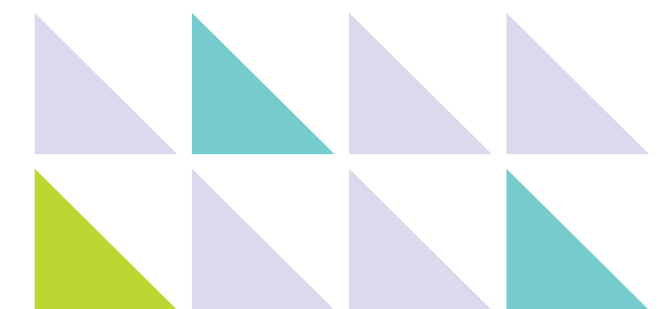


Dmitry Katson | MVP

- Business Central all-in-one:
 - architect
 - developer
 - team-leader
 - with 15+ years of experience
- Creator of [CentralQ.ai](#)



Stefano Demiliani | MVP

- MVP on Azure and Business Applications
- Long time MCT (inside and outside MS)
- Working full-time as a CTO:
 - Business Central architecture & development (started from Navision 2.60 ☹)
 - Azure serverless solutions design & development
- Author of some books about BC and Azure stuffs (that I hope you had the chance to read 😊)
- Lover of “*leaving the ERP doing the ERP*”



What you will, probably, learn today

- VS Copilot X
- chatGPT vs Azure OpenAI
- Prompt engineering
- Chat with Business Central
- Embeddings
- OpenAI SDK for .NET, Python
- Semantic Kernel



Freeing up time and focusing on what really matters



What wife can tell to the husband when she wants him to do some tasks



Certainly! Here are some phrases your wife can use to request tasks from you:

"Darling, tomorrow morning, could you drop off the dry cleaning? They said it would be ready by 10 am, and it would be wonderful to have those clothes fresh and ready for the week."

"Honey, I have a doctor's appointment tomorrow at 2 pm. Could you please take care of dinner? Maybe you could whip up one of your delicious pasta dishes?"

"Hey, Superman, tomorrow is our anniversary. Could you plan a surprise date night for us? It's always exciting to see what creative ideas you come up with."



Freeing up time and focusing on what really matters

Status

User ID

Object Type to Run

Object ID to Run

Object Caption to Run

Description

Job Queue Category Code

(There is nothing to show in this view)

D365BC Copilot Chat

Ask here ... anything?

SEND



GitHub Copilot X

- ✓ Searching for: **Github Copilot vs Github Copilot X comparison**
- ✓ Generating answers for you...

GitHub Copilot is a tool that helps developers write code by suggesting lines of code based on the context of the code they are writing. It is powered by OpenAI's GPT-3 language model, which has been trained on a massive dataset of code ¹.

GitHub Copilot X is currently a representation of GitHub's vision for the future rather than an available product offering of GitHub Copilot. It is evolving to bring chat and voice interfaces, support pull requests, answer questions on docs, and adopt OpenAI's GPT-4 for a more personalized developer experience ² ³.



GitHub Copilot X



Here is a table summarizing the differences between **GitHub Copilot** and **GitHub Copilot X**:

Feature	GitHub Copilot	GitHub Copilot X
Availability	Available now	Not yet available
Features	Suggests lines of code based on context	Evolving to bring chat and voice interfaces, support pull requests, answer questions on docs, and adopt OpenAI's GPT-4 for a more personalized developer experience
Language model	GPT-3	Evolving to adopt OpenAI's GPT-4

<https://github.com/features/preview/copilot-x>



D365BC Copilot Chat (prototype)

D365BC Copilot Chat

SEND

```
<style>textarea { border-radius: 10px; } .btn { border: none; color: white; padding: 15px 28px; border-radius: 10px; margin-top: 5px;} .success {background-color: #008089;} .success:hover {background-color: #00B7C3;}</style><textarea Id="TextAreaQUESTION" maxlength="%2" style="width:100%;height:30%;resize: none;font-size: 10pt">%1</textarea>
```

```
<BR><button class="btn success"
OnClick="window.parent.WebPageViewerHelper.TriggerCallback(document.getElementById("TextAreaQUESTION").value)">SEND</button><BR><textarea readonly
```

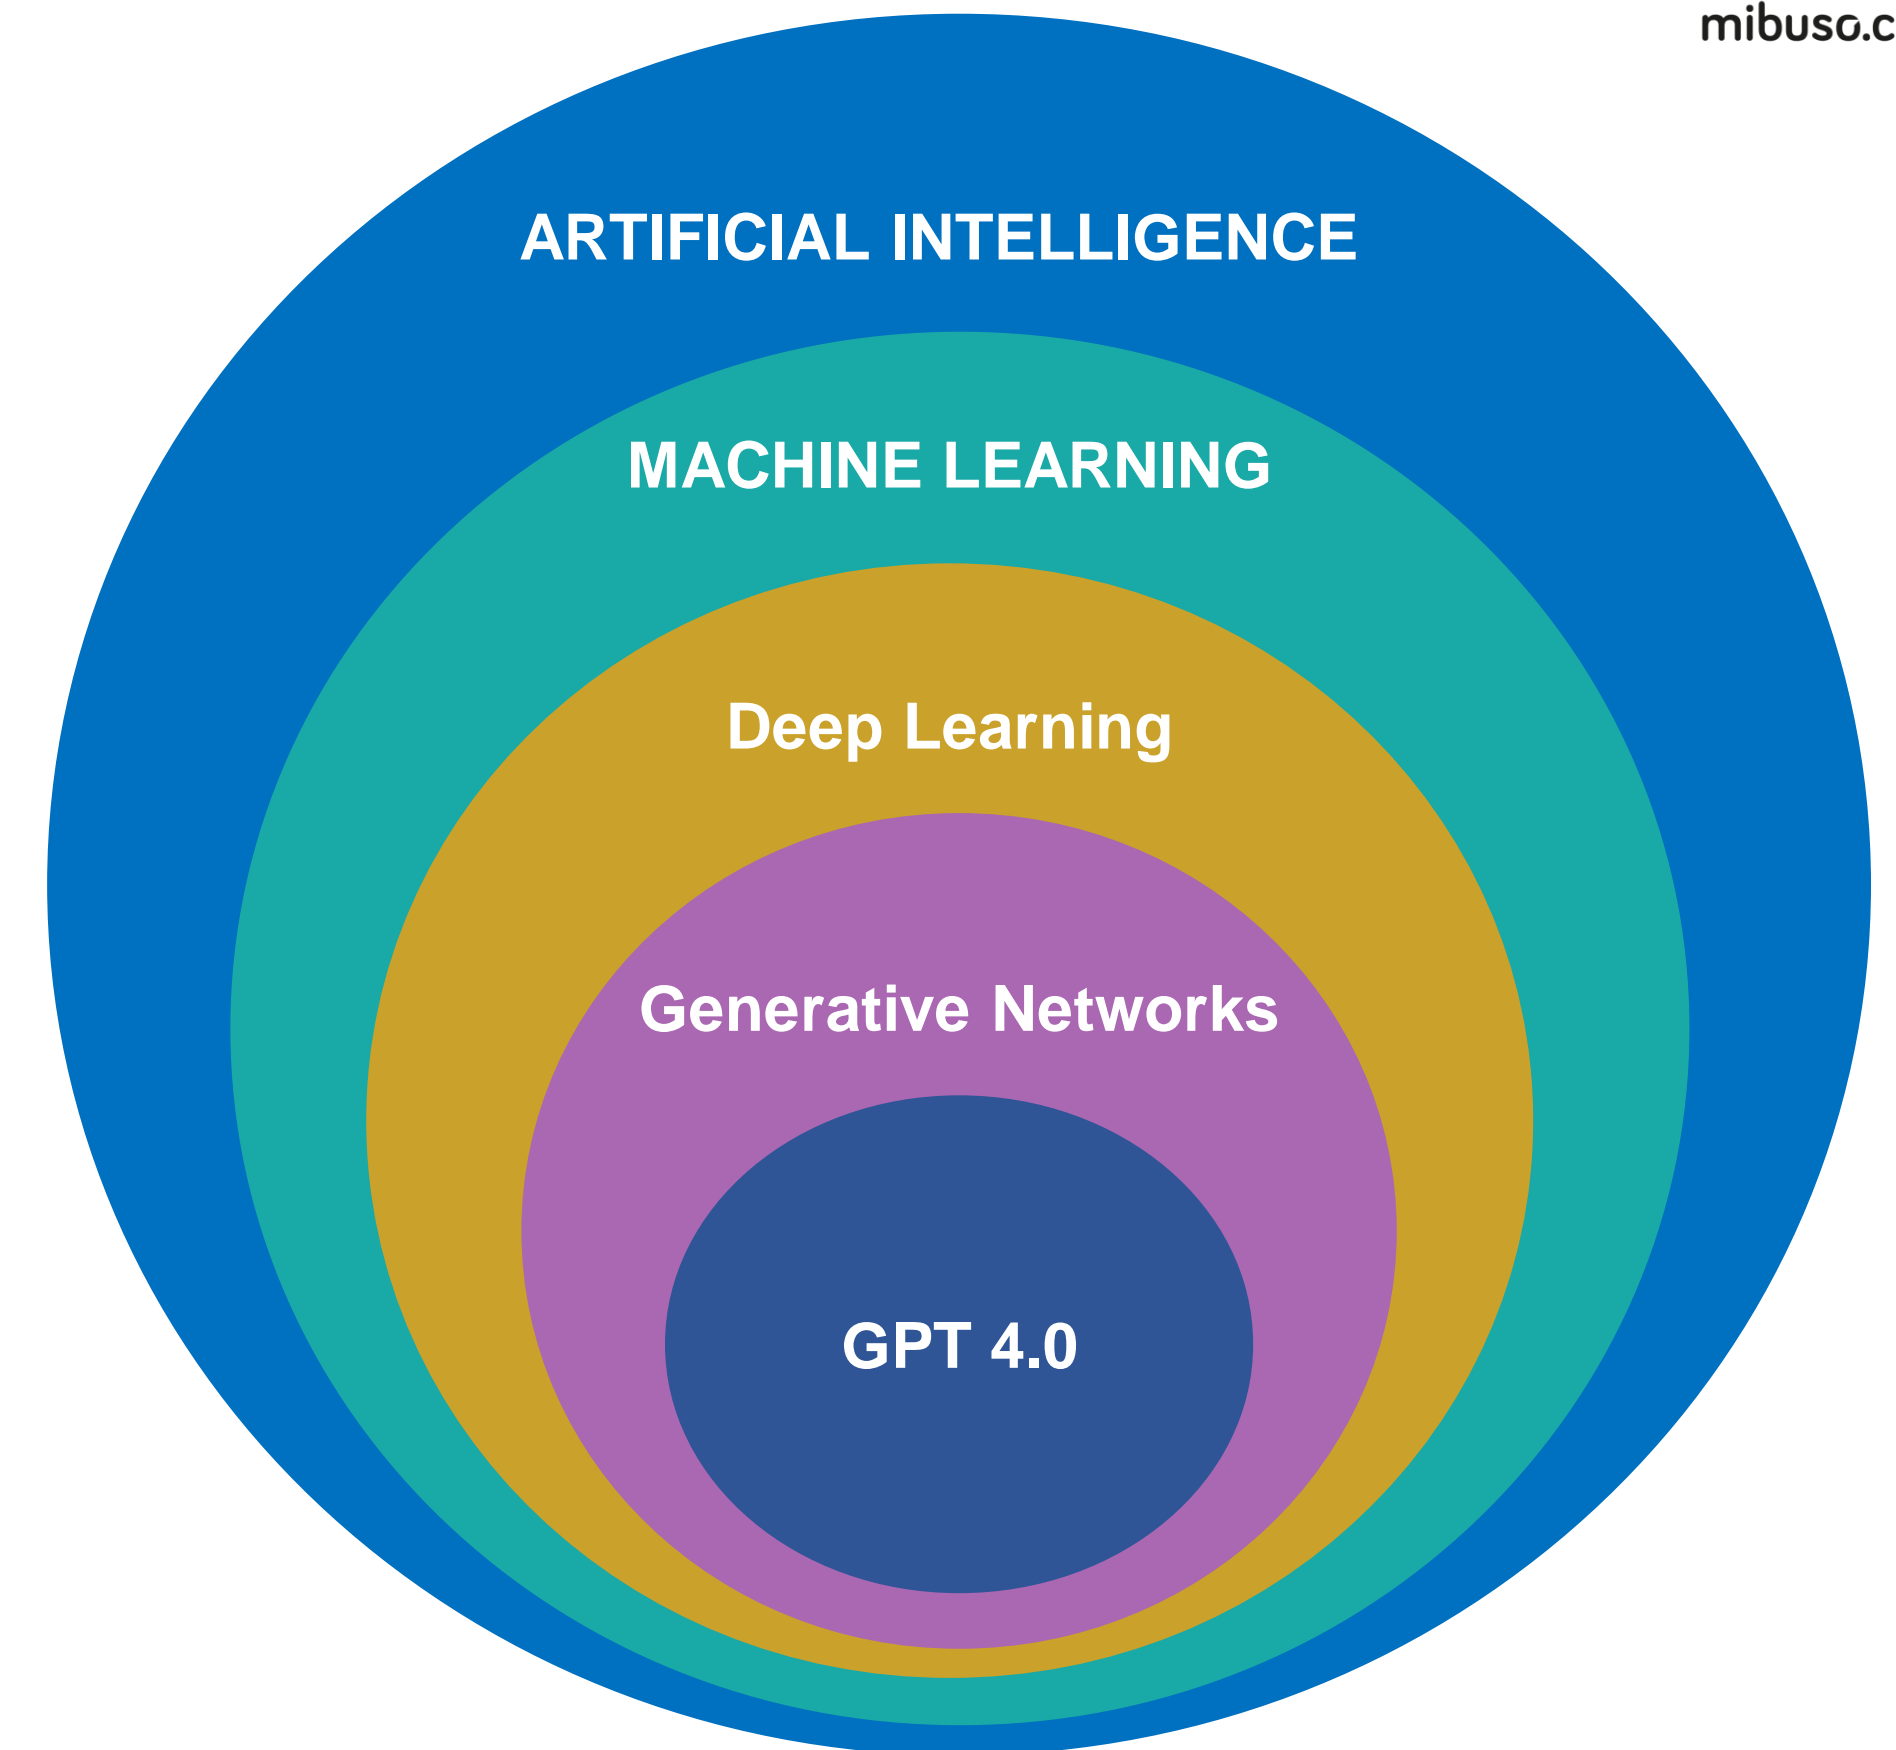
```
Id="TextAreaQUESTIONHISTORY" maxlength="%4" style="width: 80%; height: 30%; background: #f5f5f5; color: #272727; padding: 20px; padding-top: 5px; padding-bottom: 5px; box-shadow: 0px 7px 24px rgba(100, 100, 100, 0.4); border-radius: 10px; margin-top: 40px; margin-left: 19%; resize: none; border-color: transparent; font-size: 10pt; ">%3</textarea>
```

```
<textarea readonly Id="TextAreaANSWER" maxlength="%2" style="width: 80%; height:80%;resize: none; background: #ddf8f7; color: #272727; padding: 20px; padding-top: 15px; padding-bottom: 15px; box-shadow: 0px 7px 24px rgba(100, 100, 100, 0.4); border-radius: 10px; margin-top: 40px; margin-left: 10px; display: flex; flex-direction: row; align-items: flex-start; border-color: transparent;font-size: 10pt ">%3</textarea>
```



Where we now in AI field?

- **AI** - Programs with the ability to learn and reason like humans
- **ML** - Algorithms with the ability to learn without being explicitly programmed
- **DL** - Subset of machine learning in which artificial neural networks adapt and learn from vast amounts of data
- **GN** - type of deep learning that can create new, original content based on patterns and data it has learned



Azure OpenAI: overview

Azure OpenAI is a partnership between Microsoft and OpenAI.



Azure OpenAI service provides a safe and reliable way to access cutting-edge language AI models like OpenAI GPT, Codex, and DALL-E in the Azure ecosystem while safeguarding customer data.



OpenAI vs Azure OpenAI

Feature	OpenAI	Azure OpenAI
Data processing location	Most of the training of the language model is done within the US	Can be deployed within three specific regions Azure. East US, South Central US and West Europe
API access	Access to their own services	REST API access to OpenAI's models
Security capabilities	Not specified	Customers get the security capabilities of Microsoft Azure while running the same models as OpenAI
Private networking	Not specified	Yes
Regional availability	Not specified	Yes
Responsible AI content filtering	Not specified	Yes
Approval required for API access	No	Yes



How to get access

- Azure OpenAI

<https://aka.ms/oaiapply>

<https://aka.ms/oai/get-gpt4>

- OpenAI

<https://platform.openai.com/signup>

- Pricing is the same



Azure OpenAI

Region:
East US

Currency:
United States – Dollar (\$) USD

ChatGPT (gpt-3.5-turbo)	\$0.002
-------------------------	---------

GPT-4	Prompt (Per 1,000 tokens)	Completion (Per 1,000 tokens)
8K context	\$0.03	\$0.06
32K context	\$0.06	\$0.12

OpenAI

Chat

ChatGPT models are optimized for dialogue. The performance of gpt-3.5-turbo is on par with Instruct Davinci.
[Learn more about ChatGPT](#)

Model	Usage
gpt-3.5-turbo	\$0.002 / 1K tokens

GPT-4

With broad general knowledge and domain expertise, GPT-4 can follow complex instructions in natural language and solve difficult problems with accuracy.
[Learn more](#)

Model	Prompt	Completion
8K context	\$0.03 / 1K tokens	\$0.06 / 1K tokens
32K context	\$0.06 / 1K tokens	\$0.12 / 1K tokens



What is Prompt

- **Text used to generate a response** from a language model
- Should have instruction to the model describing what do you expect to receive
- Can be used to initiate a conversation, generate text in a specific style or tone, or perform a specific task such as translation or summarization
- Think of it as a **technical task to intern**

You are a world class software engineer.

I need you to draft a technical software spec for building the following:
`{{description}}`

Think through how you would build it step by step.

Then, respond with the complete spec as a well-organized markdown file.



Good Prompt Structure

- Set model into role: "You are..."
- Instructions at the beginning, separate context with ### or "" or ""
- Be specific about context, outcome, style
- Provide examples of the outcome
- Try to be short and clean in instructions, save more space to context

ChatGPT Prompting Cheat Sheet

How Prompt Engineering Works

Note: Prompt engineering designs and optimizes prompts for language models. It's important in NLP and language generation. Prompt formats guide the model and can be used for tasks like product descriptions or conversational AI. Reliable prompt formats exist, but exploring new formats is encouraged. "{your input here}" is a placeholder for text or context.

Rules of Thumb and Examples

Rule #1 – Instructions at beginning and ### or "" to separate instructions or context

✗ Rewrite the text below in more engaging language.
{your input here}

✓ Rewrite the text below in more engaging language.
Text: ""
{your input here}
""

Rule #2 – Be specific and detailed about the desired context, outcome, length, format, and style.

✗ Write a short story for kids

✓ Write a funny soccer story for kids that teaches the kid that persistence is key for success in the style of Rowling.

Rule #3 – Give examples of desired output format

✗ Extract house pricing data from the following text.
Text: ""
{your text containing pricing data}
""

✓ Extract house pricing data from the following text.
Desired format: ""
House 1 | \$1,000,000 | 100 sqm
House 2 | \$500,000 | 90 sqm
... (and so on)
""
Text: ""
{your text containing pricing data}
""

Rule #4 – First try without examples, then try giving some examples.

✓ Extract brand names from the text below.
Text: {your text here}
Brand names:

✓ Extract brand names from the texts below.
Text 1: Finxter and YouTube are tech companies. Google is too.
Brand names 2: Finxter, YouTube, Google

Text 2: If you like tech, you'll love Finxter!
Brand names 2: Finxter

Text 3: {your text here}
Brand names 3:

Rule #5 – Fine-tune if Rule #4 doesn't work

Fine-tuning improves model performance by training on more examples, resulting in higher quality results, token savings, and lower latency requests. GPT-3 can intuitively generate plausible completions from few examples, known as **few-shot learning**. Fine-tuning achieves better results on various tasks without requiring examples in the prompt, saving costs and enabling lower-latency requests.

Example Training Data

```
{ "prompt": "<input>", "completion": "<ideal output>" }  
{ "prompt": "<input>", "completion": "<ideal output>" }  
{ "prompt": "<input>", "completion": "<ideal output>" }  
...
```

Rule #6 – Be specific. Omit needless words.

✗ ChatGPT, write a sales page for my company selling sand in the desert, please write only a few sentences, nothing long and complex

✓ Write a 5-sentence sales page, sell sand in the desert.

Rule #7 – Use leading words to nudge the model towards a pattern

✗ Write a Python function that plots my net worth over 10 years for different inputs on the initial investment and a given ROI

✓ # Python function that plots net worth over 10 years for different inputs on the initial investment and a given ROI

```
import matplotlib  
  
def plot_net_worth(initial, roi):
```



New documentation from Azure OpenAI

- <https://learn.microsoft.com/en-us/azure/cognitive-services/openai/concepts/advanced-prompt-engineering>

The system message is included at the beginning of the prompt and is used to prime the model with context, instructions, or other information relevant to your use case. You can use the system message to describe the assistant's personality, define what the model should and shouldn't answer, and define the format of model responses.

The example below, shows a sample system message and the resulting model response.

System message	User	Assistant
You are an AI assistant that helps people find information and responds in rhyme. If the user asks you a question you don't know the answer to, say so.	What can you tell about me, John Doe?	Dear John, I'm sorry to say, But I don't have your name today. I'm just an AI with no knowledge in names. But without your name I can't explain. So please tell me about what you want. And I'll do my best to give you an answer that's unique.

Some other examples of system messages are:



AI Job Queue Scheduler - Prompt Structure

""""

You are Microsoft Dynamics 365 Business Central Job Queues scheduling API.

You must only use information from the context below.

You will be provided {input_text_description}.

You must always return results in a predefined {output_format} format.

{additional_output_instructions}

Context:

'''

{context_text}

'''

Example of requests in natural language and responses in JSON format:

'''

{examples_text}

'''

Answer in {format_type} FORMAT:

'''

{format_description}

'''

""""



Prompt limits

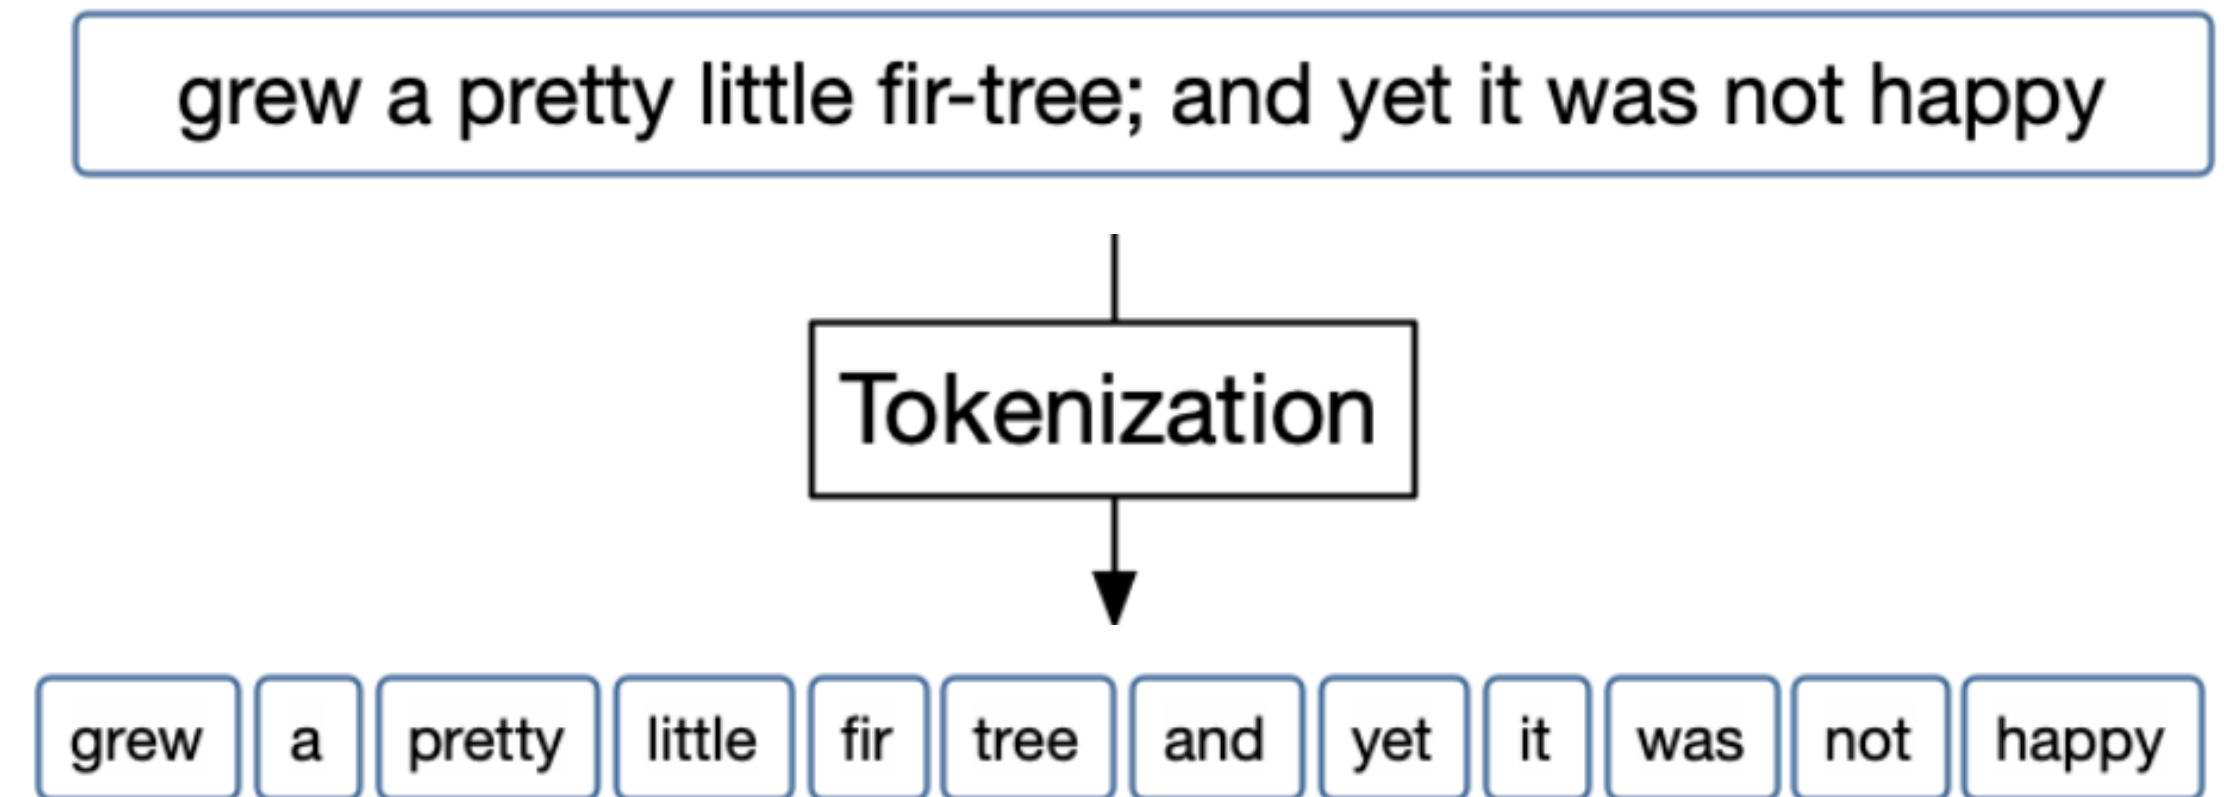
Model Name	Description	Max Tokens	Last Update
gpt-3.5-turbo	Most capable GPT-3.5 model and optimized for chat at 1/10th the cost of text-davinci-003. Will be updated with our latest model iteration 2 weeks after it is released.	4,096	Sep 2021
gpt-3.5-turbo-16k	Same capabilities as the standard gpt-3.5-turbo model but with 4 times the context.	16,384	Sep 2021
gpt-4	More capable than any GPT-3.5 model, able to do more complex tasks, and optimized for chat. Will be updated with our latest model iteration 2 weeks after it is released.	8,192	Sep 2021
gpt-4-32k	Same capabilities as the base gpt-4 mode but with 4x the context length. Will be updated with our latest model iteration.	32,768	Sep 2021

- Request + Response!



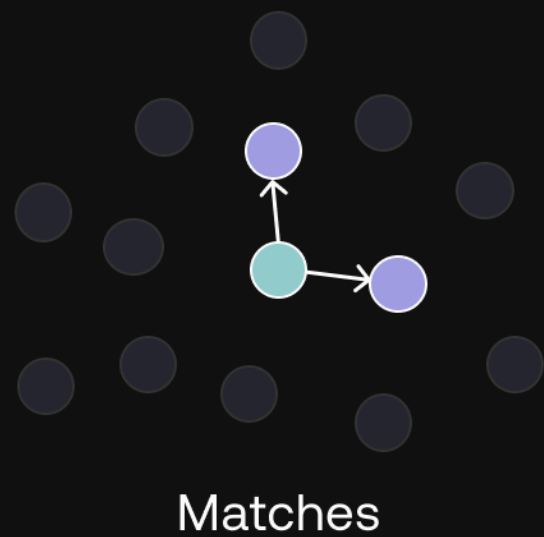
What is Token

- Token is a small unit of text that is used to represent a word or a part of a word
- Tokenization is the process of breaking down a piece of text into small units called tokens
- 1 token \approx 4 character

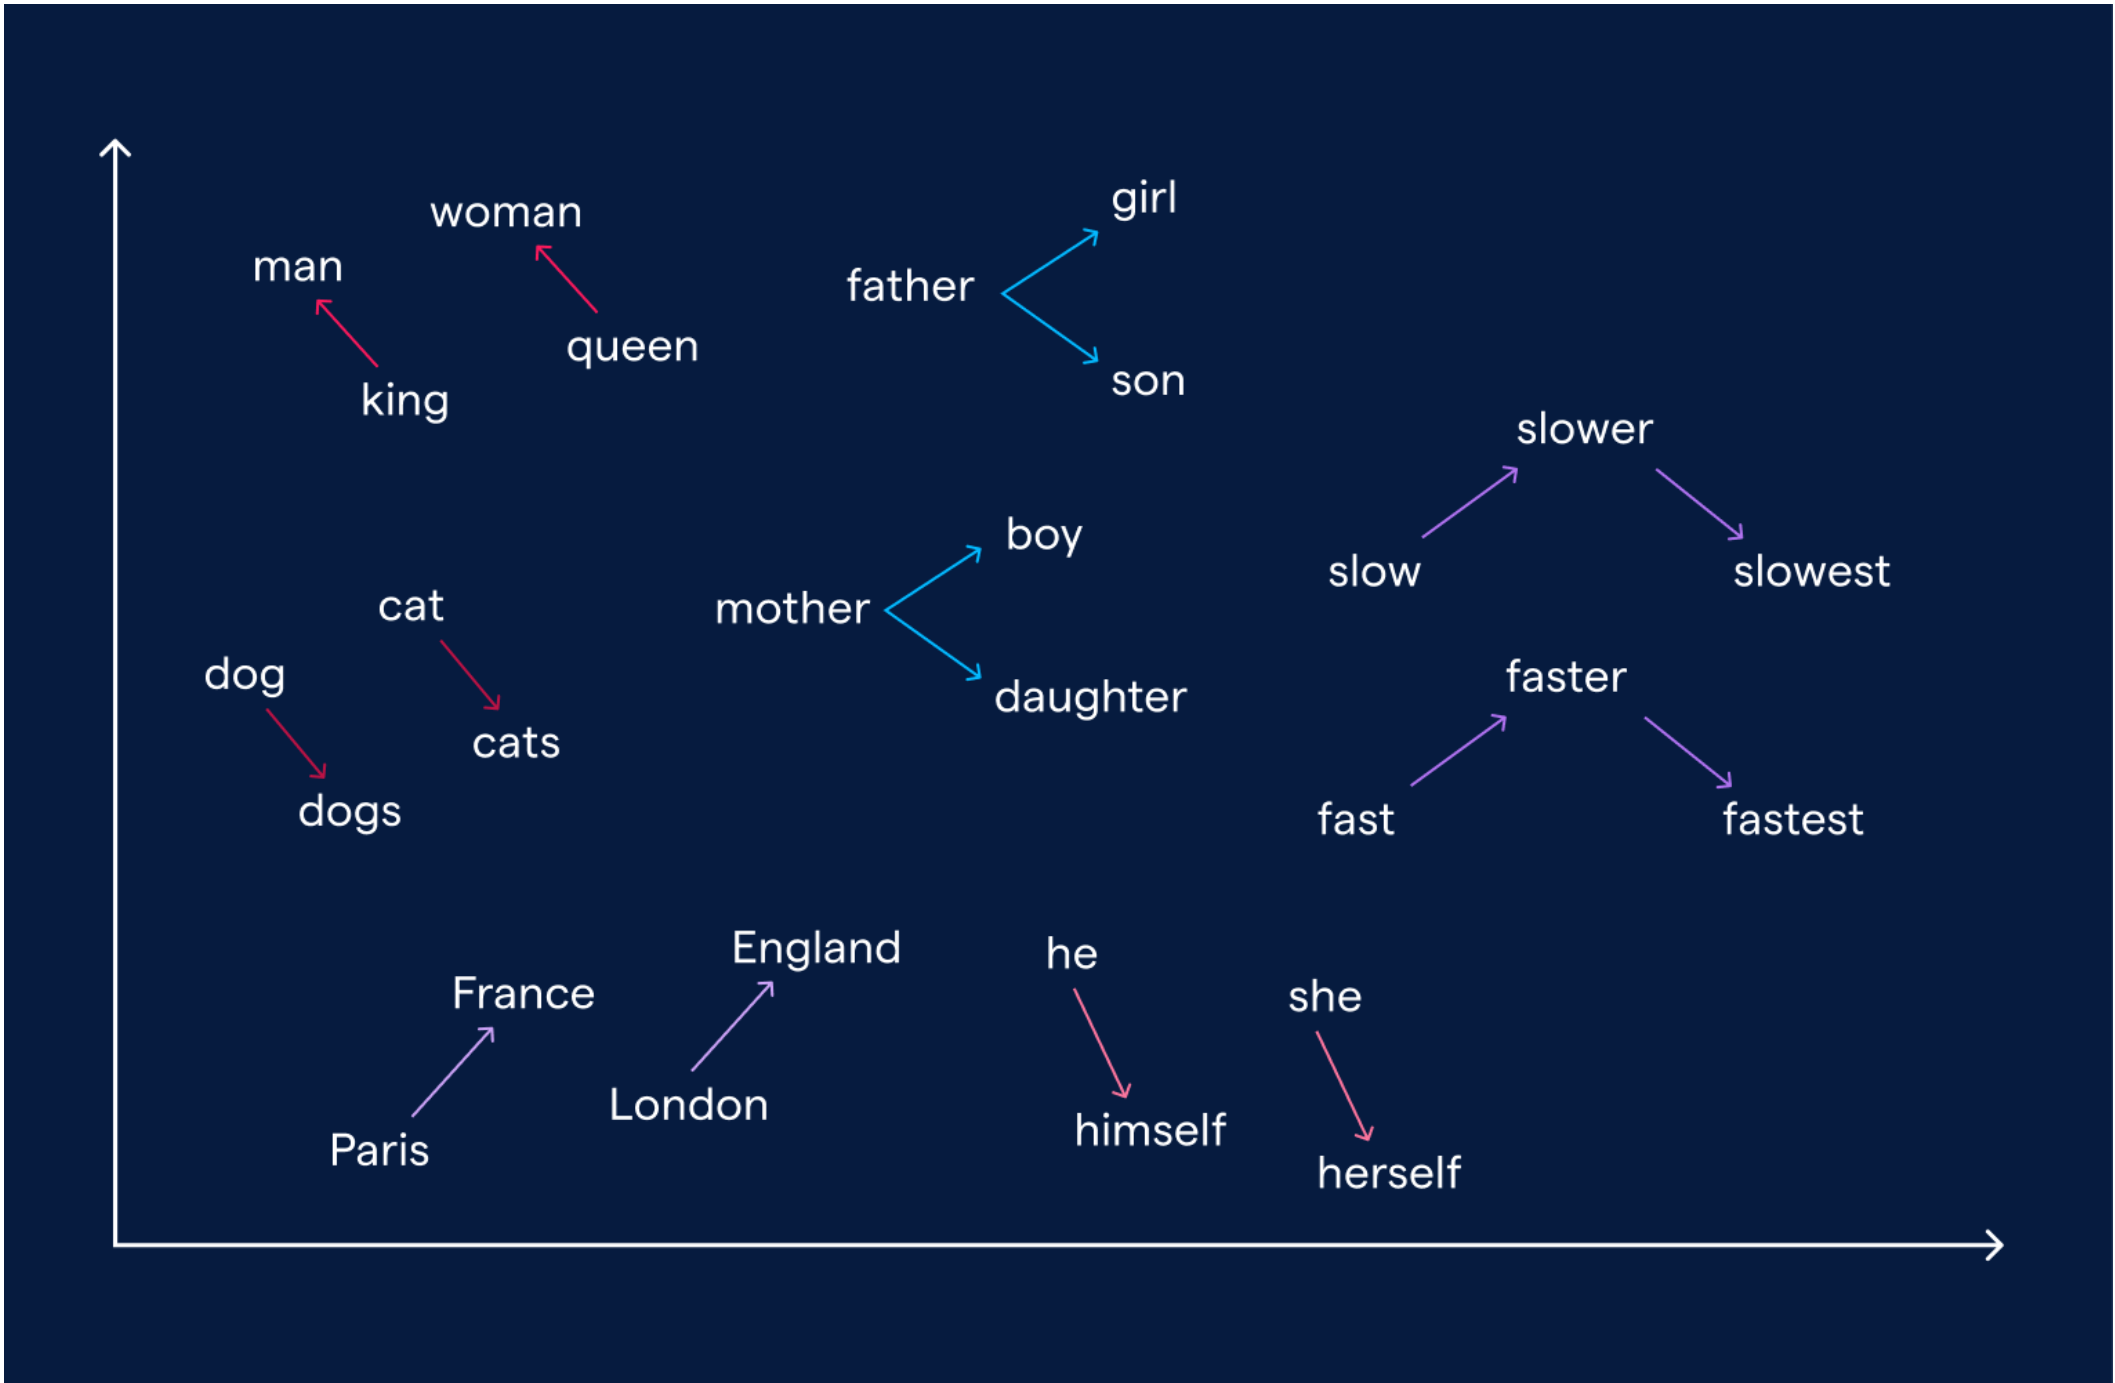
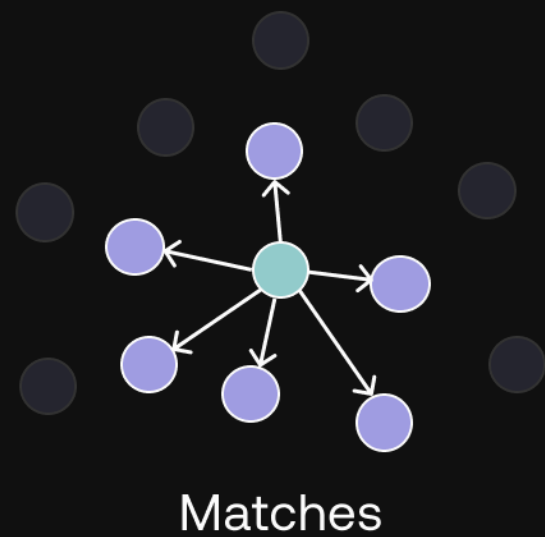


Embeddings

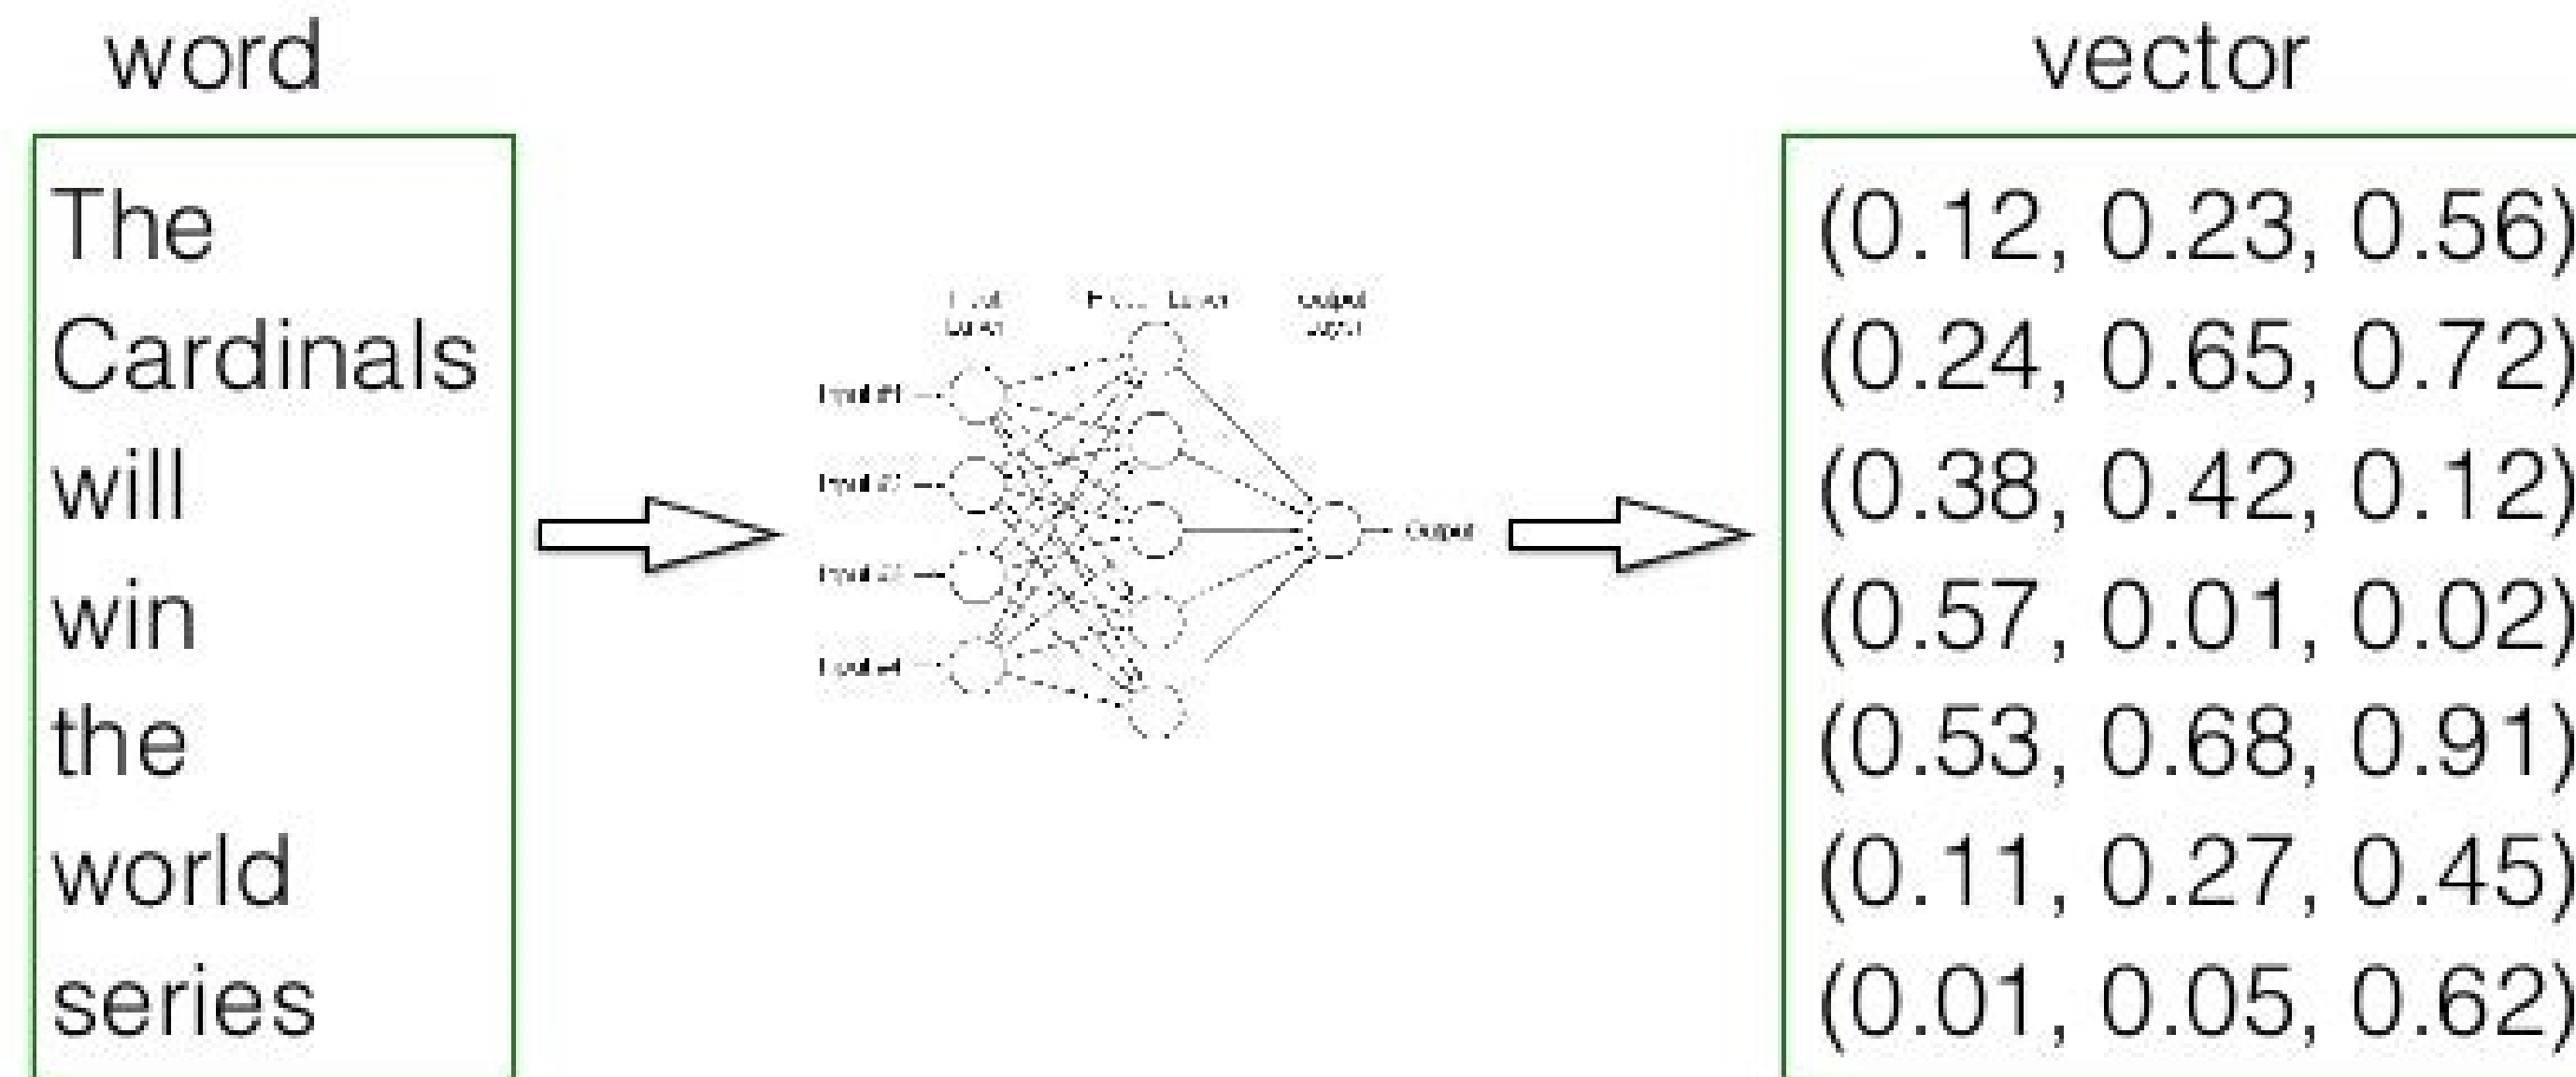
Keyword Search



Semantic Search

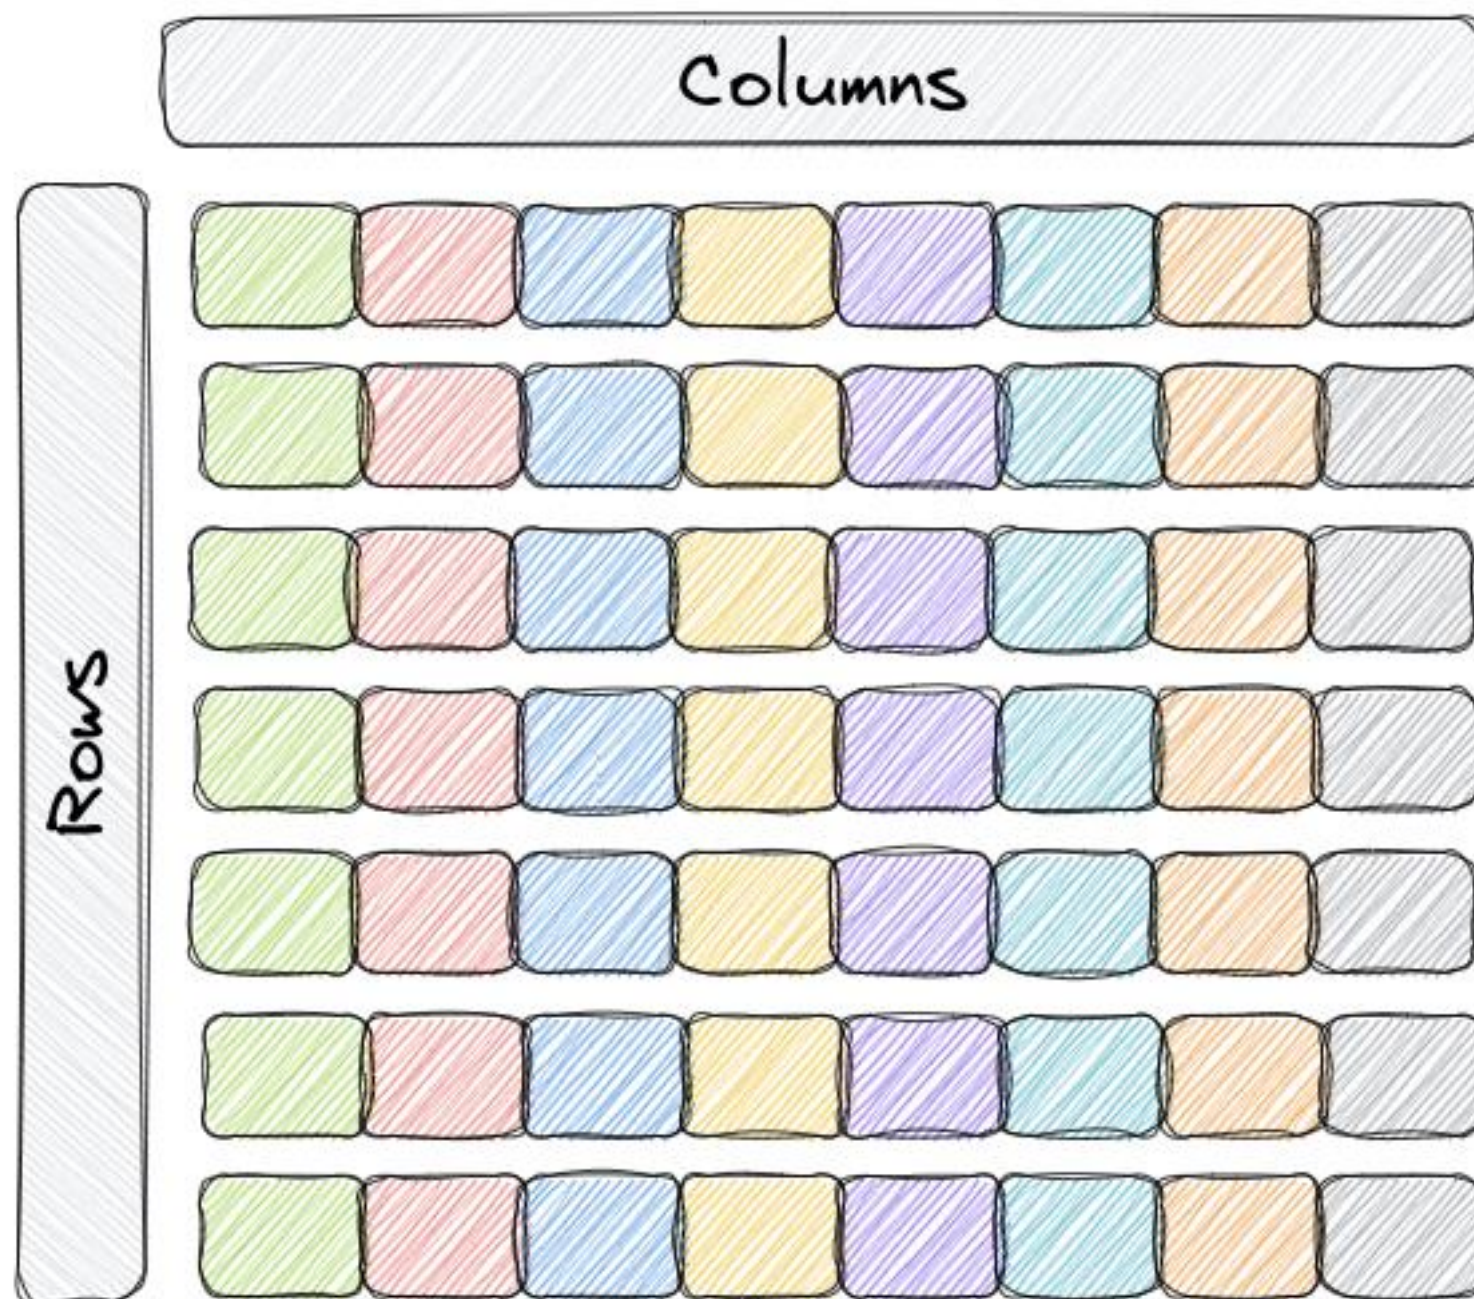


Embeddings

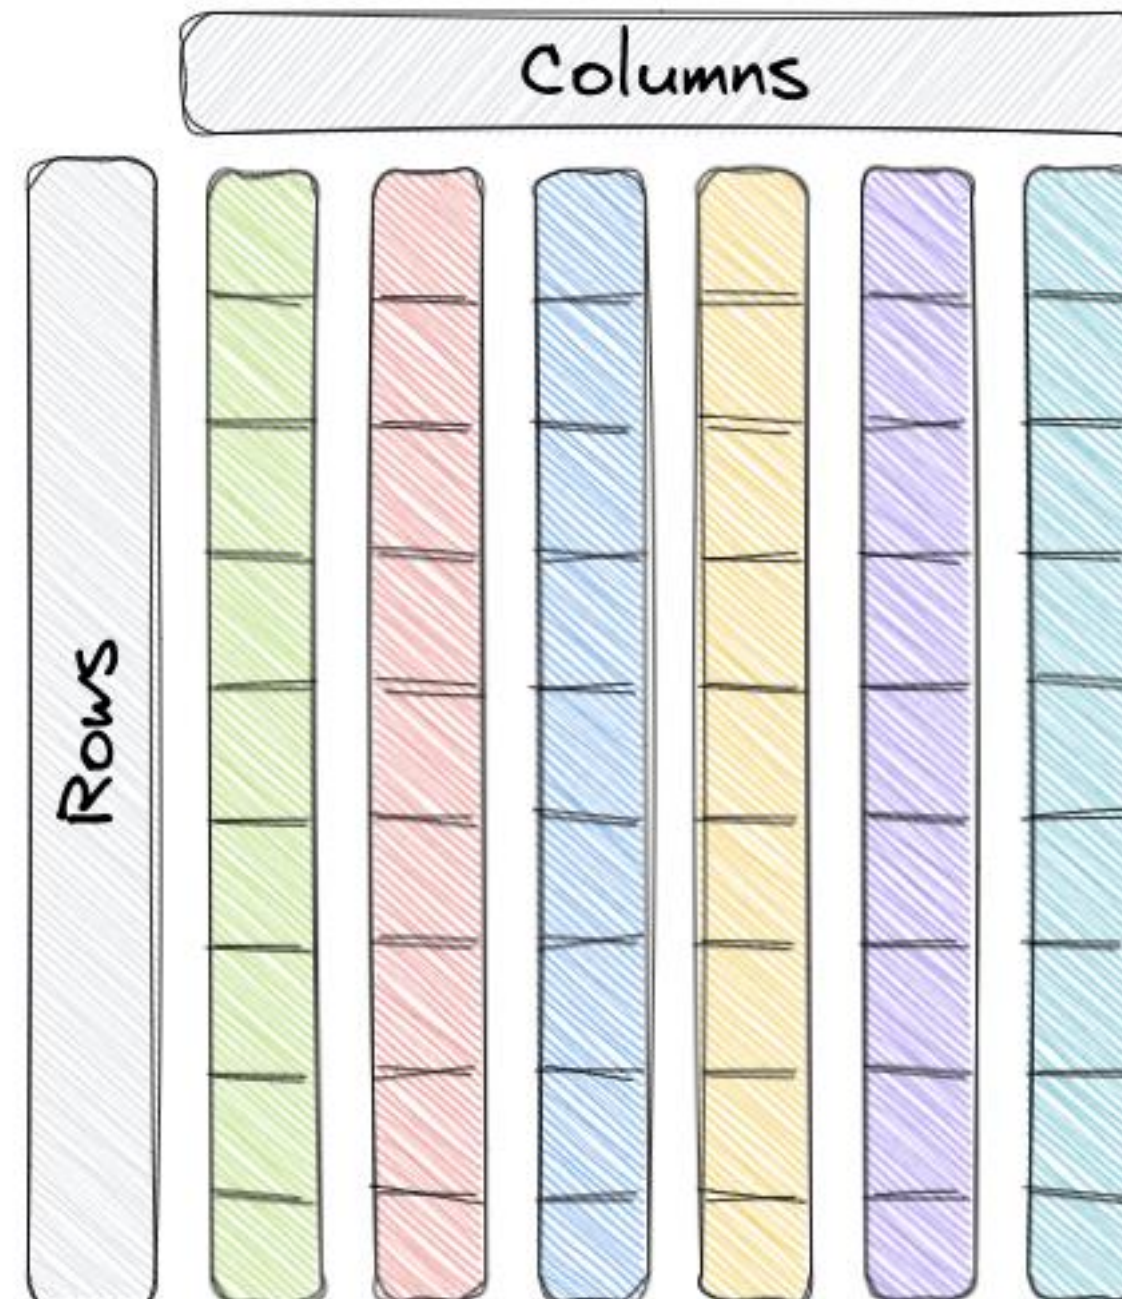


Vector Databases

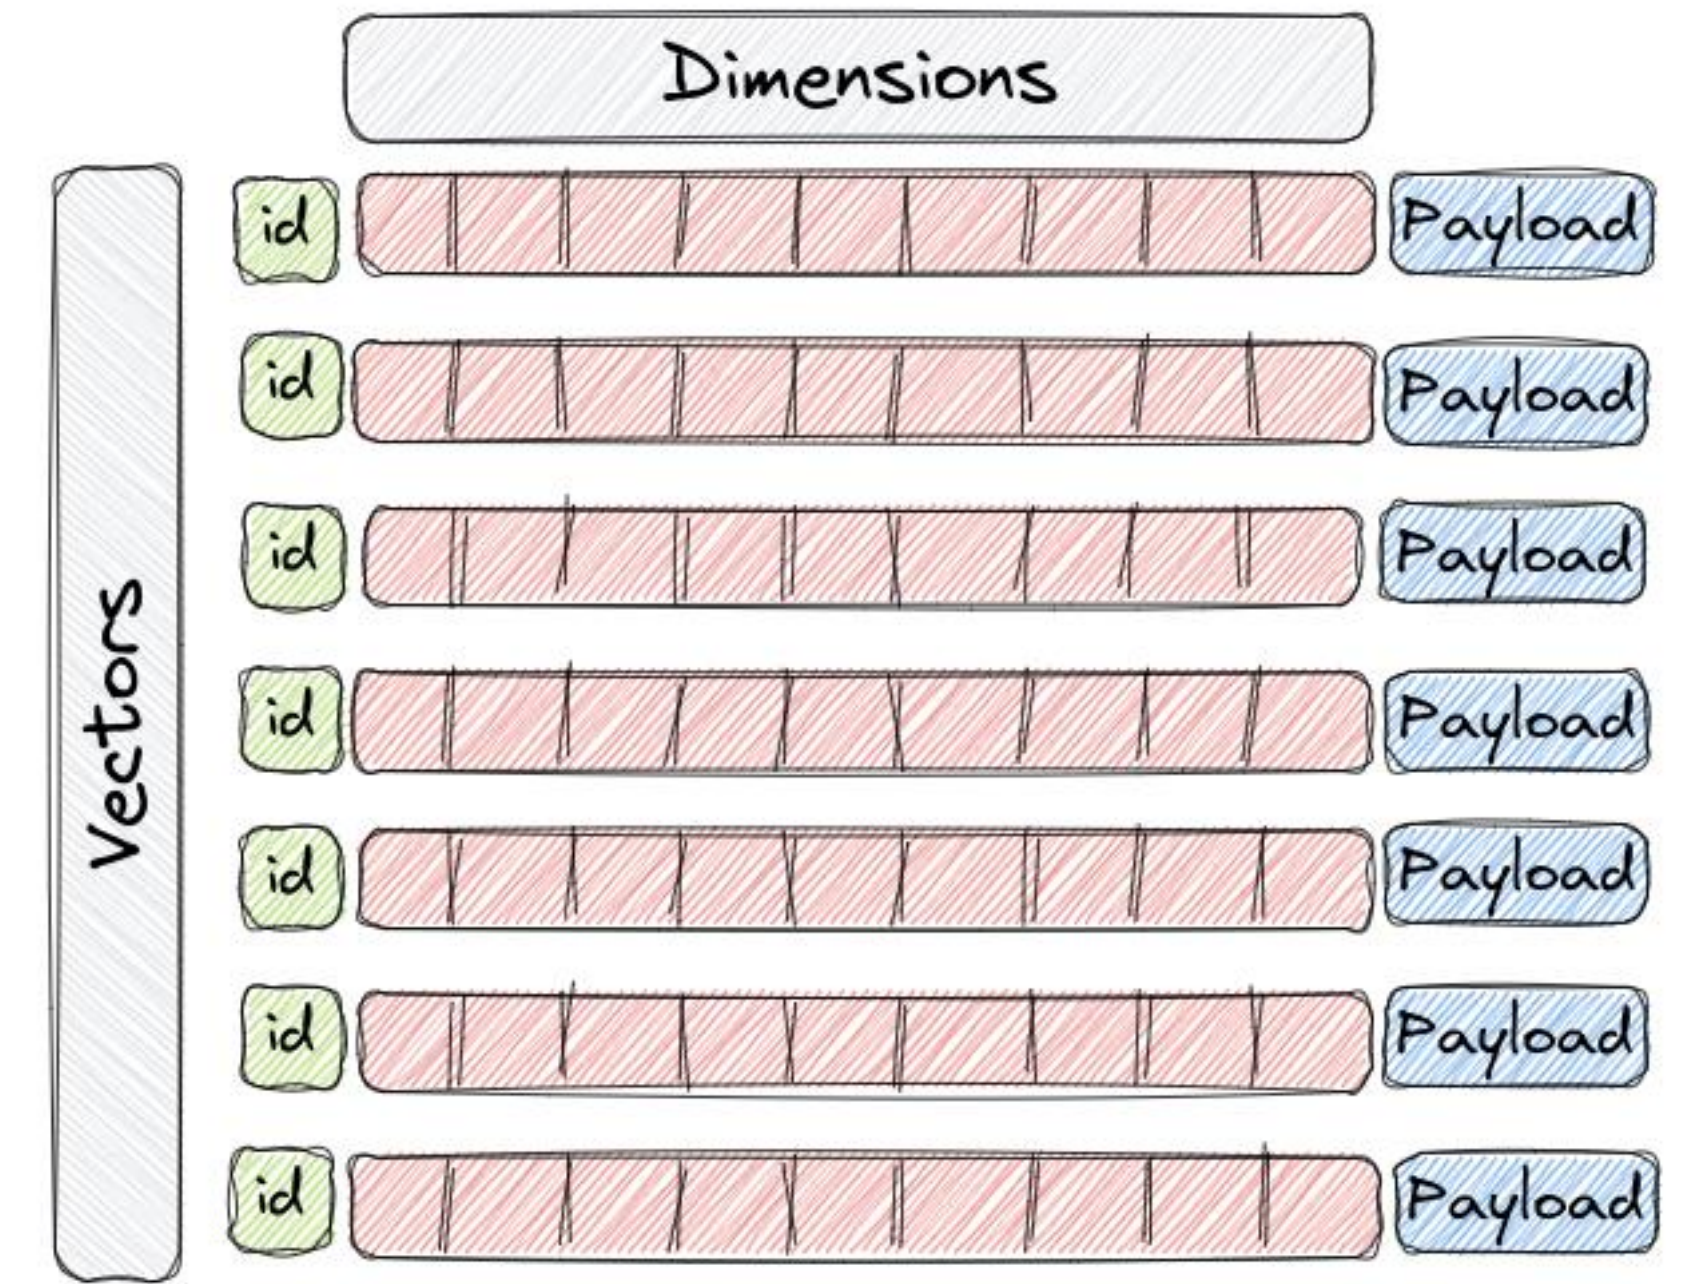
OLTP Database



OLAP Database



Vector Database



Share your BC TechDays Lifehacks

Share BCTechdays lifehacks

Anything you find valuable for others, like:

- best beer and where to get it
- how to get to the conference cheaper
- how to make friends
- etc.

1. Lifehack 1

Enter your answer

2. Lifehack 2

Enter your answer

Submit



Working with AI from a non AI-guru perspective...

- AI is expensive and unpredictable (like me 🤪) if not well used
- AI is not only ChatGPT
- AI can be used with your own data
- AI is not only Python
- To replace developers with AI, customers need to accurately describe what they want → DEVs will have a long future!



AI most common pattern...



The quality of input data reflects the quality of the output response.



Azure OpenAI on your data

Chat with your own data in different data sources in few clicks!



Azure OpenAI SDK for .NET

The **Azure OpenAI client library for .NET** (currently in beta) is an adaptation of OpenAI's REST APIs that provides an idiomatic interface and rich integration with the rest of the Azure SDK ecosystem.

It can connect to Azure OpenAI resources *or* to the non-Azure OpenAI inference endpoint, making it a great choice for even non-Azure OpenAI development.

```
dotnet add package Azure.AI.OpenAI --prerelease
```



Azure OpenAI service quotas and limits

Limit Name

Limit Value

OpenAI resources per region per Azure subscription
Default quota per model and region (in tokens-per-minute)

30
Text-Davinci-003: 120 K
GPT-4: 20 K
GPT-4-32K: 60 K
All others: 240 K

Maximum prompt tokens per request

Varies per model. For more information, see [Azure OpenAI Service models](#)

Max fine-tuned model deployments

2

Total number of training jobs per resource

100

Max simultaneous running training jobs per resource

1

Max training jobs queued

20

Max Files per resource

30

Total size of all files per resource

1 GB

Max training job time (job will fail if exceeded)

720 hours

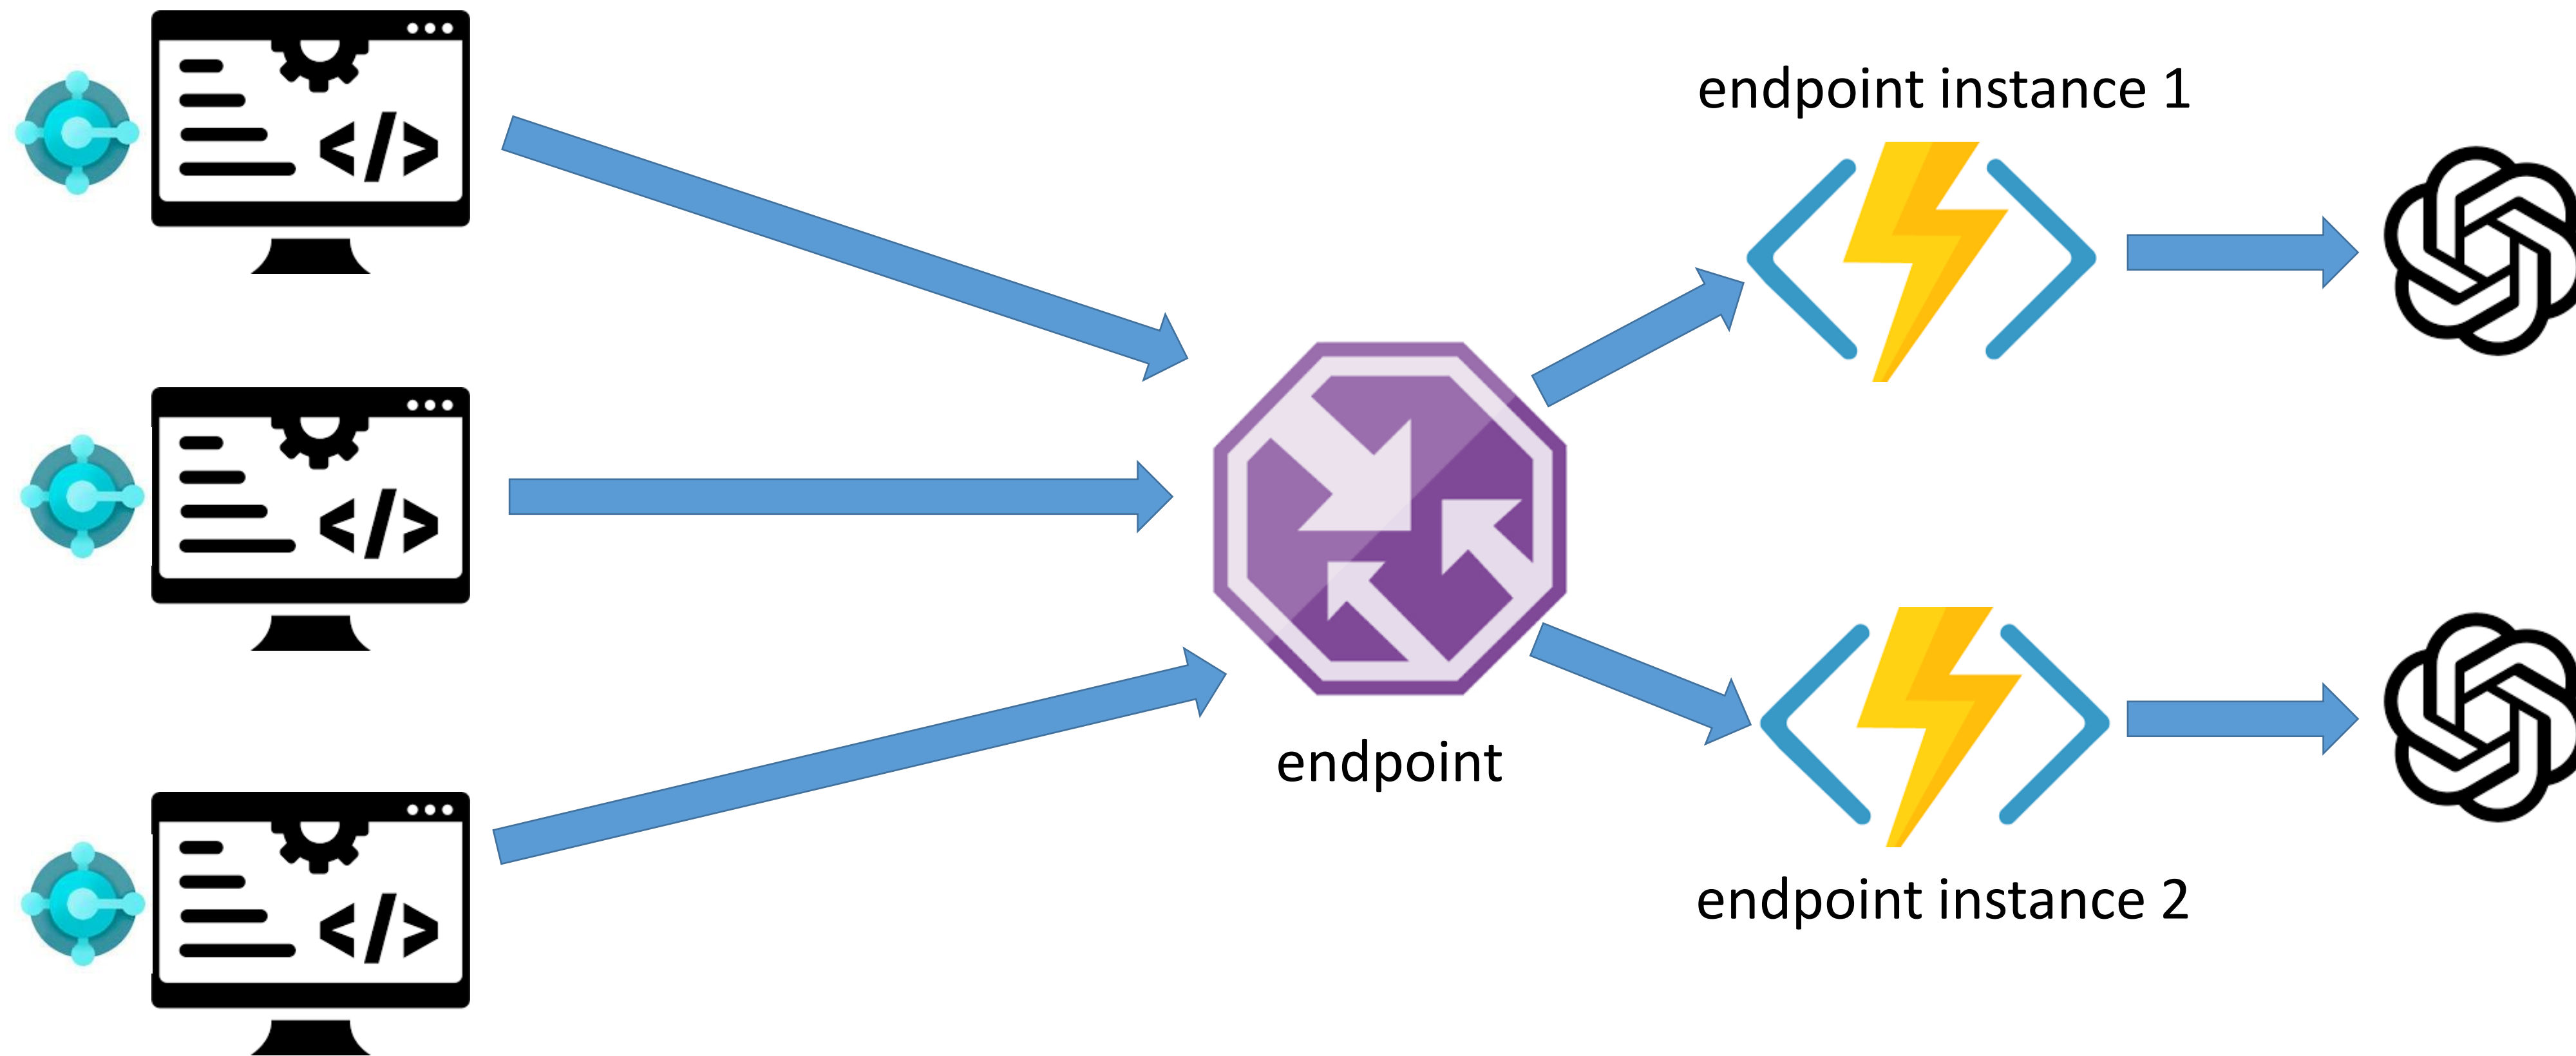
Max training job size (tokens in training file) x (# of epochs)

2 Billion

Sometimes the rate limit can be a problem!



Load balancing Azure OpenAI service



Microsoft Semantic Kernel

Semantic Kernel is a lightweight open-source orchestration SDK that lets you easily mix-and-match AI prompts with conventional programming languages like C# and Python.

Semantic Kernel has been engineered to allow developers to flexibly integrate AI into their existing apps. To do so, Semantic Kernel provides a set of *abstractions* that make it easy to create and manage prompts, native functions, memories, and connectors. You can then orchestrate these components using Semantic Kernel pipelines to complete users' requests or automate actions.

<https://aka.ms/skrepo>



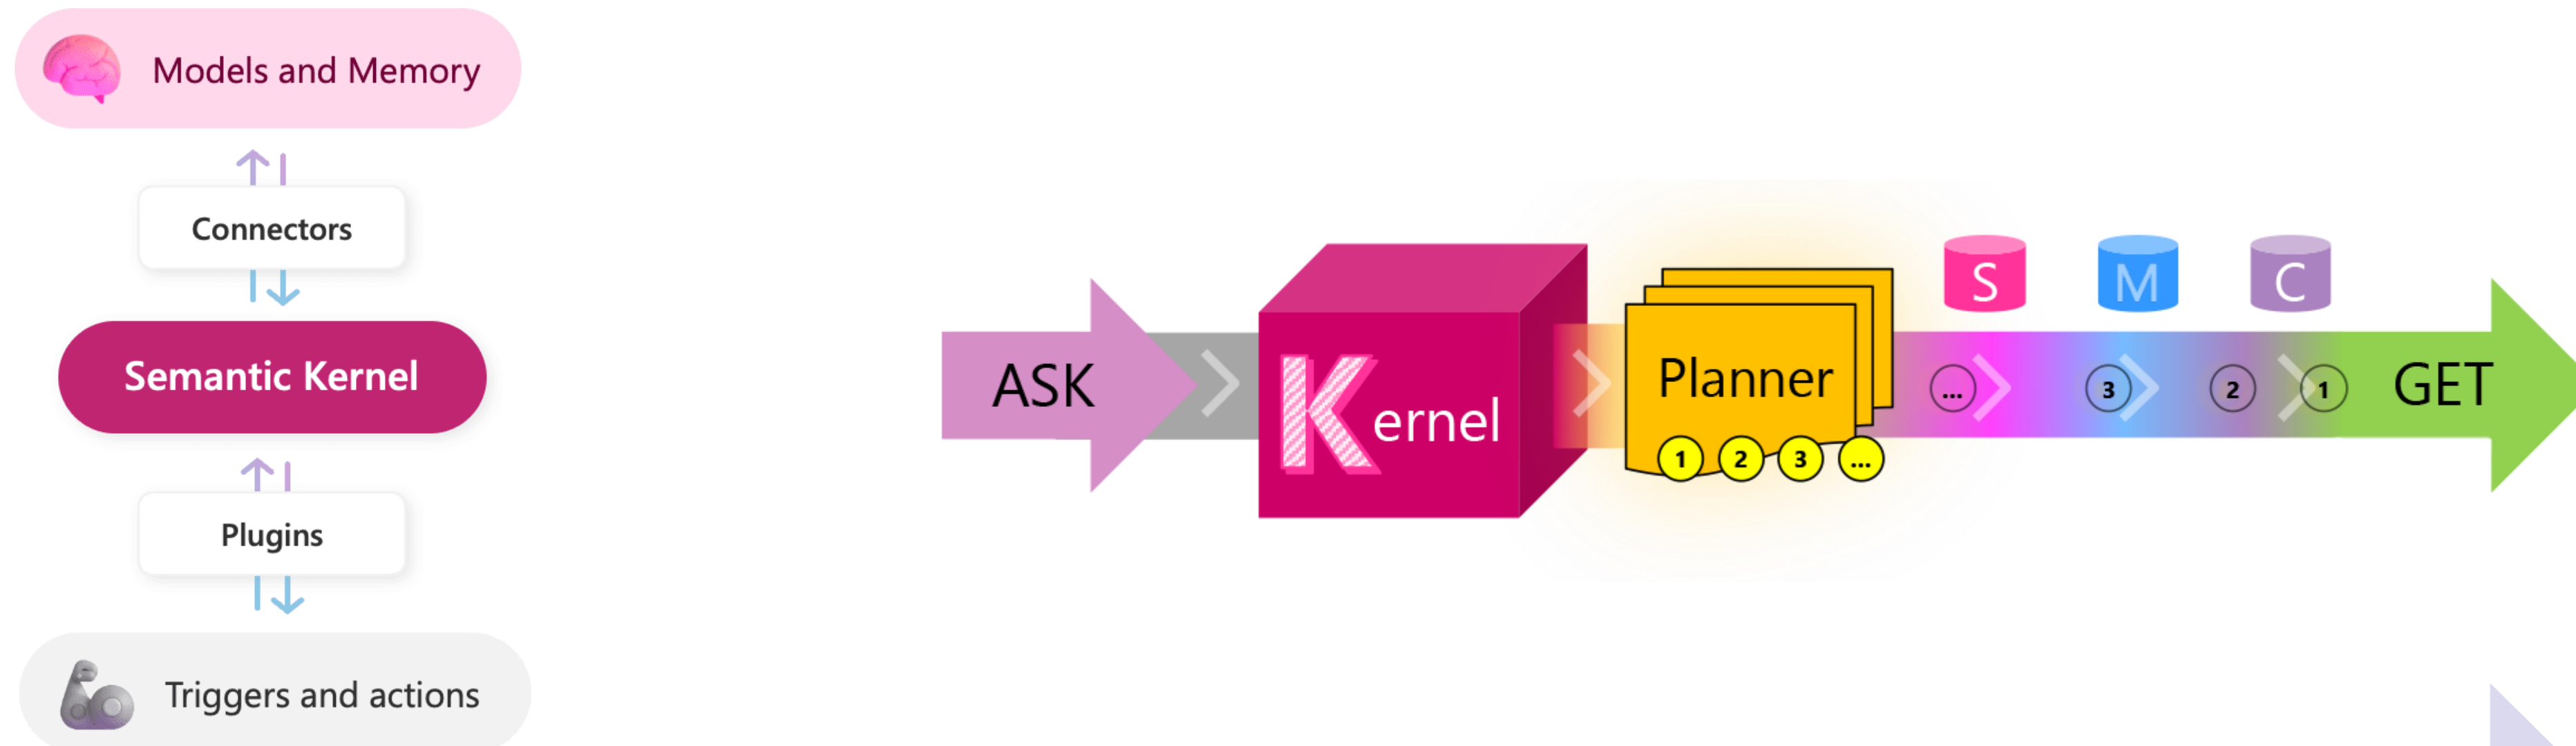
Microsoft Semantic Kernel

The Semantic Kernel understands two kinds of operation:

- **Semantic functions** - capabilities expressed as instructions in human language
- **Native functions** - additional capabilities written as code that can be executed locally

These functions are grouped together into **Skills** - and a skill can have any number of semantic or native functions.

The Kernel also has a **Planner**, which can decide which functions to chain together to complete an **Ask** (a request made to it in human language).



Microsoft Semantic Kernel: DEMO

Using **Semantic Kernel** to query customer's data (large document files etc).

- **Azure OpenAI:**
 - Embeddings
 - Completions
- **Qdrant vector database** for storing embeddings

Qdrant is a vector similarity search engine that provides a production-ready service with a convenient API to store, search, and manage points (i.e. vectors) with an additional payload.



A **vector** in this context is a mathematical representation of an object or data point, where each element of the vector corresponds to a specific feature or attribute of the object.

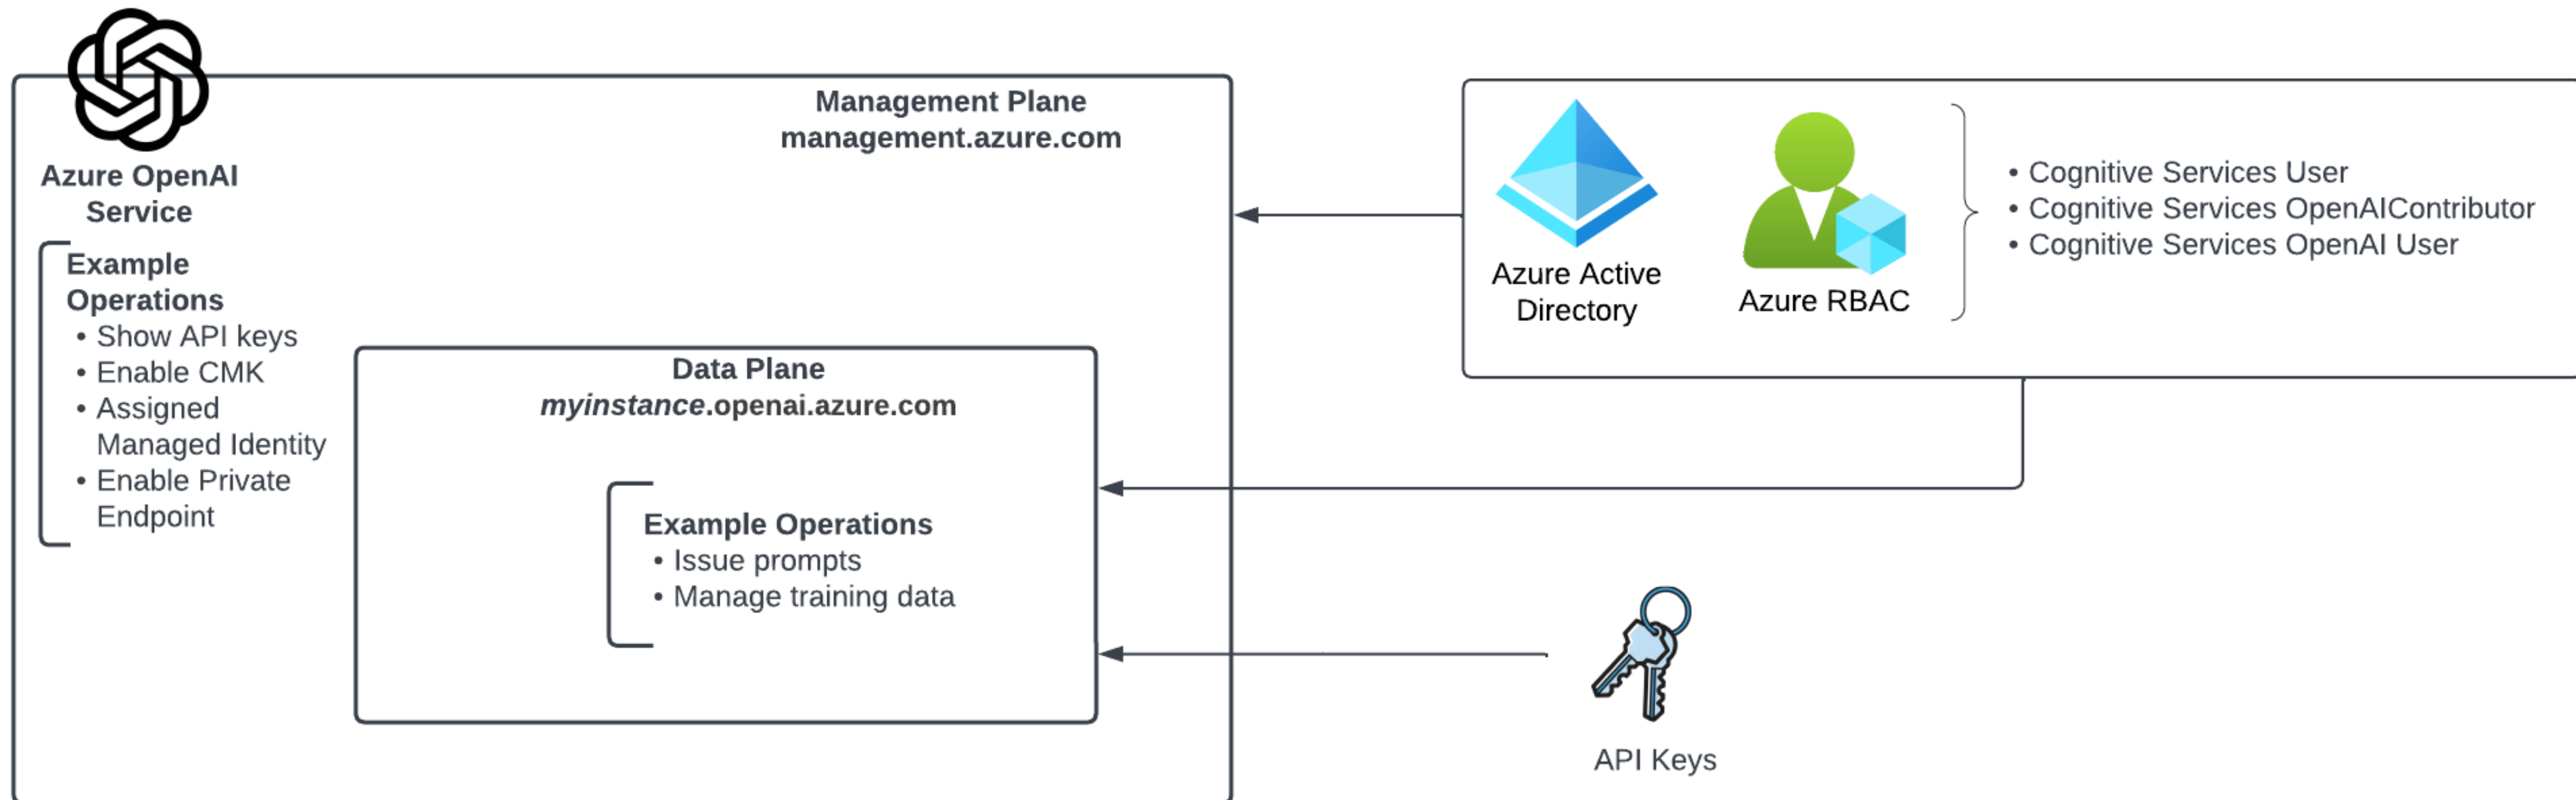
The main advantage of a vector database is that it allows for fast and accurate similarity search and retrieval of data based on their vector distance or similarity. This means that instead of using traditional methods of querying databases based on exact matches or predefined criteria, you can use a vector database to find the most similar or relevant data based on their semantic or contextual meaning.



Azure OpenAI: planes

Planes

- **Management plane:** administration of AI resources
- **Data plane:** administration of data hosted within the resource



CentralQ

AI search, fueled by the [collective knowledge](#) of Business Central community.

🔍

What is BC Techdays?

➔


Code focused


Answer 👍 🗨


BC TechDays is an annual tech conference focused on Microsoft Dynamics Business Central/NAV [4]. It is organized by Mibuso and attracts industry leaders, innovators, and experts from around the world [3]. The event is known for covering new features, best practices, and programming use cases of Business Central/NAV [4]. Attendees have the opportunity to exchange ideas, find partners, and conceive projects [4]. The conference covers topics such as cloud, Power Platform, code optimization, performance, and telemetry [4].


Sources (4)

☰ View List

-  learn.microsoft.com

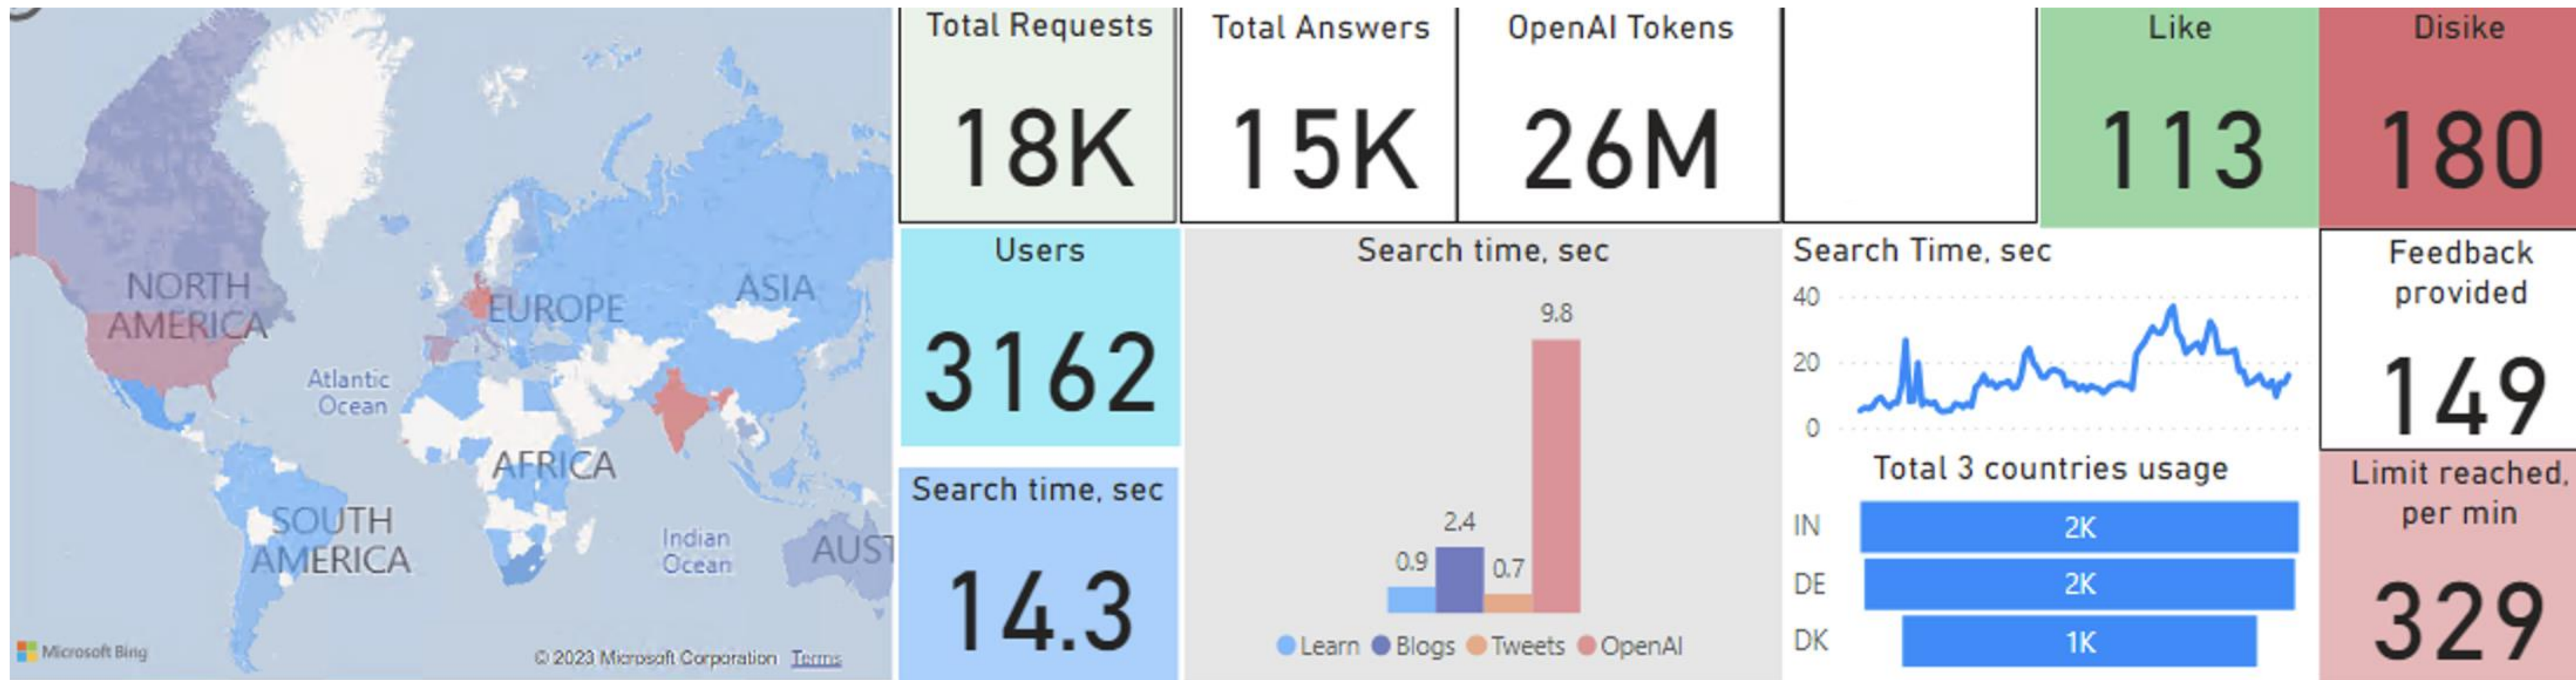
 learn.microsoft.com

 tharangac.com

 techspheredynamics.com

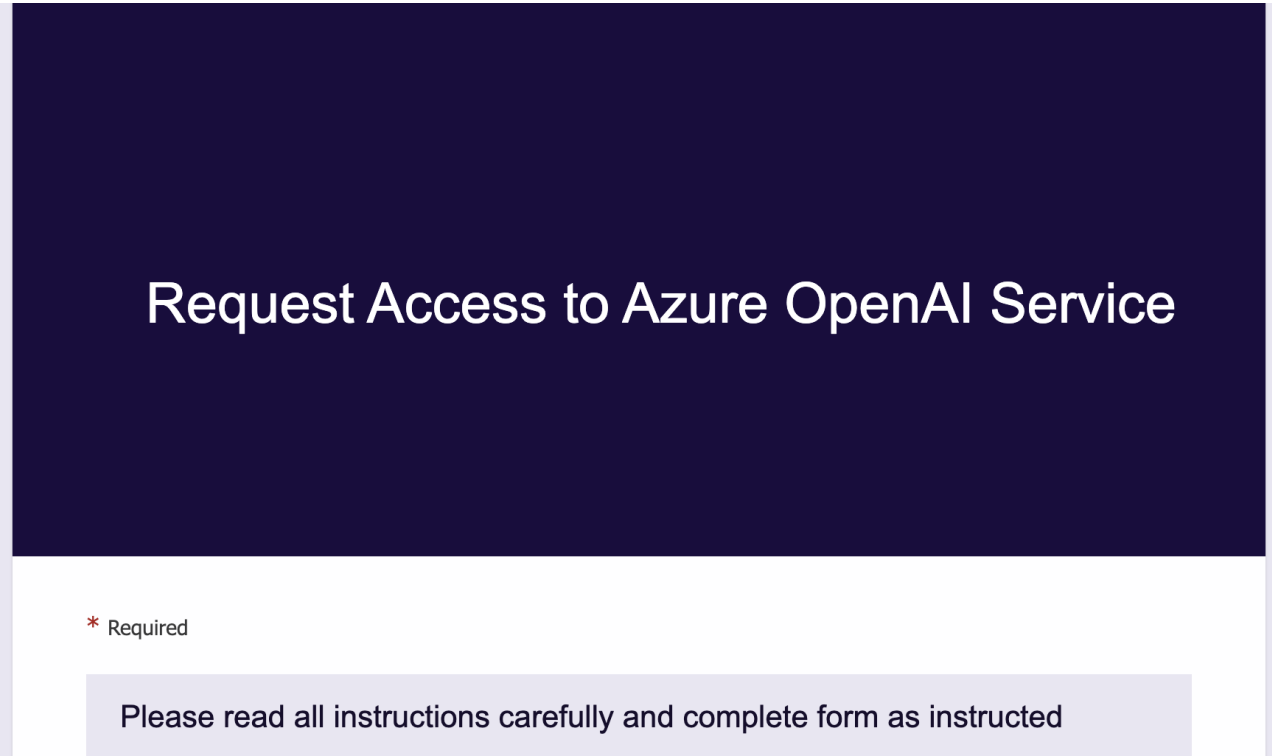


CentralQ telemetry



search	answer
What a service item is?	In Business Central, a service item refers to equipment or items that require service as a consultancy hour, for limited business support [3]. When creating a service order, link a service item to an item in inventory or a service item group [1]. There are several ways to create a service item on the Service Items page, creating them as part of another process, or when shipping items [1]. To work with service items, you must set up service item groups.
does BC automatically copy comments from Customer Card to Sales order document?	Yes, you can add comments to both Customer Card and Sales Order Document in Microsoft Dynamics 365 Business Central. When you add comments to an ongoing Sales Order Document, the comments are transferred to the resulting posted sales order document. However, comments are not printed or output to reports or externally-facing documents. To add comments to a sales order document, you can choose the Comments action, which opens the Comment Sheet page where you can add comments to individual lines [1]. The steps to add comments to a sales order document are similar to all other cards and documents [1].
does BC automatically copy comments from Customer to Sales order?	Yes, comments on ongoing documents in Microsoft Dynamics 365 Business Central are transferred to the resulting posted sales order document. For example, a comment on a sales order is transferred to a resulting posted sales order document. Comments are not printed or output to reports or externally-facing documents [1].

Actions for you to make your app AI-powered



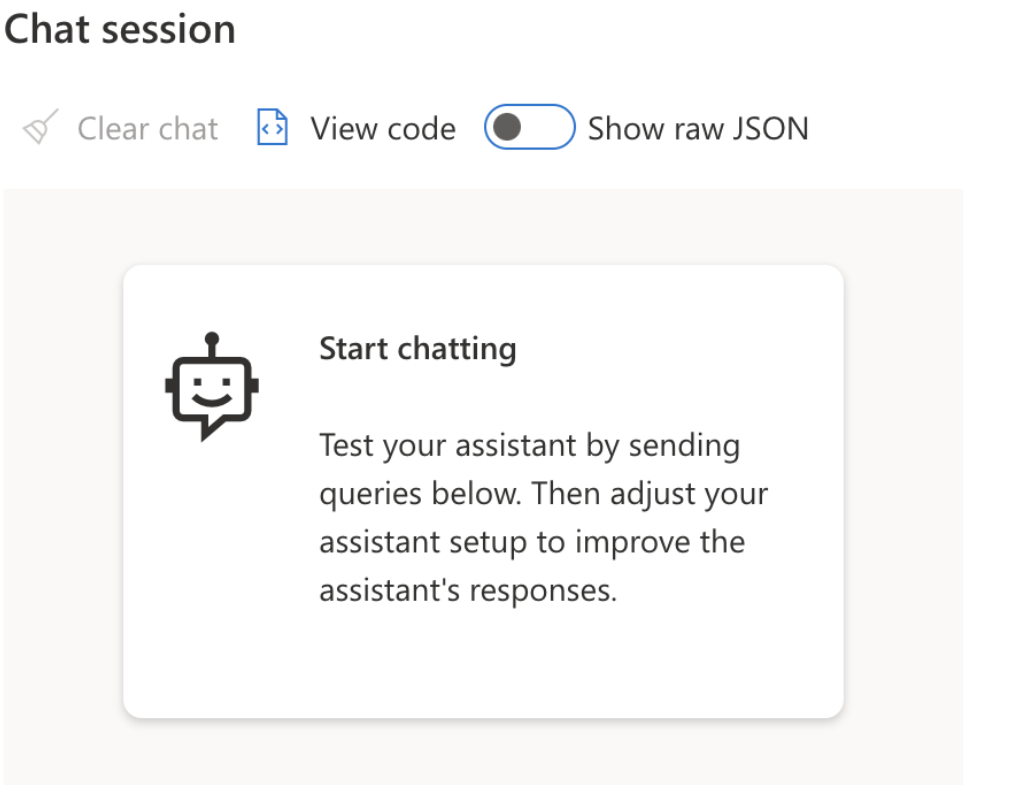
You are a world class software engineer.

I need you to draft a technical software spec for building the following:

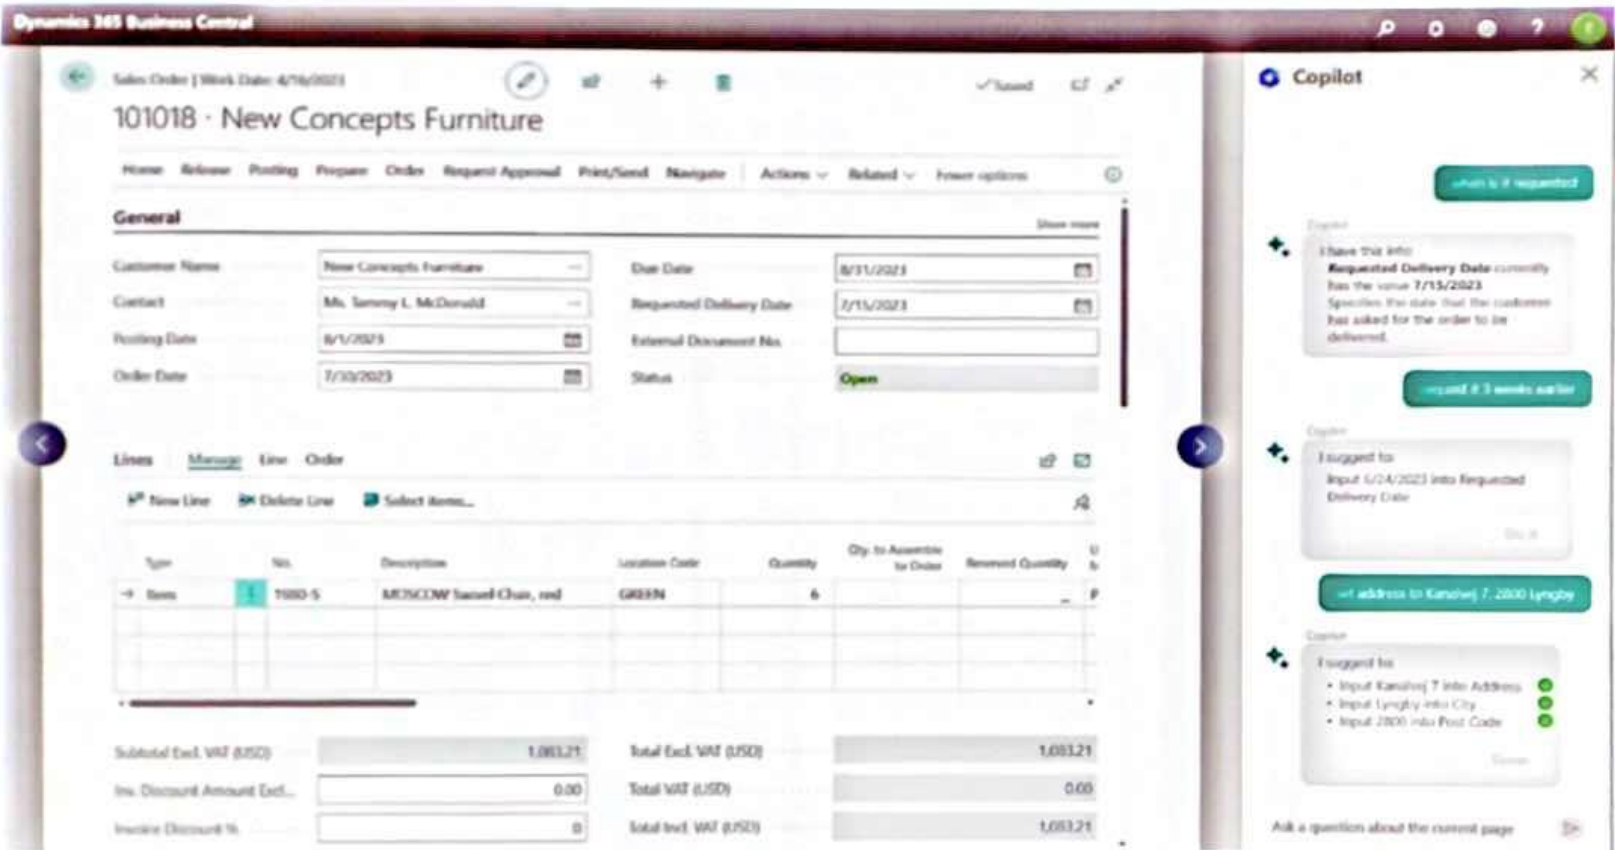
```
{{description}}
```

Think through how you would build it step by step.

Then, respond with the complete spec as a well-organized markdown file.



- Register on the Azure OpenAI/OpenAI
- Learn Prompts
- Experiment in the playground



- Think about embed scenarios, experiment and have Fun!



... and maybe one day

We will come from this



Dynamics 365 Business Central

Sales Order

S-ORD101001 · Adatum Corporation

Home

Prepare

Print/Send

Request Approval

Order

Report

More options

Post...

Release

Create Warehouse Shipment

Create Inventory Put-away/Pick...

Archive Document

General

Customer NameAdatum Corporation

ContactRobert Townes

Posting Date02/04/2023

Order Date02/04/2023

Due Date02/05/2023

Requested Delivery Date03/04/2023

External Document No.

StatusOpen

Lines

Manage

New Line

Delete Line

Select items...

Type	No.	Item Reference No.	Description	Location Code	Quantity	Qty. to Assemble to Order	Reserved Quantity	Unit of Measure Code	Unit Price Excl. Tax	Tax Area Code	Tax Group Code	Lir
→ Item	1996-S		ATLANTA Whiteboard, base		12			PCS	1,397.30	ATLANTA, GA	FURNITURE	

Subtotal Excl. Tax (USD)

16,767.60

Total Excl. Tax (USD)

16,767.60

Inv. Discount Amount Excl. Tax (USD)

0.00

Total Tax (USD)

1,006.06

Invoice Discount %

0

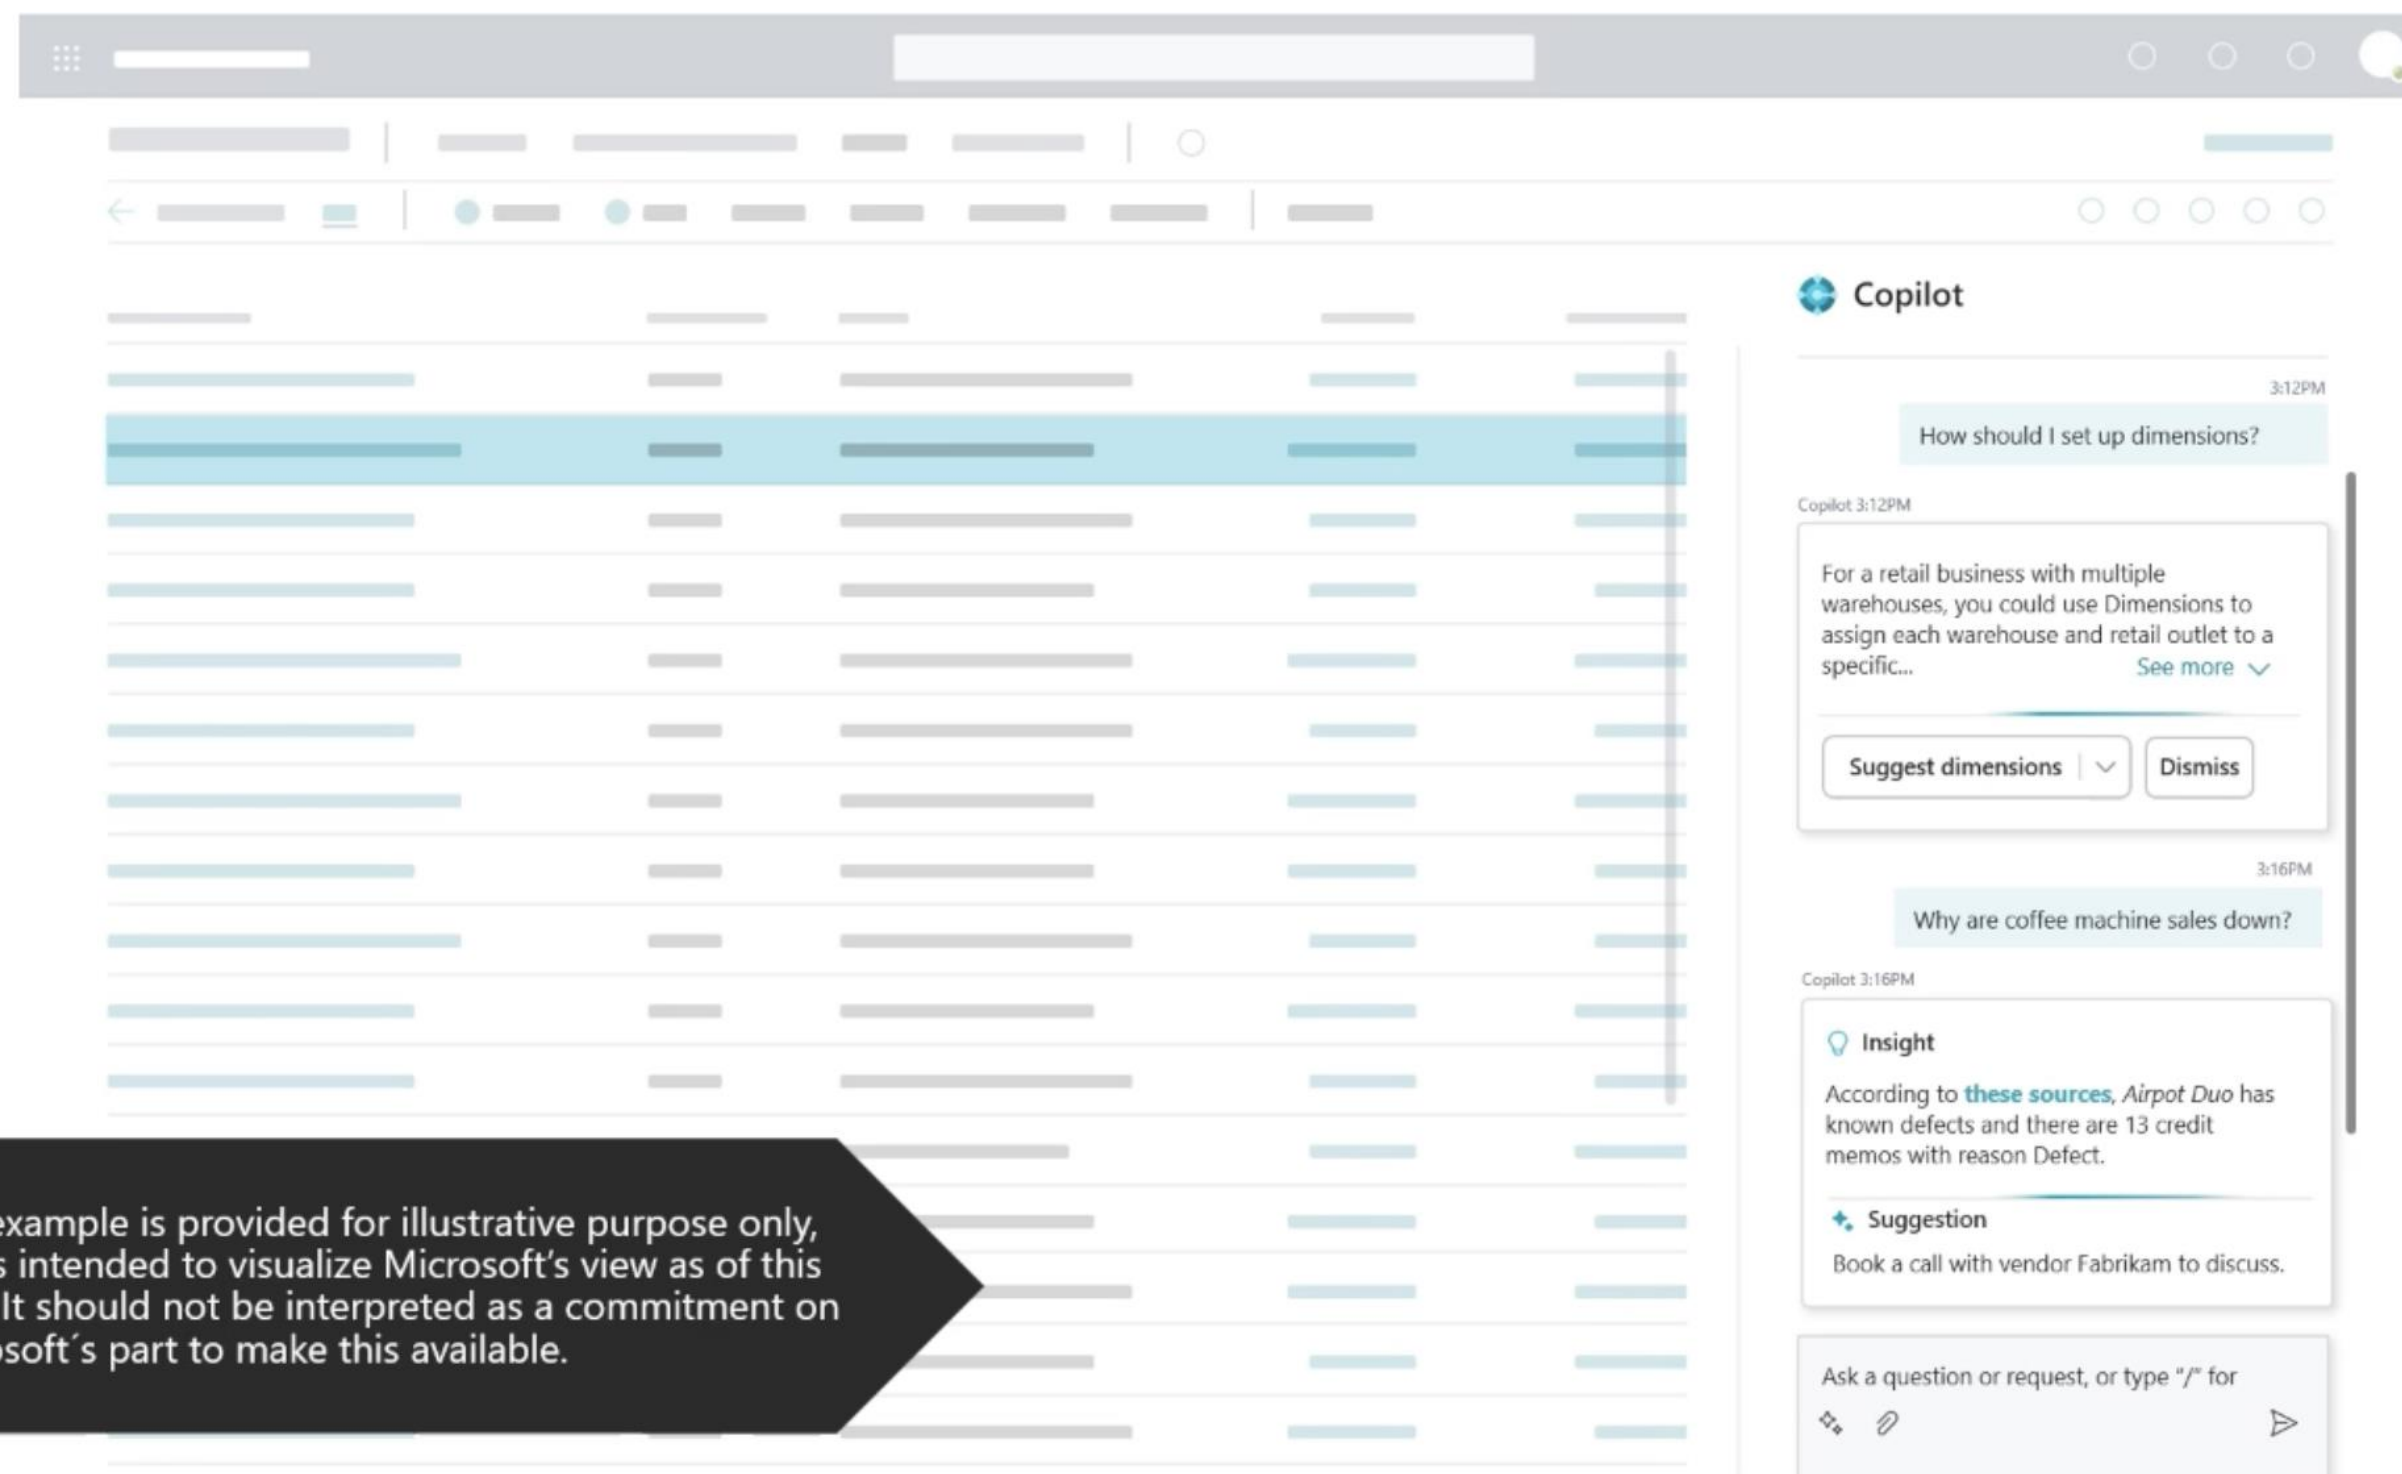
Total Incl. Tax (USD)

17,773.66

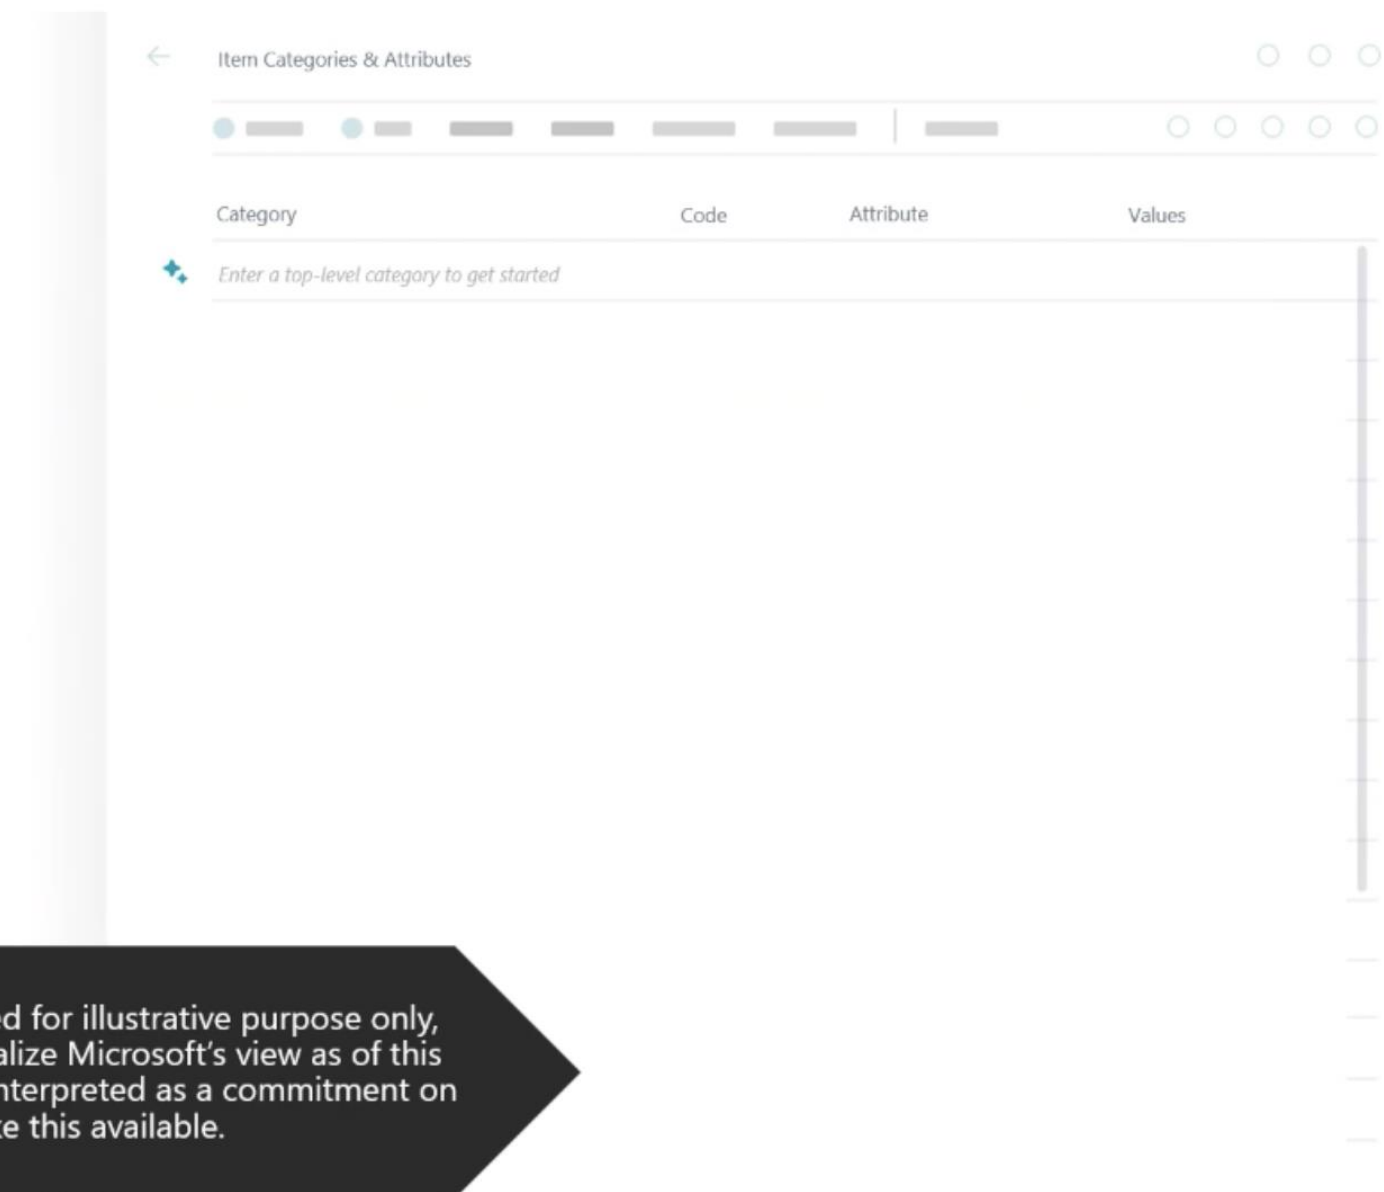


We will come from this

Through this



This example is provided for illustrative purpose only, and is intended to visualize Microsoft's view as of this date. It should not be interpreted as a commitment on Microsoft's part to make this available.

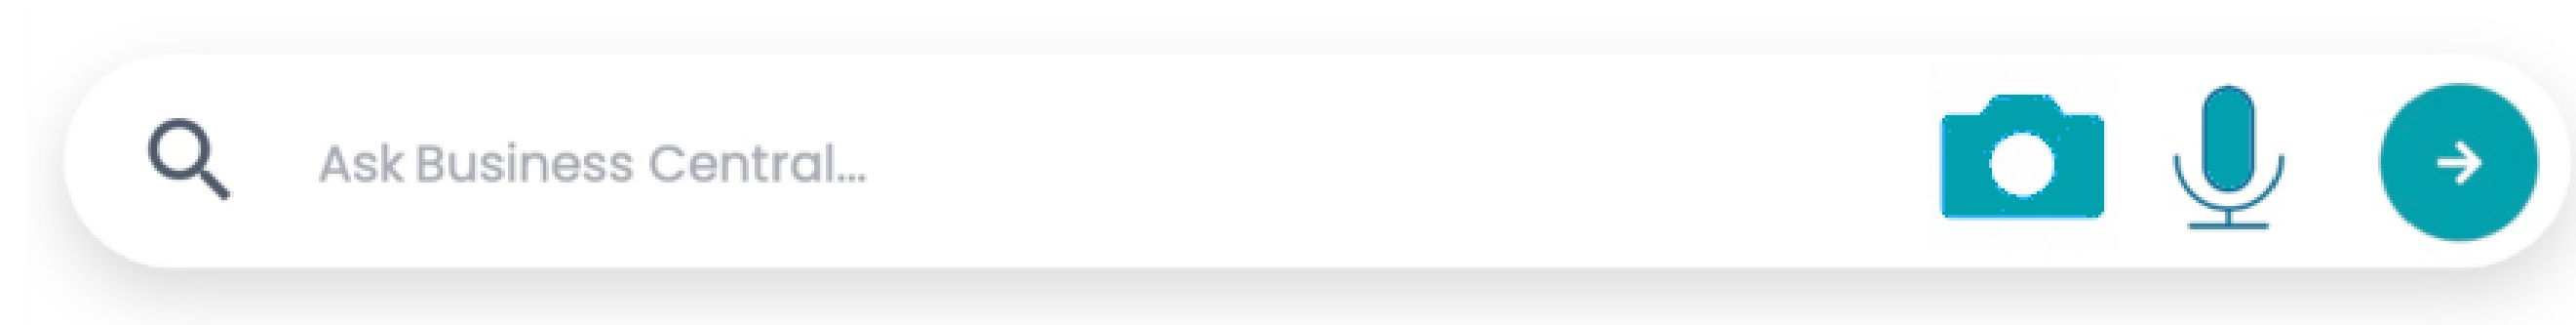


This example is provided for illustrative purpose only, and is intended to visualize Microsoft's view as of this date. It should not be interpreted as a commitment on Microsoft's part to make this available.



Through this

Into this ;)





Dmitry Katson | MVP
@DmitryKatson



Stefano Demilliani | MVP
@Demilliani



Q&A

Any Questions?



Thank
You!

