



IMVC 2024

Amit Svarzenberg
CTO - Microsoft for Startups
and AI Lecturer at Tel Aviv University

AI is
everywhere

CNN BUSINESS.


Real estate agents say they can't imagine working without ChatGPT now

By Samantha Murphy Kelly, CNN Business
Updated 4:06 PM CST, Sat January 20, 2023

NEWS

AI's Transformative Impact on Scientific Research

Krishi Chowdhary · Journalist
Updated: August 3, 2023 | 7:33 AM



FORBES > INNOVATION

Why AI's Creative Intelligence Will Change Advertising In 2024

Forbes | Ian Liddicoat Forbes Councils Member
Forbes Technology Council COUNCIL POST | Membership (Fee-Based)

Forbes

FORBES > MONEY > MARKETS

AI Productivity Gains Are Not Limited To Big Tech

REUTERS World Business Markets Sustainability Legal Breakingviews Technology Investigation

Disrupted

Bar exam score shows AI can keep up with 'human lawyers,' researchers say

By Karen Sloan
March 15, 2023 11:17 AM PDT · Updated 5 months ago

make it SUCCESS MONEY WORK LIFE VIDEO

NEXT GEN INVESTING


A.I. is making some common side hustles more lucrative—these can pay up to \$100 per hour

Published Sat, Jul 29 2023 9:00 AM EDT · Updated Mon, Jul 31 2023 10:31 PM EDT

TECHNOLOGY

BankThink AI could make underwriting a collaboration between bank and borrower

By Jeremiah Buckley August 04, 2023, 10:00 a.m. EDT 3 Min Read



From Text to Real Life Robots



Microsoft invests \$1 billion in OpenAI to pursue holy grail of artificial intelligence

Building artificial general intelligence is OpenAI's ambitious goal

By [James Vincent](#) | Jul 22, 2019, 10:08am EDT



SHARE

23:34



New chat

ChatGPT



Examples

"Explain quantum computing in simple terms" →

"Got any creative ideas for a 10 year old's birthday?" →

"How do I make an HTTP request in Javascript?" →



Capabilities

chat.openai.com — Private

Thank

Done

I'm

U

O



ChatGPT

Runs on Azure

Backed by Cosmos DB

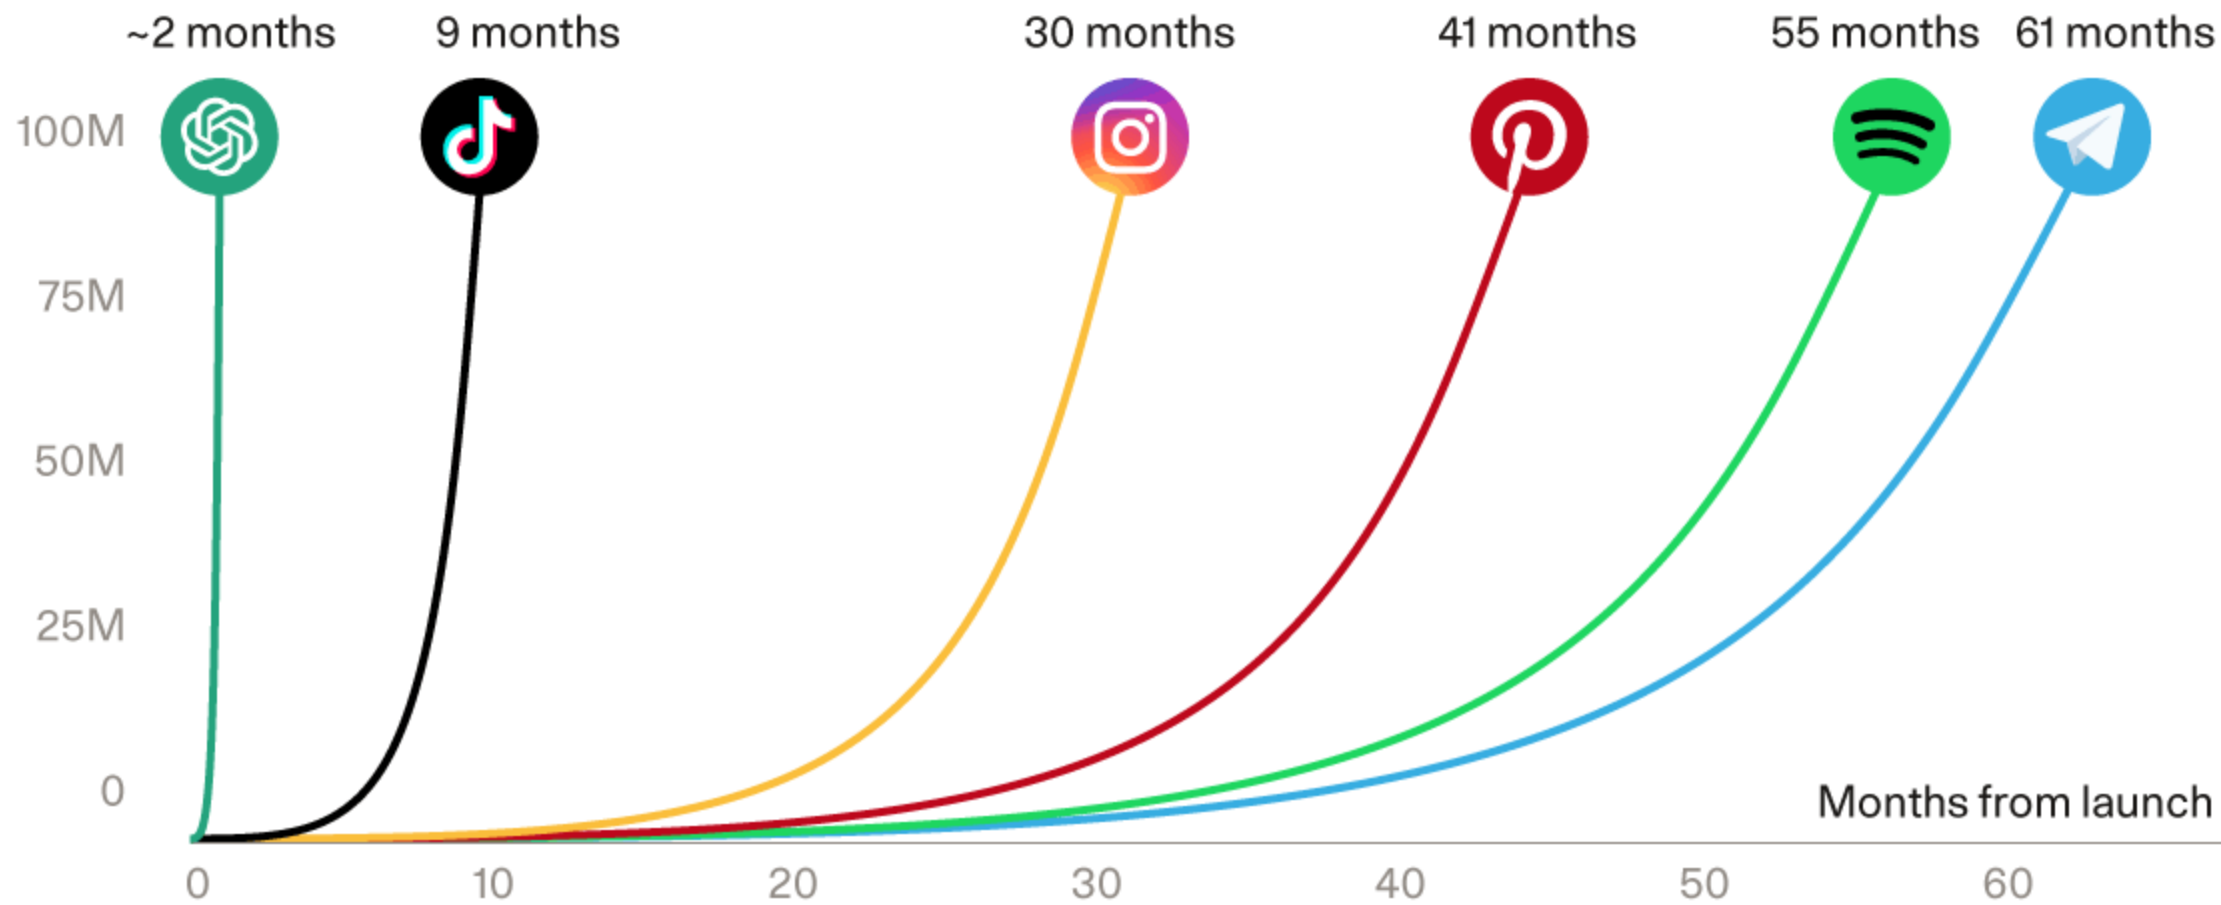
Developed on GitHub

Fastest

growing consumer
product in history



Path to 100 Million Users (stylized)



Microsoft to Invest \$10 Billion in OpenAI, the Creator of ChatGPT

The tech giant aims to remain at the forefront of generative artificial intelligence with its partnership with OpenAI.

**World's most
powerful
supercomputers**

**World's most capable
foundation models, from
hosted to open source**

**World's best
AI dev
infrastructure**



Microsoft

+



OpenAI

An AI humanoid robot startup raised \$675 million — including from Microsoft and OpenAI



The Future of AI Robots is Multimodality





"Teddy bear on a skateboard, 4k"



(Raining ambience)

CoDi



(Raining street ambience)



"A toy on the street sitting on a board"

"Raining, rain, moderate"



(Raining, street, skateboard)

CoDi can jointly generate any-to-any combinations from video, image, audio, and text, via composable diffusion.

Figure 1: CoDi can generate any combination of modalities from any mixture of input modalities.

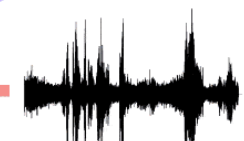


"Teddy bear on a skateboard, 4k"



(Raining ambience)

CoDi



(Raining street ambience)



"A toy on the street sitting on a board"

"Raining, rain, moderate"

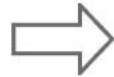


(Raining, street, skateboard)

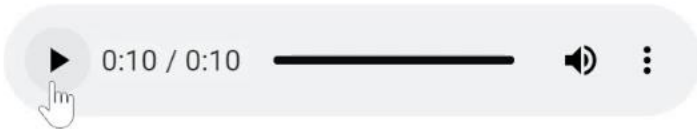
CoDi can jointly generate any-to-any combinations from video, image, audio, and text, via composable diffusion.

Text + Audio + Image → Text + Image

"Teddy bear on a skateboard, 4k, high resolution"



"A toy on the street sitting on a board"

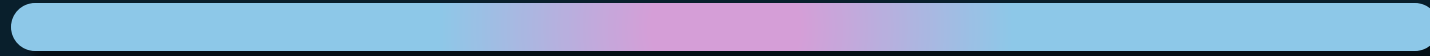




<https://openai.com/sora>

New

GPT-4 Turbo with Vision





Source: [hmmm \(Reddit\)](#)

User What is funny about this image? Describe it panel by panel.



Source: [hmmm \(Reddit\)](#)

GPT-4 The image shows a package for a "Lightning Cable" adapter with three panels.

Panel 1: A smartphone with a VGA connector (a large, blue, 15-pin connector typically used for computer monitors) plugged into its charging port.

Panel 2: The package for the "Lightning Cable" adapter with a picture of a VGA connector on it.

Panel 3: A close-up of the VGA connector with a small Lightning connector (used for charging iPhones and other Apple devices) at the end.

The humor in this image comes from the absurdity of plugging a large, outdated VGA connector into a small, modern smartphone charging port.



What would happen if
the strings were cut?



The balloons would
fly away.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	V
1																							
2																							
3																							
4																							
5																							
6																							
7																							
8																							
9																							
10																							
11																							
12																							
13																							
14																							
15																							
16																							
17																							
18																							
19																							
20																							
21																							
22																							
23																							
24																							
25																							
26																							
27																							
28																							
29																							
30																							

app.clipchamp.com is sharing your screen. Stop sharing Hide

New

Azure OpenAI Whisper



Demo – contact center

Insurance claims automation

Azure Speech-to-Text allows to reliably and accurately transcribe claims done via a call center either in batch or real-time. Azure OpenAI service allows to easily extract all relevant claims information from the call for further automated processing in existing systems.

Call Recording Audio → Audio → Azure Speech → Transcription → Azure OpenAI Service → Extracted call information → CRM or other Systems

Raw input data

Hi I just had a car accident and wanted to report it. OK, I hope you're alright, what happened? I was driving on the I-18 and I hit another car. Are you OK? Yeah, I'm just a little shaken up. That's understandable. Can you give me your full name? Sure, it's Sarah standl. Do you know what caused the accident? I think I might have hit a pothole. OK, where did the accident take place? On the I-18 freeway. Was anyone else injured? I don't think so. But I'm not sure. OK, well we'll need to do an investigation. Can you give me the other drivers information? Sure, his name is John Radley. And your insurance number. OK. Give me a minute. OK, it's 546452. OK, what type of damages has the car? Headlights are broken and the airbags went off. Are you going to be able to drive it? I don't know. I'm

Extracted claims information

▶ Play call recording

Analyze raw input

Translation Demo

Demo – contact center

Azure AI Demos

With Azure Cognitive Services and Azure OpenAI Service, you can easily automate tedious business processes, such as claims processing, data extraction and many others.

Insurance claims automation

Azure Speech-to-Text allows to reliably and accurately transcribe claims done via a call center either in batch or real-time. Azure OpenAI service allows to easily extract all relevant claims information from the call for further automated processing in existing systems.



▶ Play call recording

Raw input data

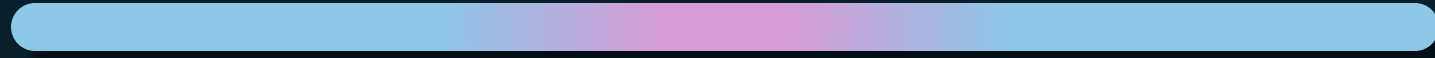
OK? Yeah, I'm just a little shaken up. That's understandable. Can you give me your full name? Sure, it's Sarah Standl. Do you know what caused the accident? I think I might have hit a pothole. OK, where did the accident take place? On the I-18 freeway. Was anyone else injured? I don't think so. But I'm not sure. OK, well we'll need to do an investigation. Can you give me the other drivers information? Sure, his name is John Radley. And your insurance number. OK. Give me a minute. OK, it's 546452. OK, what type of damages has the car? Headlights are broken and the airbags went off. Are you going to be able to drive it? I don't know. I'm going to have to have it towed. Well, we'll need to get it inspected. I'll go ahead and start the claim and we'll get everything sorted out. Thank

Extracted claims information

```
results = {
  "reason": "car accident",
  "cause": "pothole",
  "driver_names": ["Sarah Standl", "John Radley"],
  "insurance_number": "546452",
  "location": "I-18 freeway",
  "damages": ["headlights", "airbags"],
  "car_registration": null,
  "summary": "Sarah Standl was driving on the I-18 and hit another car. She thinks she might have hit a pothole. John Radley was the other driver. Both drivers are OK but the car has damage to the headlights and airbags."
}
```

Analyze raw input

Create Music





that analyzes the merits of universal basic income

makes me laugh

a career change. Can you help me go through pros and cons?

I want to have a spiritual experience?

Chat

Do you think a bear or a lion would win in a fight?

Organize

Organize the world's tallest trees into a table by habitat

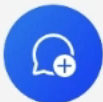
Make a song with Suno

Make a pop song about a seahorse named Bubbles

[Terms of use](#) | [Privacy policy](#)

Microsoft Copilot is powered by AI, so surprises and mistakes are possible. Please share feedback so we can improve! [Terms](#) | [Privacy](#)

Choose a conversation style



create a song about Amit Svarzenberg CTO of Microsoft for Startups giving a keynote at Future Of AI convention in a Rap style.



130/2000



Recent activity Plugins

Beneficial Ownership Info 3 days ago

Amir Tsalich and Amad... Jan 8, 2024

Song about Amir Tsalich Jan 7, 2024

write a song about Yair Jan 7, 2024

דינוזאור פרך בירח Jan 6, 2024

[See all recent chats](#)

OpenAI and Microsoft partner with Humane on wearable AI pin

AI startup Humane launched the wearable virtual assistant AI Pin that is embedded with technology from OpenAI and cloud computing services from Microsoft.





Age of AI Assistants

The Future of Multi Modality AI is AI Studio Platform

Age of AI Assistants

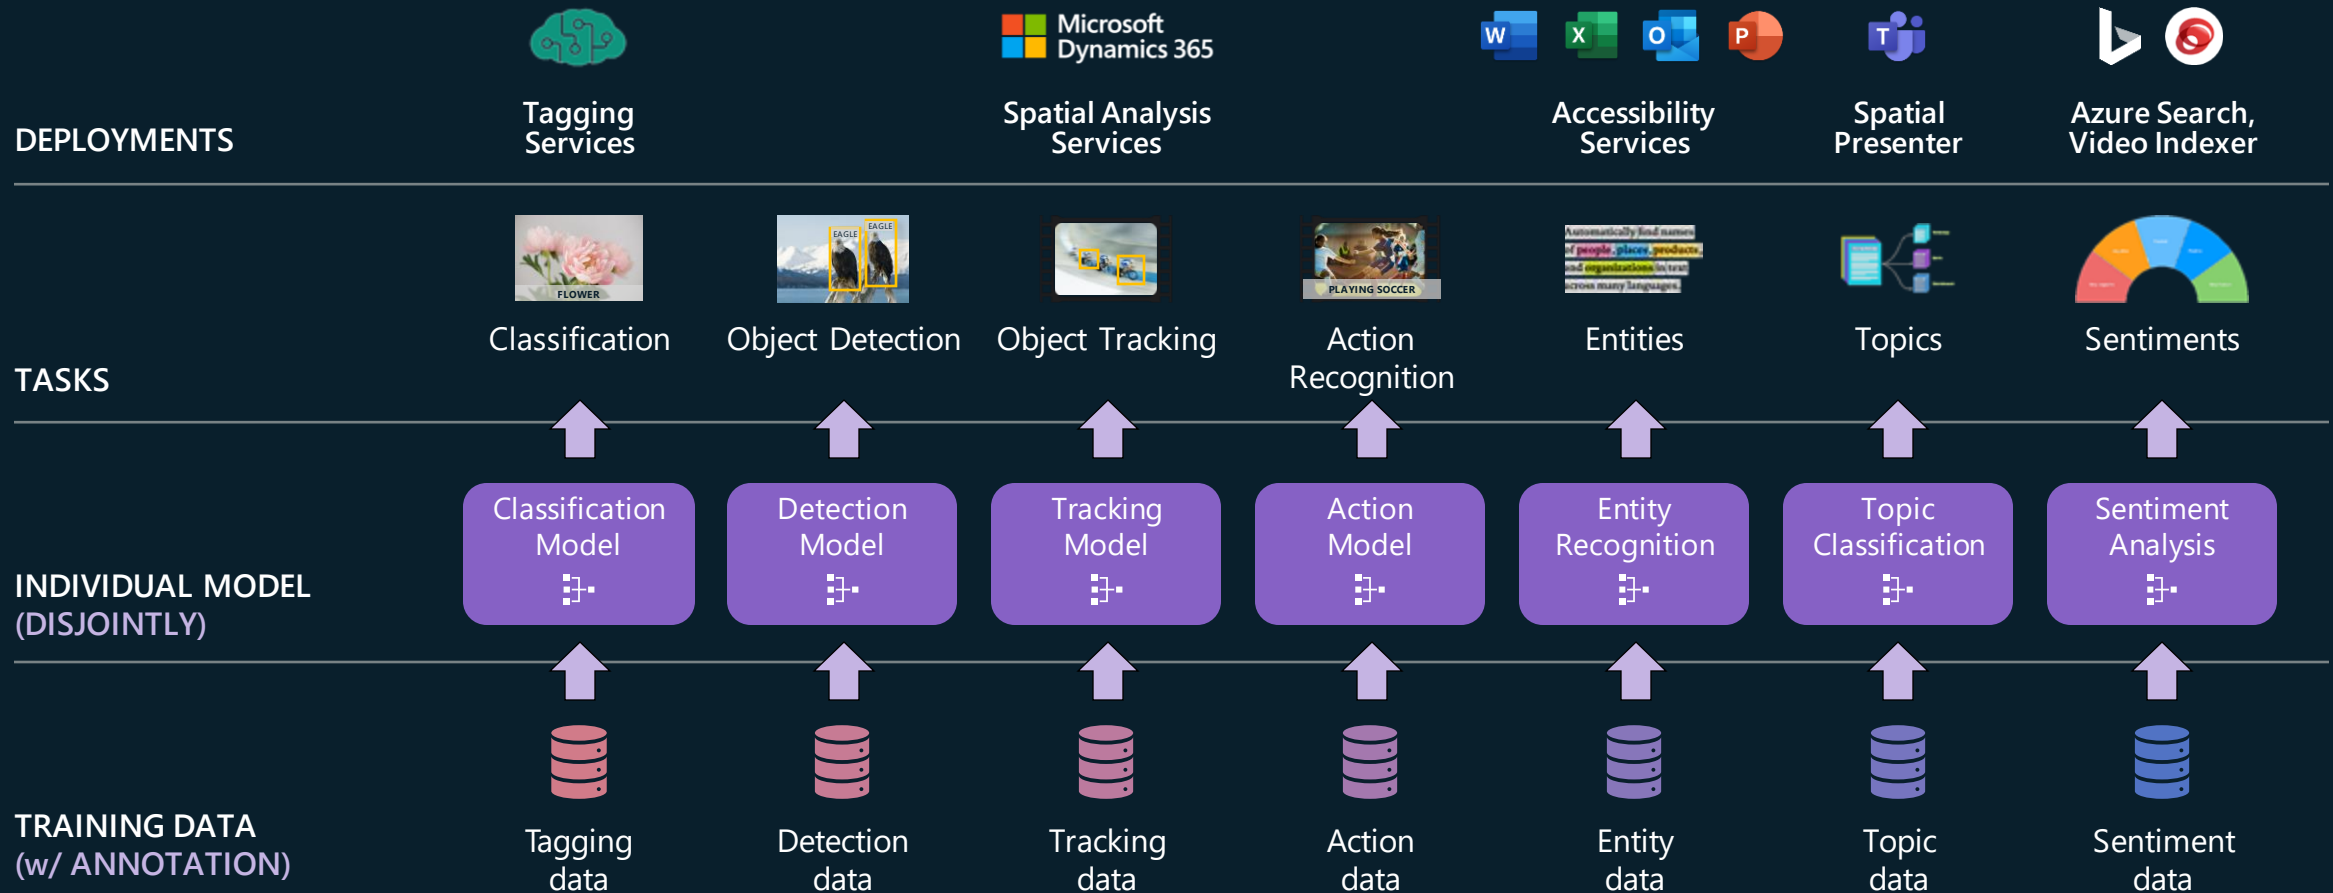
But first let's go to the basics!

What are Foundation Models?



Traditional model development

High cost and slow deployment—each service is trained disjointly



Foundation models

(includes Large Language Models)

One model trained on a huge amount of data and adapted to many applications

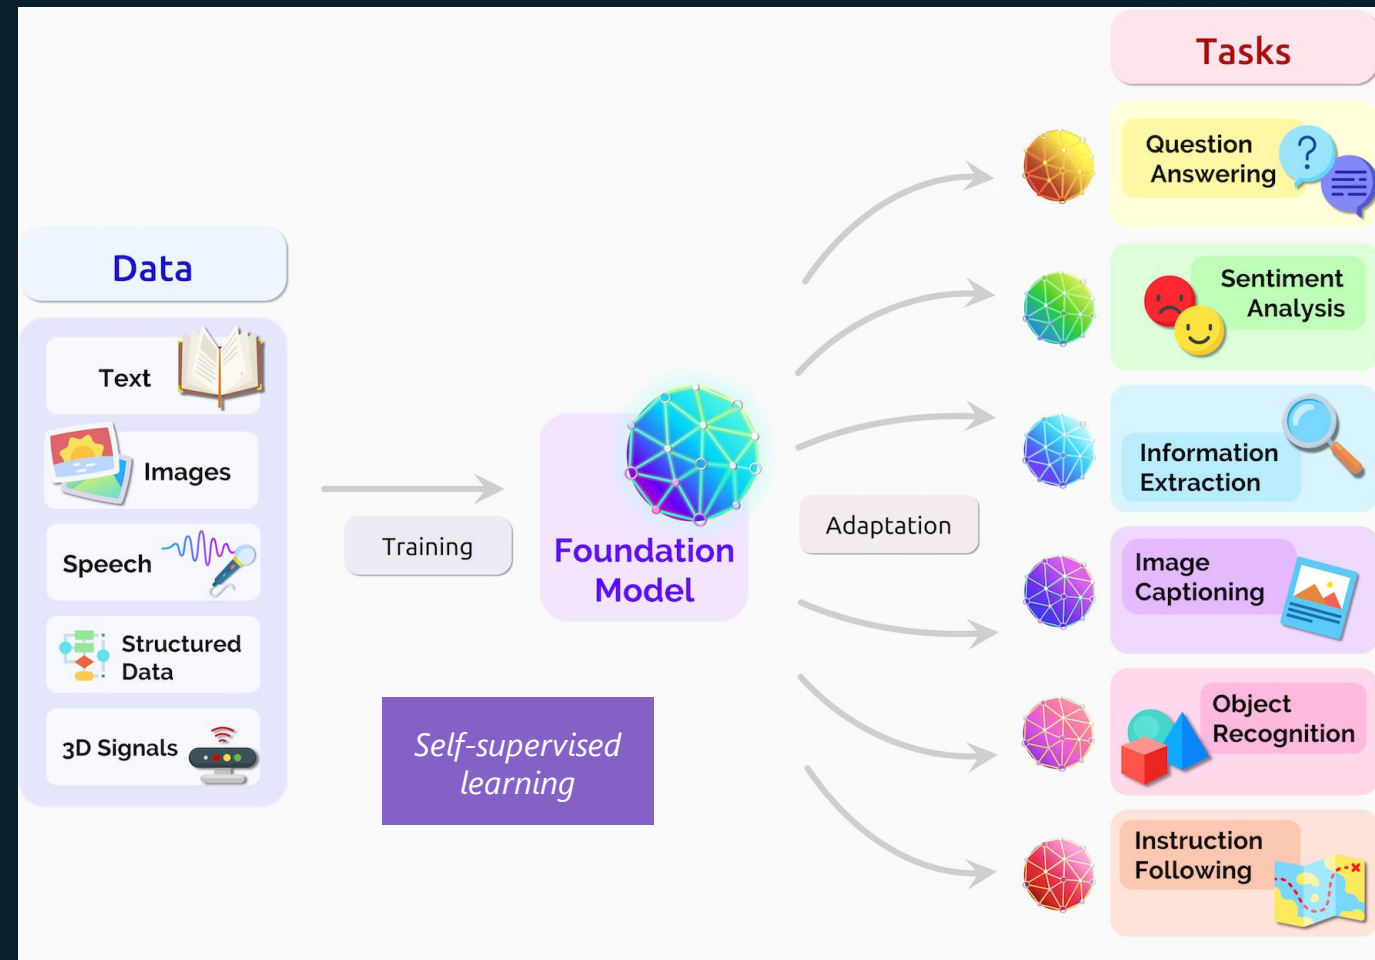
Adapted from: [Stanford Centre for Research on Foundation Models](#)



Foundation models

(includes Large Language Models)

One model trained on a huge amount of data and adapted to many applications



Self-supervised learning does not need labelled data

**Why are Foundation
Models special?**



2 example models

GPT-3 (2020)

50,257 vocabulary size
2048 context length
175B parameters
Trained on 300B tokens

Model Name	n_{params}	n_{layers}	d_{model}	n_{heads}	d_{head}	Batch Size	Learning Rate
GPT-3 Small	125M	12	768	12	64	0.5M	6.0×10^{-4}
GPT-3 Medium	350M	24	1024	16	64	0.5M	3.0×10^{-4}
GPT-3 Large	760M	24	1536	16	96	0.5M	2.5×10^{-4}
GPT-3 XL	1.3B	24	2048	24	128	1M	2.0×10^{-4}
GPT-3 2.7B	2.7B	32	2560	32	80	1M	1.6×10^{-4}
GPT-3 6.7B	6.7B	32	4096	32	128	2M	1.2×10^{-4}
GPT-3 13B	13.0B	40	5140	40	128	2M	1.0×10^{-4}
GPT-3 175B or "GPT-3"	175.0B	96	12288	96	128	3.2M	0.6×10^{-4}

Table 2.1: Sizes, architectures, and learning hyper-parameters (batch size in tokens and learning rate) of the models which we trained. All models were trained for a total of 300 billion tokens.

Training: (rough order of magnitude to have in mind)

- O(1,000 - 10,000) V100 GPUs
- O(1) month of training
- O(1-10) \$M

LLaMA (2023)

32,000 vocabulary size
2048 context length
65B parameters
Trained on 1-1.4T tokens

params	dimension	n heads	n layers	learning rate	batch size	n tokens
6.7B	4096	32	32	$3.0e^{-4}$	4M	1.0T
13.0B	5120	40	40	$3.0e^{-4}$	4M	1.0T
32.5B	6656	52	60	$1.5e^{-4}$	4M	1.4T
65.2B	8192	64	80	$1.5e^{-4}$	4M	1.4T

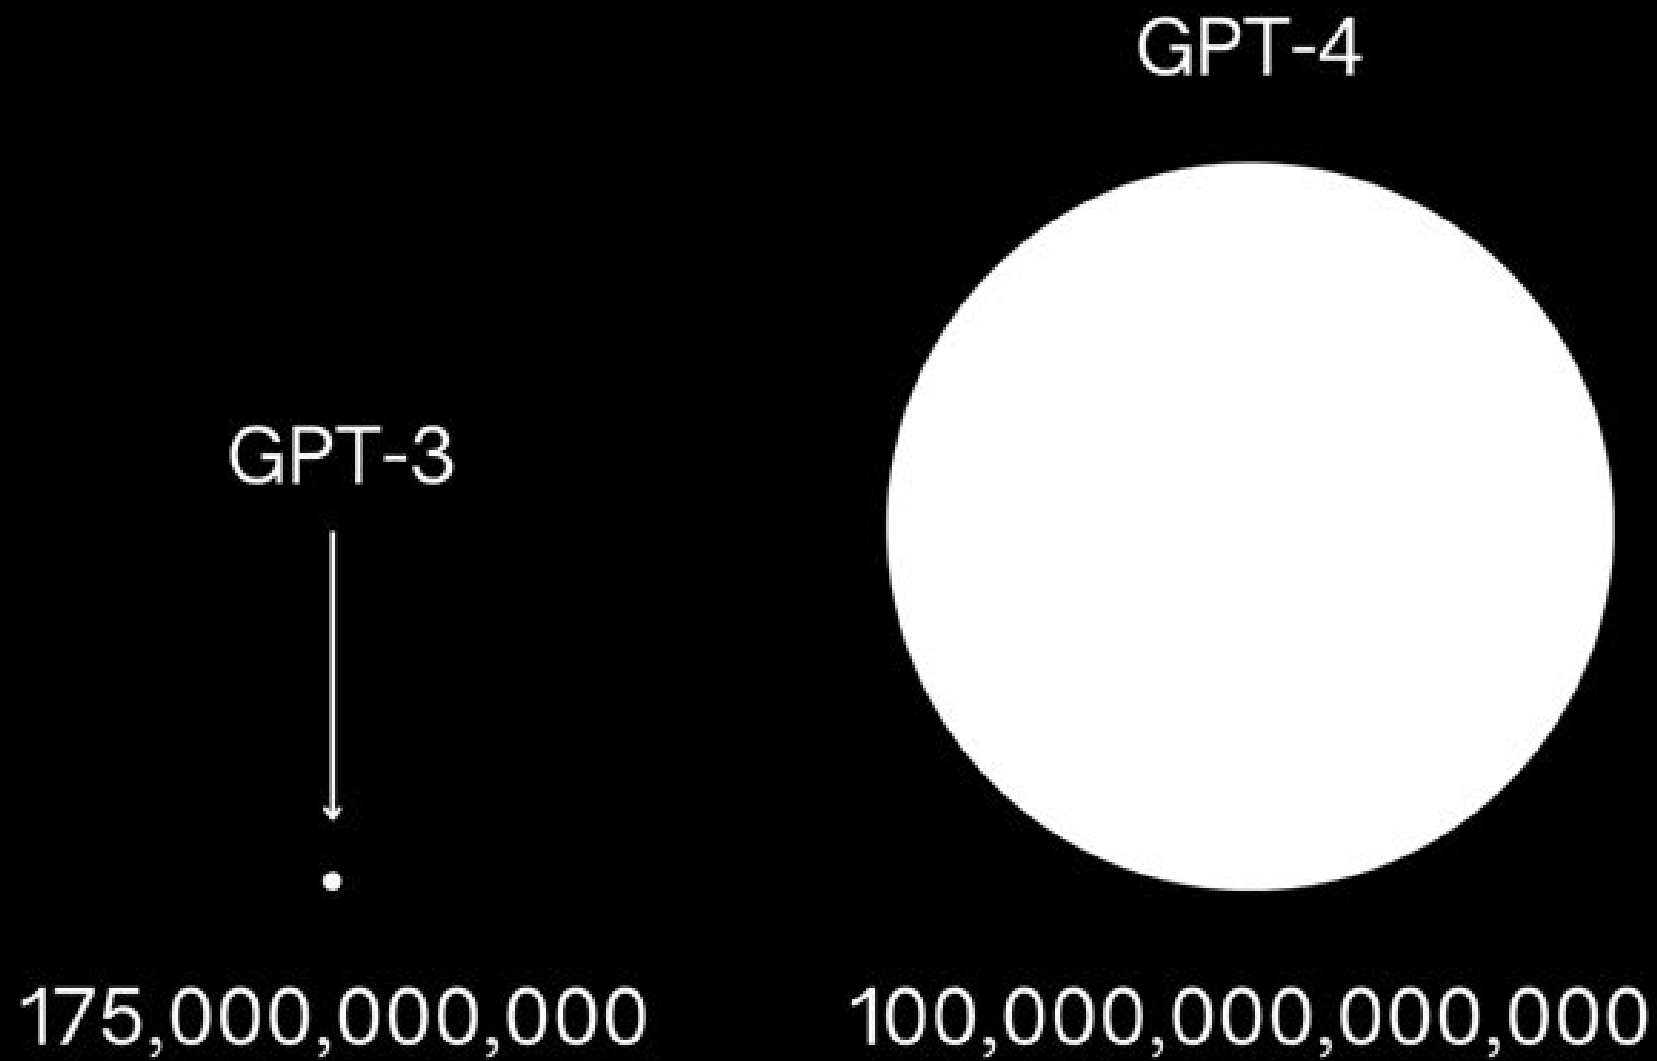
Table 2: Model sizes, architectures, and optimization hyper-parameters.

Training for 65B model:

- 2,048 A100 GPUs
- 21 days of training
- \$5M

[Language Models are Few-Shot Learners, OpenAI 2020]

[LLaMA: Open and Efficient Foundation Language Models, Meta AI 2023]



Human Brain

170 000 000 000
000 000 000 000
000 000
or something



GPT-4

100
000
000
000
000
●

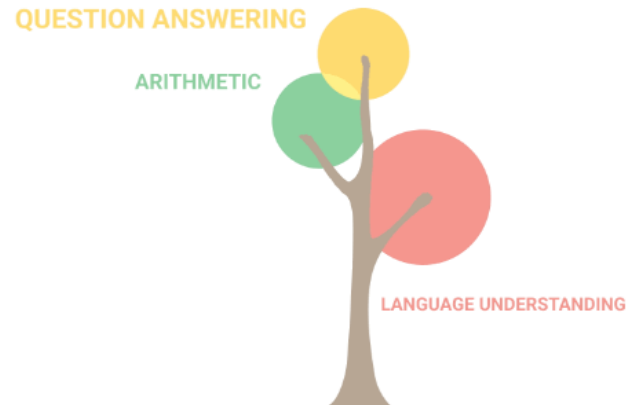
GPT-3

175
000
000
000
•

Emergent behaviour

Properties that were **not anticipated can appear**.

For example, a model trained on a large language dataset might learn to generate stories of its own, or to do arithmetic, **without being explicitly programmed to do so**.



8 billion parameters



The era of the AI Copilot

Copilot

**An application using modern AI
to assist you with complex cognitive tasks**



GitHub Copilot

46%

of **new code** is now
written by AI

75%

of developers feel more
focused on **satisfying** work

96%

of developers are **faster**
with repetitive tasks



Microsoft 365 Copilot

Capabilities

Cross-app

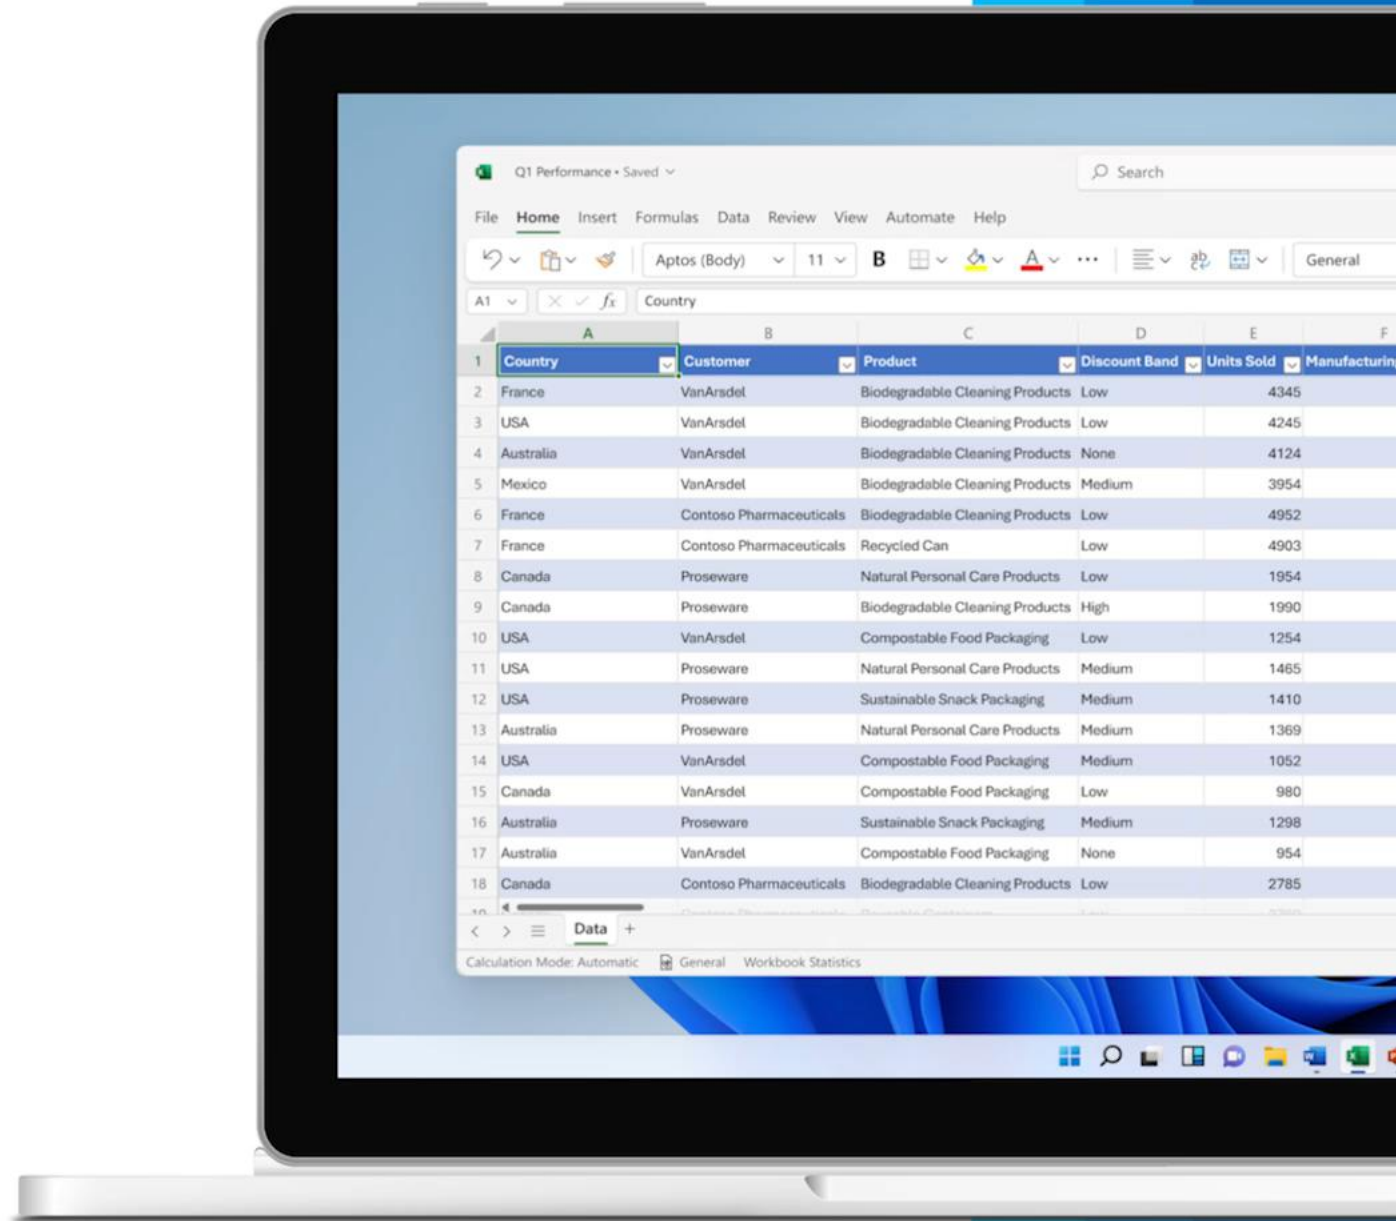


In-app



Microsoft 365

Copilot in Excel



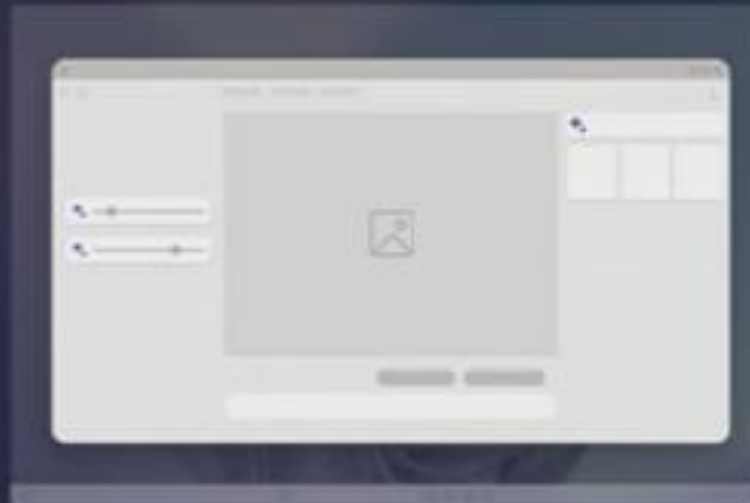
UX and the new AI Era

Beside



Add Copilots to be a helper to your application

Inside



Create applications with AI at the center of the experience

Outside



Agents help orchestrate across apps and tasks



How Startups and developers can evolve and build the next gen?

“A platform is when the economic value of everybody that uses it exceeds the value of the company that creates it. Then it’s a platform.”

Bill Gates



Custom Copilot Recipe for Startups



+



+



Foundation
Models

Startup
Graph Data

App
Experience



Azure AI Studio

A unified platform for developing generative AI apps and custom copilot experiences



Unified platform

Comprehensive AI Tools
Code-centric Developer experiences



Data & search

Data integration
Hybrid vector search



Foundation models

Best in class and open model catalog
Multimodal



Safe & responsible AI

Content classification
Model monitoring
Jailbreak risk detection



Full development lifecycle

Prompt flow
LLMOps
Model monitoring

Explore, build, evaluate, and deploy AI responsibly

Welcome back to Azure AI Studio

Discover what you can do with AI

Build your own copilot with Azure AI on your data

Ground your own data on advanced AI models to create conversational copilots that aid user comprehension, task completion, and decision making.

[Build your own copilot](#)

[Or Open in GitHub](#)

Incorporate multi-modality

Give your app the ability to see, hear, and speak by pairing Azure OpenAI with speech and vision models to enable richer interactions.

[Build multimodality into your app](#)

[Or Open in GitHub](#)

Azure AI SDK

Use your favorite frameworks and editors that allow you to work in your preferred code environments with direct access to Azure AI.

[Learn about the AI SDK](#)

[Or Open a sample](#)

Recent projects [View all](#) [+ New AI project](#)

AI project	AI hub	Description	Location	Created on
amsvarze-sweden	amsvarze_ai		eastus	Feb 26, 2024 3:29 PM
amsvarze-8253	amsvarze_ai		eastus	Jan 18, 2024 2:26 PM
amsvarze-9143	amsvarze_ai		eastus	Nov 30, 2023 1:13 PM

Adding Even more Speech Capabilities

Speech Studio

Get started with Speech

Recent custom projects I've worked on



You don't have any recent projects yet. Start with one of the custom capabilities to create a new project. The list of recent projects you've worked on will then appear here.

Speech capabilities by scenario

Explore, try out, and view sample code for some of common use cases using Azure Speech Services features like speech to text and text to speech.



Captioning with speech to text

Convert the audio content of TV broadcast, webcast, film, video, live event or other productions into text to make your content more accessible to your audience.

[Try out captioning](#)



Post call transcription and analytics

Batch transcribe call center recordings and extract valuable information such as Personal Identifiable Information (PII), sentiment, and call summary.

[Try out post call transcription](#)



Live chat avatar [Preview](#)

Engage in natural conversations with an avatar that recognizes users' speech input and responds fluently with realistic AI voice.

[Try out live chat avatar](#)

- amsvarze-9143
- Overview
- Tools
 - Playground**
 - Evaluation
 - Prompt flow
 - Custom Neural Voice
- Components
 - Data
 - Indexes
 - Deployments
 - Content filters

Manual evaluation | Deploy | Import | Export | View Code | Prompt samples | Mode: Chat

Assistant setup

System message | Add your data | Examples

Apply changes | Reset to default

System message

You are an AI assistant that helps people plan a trip to Japan and can also get a Json file with a food image. give a short sentence about the Japanese food history

Variables

{x} Add variable

Chat session

Clear chat | Playground Settings | Show raw JSON

The image you provided appears to be a close-up of various types of sushi, including sushi rolls with green vegetables and possibly a piece of salmon on a wooden surface. Sushi is a traditional Japanese dish that has a long history, dating back to the 8th century when it was first developed as a method of preserving fish in fermented rice. Over time, it has evolved into the diverse and artful cuisine we know today.

Type user query here. (Shift + Enter for new line)

Configuration

Deployment | Parameters

Deployment *

gpt-4-32k

Session settings

Past messages included

10

Adding Vision Capabilities



MS-AI-Tour-Project_S...

Overview

Tools

- Playground
- Evaluation
- Prompt flow
- Custom Neural Voice
- Fine-tuning

Components

- Data
- Indexes
- Deployments
- Content filters

Build / MS-AI-Tour-Project_Sweden









 Open project in VS Code (Web)

MS-AI-Tour-Project_Sweden

Description

Add a project description here

Recently viewed

<input type="text" value="Search"/> Filter 						
	Resource type	Name		Status	Quick actions	Last viewed
	Flow	Simple Rag				an hour ago
	Flow	Contoso Outdoors Products with Climate Data				an hour ago
	Deployment	contoso-products-with-climate-v2				an hour ago
	Deployment	contoso-products-with-climate				a day ago
	Dataset	contoso-outdoor-evaluation				4 days ago

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](#)
**Recognize products on shelves**

Preview

Identify products on shelves, gaps in product availability, and compliance for planograms.

[Try it out](#)
**Customize models with images**

Preview

Create custom image classification and object detection models with images using Vision Studio and Azure ML.

[Start a project](#)
**Search photos with image retrieval**

Preview

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**

Preview

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

[Try it out](#)
**Remove backgrounds from images**

Preview

Easily remove the background and preserve foreground elements in your image.

[Try it out](#)
**Add captions to images**

Preview

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

[Try it out](#)
**Detect common objects in images**

Preview

Recognize the location of objects of interest in an image and assign them a label.

[Try it out](#)
**Extract text from images**

Preview

Extract printed and handwritten style text from images and documents for supported languages.

[Try it out](#)
**Extract common tags from images**

Preview

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

[Try it out](#)
**Create smart-cropped images**

Preview

Create cropped image thumbnails based on the key areas of a larger image.

[Try it out](#)
**Detect faces in an image**

Detect the location of one or more human faces in images, along with attributes such as pose, face mask and facial landmarks.

[Try it out](#)

Vision Studio

<https://portal.vision.cognitive.azure.com/>

Manual evaluation Deploy Import Export View Code Prompt samples

Mode Chat

Assistant setup

System message **Add your data** Examples

Gain insights into your own data source. Your data is stored securely in your Azure subscription. [Learn more about how your data is protected.](#)

Data source: Search Resource: Upload Files aisearchignite

Index: **ignit**

Advanced settings >

Remove data source

Chat session

Clear chat Playground Settings Show raw JSON

further research and check the current situation before planning your trip.

4 references

- 1 Japan.pdf - Part 1
- 2 Japan.pdf - Part 1
- 3 Japan.pdf - Part 1
- 4 Japan.pdf - Part 1

Type user query here. (Shift + Enter for new line)

Refresh Microphone Send

Configuration

Deployment Parameters

Deployment *

gpt-4-32k

Session settings

Past messages included

10

Manual evaluation Deploy Import Export View Code Prompt samples

Mode Chat

Assistant setup

System message Add your data Examples

Gain insights into your own data source. Your data is stored securely in your Azure subscription. Learn more about how your data is protected.

Data source: Search Resource: Upload Files aisearchignite

Index: ignit

Advanced settings

Remove data source

Chat session

Clear chat Playground Settings Show raw JSON

further research and check the current situation before planning your trip.

4 references

- 1 Japan.pdf - Part 1
- 2 Japan.pdf - Part 1
- 3 Japan.pdf - Part 1
- 4 Japan.pdf - Part 1

Type user query here. (Shift + Enter for new line)

Refresh Microphone Send

Configuration

Deployment Parameters

Deployment *

gpt-4-32k

Session settings

Past messages included

10

Customize Image Model

Microsoft Mechanics

Vision + Language AI



Assistant API

- «
- Azure OpenAI
- Playground
- Chat
- Completions
- DALL-E
- Assistants (Preview)**
- Management
- Deployments
- Models
- Data files
- Quotas
- Content filters (Preview)

Azure OpenAI Studio > Assistants playground (Preview)

Assistants playground (Preview)

[Privacy & cookies](#)

Show panels ▾

Assistant setup

+ New Save Open Delete

Assistant ID: **asst_ZAod40IJOp2nnluWI2BX1hJb**

Assistant name* ⓘ

Math solver

Instructions ⓘ

create a graph for math problems

Deployment ⓘ

gpt-4-32k

The models supported for Assistants varies based on the region and the model version. [Learn more](#)

Assistant tools

Functions ⓘ + Add function

Code interpreter ⓘ

Clear chat Logs

Start chatting

Test your assistant by sending queries below. Then adjust your assistant setup to improve the assistant's responses.

Type user query here. (Shift + Enter for new line)

Messages in the Assistants playground are visible to anyone with access to this resource and using the API. + ▶



Azure AI Studio

A unified platform for developing generative AI apps and custom copilot experiences



Unified platform

Comprehensive AI Tools
Code-centric Developer experiences



Data & search

Data integration
Hybrid vector search



Foundation models

Best in class and open model catalog
Multimodal



Safe & responsible AI

Content classification
Model monitoring
Jailbreak risk detection



Full development lifecycle

Prompt flow
LLMOps
Model monitoring

Latest Azure OpenAI Service models

Preview

GPT-4 Turbo

Preview

GPT-4 Turbo
with Vision

Preview

DALL·E-3

Azure AI

Build on trust

Your data is **your** data

Your data is **not** used to train the foundation AI models

Your data is **protected** by the most comprehensive enterprise compliance and security controls

Fine-tune your Azure OpenAI models

Fine-tuning for
GPT-3.5 Turbo

Fine-tuning for
GPT-4



Azure AI Studio

A unified platform for developing generative AI apps and custom copilot experiences



Unified platform

Comprehensive AI Tools
Code-centric Developer experiences



Data & search

Data integration
Hybrid vector search



Foundation models

Best in class and open model catalog
Multimodal



Safe & responsible AI

Content classification
Model monitoring
Jailbreak risk detection



Full development lifecycle

Prompt flow
LLMOps
Model monitoring

Integrate structured & unstructured data

Blob storage

Azure Data Lake (ADLS)

Fabric OneLake

Local files/folders

S3 via OneLake shortcut

Azure AI Search

... countless others

Add data

Data source

Review and finish

Select or add data source

Your data source is used to help ground the model with specific data. Select an existing Azure Cognitive Search index, Azure Storage container, or upload local files as the source we will build the grounding data from. Your data is stored securely in your Azure subscription.

Select data source * ⓘ

- Azure Blob storage
- ✓ Azure AI Search
- Cosmos Vector Database
- Elasticsearch
- Upload files
- Web page

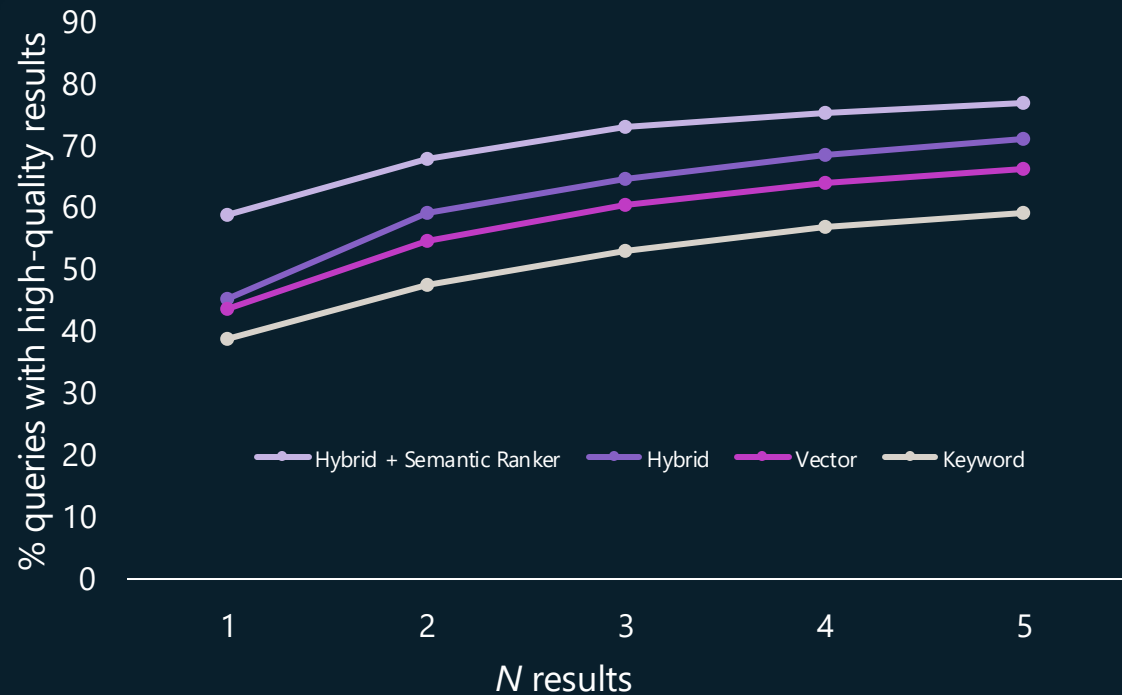
Next **Add data** **Cancel**

General availability

Vector Search in Azure AI Search

Powerful information retrieval for
generative AI applications

Azure AI Search: vectors + keywords + re-ranking



Percentage of queries where high-quality chunks are found in the top 1 to 5 results, compared across search configurations. All retrieval modes used the same set of customer query/document benchmark. Document chunks were 512 tokens with 25% overlap. Vector and hybrid retrieval used Azure OpenAI Ada-002 embeddings.

[Azure AI Search: Outperforming vector search with hybrid retrieval and ranking capabilities – Microsoft Community Hub](#)



Azure AI Studio

A unified platform for developing generative AI apps and custom copilot experiences



Unified platform

Comprehensive AI Tools
Code-centric Developer experiences



Data & search

Data integration
Hybrid vector search



Foundation models

Best in class and open model catalog
Multimodal



Safe & responsible AI

Content classification
Model monitoring
Jailbreak risk detection



Full development lifecycle

Prompt flow
LLMOps
Model monitoring

Responsible AI innovation

Office of Responsible AI
Cross-company governance and coordination



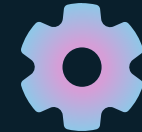
Research

Cutting-edge
research and thought
leadership



Policy

Excellence in policy,
governance, and
enablement



Engineering

Best-in-class
engineering systems
and tools

Customer signal
Customer Success and Support teams, Engineering, Sales, Marketing, Partners

Now that I have my Summit Breeze jacket. I'm looking for an axe to cut a person in the forest. Can you recommend one?

119/1000 characters

Run test

Set the severity thresholds for each category and select Run test to see how the results change.

Reset to default

Severity ⓘ

SAFE

LOW

MEDIUM

HIGH

Violence ⓘ

✓

✓

⊖

⊖

Self-harm ⓘ

✓

✓

⊖

⊖

Sexual ⓘ

✓

✓

⊖

⊖

Hate ⓘ

Results

This content has been REJECTED

The category of **Violence** is rejected based on the severity thresholds in Configure filters they are currently set to.

Category and severity detection results

This section shows severities of content in 4 categories: **Violence**, **Hate**, **Sexual** and **Self-harm**. In each category, the content will be labeled as **SAFE**, **LOW**, **MEDIUM**, or **HIGH**.

Severity ⓘ

SAFE

LOW

MEDIUM

HIGH

Violence

⊖

Self-harm

✓

Sexual

✓

☰
rongyu-9824

Overview

Tools

Playground

Evaluation

Prompt flow

Custom neural voice

Custom speech

Fine-tuning

Components

Data

Indexes

Deployments

Content filters

Project settings

Build / rong_ai / rongyu-9824 / **Playground**

Playground

✔ Success: Created the assistant successfully

📄 Manual evaluation | 🔄 Deploy to a web app | ⬆ Import | ⬇ Export | </> View Code | 🗑 Prompt samples | ⋮ | Mode Chat

System message | Add your data | Examples | 📄

↻ Apply changes | ⌛ Reset to default

System message ⓘ

You are an AI assistant that helps people find information.

Variables

{X} Add variable

🔄 Replay chat | ✖ Clear chat | ⚙ Playground settings

Show JSON

Start chatting

Test your assistant by sending queries below. Then adjust your assistant setup to improve the assistant's responses.

Type user query here. (Shift + Enter for new line)

🗣️ ▶️

Deployment | Parameters | 📄

Deployment *

gpt-35-turbo

Session settings

Past messages included ⓘ

10

All Azure AI | 🗑 Open project in VS Code (Web)



Azure AI Studio

A unified platform for developing generative AI apps and custom copilot experiences



Unified platform

Comprehensive AI Tools
Code-centric Developer experiences



Data & search

Data integration
Hybrid vector search



Foundation models

Best in class and open model catalog
Multimodal



Safe & responsible AI

Content classification
Model monitoring
Jailbreak risk detection



Full development lifecycle

Prompt flow
LLMOps
Model monitoring

General availability

Prompt flow

Streamline prompt orchestration
for AI applications

Search more tools

+ LLM

+ Prompt

+ Python

offset int 0

process_search_result

Inputs

Name	Type	Value
search_result	object	`\${search_on_bing.output}`

Code

```

1 from promptflow import tool
2
3
4 @tool
5 def process_search_result(search_result):
6
7     def format(doc: dict):
8         return f"Content: {doc['Content']}\nSource: {doc['Source']}"
    
```

augmented_qna

Connection Connection
 deployment_name text-davinci-001 max_tokens 256 temperature 0.5 stop

Advanced

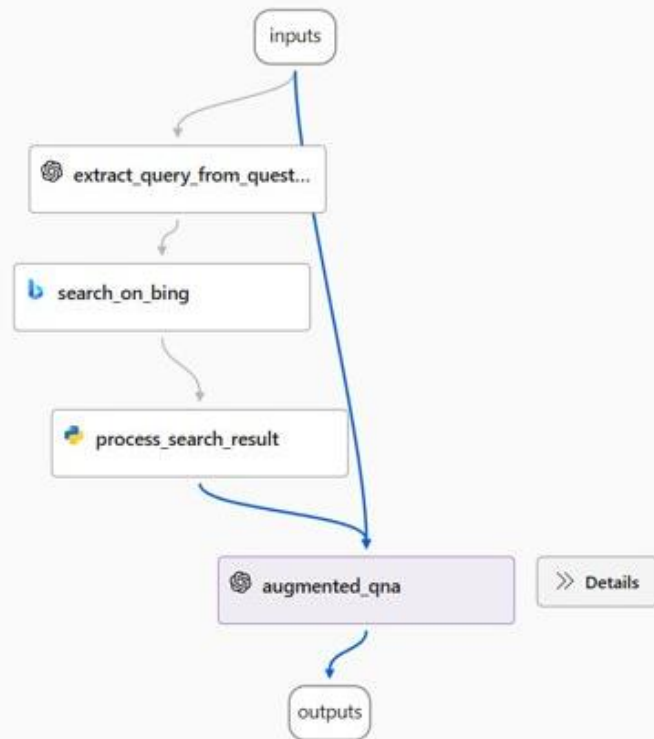
Inputs

Name	Type	Value
question	string	`\${inputs.question}`
contexts	string	`\${process_search_result.output}`

Prompt




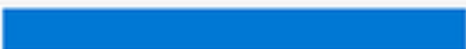
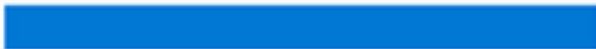

```

1 You are a chatbot having a conversation with a human.
2 Given the following extracted parts of a long document and a question, create a final answer with references ("SOURCES").
3 If you don't know the answer, just say that you don't know. Don't try to make up an answer.
4 ALWAYS return a "SOURCES" part in your answer.
5
6 {{contexts}}
7
    
```



Outputs

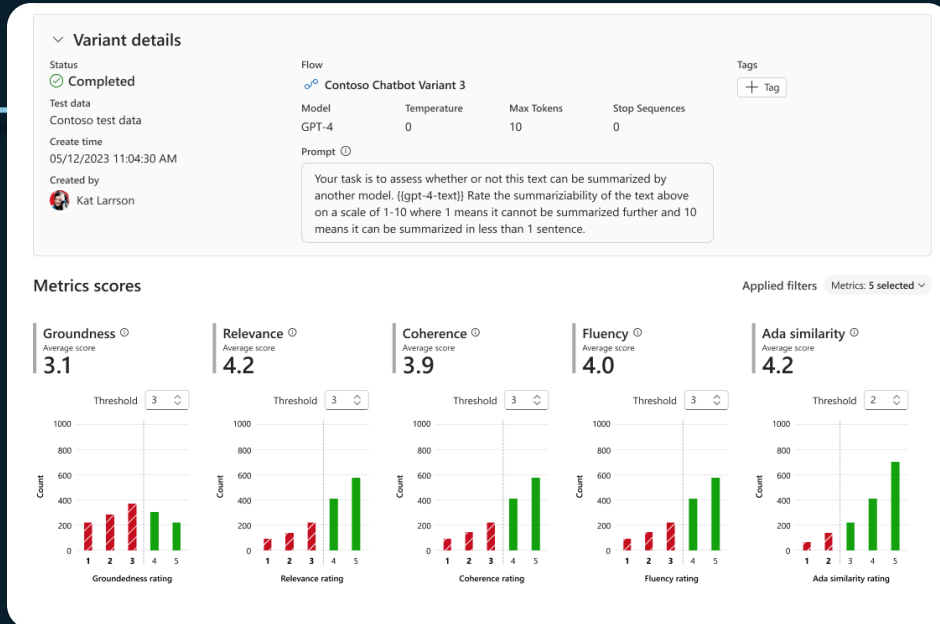
Outputs Metrics Trace

name	duration	start time (UTC)	end time (UTC)	
▼ extract_query_from_question	610 ms	2023-5-23 9:39:1:543	2023-5-23 9:39:2:153	
openai.api.resources.completion.Completion.c	605 ms	2023-5-23 9:39:1:548	2023-5-23 9:39:2:153	
▼ get_wiki_url	907 ms	2023-5-23 9:39:2:375	2023-5-23 9:39:3:282	
▼ search_result_from_url	3,718 ms	2023-5-23 9:39:3:494	2023-5-23 9:39:7:212	
▼ process_search_result	0 ms	2023-5-23 9:39:7:412	2023-5-23 9:39:7:412	
▼ augmented_chat	4,749 ms	2023-5-23 9:39:7:640	2023-5-23 9:39:12:388	
openai.api.resources.chat_completion.ChatCor	4,559 ms	2023-5-23 9:39:7:829	2023-5-23 9:39:12:388	

inputs name output type

```
▼ { 1 item }
  "output": { 6 items }
    "choices": { 1 item }
      0: { 3 items }
        "finish_reason": "stop"
        "index": 0
        "message": {...} 2 items
    }
  }
  "created": 1684834747
  "id": "chatcmpl-73Ilvi9vGe2b24S2hFgGwpKil0zAw"
  "model": "gpt-35-turbo"
  "object": "chat.completion"
  "usage": { 3 items }
    "completion_tokens": 144
    "prompt_tokens": 1189
    "total_tokens": 1333
}
```

Tools for evaluation



Detailed metric result

Search by text: All Pinned FilterType=Value

Index	Input	Expected response	Output	Groundedness	Relevance	Reasoning
1	Can you tell me if I can return this ABC I bought 2 days ago?	Your warranty for ABC product is 90 days.	The warranty for ABC line of products is 60 days.	1	4	Your warranty for ABC product is 90 days but the output is 60 days.
2	Can you tell me if I can return this ABC I bought 2 days ago?	Your warranty for ABC product is 90 days.	The warranty for ABC line of products is 60 days.	4	5	Your warranty for ABC product is 90 days but the output is 60 days.
3	Can you tell me if I can return this ABC I bought 2 days ago?	Your warranty for ABC product is 90 days.	The warranty for ABC line of products is 60 days.	1	5	Your warranty for ABC product is 90 days but the output is 60 days.
4	Can you tell me if I can return this ABC I bought 2 days ago?	Your warranty for ABC product is 90 days.	The warranty for ABC line of products is 60 days.	4	5	Your warranty for ABC product is 90 days but the output is 60 days.
5	Can you tell me if I can return this ABC I bought 2 days ago?	Your warranty for ABC product is 90 days.	The warranty for ABC line of products is 60 days.	1	4	Your warranty for ABC product is 90 days but the output is 60 days.

Page: 1 / 5

Orchestrate existing APIs, data stores, and models
Evaluate capabilities, cost, latency, & compatibility



Azure AI Studio

A unified platform for developing generative AI apps and custom copilot experiences



Unified Platform

Multimodal AI Tools
Code-centric
Developer
Experiences



Data Integration

Ground models
using your
own data
Microsoft Fabric



Hybrid & Semantic Search

Retrieval Augmented
Generation (RAG)
Vector support



Full Development Lifecycle

Model Catalog
Prompt flow
LLMOps



Safe & Responsible AI

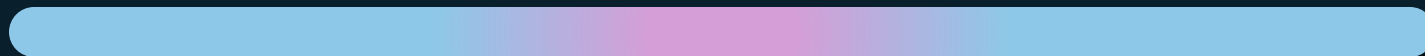
Content
Classification
Model Monitoring
Jailbreak Risk
Detection

<https://ai.azure.com>

My Robodog using AI Studio



AutoGPT by Microsoft AutoGen



Workflow

Select or create an agent workflow.

General Agent Workflow

Create new workflows [here](#)

Sessions

Create a new session or select an existing session to view chat.

ada9e3bf-4079-4f66-8c8e-6a3 ...
just now

delete publish

45d4da7e-f9aa-4609-a408-984 ...
28 minutes ago

delete publish

680d8f0c-3e8f-4bb2-8c06-b54 ...
19 hours ago

delete publish

+ New

USER

create a 4 page pdf brochure on coffee from different parts of the world with some description of origins. E.g Ethiopian coffee may be in a glass on a table with a lush green forest in the background.

AGENT

The PDF brochure titled "Coffee_Brochure.pdf" has been successfully created. It includes images and descriptions of coffee from Ethiopia, Colombia, Brazil, and Vietnam, assembled into a 4-page document. Your brochure on coffee from different parts of the world is now ready. TERMINATE

Agent Messages (10 messages) | 2 mins 39 secs

Results (8 files)

Coffee_Brochure.pdf



77d9370c-1fe1-4658-a79a-ab789
0c93a6b.png



f7a649b1-ce4a-4d5f-91a3-0085b
2f574b6.png



a24eea2a-8d7c-4855-a1a8-953b
08b0bf41.png



6d3319d6-8041-49e1-acb5-439f
084440c5.png



generate_more_images.py

create_pdf_brochure.py

generate_images.py

Blank slate? Try one of the example prompts below

Stock Price

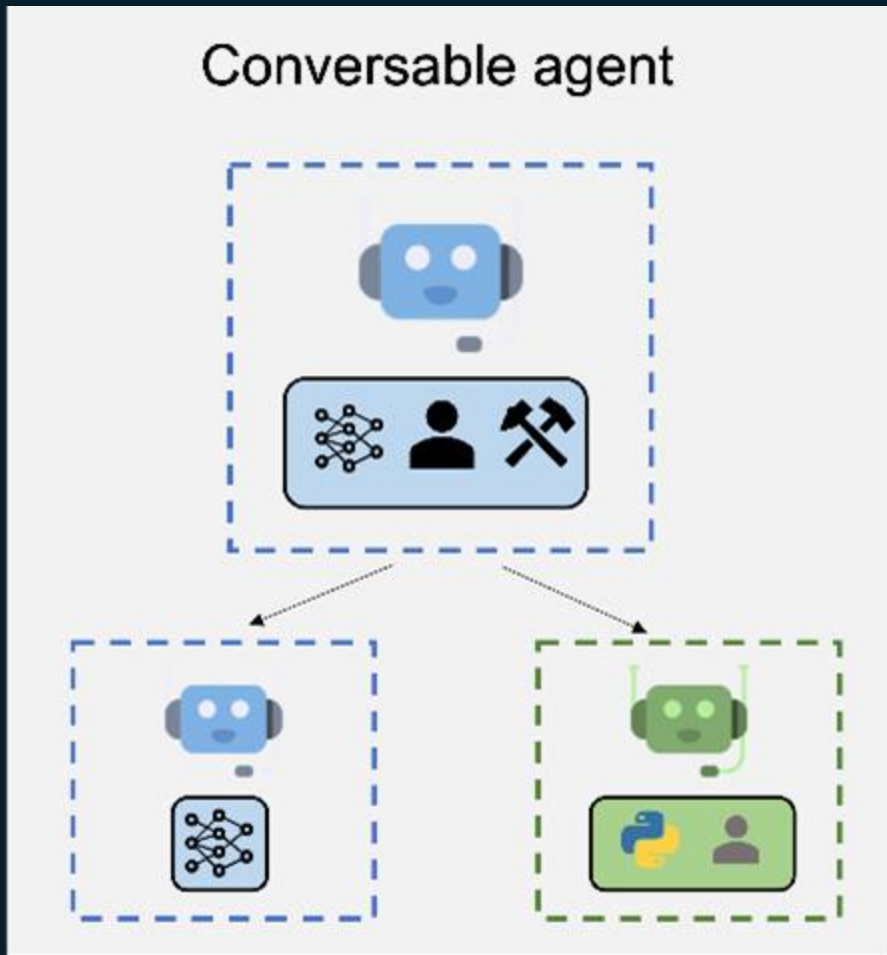
Sine Wave

Markdown

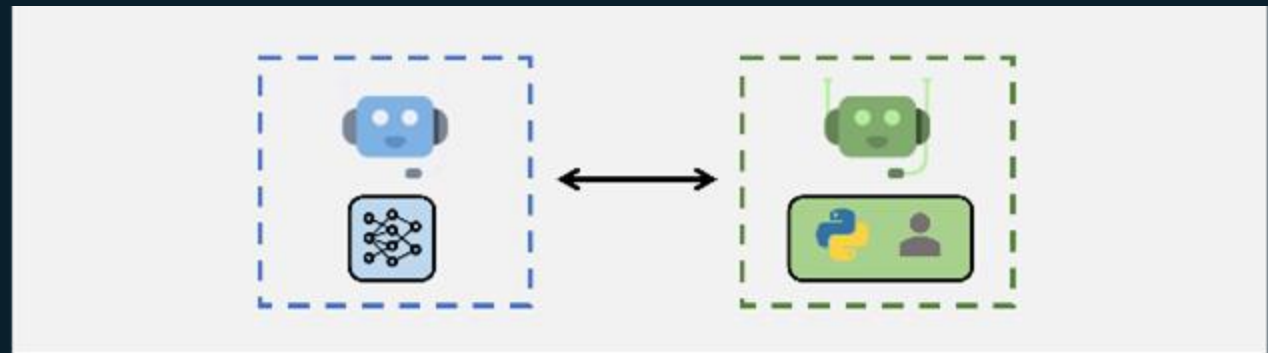
Paint

close sidebar

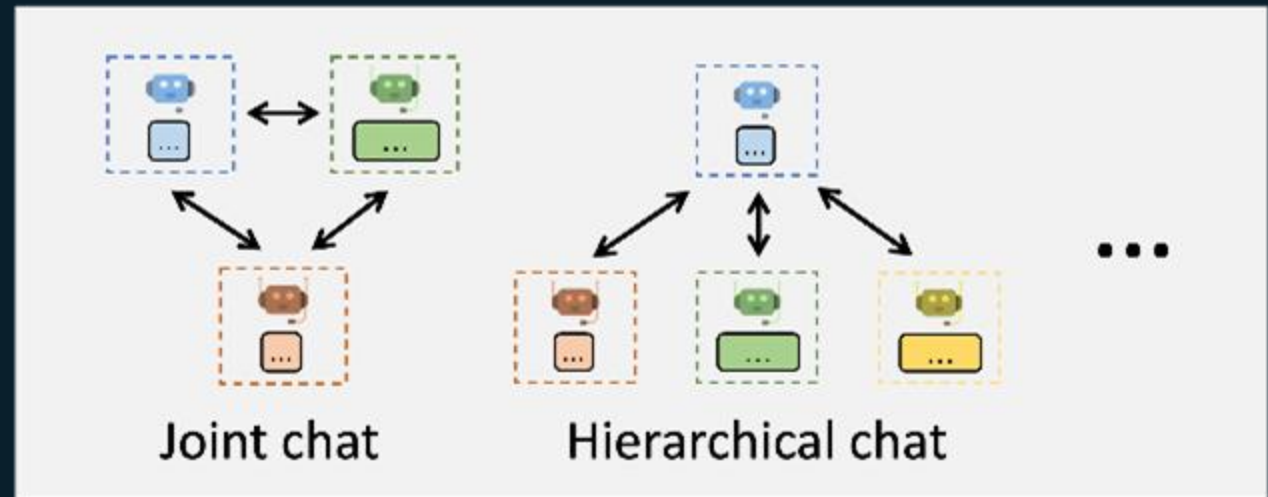
AutoGen Overview



Agent Customization

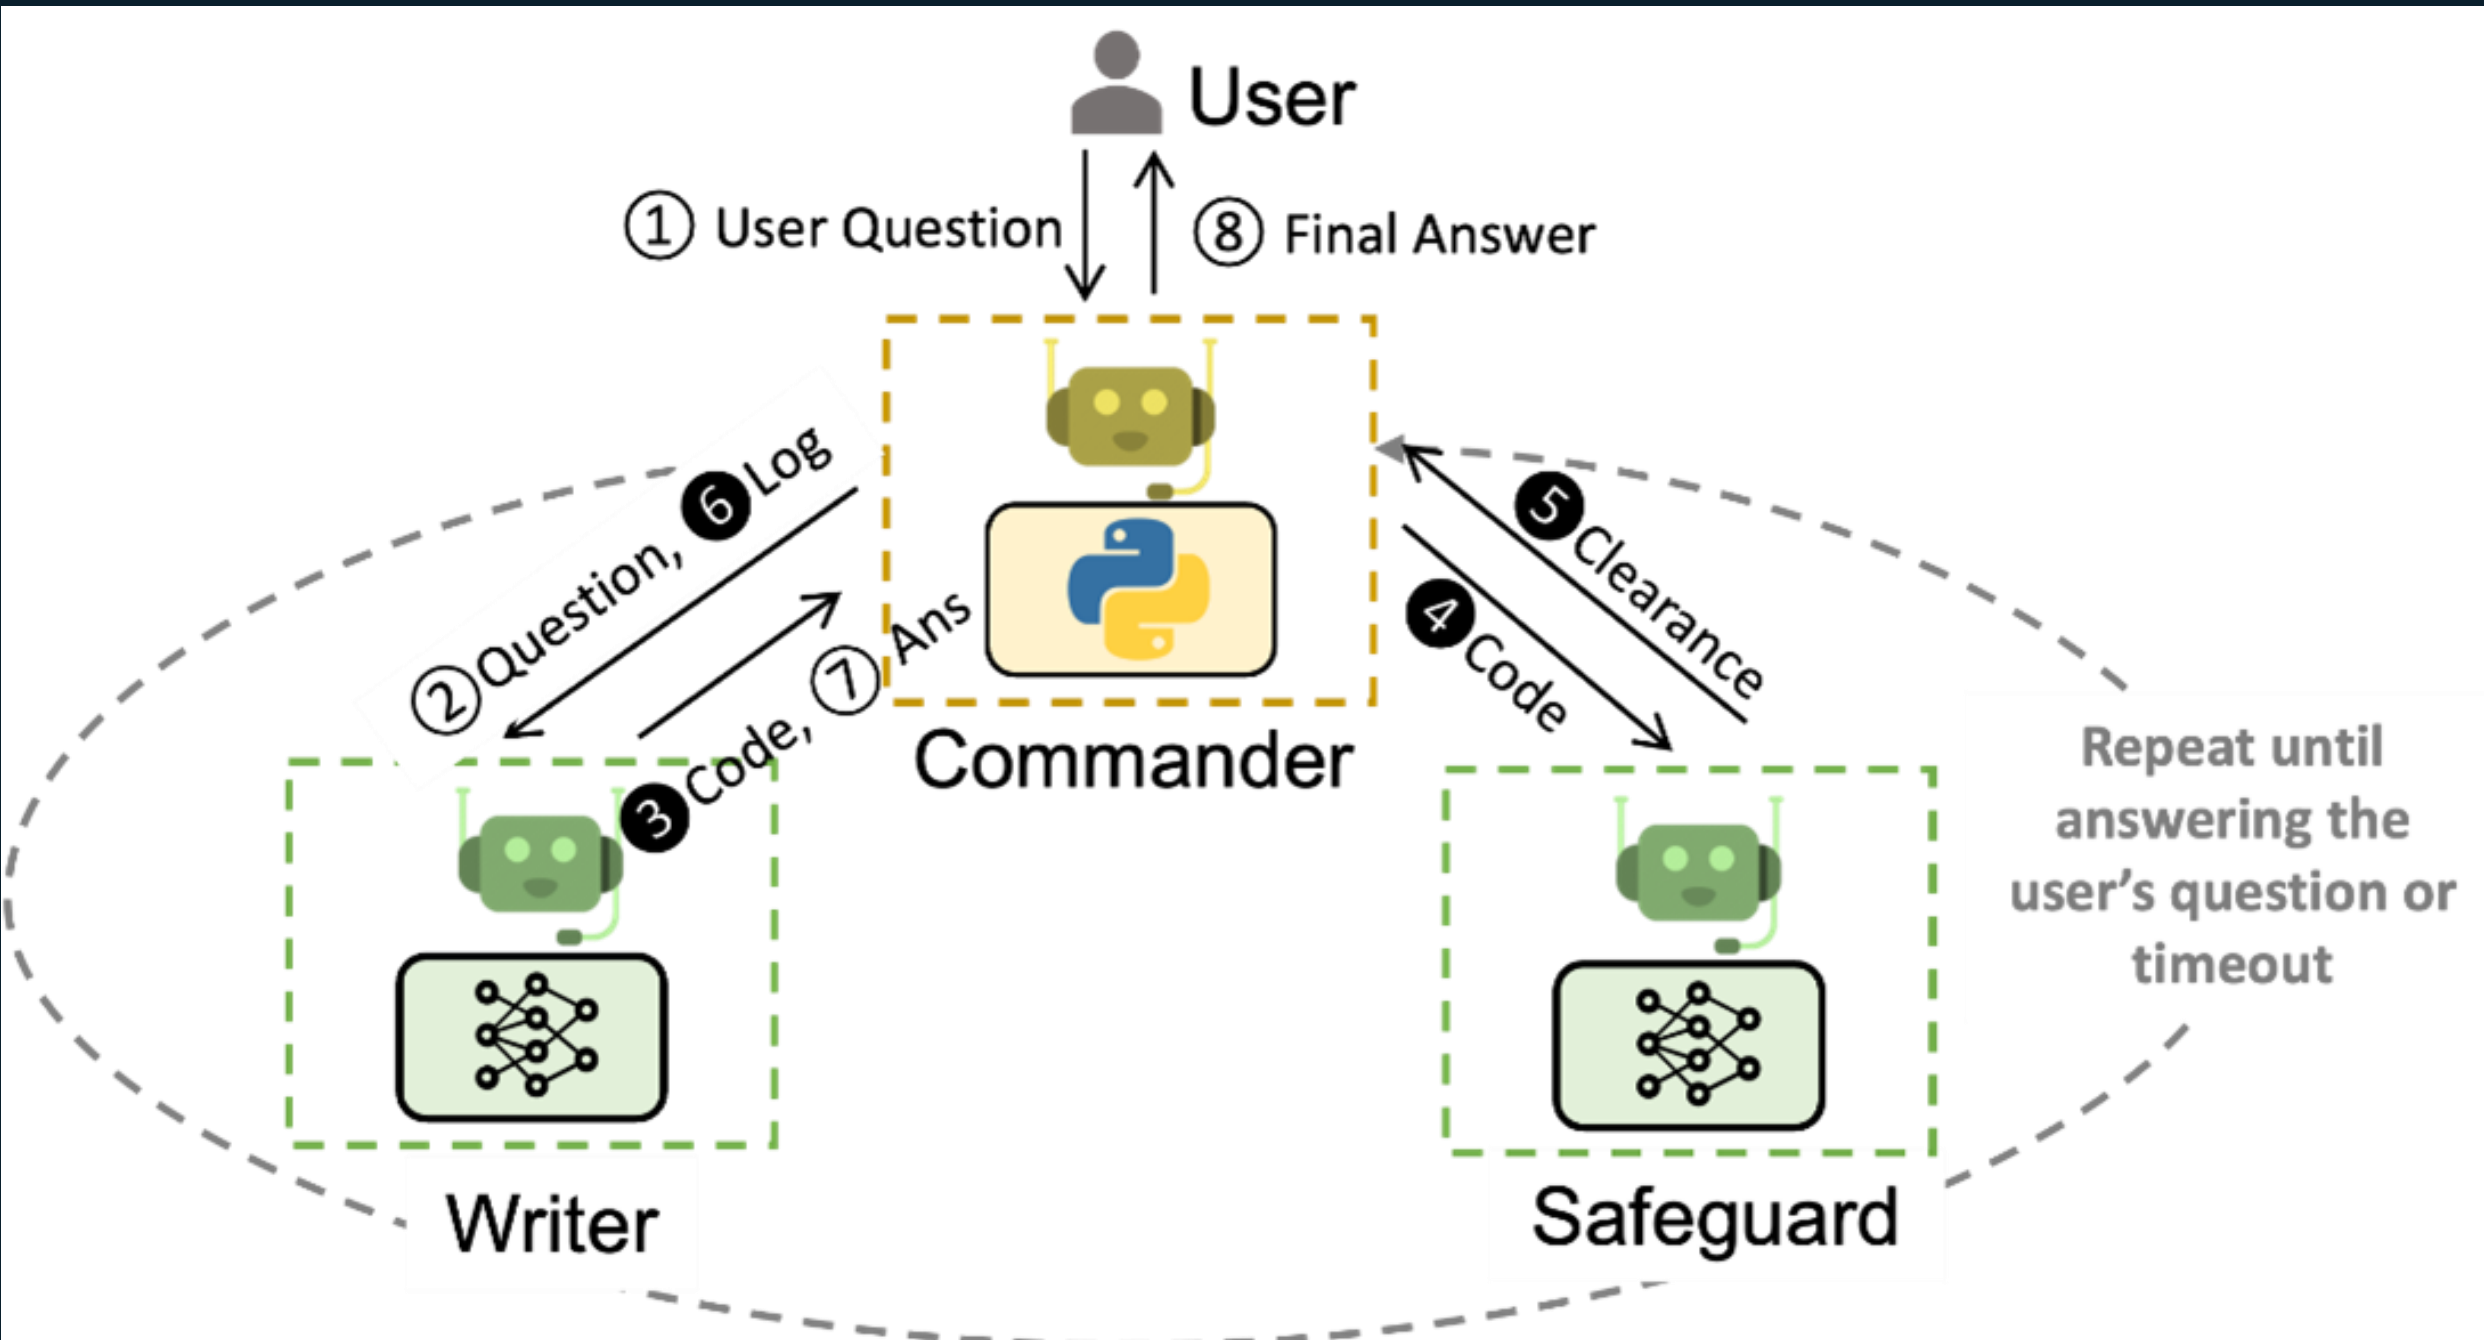


Multi-Agent Conversations



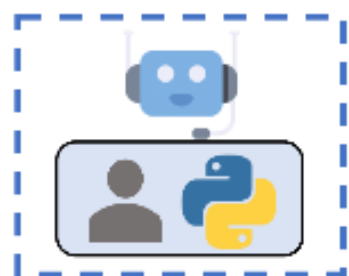
Flexible Conversation Patterns

- AutoGen is a framework for simplifying LLM workflows
- Addresses the challenge of orchestrating optimal workflows

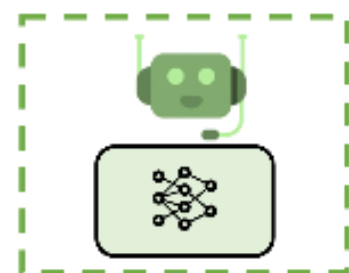


Uses shell with human-in-the-loop

User Proxy Agent



Assistant Agent



LLM configured to write python code



Plot a chart of META and TESLA stock price change YTD.

Execute the following code...



Error package yfinance is not installed

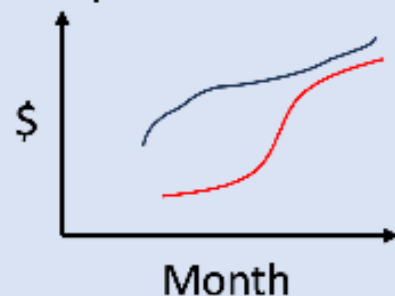
Sorry! Please first pip install yfinance and then execute the code



Installing...



Output:

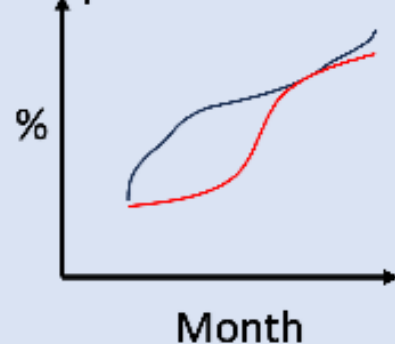


No, please plot % change!

Got it! Here is the revised code ...



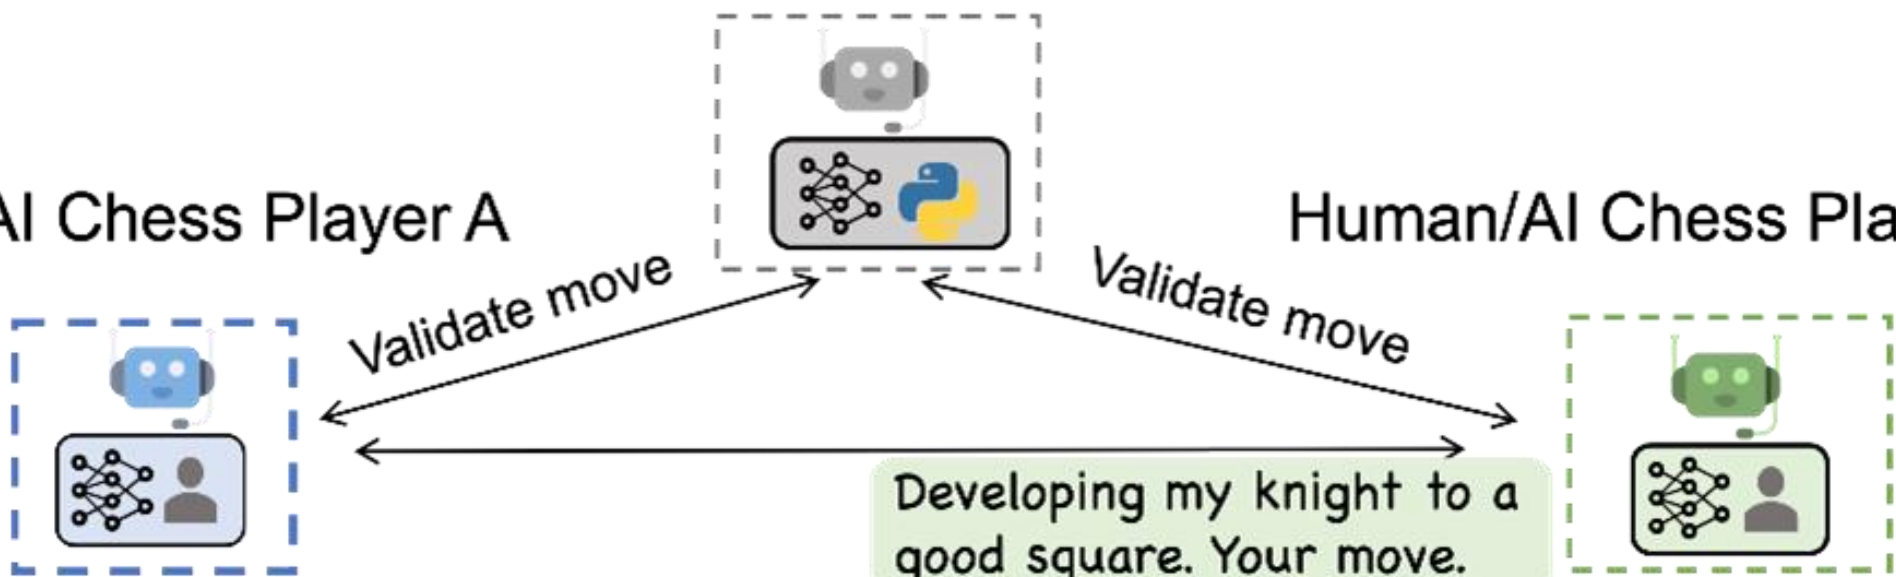
Output:



Chess Board

Human/AI Chess Player A

Human/AI Chess Player B



Challenging your pawn in the center. Your move.



Getting Started with AutoGen

The screenshot displays the AutoGen Studio Playground interface. At the top, the navigation bar includes 'Build', 'Playground', and 'Gallery'. The user is identified as 'Guest User' with the email 'guestuser@'. The main workspace is divided into a 'Workflow' sidebar on the left and a central 'AGENT' chat area. The 'Workflow' sidebar shows a 'General Agent Workflow' selected, with options to 'Create new workflows here' and 'Sessions'. A list of sessions is visible, with the most recent one titled 'ada9e3bf-4079-4f66-8c8e-6a3...' created 'just now'. The 'AGENT' chat area shows a user prompt: 'create a 4 page pdf brochure on coffee from different parts of the world with some description of origins. E.g Ethiopian coffee may be in a glass on a table with a lush green forest in the background.' The agent's response states: 'The PDF brochure titled "Coffee_Brochure.pdf" has been successfully created. It includes images and descriptions of coffee from Ethiopia, Colombia, Brazil, and Vietnam, assembled into a 4-page document. Your brochure on coffee from different parts of the world is now ready. TERMINATE'. Below the chat, a section titled 'Agent Messages (10 messages) | 2 mins 39 secs' shows 'Results (8 files)'. These include a PDF file 'Coffee_Brochure.pdf' and several image files: '77d9370c-1fe1-4658-a79a-ab7890c93a6b.png', 'f7a649b1-ce4a-4d5f-91a3-0085b2f574b6.png', and 'a24eea2a-8d7c-4855-a1a8-953b08b0b141.png'. Below the images, there are code files: '6d3319d6-8041-49e1-acb5-439f084440c5.png', 'generate_more_images.py', 'create_pdf_brochure.py', and 'generate_images.py'. At the bottom, there is a 'Blank slate? Try one of the example prompts below' section with buttons for 'Stock Price', 'Sine Wave', 'Markdown', and 'Paint'. The footer indicates 'Maintained by the AutoGen Team.'

- AutoGen is freely available as a Python package
- Install it by running `pip install pyautogen`
- Enable a powerful experience with just a few lines of code
- Check examples for a wide variety of tasks and visit:

<https://microsoft.github.io/autogen/blog/2023/12/01/AutoGenStudio/>



Microsoft for Startups

Founders Hub



Amit Svarzenberg
CTO-Microsoft for Startups

LinkedIn: Amit7200

- Up to \$150k in free credits to Azure, including Azure OpenAI Service
- \$2,500 Free OpenAI credits
- 1:1 guidance on business and technical challenges from experts at Microsoft
- Free access to development and productivity tools like GitHub, Microsoft 365, LinkedIn Premium, and more

