
I want unlimited enterprise-grade apps, with the security of Snowflake.
[Start trial →](#)



<https://ahmedapp1.streamlit.app/>



You can adjust app settings, including your public name.



Computer Vision

Microsoft Azure Search resources, services, and docs (G+)

Home > Create Computer Vision

Basics Network Identity Tags Review + create

Boost content discoverability, accelerate text extraction, and create products that more people can use by embedding vision capabilities in your apps. Use visual data processing to label content (from objects to concepts), extract printed and handwritten text, recognize familiar subjects like brands and landmarks, and moderate content. No machine learning expertise is required.

[Learn more](#)

Project Details

Subscription *

Resource group * [Create new](#)

Instance Details

Region

Name *

Pricing tier *

[View full pricing details](#)

Responsible AI Notice

[Previous](#) [Next](#) [Review + create](#)

Microsoft Azure Search resources, services, and docs (G+) Copilot

Home > Azure AI services | Computer vision >

All-ncstance1

Computer vision

Search Delete

Overview Get Started Monitoring

- Activity log
- Access control (IAM)
- Tags
- Diagnose and solve problems
- Resource visualizer
- Resource Management
- Security
- Monitoring
- Automation
- Help

Get started with your resource in Vision Studio

Try out all Computer Vision features and build your own custom models

[Go to Vision Studio](#)

Keys and endpoint

These keys are used to access your Azure AI services API. Do not share your keys. Store them securely—for example, using Azure Key Vault. We also recommend regenerating these keys regularly. Only one key is necessary to make an API call. When regenerating the first key, you can use the second key for continued access to the service.

[Show Keys](#)

KEY 1

KEY 2

Location/Region



Try Azure AI Content Understanding in AI Foundry

Turn unstructured documents, images, video, and audio into structured data with the new Generative AI-powered Content Understanding service, built by the same team that created the Vision service. ([Learn more](#))

Visit AI Foundry

Vision Studio



Video Retrieval and Summary Preview

Generate a brief summary of the main points shown in video. Locate specific keywords and jump to the relevant section.

Try it out



Search photos with image retrieval

Retrieve specific moments within your photo album. For example, you can search for: a wedding you attended last summer, your pet, or your favorite city.

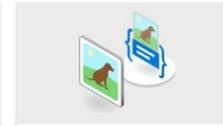
Try it out



Add dense captions to images

Generate human-readable captions for all important objects detected in your image.

Try it out



Add captions to images

Generate a human-readable sentence that describes the content of an image.

Try it out



Detect common objects in images

Recognize the location of objects of interest in an image and return them as labels.



Extract common tags from images

Use an AI model to automatically assign one or more labels to an image.



Create smart-cropped images

Create cropped image thumbnails based on the most important features in an image.



Detect faces in an image

Detect the location of one or more human faces in an image, along with whether each

Vision Studio > Extract text from images

[View documentation](#) [View SDK reference](#) [Use the REST API](#) [View supported languages](#)

Try it out

I acknowledge that this demo will incur usage to my account.

Use one of your own files or choose from a sample below.

Drag and drop a file here or
Browse for a file
or
Take a photo



Select an Azure resource

To access Vision Studio, you need a Cognitive Services resource. You can use an existing resource or create a new one. You can change the resource you use at any time. [Learn more about resources in Azure](#)

Azure directory

ahmedabdulwahedoutlook179.onmicrosoft.com

Subscription *

MSDN533

Azure Resources * ⓘ

All-ntance1

Create a new resource

Confirm

Cancel

Step-by-Step publish setup on Ubuntu on Azure

Step 1: Install Required Packages

SSH into your Ubuntu VM and run:

```
sudo apt update && sudo apt install -y python3-pip python3-venv  
pip3 install flask requests
```

Step 2: Project Structure

```
mkdir ocr_web_app  
cd ocr_web_app  
python3 -m venv venv  
source venv/bin/activate  
pip install flask requests  
mkdir templates uploads  
touch app.py templates/index.html
```

Step 3: app.py – The Flask Web App

Paste the following into app.py:

```
from flask import Flask, request, render_template  
import os  
import requests  
app = Flask(__name__)  
UPLOAD_FOLDER = 'uploads'  
app.config['UPLOAD_FOLDER'] = UPLOAD_FOLDER  
# Replace with your Azure credentials  
AZURE_KEY = "YOUR_AZURE_KEY"  
AZURE_ENDPOINT = "https://YOUR_REGION.api.cognitive.microsoft.com"  
OCR_URL = AZURE_ENDPOINT + "/vision/v3.2/ocr"  
@app.route("/", methods=["GET", "POST"])  
def upload_file():  
    extracted_text = ""  
    if request.method == "POST":  
        file = request.files['file']  
        if file:  
            filepath = os.path.join(app.config['UPLOAD_FOLDER'], file.filename)  
            file.save(filepath)  
            with open(filepath, 'rb') as image_data:  
                headers = {  
                    'Ocp-Apim-Subscription-Key': AZURE_KEY,  
                    'Content-Type': 'application/octet-stream'  
                }  
                response = requests.post(OCR_URL, headers=headers, data=image_data)  
                response.raise_for_status()  
                result = response.json()  
                for region in result.get("regions", []):
```

```
        for line in region["lines"]:
            extracted_text += " ".join([word["text"] for word in
line["words"]]) + "\n"
            return render_template("index.html", extracted_text=extracted_text)
if __name__ == "__main__":
    app.run(host="0.0.0.0", port=5000)
```

Step 4: templates/index.html

Paste this into templates/index.html:

```
<!DOCTYPE html>
<html>
<head>
    <title>OCR Web App</title>
</head>
<body>
    <h2>Upload Image for OCR</h2>
    <form method="POST" enctype="multipart/form-data">
        <input type="file" name="file" required>
        <button type="submit">Extract Text</button>
    </form>
    {% if extracted_text %}
    <h3>Extracted Text:</h3>
    <pre>{{ extracted_text }}</pre>
    {% endif %}
</body>
</html>
```

Step 5: Run the App

```
python3 app.py
```

Then open in your browser:

```
http://<your-azure-vm-public-ip>:5000
```

Step 6: Open Port 5000 in Azure

1. Go to [Azure Portal](#)
2. Find your VM → *Networking*
3. Add **inbound port rule**:
 - Port: 5000
 - Protocol: TCP
 - Action: Allow

Azure OpenAI

Microsoft Azure

Home > Azure AI services | Azure OpenAI >

Create Azure OpenAI ...

⚠ Changes on this step may reset later selections you have made. Review all options prior to deployment.

Azure OpenAI Service provides access to OpenAI's powerful language models, including all the latest OpenAI models. These models can be easily adapted to your specific tasks, including but not limited to content generation, summarization, image understanding, semantic search, and natural language to code translation. Top use cases include Call Centers, Virtual Assistants, Accessibility, Content Generation, and Code Development. The service also features the Assistants API, Fine Tuning capabilities and many ways to connect your data to the service for conversational experiences. The service can be scaled through Standard (tokens) and Provisioned (PTUs) deployment types.

[Learn more](#)

Project Details

Subscription * ⓘ

Resource group * ⓘ [Create new](#)

Instance Details

Region ⓘ

Name * ⓘ

Pricing tier * ⓘ

[View full pricing details](#)

Microsoft Azure Copilot ahmed_abdu AHMED TENANT

Home > Microsoft.CognitiveServicesOpenAI-20250507005434 | Overview >

AhmedAiRG

Resource group

Overview

Essentials

Subscription (move) : [MSDN533](#) Deployments : [2 Succeeded](#)

Subscription ID : 6bde611b-7c73-4f7d-8942-8d3baa229533 Location : East US

Tags (edit) : [Add tags](#)

Resources Recommendations

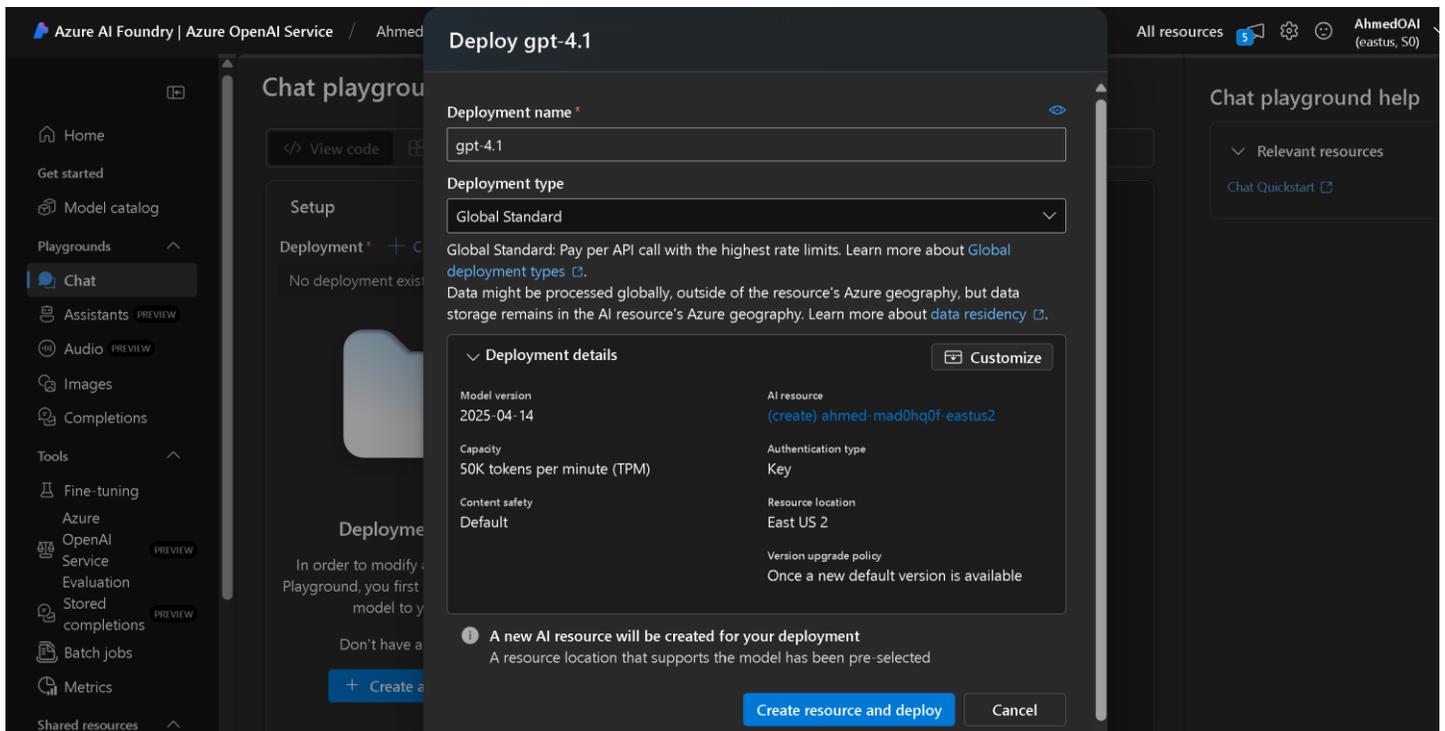
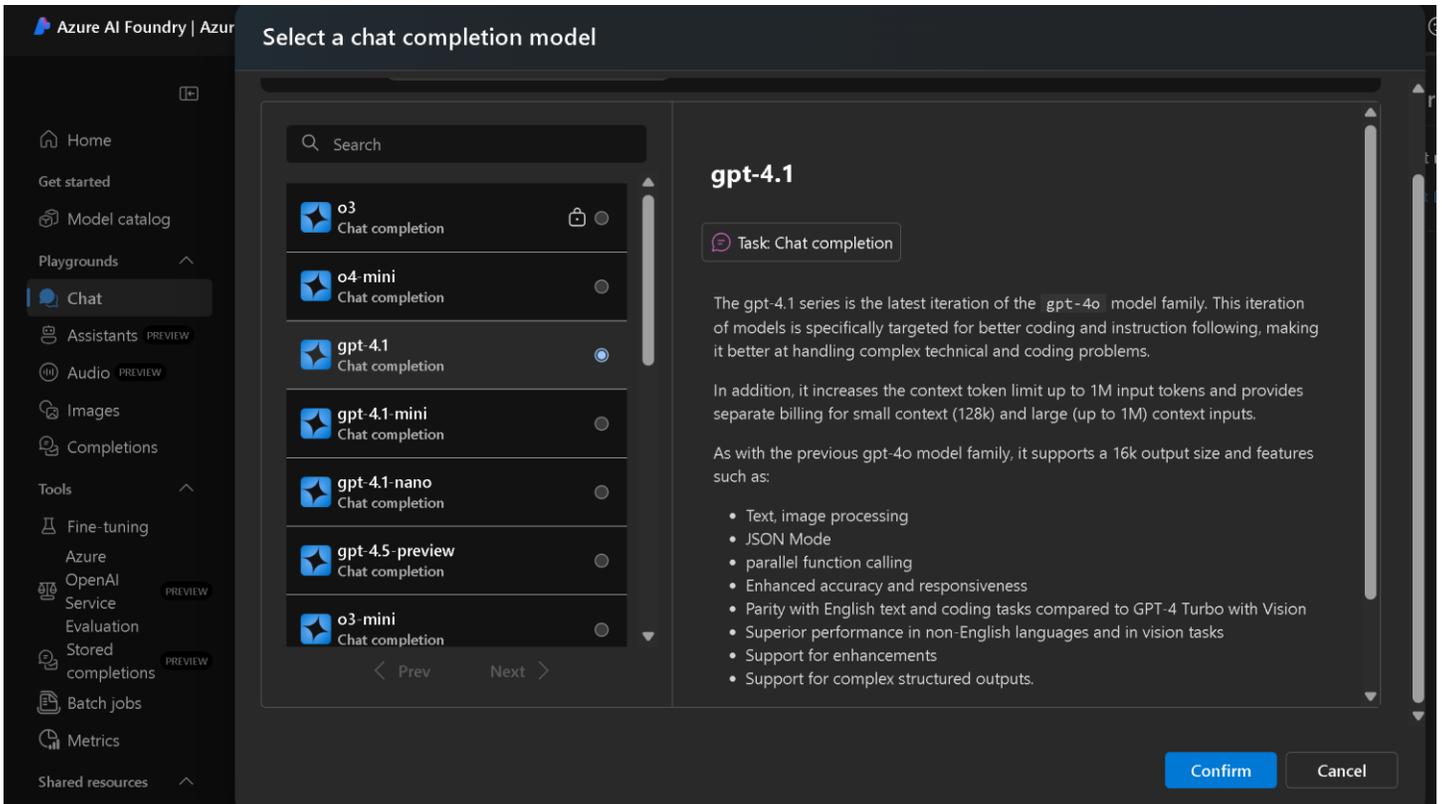
Filter for any field... Type equals all Location equals all Add filter

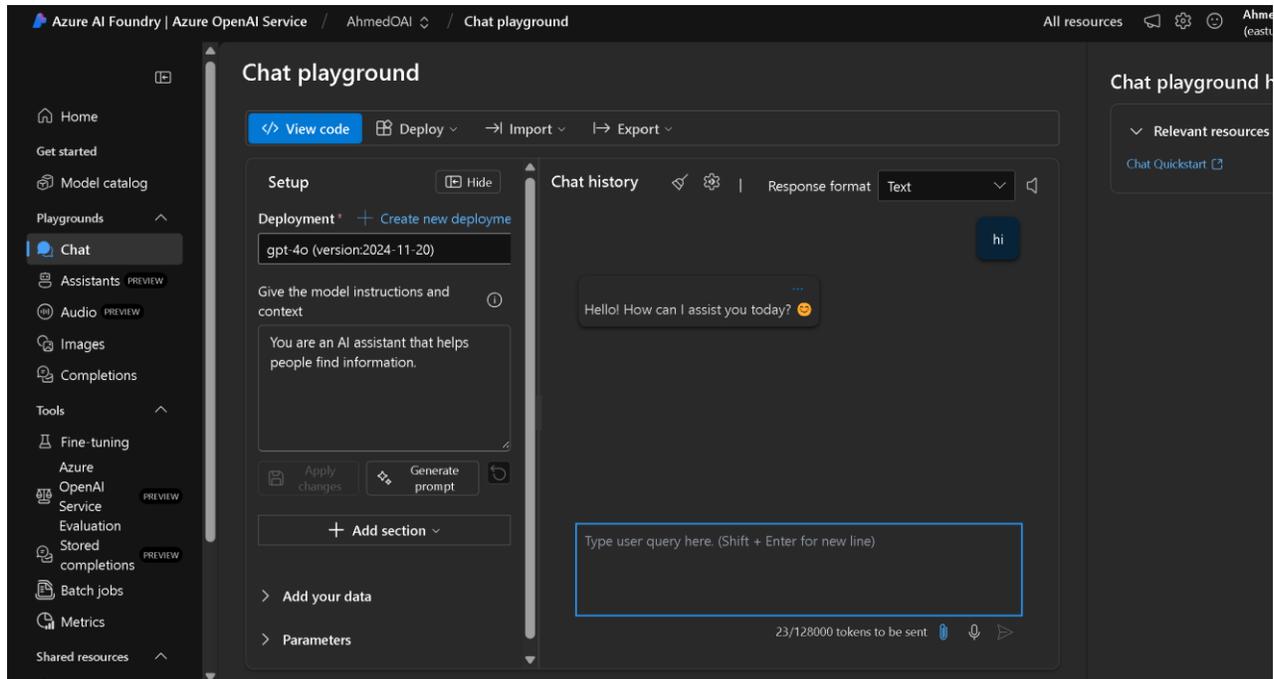
Showing 1 to 2 of 2 records. Show hidden types

<input type="checkbox"/>	Name ↑↓	Type ↑↓	Location ↑↓
<input type="checkbox"/>	AhmedOAI	Azure OpenAI	East US
<input checked="" type="checkbox"/>	All-instance1	Computer vision	East US

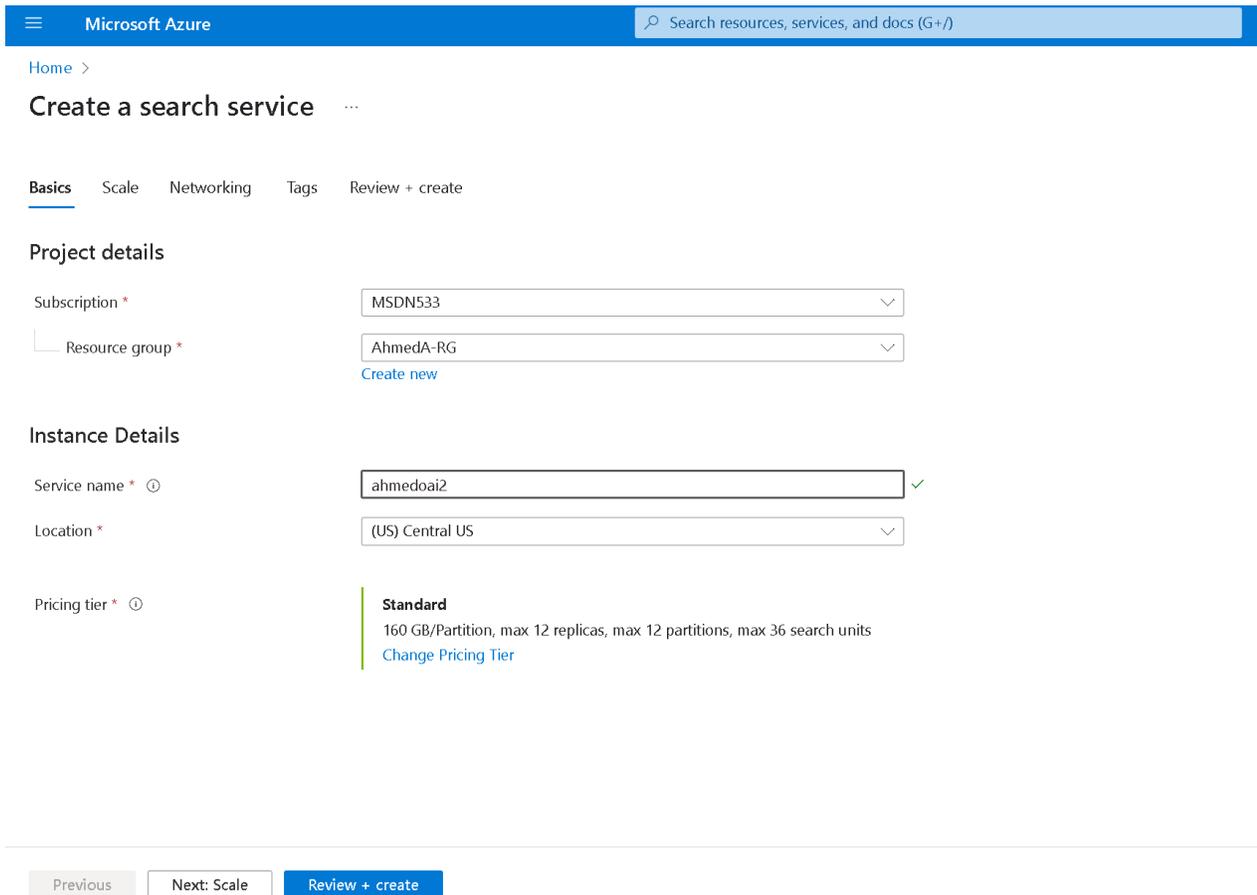
The screenshot shows the Microsoft Azure portal interface. At the top, there's a search bar and navigation icons. Below that, the breadcrumb trail reads: Home > Microsoft.CognitiveServicesOpenAI-20250507005434 | Overview > AhmedAIRG >. The main heading is 'AhmedOAI' with 'Azure OpenAI' underneath. A search bar is present, and there are links for 'Go to Azure AI Foundry portal' and 'Delete'. A left-hand navigation menu includes 'Overview', 'Activity log', 'Access control (IAM)', 'Tags', 'Diagnose and solve problems', 'Resource visualizer', 'Resource Management', 'Security', 'Monitoring', 'Automation', and 'Help'. The 'Essentials' section displays resource details: Resource group (move) : AhmedAIRG, Status : Active, Location : East US, Subscription (move) : MSDN533, Subscription ID : 6bde611b-7c73-4f7d-8942-8d3baa229533, and Tags (edit) : WebApp : Test. On the right, it shows API Kind : OpenAI, Pricing tier : Standard, Endpoints : [Click here to view endpoints](#), and Manage keys : [Click here to manage keys](#). Below this, there are tabs for 'Get Started', 'Develop', and 'Monitor'. A large heading reads 'Build your own secure copilot and generative AI applications with Azure OpenAI Service'. The text below explains: 'Deploy an Azure OpenAI model and start making API calls. Connect your own data, call functions, and improve workflow with Azure OpenAI language, image and speech models. You can access the service through REST APIs, Python SDK, or our web-based interface in the Azure AI Foundry portal. [Learn More](#)'. A blue button labeled 'Explore Azure AI Foundry portal' is at the bottom right.

The screenshot shows the 'Chat playground' interface within the Azure AI Foundry. The breadcrumb trail is: Azure AI Foundry | Azure OpenAI Service / AhmedOAI / Chat playground. The left sidebar contains a navigation menu with 'Home', 'Get started', 'Model catalog', 'Playgrounds', 'Chat', 'Assistants (PREVIEW)', 'Audio (PREVIEW)', 'Images', 'Completions', 'Tools', 'Fine-tuning', 'Azure OpenAI Service (PREVIEW)', 'Evaluation', 'Stored completions (PREVIEW)', 'Batch jobs', 'Metrics', and 'Shared resources'. The main area is titled 'Chat playground' and has a toolbar with 'View code', 'Deploy', 'Import', and 'Export'. Below the toolbar, there are two panels: 'Setup' and 'Preview'. The 'Setup' panel shows 'Deployment * + Create new deployment' and 'No deployment exists'. It features a large blue folder icon with a white plus sign. Below this, it says 'Deployment needed' and 'In order to modify and interact with the Playground, you first need to deploy a base model to your project.' A button 'Don't have a deployment?' is followed by a blue button '+ Create a deployment'. The 'Preview' panel is currently empty.





Add your data



Select or add data source

Your data source is used to ground the generated results with your data. Select an existing data source or create a new data connection with Azure Blob Storage, databases, search, URLs, or local files as the source the grounding data will be built from.

[Learn more about data privacy and security in Azure AI.](#)

Select data source *

Upload files (preview)

Subscription *

MSDN533

Select Azure Blob storage resource ① *

ahmedstorage0001

[Create a new Azure Blob storage resource](#)

✔ Cross-origin resource sharing (CORS) is turned on for this resource.

Select Azure AI Search resource (free tier not supported) ① *

ahmedoai2

[Create a new Azure AI Search resource](#)

Enter the index name ① *

ahmed

Using Azure AI Search will incur usage to your account. [View Pricing](#)

Add vector search to this search resource. ①

Upload files

Select which files to add. Files will be stored in your Azure Blob Storage and indexed by the Cognitive Search resource created or selected in the previous step.

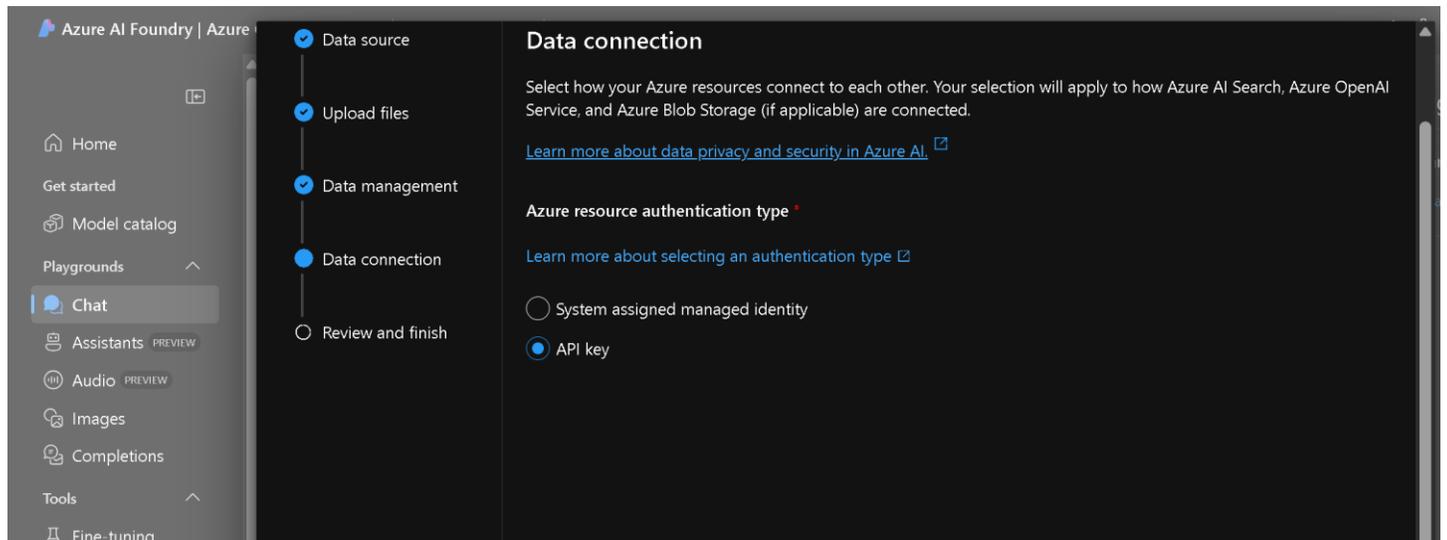
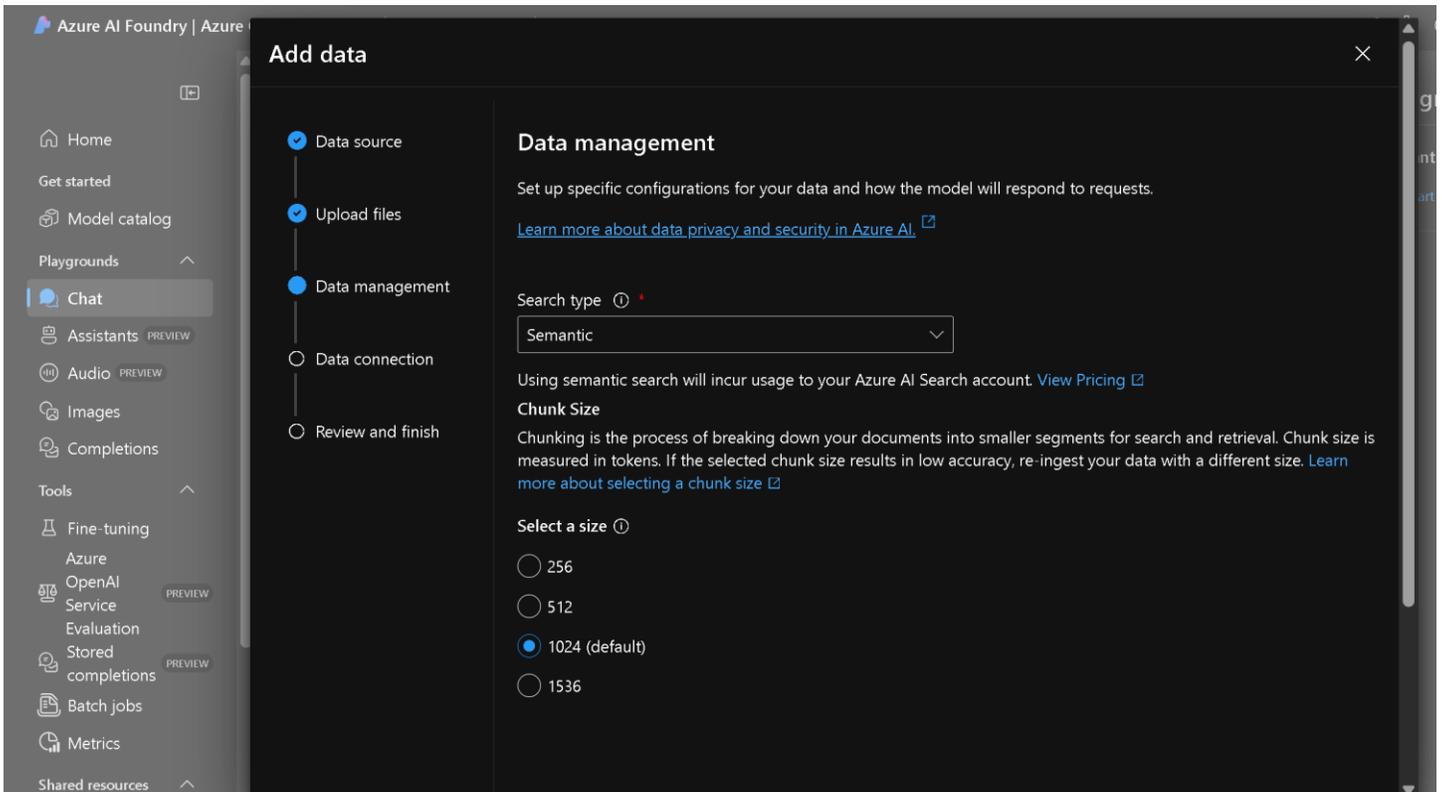
[Learn more about data privacy and security in Azure AI.](#)

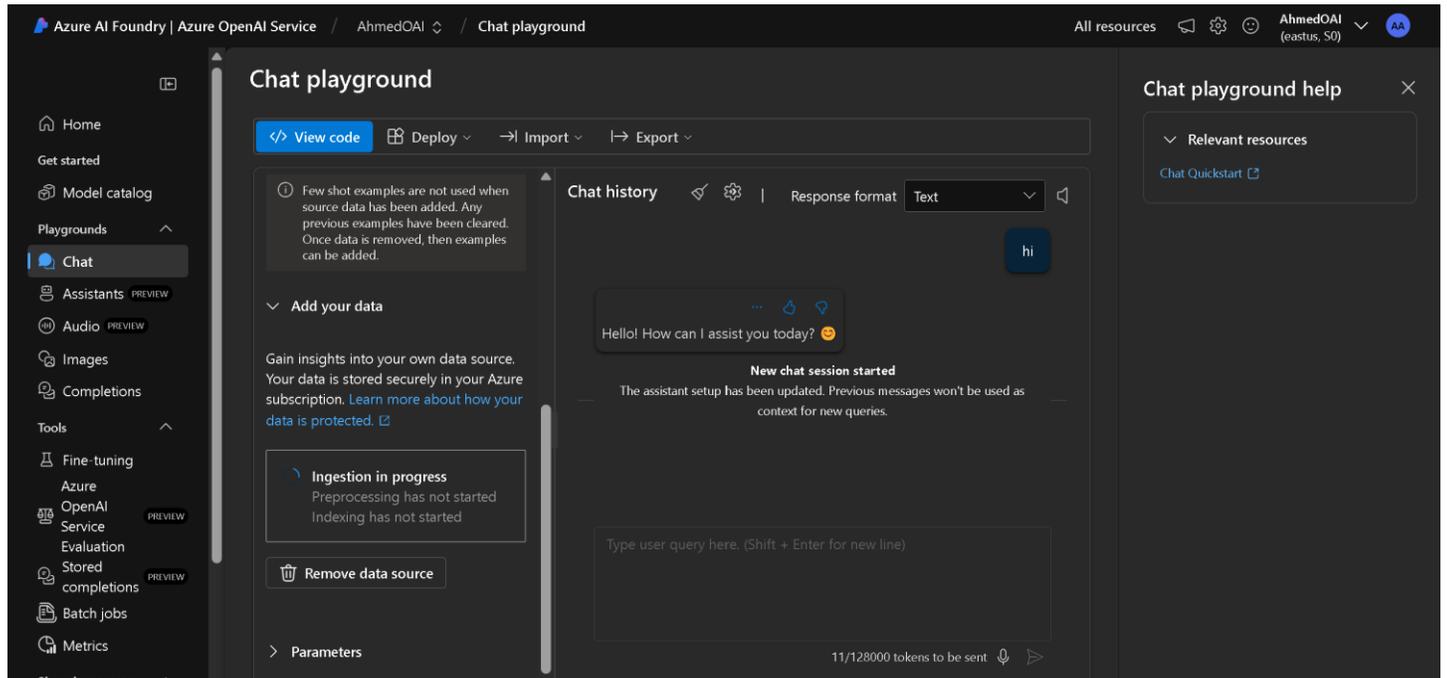
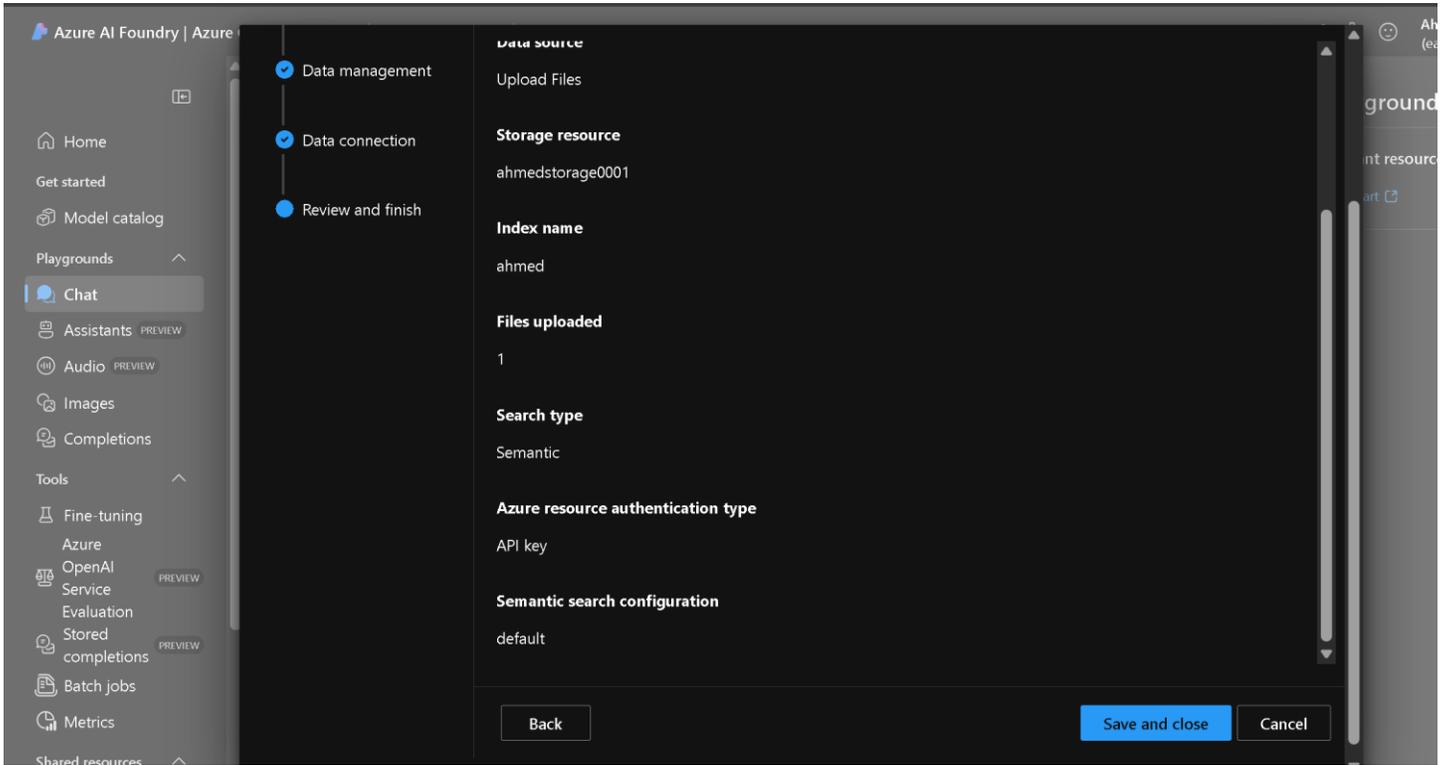
Drag and drop
or
Browse for a file

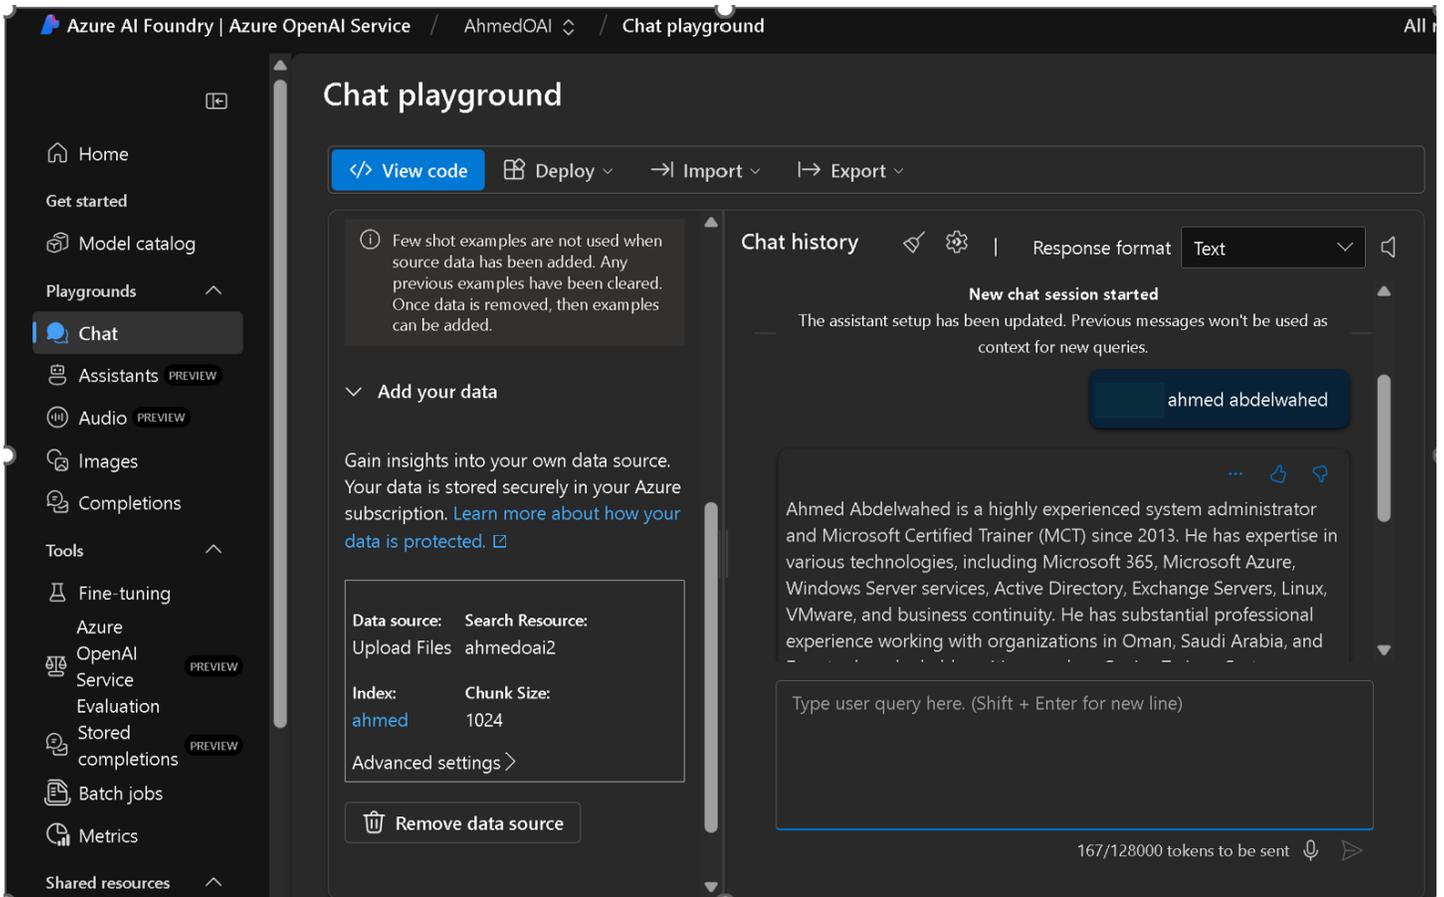
(.txt, .md, .html, .pdf, .docx, .pptx)
16 MB size limit

File name	Type	Size	Status
Ahmed Abdelwahe	PDF	143.82 KB	Uploaded

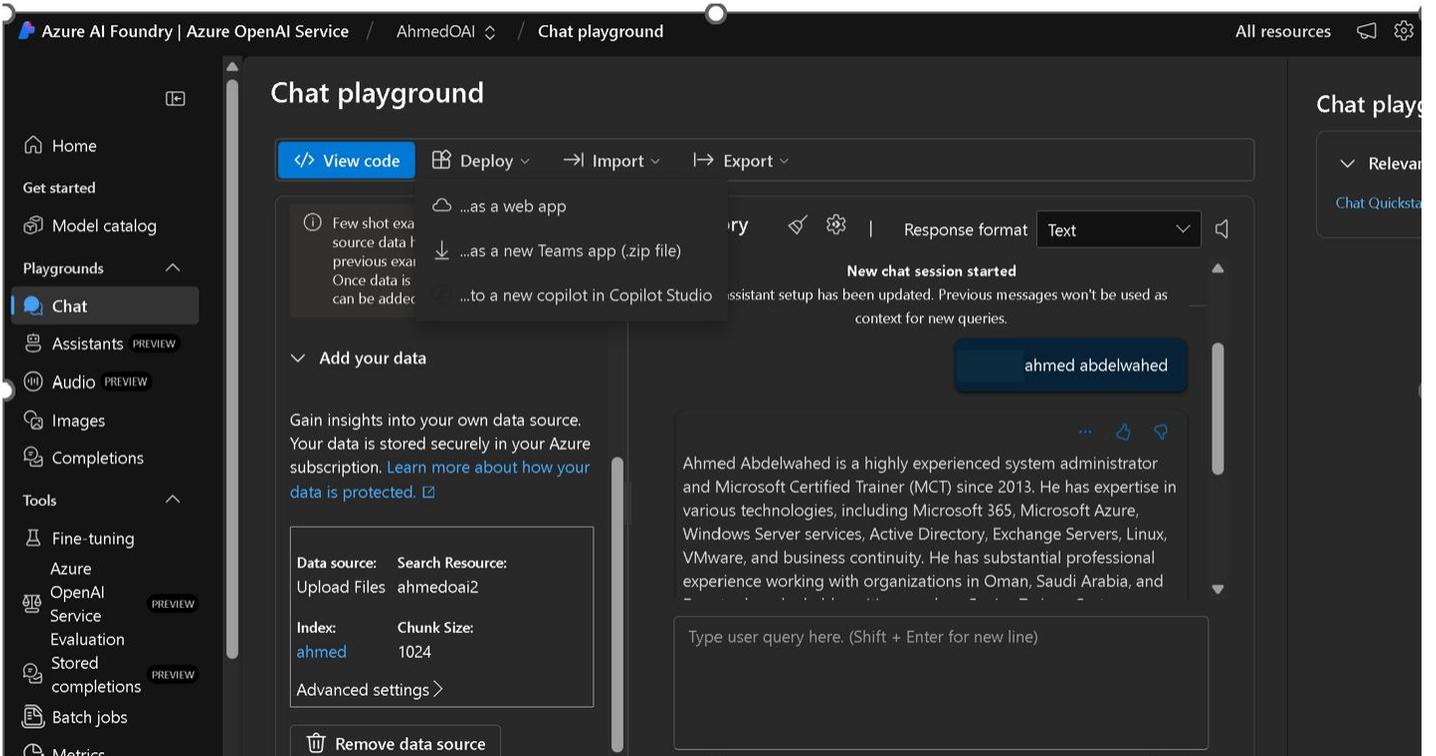
✔ Your files were successfully uploaded.

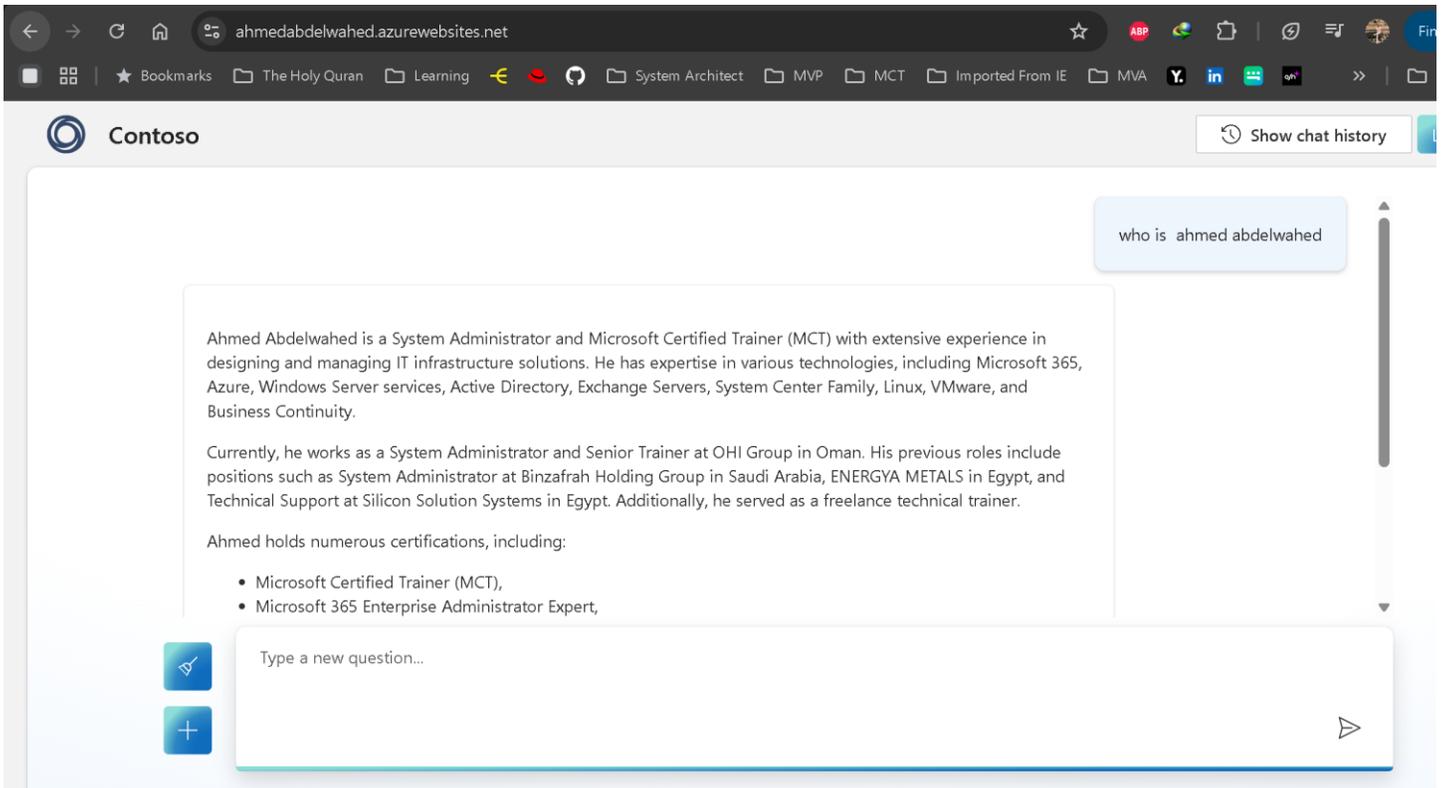
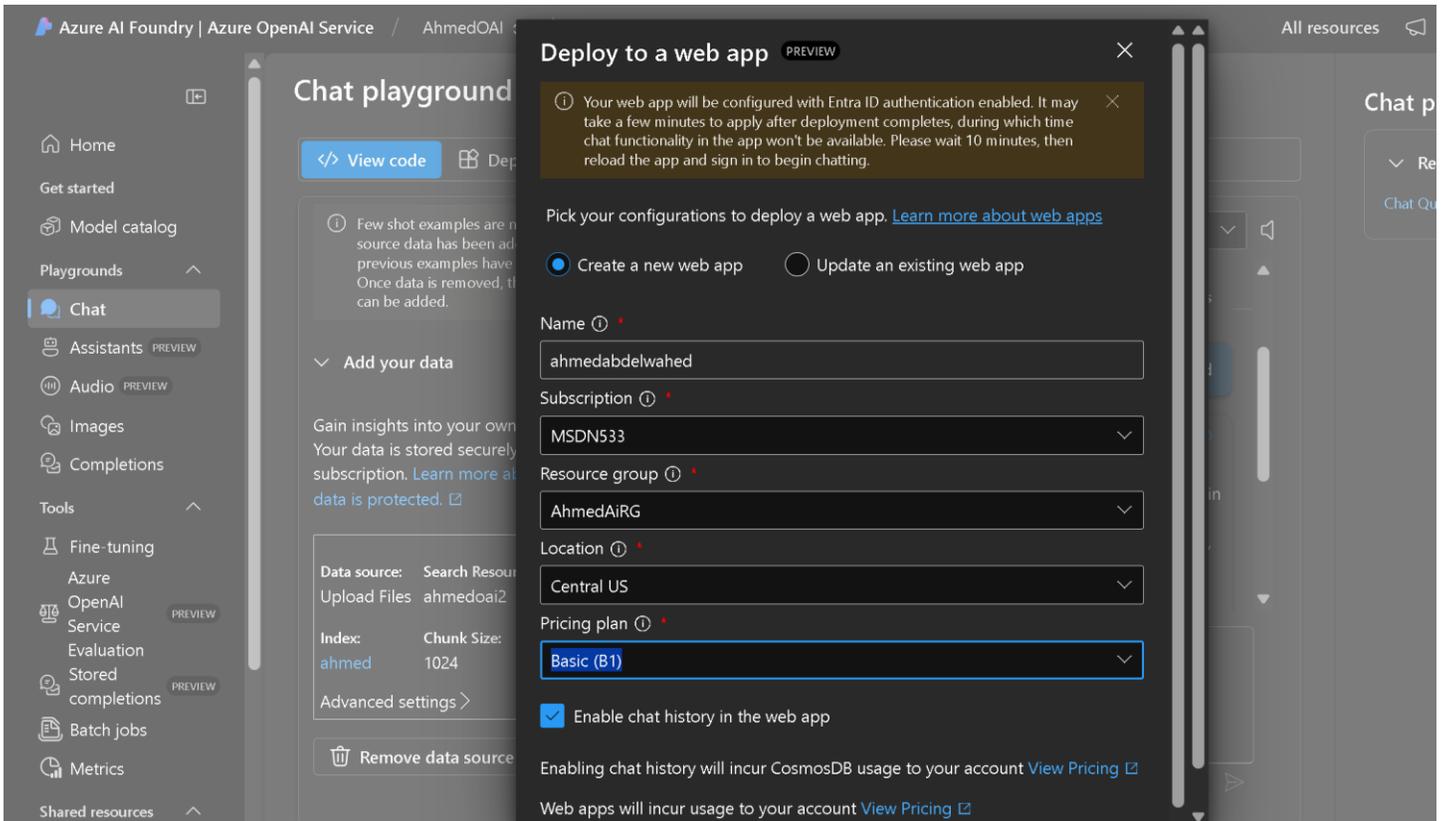




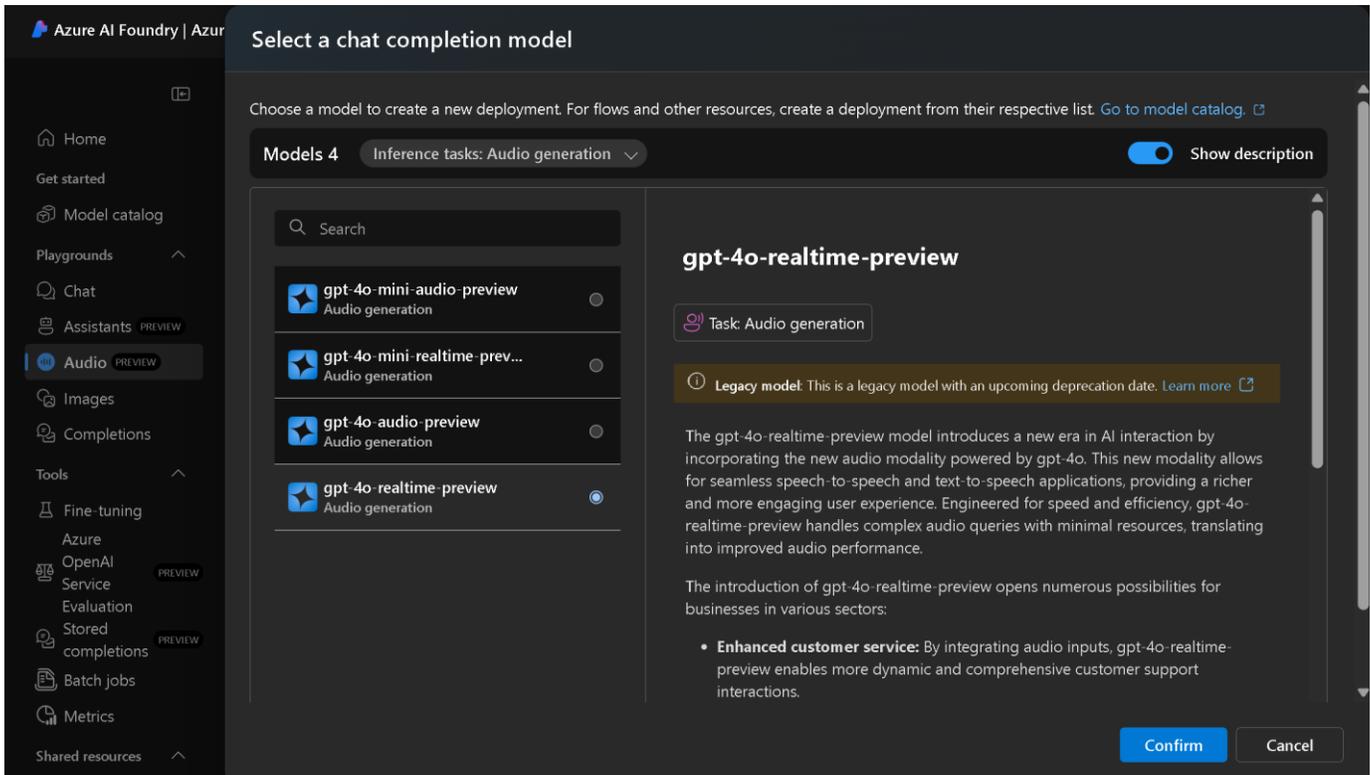
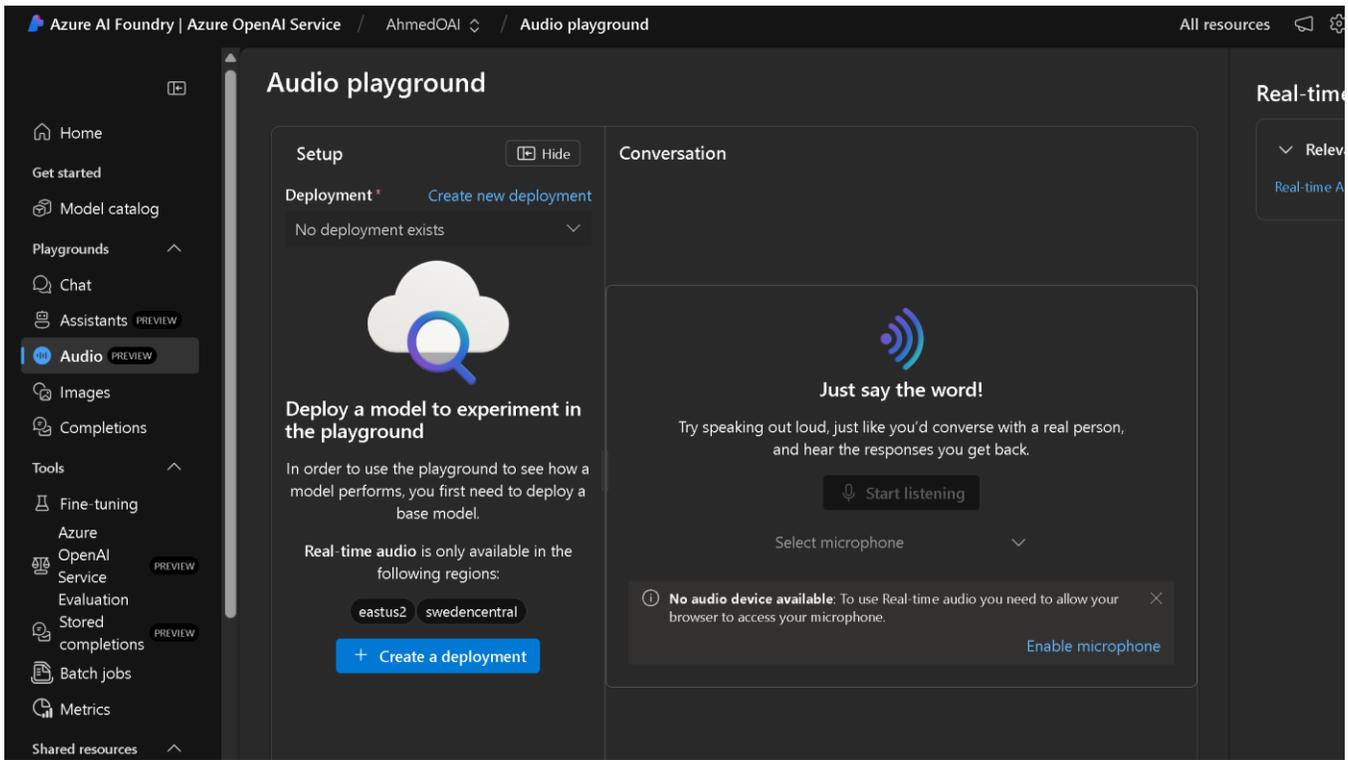


Deploy web app for your Azure OpenAI





Audio



Deploy gpt-4o-realtime-preview

Deployment name *
gpt-4o-realtime-preview

Deployment type
Global Standard

Global Standard: Pay per API call with the highest rate limits. Learn more about [Global deployment types](#).

Data might be processed globally, outside of the resource's Azure geography, but data storage remains in the AI resource's Azure geography. Learn more about [data residency](#).

Deployment details Customize

Model version 2024-12-17	AI resource (change) ahmed-mad0hq0f-eastus2
Capacity 1K tokens per minute (TPM)	Resource location East US 2
Content safety DefaultV2	Version upgrade policy Once a new default version is available

The model will be deployed to the selected resource
A different resource has been pre-selected for deploying this model because the current resource either lacks sufficient quota or is located in a region that does not support this model.

Deploy to selected resource Cancel

Audio playground

Setup Hide

Deployment * Create new deployment

No deployment exists

Deploy a model to experiment in the playground

In order to use the playground to see how a model performs, you first need to deploy a base model.

Real-time audio is only available in the following regions:

eastus2 swedencentral

+ Create a deployment

Conversation

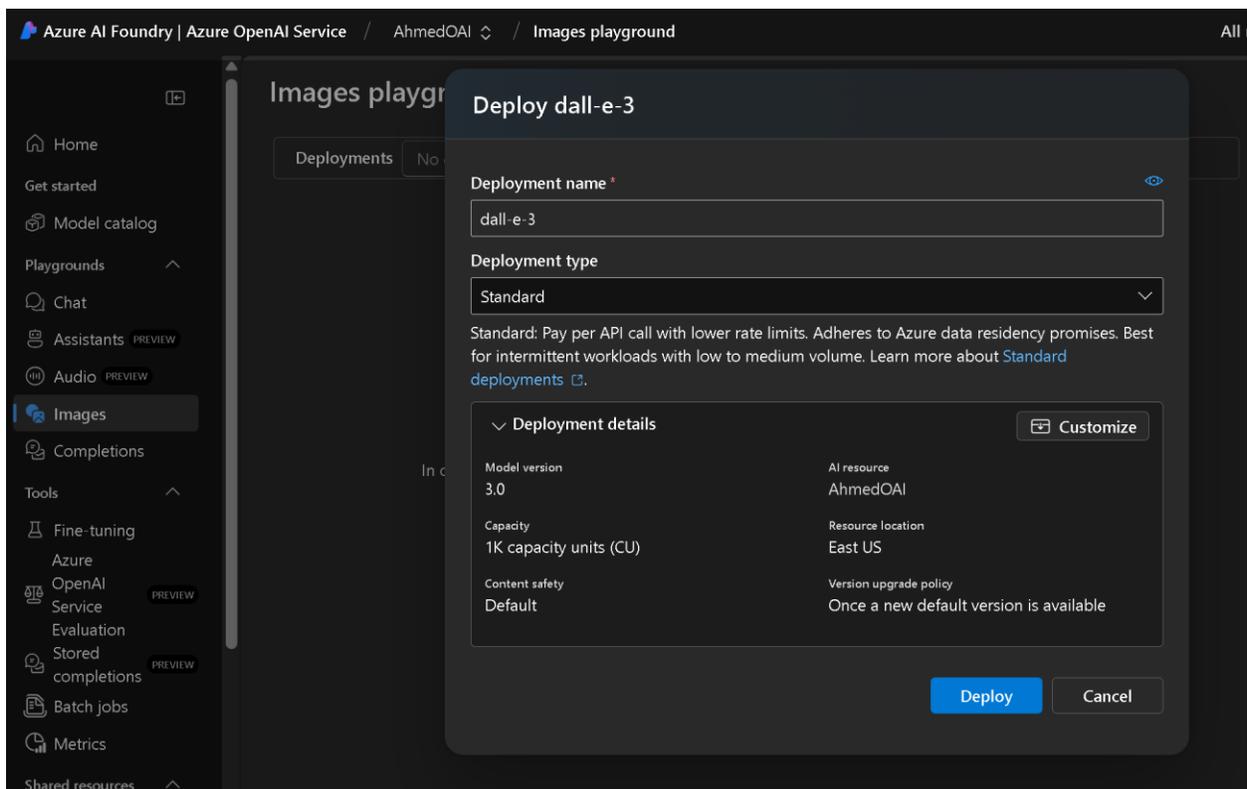
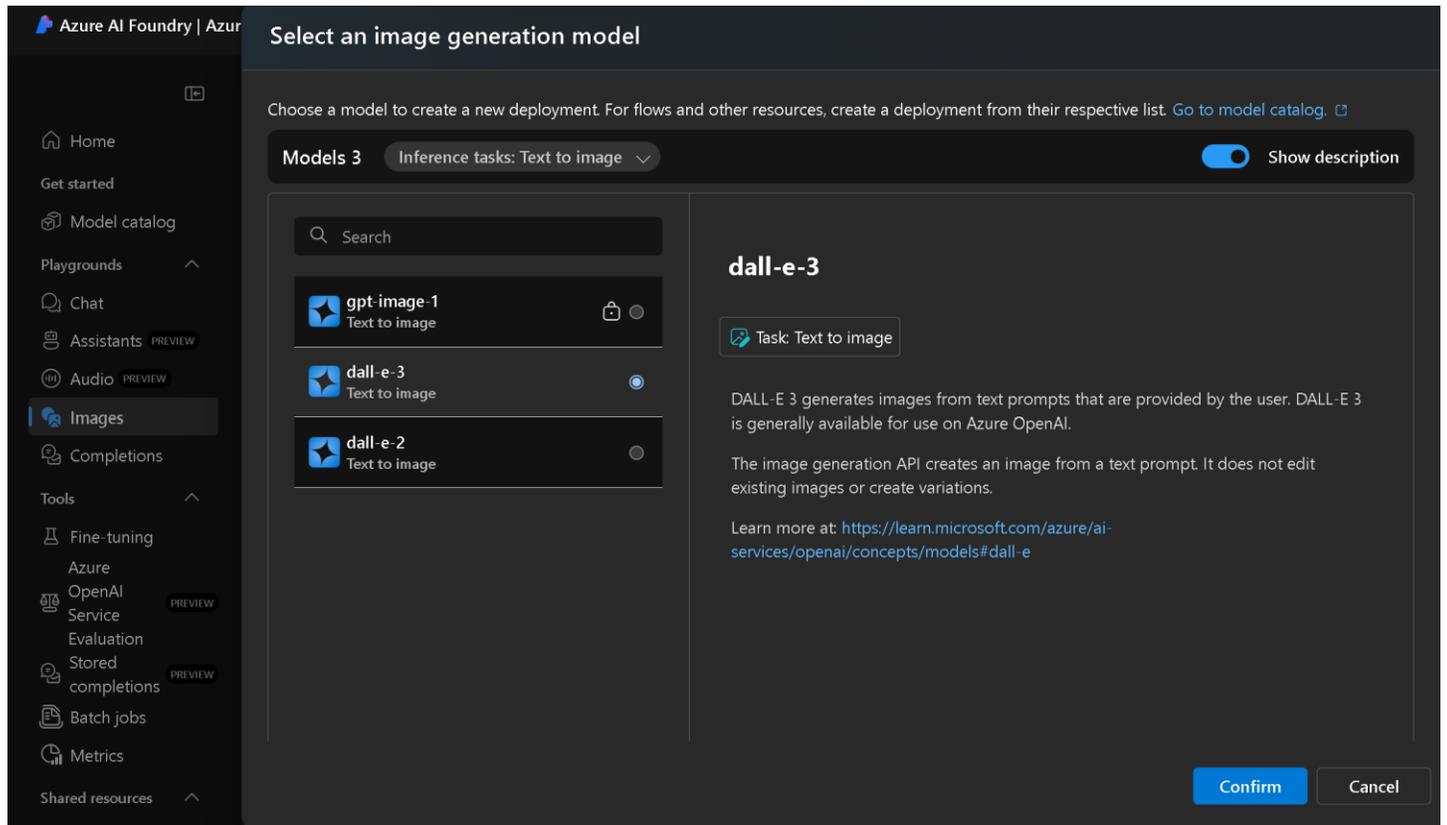
Just say the word!

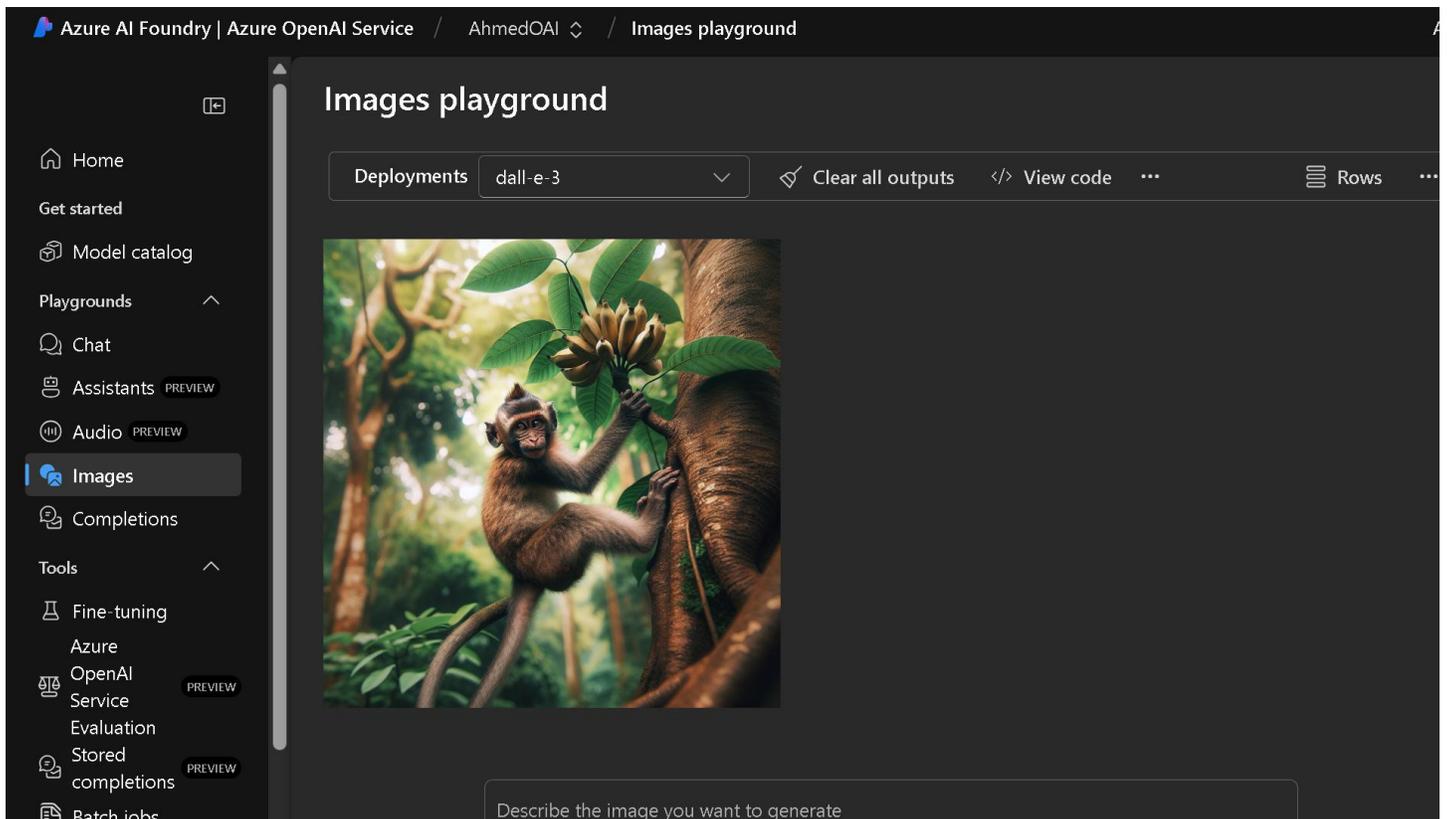
Try speaking out loud, just like you'd converse with a real person, and hear the responses you get back.

Start listening

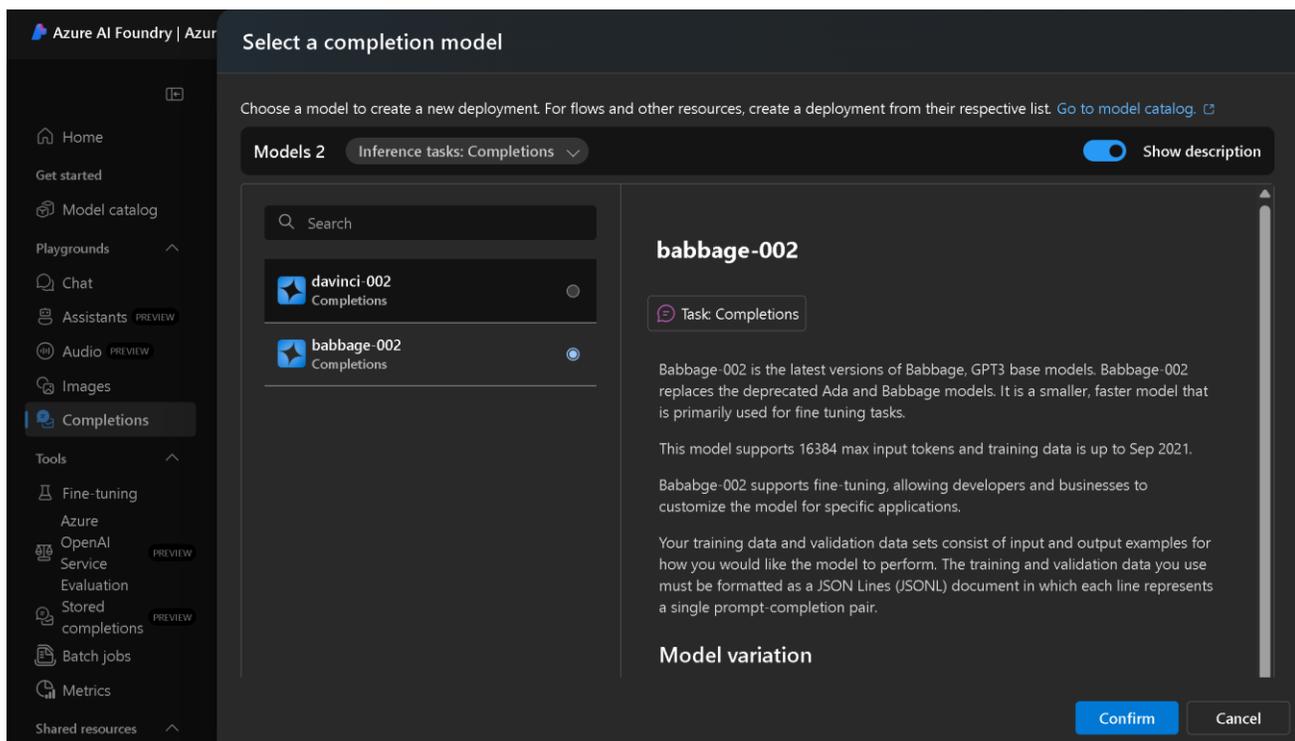
Default - Microphone Array (Intel® Smart Sound Technology for Digital Microphones)

Image





Completion



Deploy babbage-002

Deployment name *
babbage-002

Deployment type
Standard

Standard: Pay per API call with lower rate limits. Adheres to Azure data residency promises. Best for intermittent workloads with low to medium volume. [Learn more about Standard deployments](#).

Deployment details Customize

Model version	AI resource
1	(create) ahmed-mad2oz6k-swedencentral
Capacity	Authentication type
120K tokens per minute (TPM)	Key
Content safety	Resource location
Default	Sweden Central
	Version upgrade policy
	Once a new default version is available

i A new AI resource will be created for your deployment
A resource location that supports the model has been pre-selected

[Create resource and deploy](#) [Cancel](#)