

Real-World Gen-AI with SAP

Generative AI

AI adoption is accelerating in the industry

Generative AI could add **\$2.6 trillion to \$4.4 trillion** incremental value annually to the global Economy

Existing productivity value from previous generation of AI and analytics

WIDESPREAD ADOPTION, DESPITE EARLY DAYS

1/3

companies globally are using Generative AI regularly in at least one business function

40%

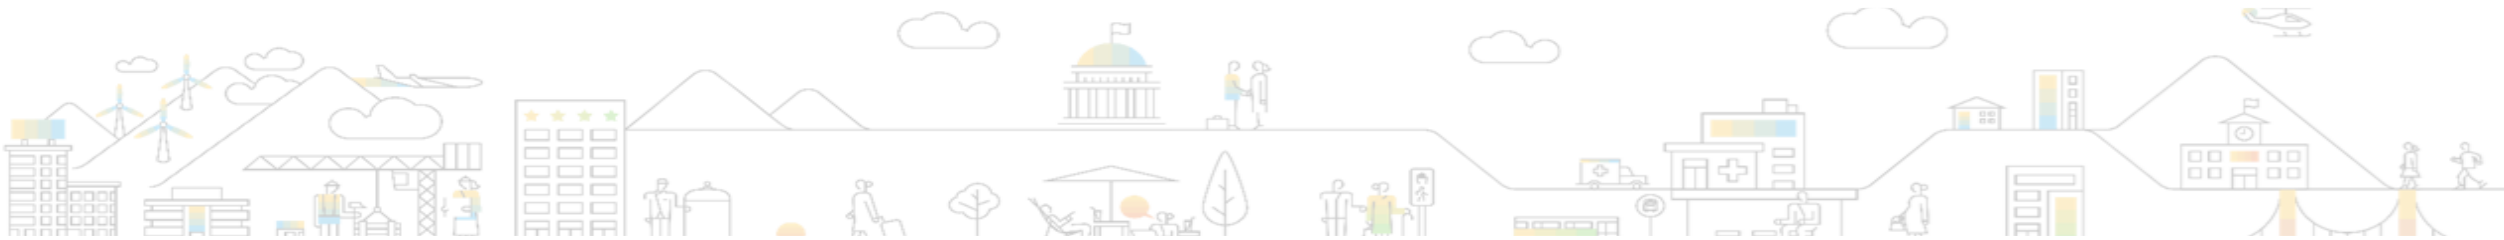
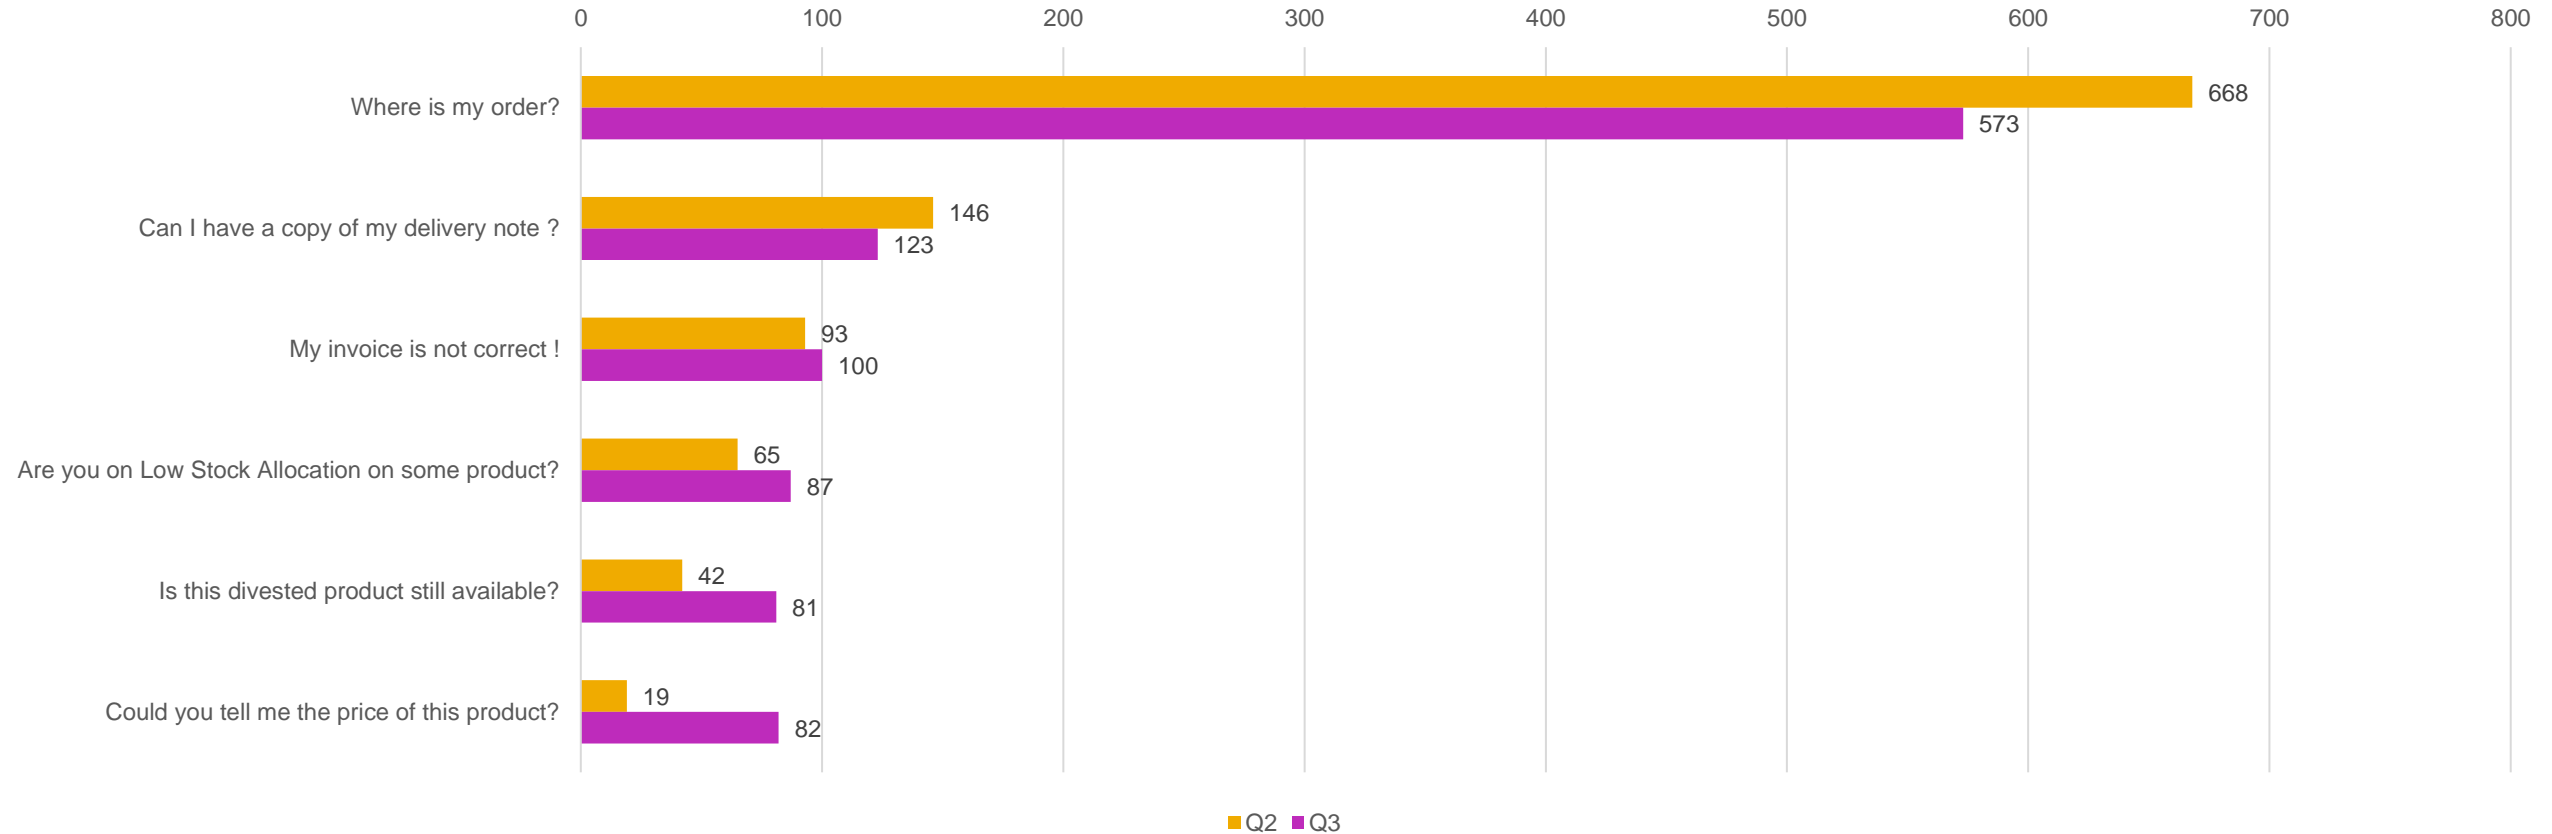
organizations will increase their investment in AI overall because of advances in Generative AI

60%

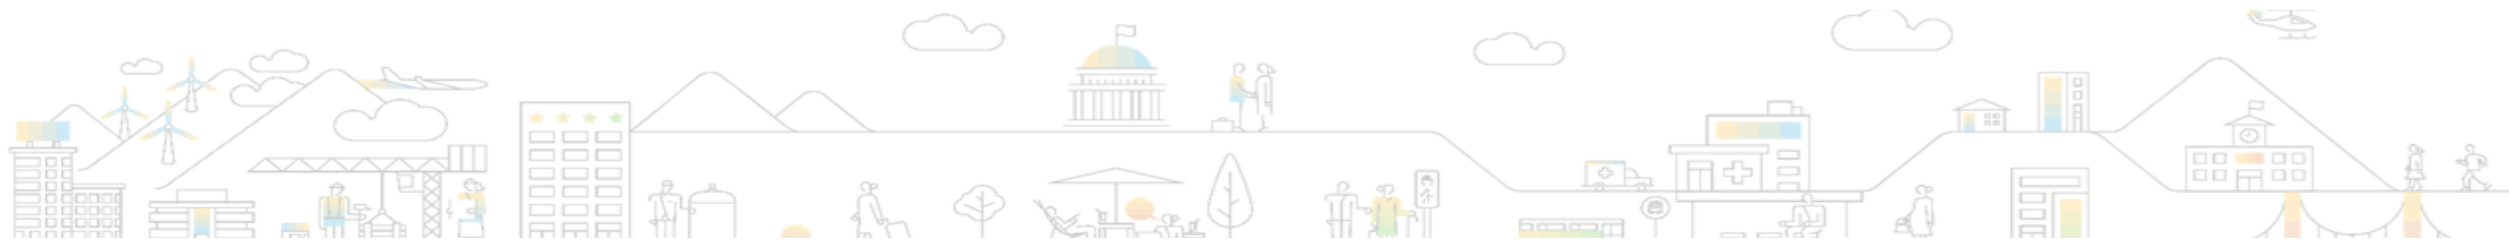
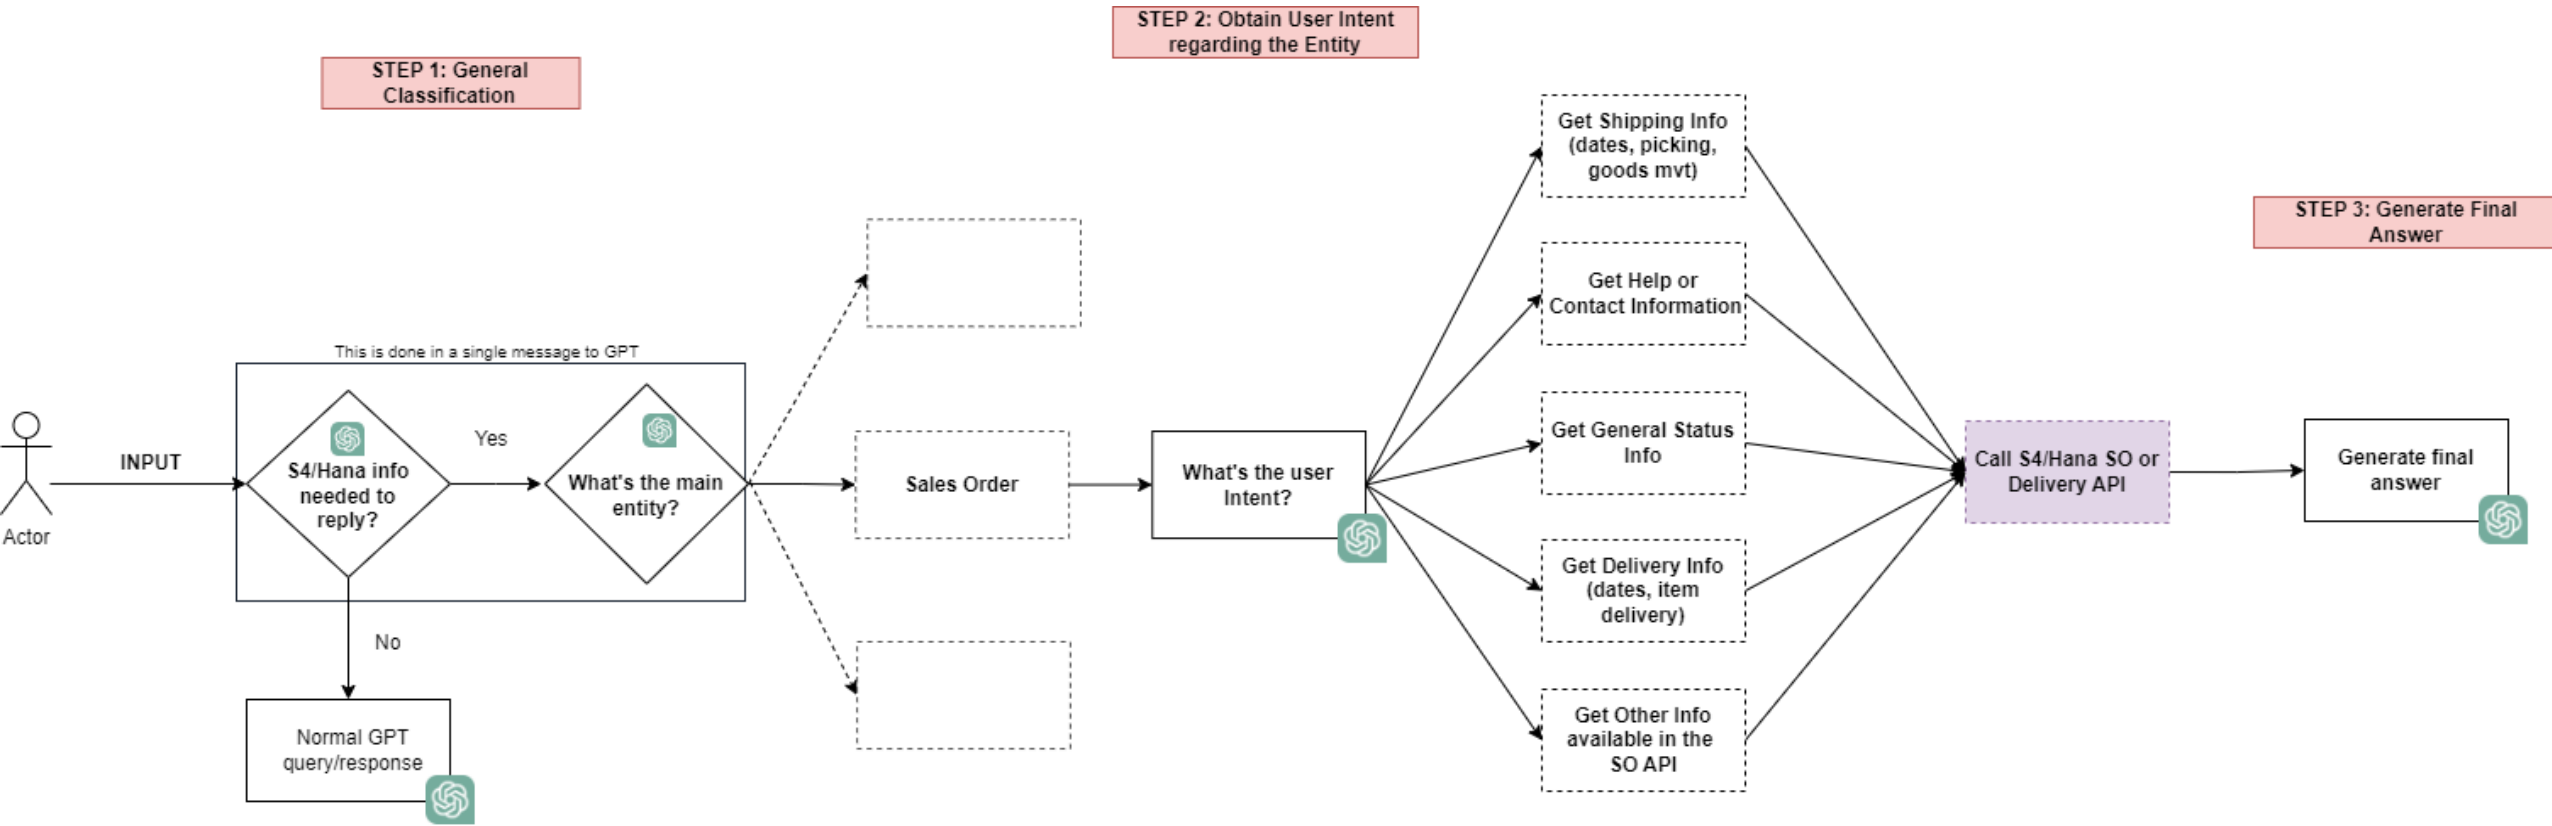
organizations with AI adoption are using Generative AI

What issues are we solving for our customers?

Inbound Communication



Our approach



OpenAI chat integration

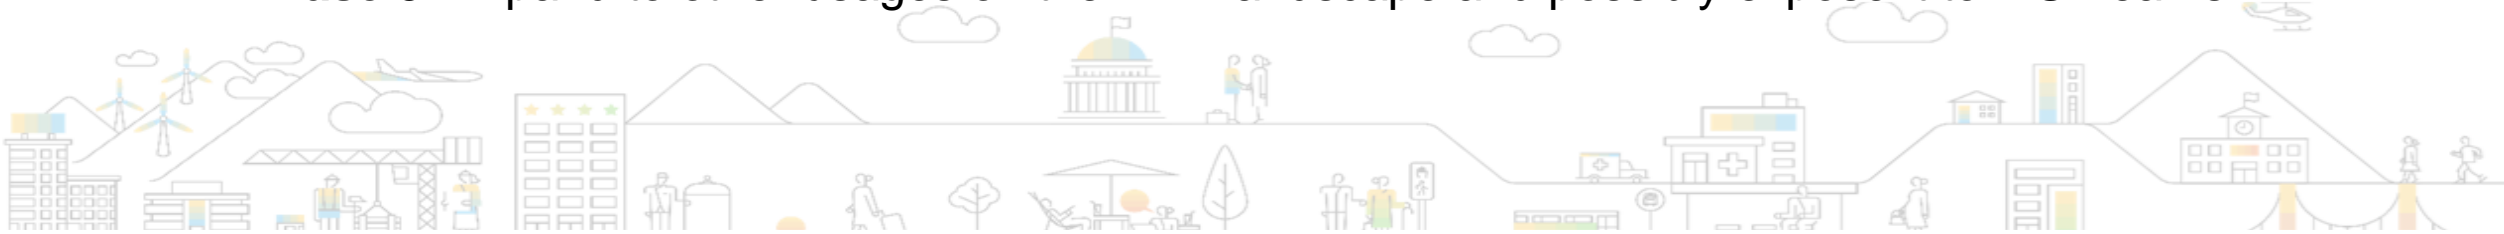
ChatGPT-like experience for business users

Motivation:

- SAP deprecating SAP CAI, presents as an opportunity to expand on new technologies for Chat assistance.
- Azure OpenAI allows fine-tune models that avoid ChatGPT problems regarding general purpose data and privacy of business data.

Plan:

- **Phase 1:** Azure OpenAI API just as a general-purpose chat bot. Demonstrate connectivity from Work Zone and Fiori.
- **Phase 2:** Use a fine-tuned model (all related customer/vendor oData), to facilitate a chatbot for external facing Customer and Vendor portal
- **Phase 3:** Expand to other usages on the BTP landscape and possibly expose it to MS Teams.



Chat history

+

<>

Today

Please provide the ID of ... 6 responses >

67%


-

+

Reset




Can you confirm the delivery date for my sales order?




Please provide the ID of the Sales Order to retrieve information from SAP.

...




6830




The confirmed delivery date for Sales Order ID 6830 is Thu Jul 09 1998. The order contains 49 items to be delivered and some of the items listed in the API response do not have delivery dates at the moment.

...




can you expand your answer?



The order consists of 49 items to be delivered. The confirmed delivery date is set for Thu Jul 09 1998. However, there are several items that do not have delivery dates yet, but they are not blocked. These items include various creams, ointments, lotions, powders, and tablets. Would you like me to provide a list of the items without delivery dates?

...




yes




"What is the delivery date and number of items for my order?"

[A]: "Your order has 49 items to be delivered and the confirmed delivery date is set for Thu Jul 09 1998. However, the following items have no delivery dates at the moment (but are not blocked): 490 (GRANEODIN OINT (1TUB...), 480 (TRIADCORTYL OINT (1T...), 470 (TRIADCORTYL CRM (1TU...), 460 (NYSTAN OINT (1TUBEX3...), 450 (NYSTAN CRM (1TUBEX3...), 440 (ECOSTATIN CRM 0.1% (...), 430 (ECONACORT CRM 0.1% (...), 440 (ECOSTATIN CRM 0.1% (...), 430 (ECONACORT CRM 0.1% (...), 420 (KERI THERAPEUTIC LOT...), 410 (KERI THERAPEUTIC LOT...), 400 (ALPHA KERI BATH OIL ...), 390 (ALPHA KERI BATH OIL ...), 380 (NYSTAN PASTILLES (1P...), 390 (ALPHA KERI BATH OIL ...), 380 (NYSTAN PASTILLES (1P...), 370 (NYSTAN R/MIX SUSP (1...), 360 (FUNGILIN LOZ 10MG (6...), 350 (KENALOG SUS-INJ 40MG...), 340 (KENALOG SUS-INJ 40MG...), 330 (TRIADCORTYL OTIC (1T...), 340 (KENALOG SUS-INJ 40MG...), 330 (TRIADCORTYL OTIC (1T...), 320 (ADCORTYL-IN-ORABASE ...), 310 (PREGESTIMIL PWD (1CA...), 300 (PROSOBEE PWD (1CANX4...), 290 (NUTRAMIGEN PWD (1CAN...), 280 (DUTONIN TAB 200MG (4...), 290 (NUTRAMIGEN PWD (1CAN...), 280 (DUTONIN TAB 200MG (4...), 270 (DUTONIN TAB 100MG (4...), 260 (BUSPAR TAB 10MG (1BT...), 250 (NYSTAN VAG CRM (1TUB...), 240 (NYSTAN TAB (1BLPX56)...), 230 (FUNGILIN TAB 100MG (...), 240 (NYSTAN TAB (1BLPX56)...), 230 (FUNGILIN TAB 100MG (...), 220 (VELOSEF PSUS 250MG/5...), 210 (VELOSEF CAP 250MG (2...), 200 (VELOSEF CAP 500MG (2...), 190 (BAXAN PSUS 500MG/5ML...), 180 (BAXAN PSUS 250MG/5ML...), 190 (BAXAN PSUS 500MG/5ML...), 180 (BAXAN PSUS 250MG/5ML...), 170 (MEGACE TAB 160MG (3B...), 160 (HYDREA CAP 500MG (10...), 150 (FLORINEF TAB 0.1MG (...), 140 (LIPOSTAT RM TAB 20MG...), 130 (LIPOSTAT RM TAB 10MG...), 140 (LIPOSTAT RM TAB 20MG...), 130 (LIPOSTAT RM TAB 10MG...), 120 (STARIL TAB 20MG (2BL...), 110 (STARIL TAB 10MG (2B...), 100 (CAPOZIDE TAB 50/25MG...), 90 (BOOTS CAPOTEN TAB 50...), 80 (BOOTS CAPOTEN TAB 25...), 90 (BOOTS CAPOTEN TAB 50...), 80 (BOOTS CAPOTEN TAB 25...), 70 (BOOTS CAPOTEN TAB 25...), 60 (BOOTS CAPOTEN TAB 12...), 50 (ACEPRIL TAB 50MG (1B...), 40 (ACEPRIL TAB 12.5MG (...), 30 (LIPOSTAT RM TAB 40MG...), 40 (ACEPRIL TAB 12.5MG (...), and 30 (LIPOSTAT RM TAB 40MG...)."

...



and what is the status?



The overall delivery status for all items in the sales order API response is Partially delivered. The overall processing status for all items is Partially processed.

...





hey



Hello, how can I assist you today?



are you able to confirm the delivery date for my sales order?



Please provide the ID of the Sales Order to retrieve information from SAP.



6445



"Yes, according to our records, the delivery date for Sales Order ID 6445 is scheduled for Thursday, July 9th, 1998. The order contains one item with the ID of 490."



who can I contact for more information?



For more information on this topic, please contact customerinfo.es@bms.com.



how much is 3+6?



The answer is 9.




who created the sales order that I mentioned before?



The user with ID 204805 created the sales order that was mentioned before.




BMS Chat (OpenAI)
Chat GPT




Other

Sales Orders


 0

Process Automation Apps


Process Trigger
Request Form





Visibility Scenario
Instances





My Inbox


 0


ChatBot  Open App

 Hello!

 Hello! How can I assist you today?

 Find contact information for our supplier ABC in SAP

 Supplier ABC's contact information in our SAP system is as follows:
-Supplier Name: ABC Pharmaceuticals
-Contact Name: John Smith
-Email: john.smith@abcpharma.com
-Phone: +123-456-7890
-Address: 123 Supplier Lane, Supplier City, Supplier Country



Chat with me!

Click on me!

OpenAI chat integration

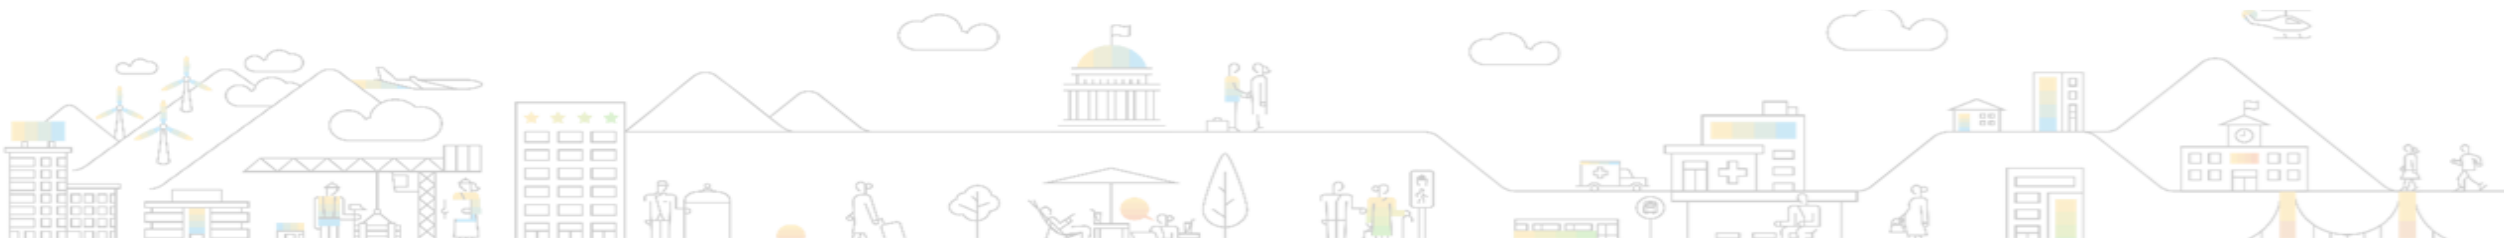
Main challenges and approach

Challenges:

- Security and Prompt injection. Retrieve unauthorized data
- Hallucinations of Large Language Models
- Accurate Fine-tuning and prompt-engineering

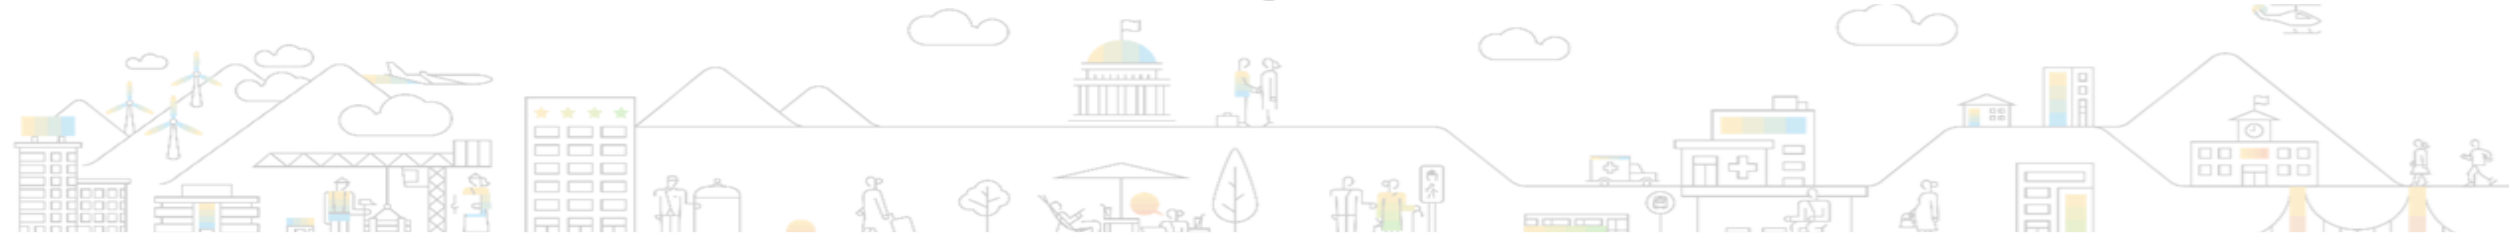
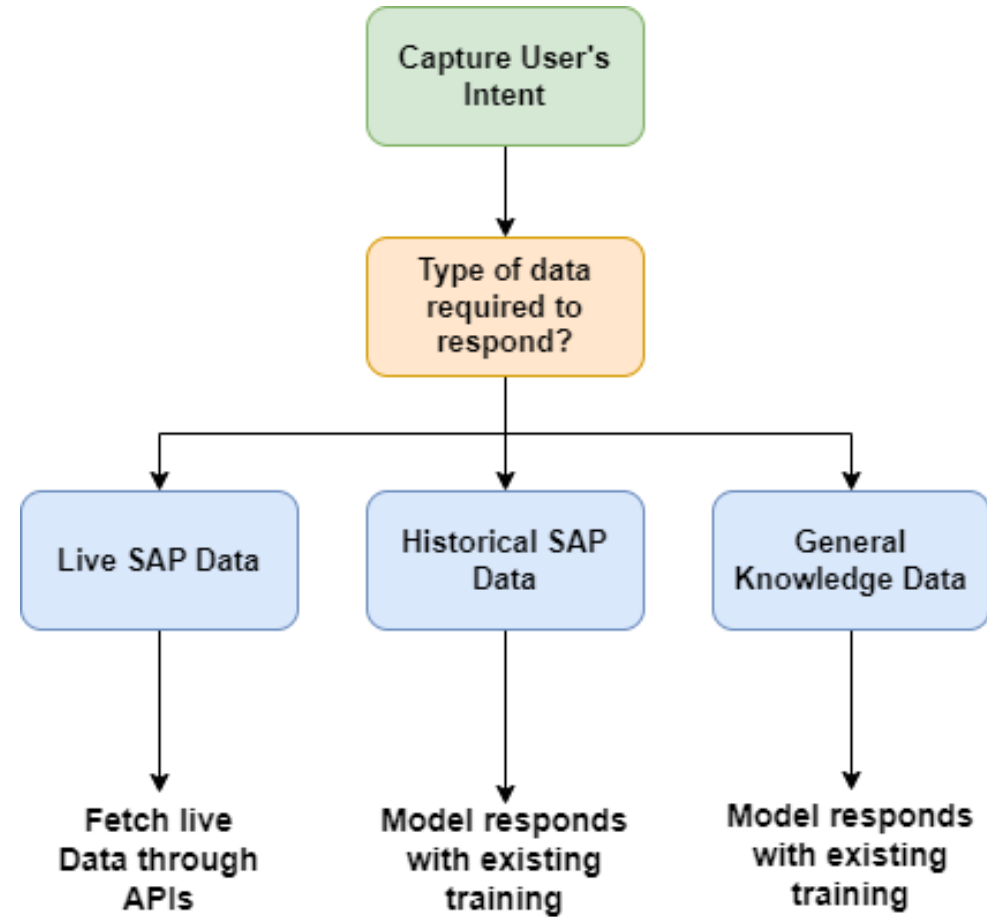
Approach:

- Capture user's intent
- Call APIs when Live Data is required to respond
- Use fine-tuned training and custom vector databases for responding historical questions
- **Key concept: Prompt-Engineering**
- Start with specific use cases and expand from there



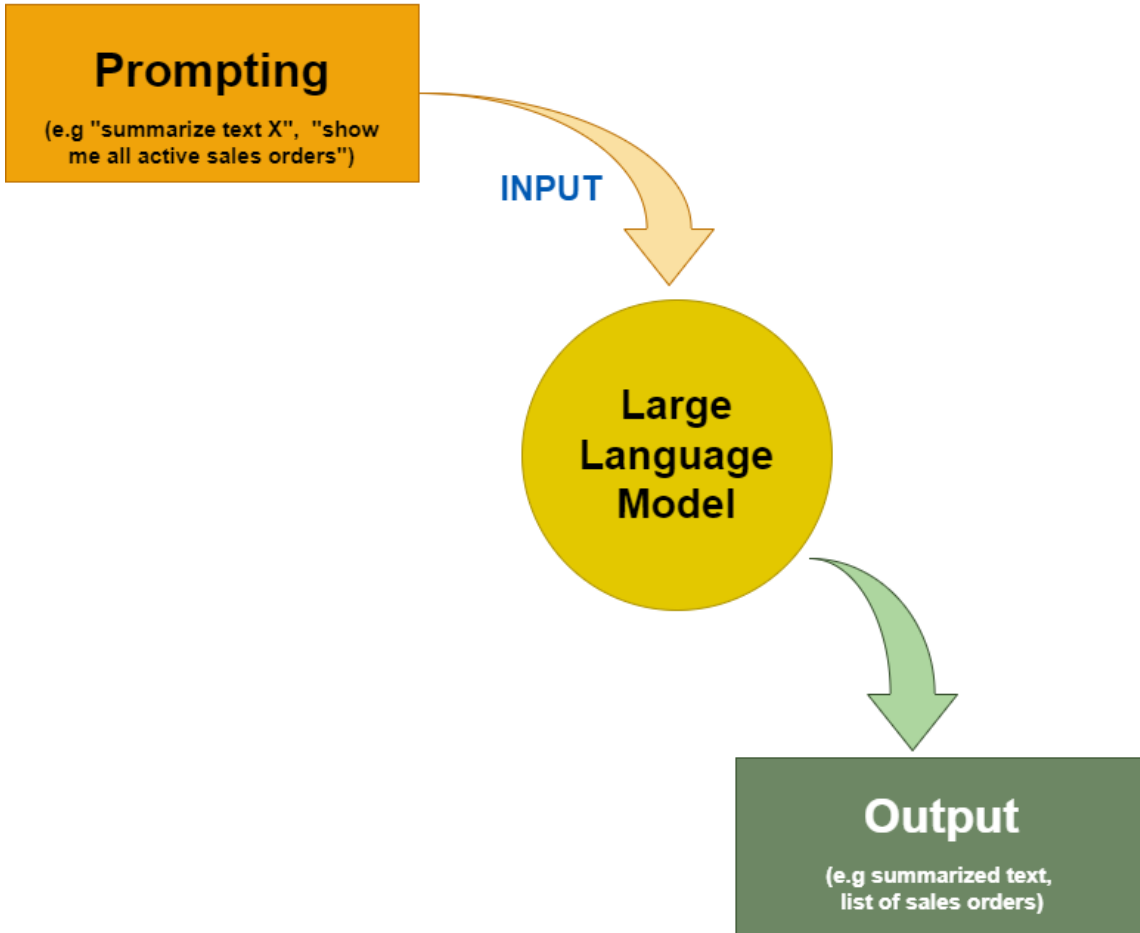
OpenAI chat integration

Main challenges and approach



Prompt-Engineering

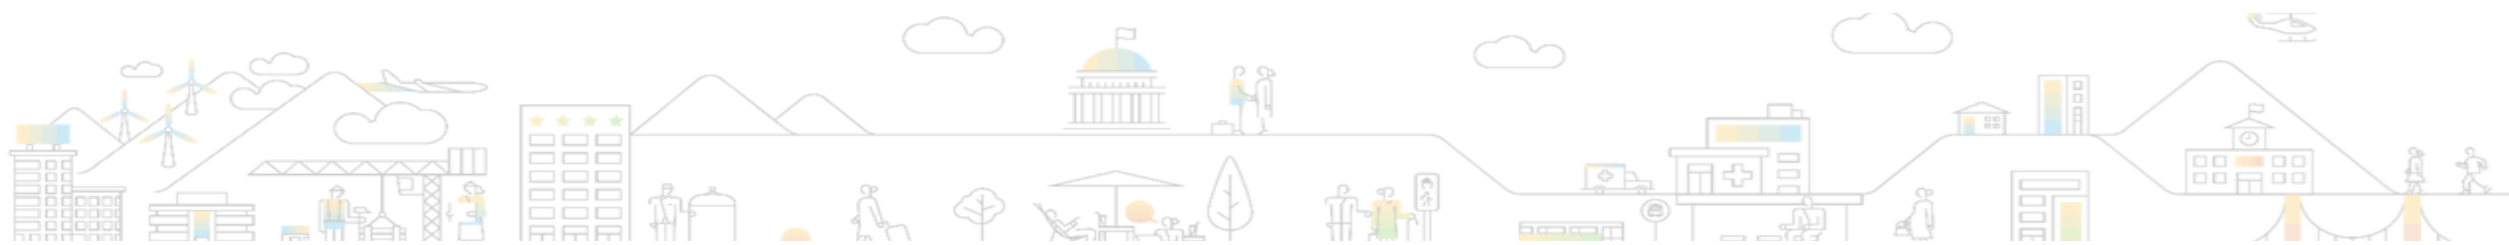
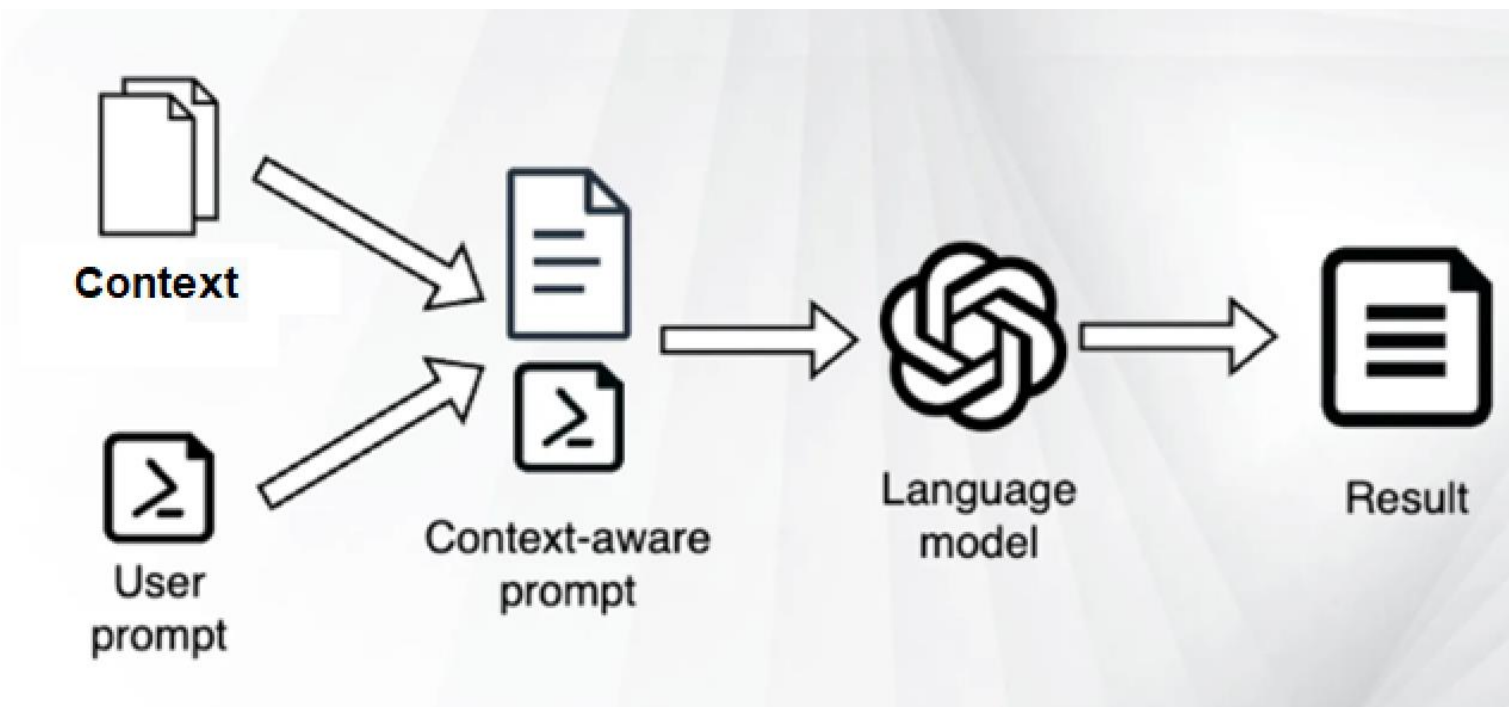
Main challenges and approach



- LLM models provide the most probable response according to training and the given input, but they don't provide an accuracy value assigned to the response given, which might confuse users that take the result as 100% accurate.
- The quality of the inputs determine the quality of the output
- **Design effective prompts** increases the likelihood that the model will return an accurate response
- One methodology is the **Chain-of-thought** prompting, in which the initial prompt is added with several additional context texts that describe the rational process that the LLM should replicate.



Improving the prompt



Chain-of-thought (COT)

Prompt-Engineering / approaches to improve response

Normal Input (Zero shot):

Prompt

If John has 5 pears, then eats 2, and buys 5 more, then gives 3 to his friend, how many pears does he have?



Output

John has 8 pears. ❌ incorrect

Full COT

Prompt

Example Q: Joe has 23 apples. If they used 20 to make lunch and bought 6 more, how many do they have?

Example A: Joe had 23 apples originally. They used 20, so they had $23-20=3$. Then they bought 6 more so they have $3+6=9$. The answer is 9

Q: If John has 5 pears, then eats 2, and buys 5 more, then gives 3 to his friend, how many pears does he have?



Output

John starts with 5 pears. He eats 2 pears, leaving him with 3 pears. He buys 5 more pears, giving him a total of 8 pears. He gives 3 pears to his friend, leaving him with only 5 pears. ✅ correct

Zero shot COT (Kojima et al.)

Prompt

If John has 5 pears, then eats 2, and buys 5 more, then gives 3 to his friend, how many pears does he have?

Let's think step by step.



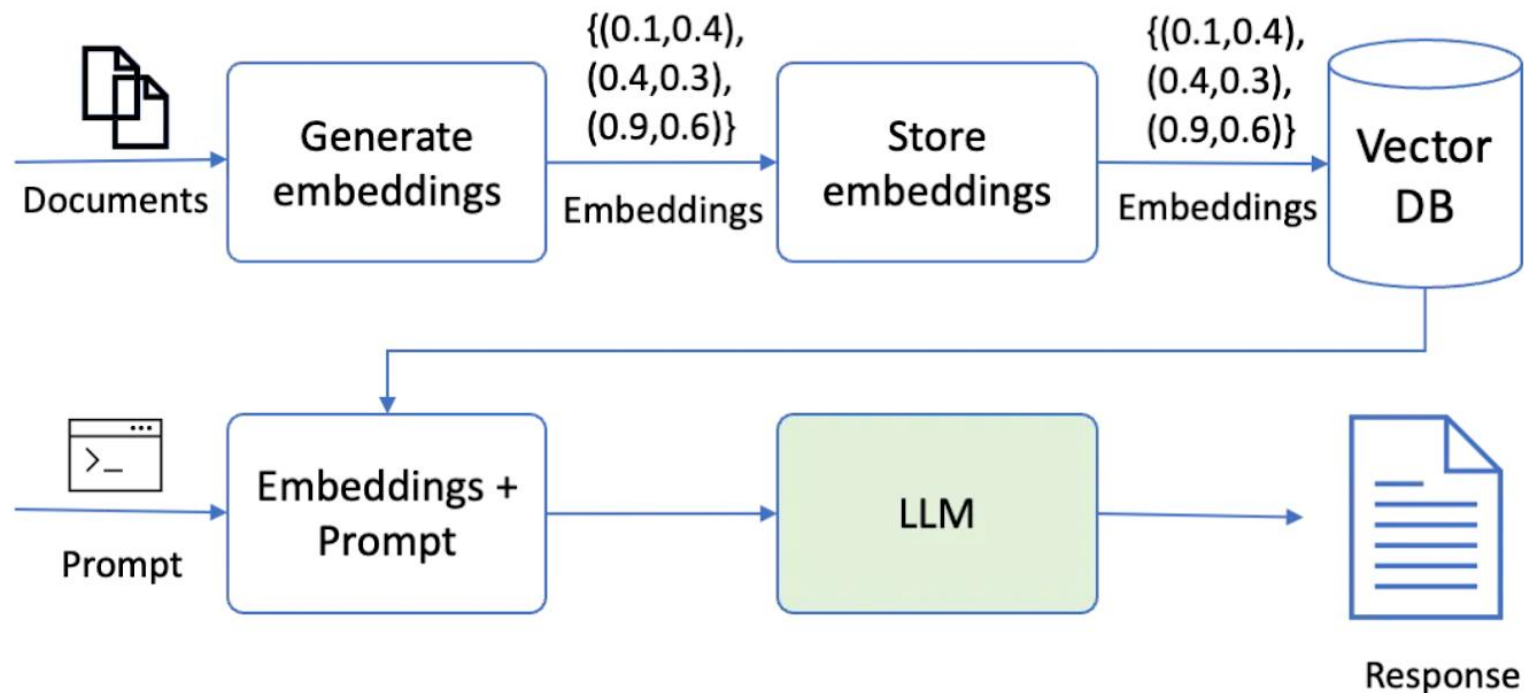
Output

John starts with 5 pears. He eats 2 pears, leaving him with 3 pears. He buys 5 more pears, giving him a total of 8 pears. He gives 3 pears to his friend, leaving him with only 5 pears. ✅ correct

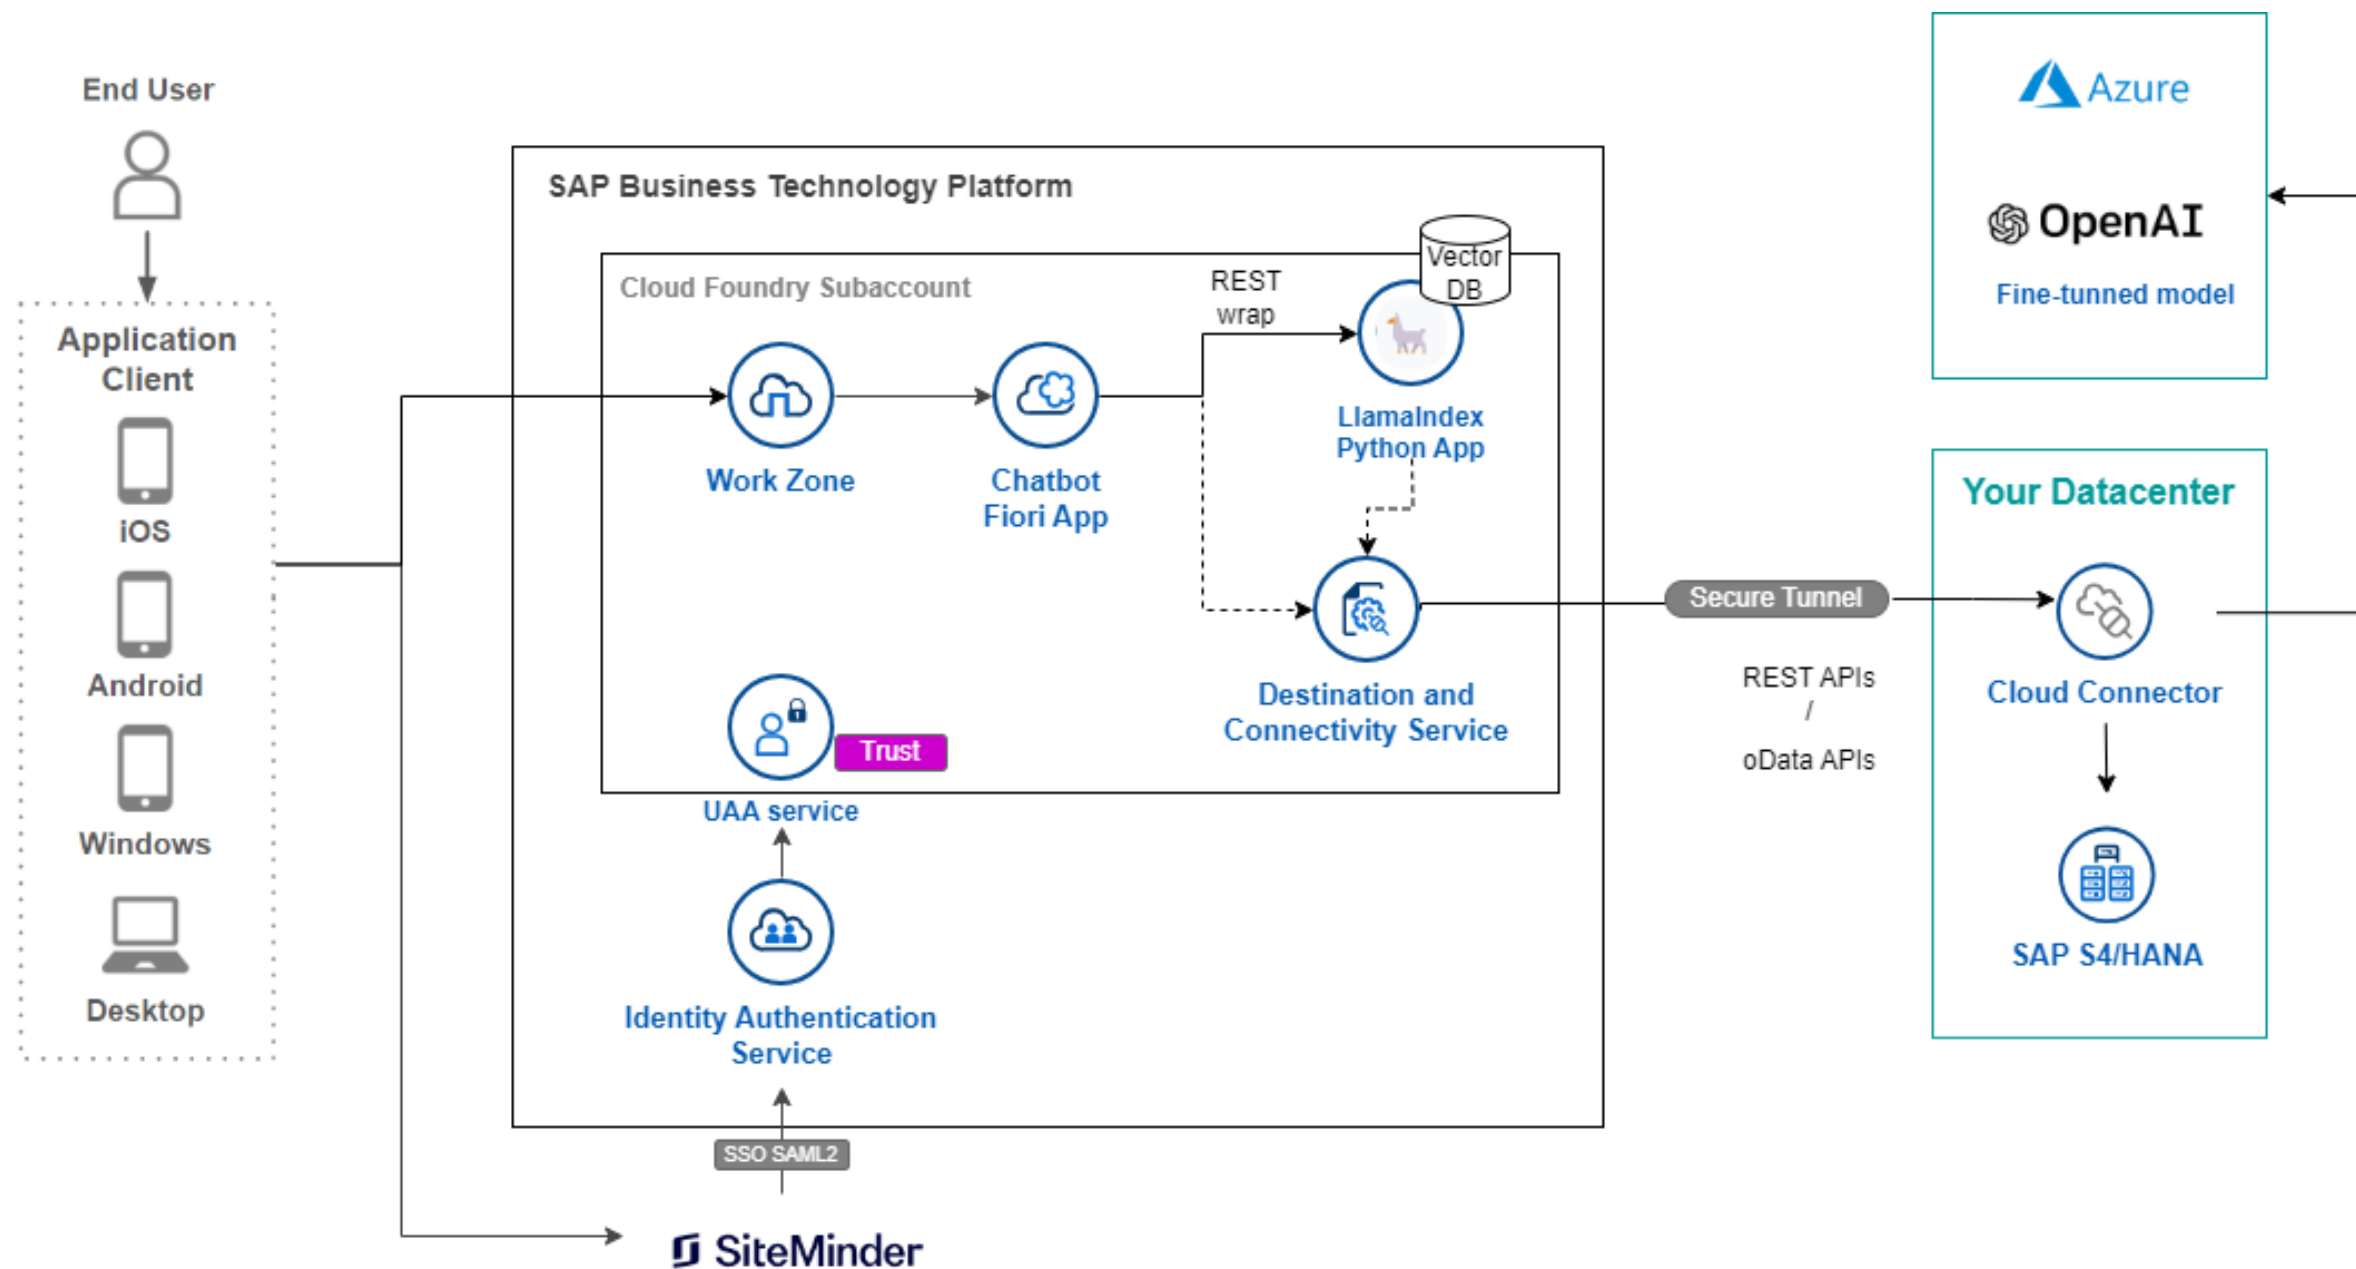
Vector Databases

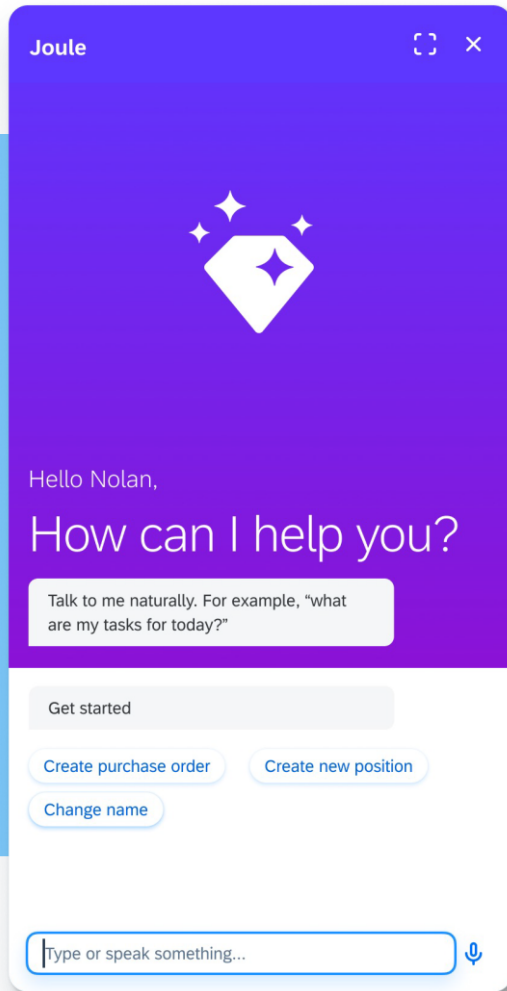
What if the “context” was already converted to embeddings and stored somewhere?

Embeddings are representations of text as real number vectors that contain the encoded context and meaning of the text.



This concept is known as **Retrieval Augmented Generation (RAG)**





Joule in your SAP applications

Joule is included in
your SAP applications,
starting with SAP Start
and SAP SuccessFactors,
with more to come...

Access SAP SuccessFactors and SAP S/4HANA Cloud, public edition from Joule in SAP Start

Give employees and managers an AI assistant that enhances their experience while reducing manual HR efforts

© 2023 SAP SE or an SAP affiliate company. All rights reserved.





Workspace

SAP AI Core Administr...

Generative AI Hub

Model Grounding

Prompt Editor

Prompt Management

Generative AI / Prompt Management

Prompt Management

Filters

All

Favourites

Search

Created On

Today

Last 7 days

Model

Falcon-40b

gpt-3.5-turbo

gpt-4

Collection

Prompts (X)

Create



Name

Model

Collection

Tags

Created On

Last Modified

☐ **Dynamic Table**
Versions: 1

☆ gpt-4

Data retrieval

Yesterday, 4:14:09 PM

Yesterday, 4:14:09 PM

☐ **HANA Views**
Versions: 1

★ gpt-4

☐ **HANA SQL Lag**
Versions: 4

☆ Falcon-40b

Grounding

☐ **HANA Calc View**
Versions: 1

☆ gpt-4

Optimization

PREVIEW



SAP AI Core

Generative AI Hub

Workspace

SAP AI Core Administr...

Generative AI Hub

Model Grounding

Prompt Editor

Prompt Management

Prompt Editor

Name: HANA SQL Lag

Save

Create New

Select

Prompt

Message

Instruction:
You are an SAP HANA Cloud expert.
You are provided multiple content items that are related to the prompt you have to answer.
Use the following pieces of content to answer the user's question.

User:
I want to calculate a shortest path. How do I do that?

Context:

Model

Selected Model:

Falcon-40b

Parameters (2)

Frequency_penalty:
Fine-tune the model's behaviour regarding the repetition in the responses.

0.5 (0)

Max_tokens:

Limit the length of the response generated by the model.

500 (500)

Response

The SHORTEST_PATH function is implemented as a built-in function in SAP HANA Cloud. It uses the graph search algorithms provided by the SAP HANA Graph Library.

The implementation of weight function depends on the type of the function parameter. If the parameter is an 'EDGE' object, the function should connect to the edge's begin vertex and check the edge's weight and/or the edge's parent edge's weight. If the parameter is a 'VERTEX' object, the function should check the vertex's degree and/or its in-degree.

The function may use the 'begin' or 'double' data type for the function's result. If the weight function returns one of these data types, if the function returns an 'integer' data type, it is converted to a 'begin'. If the function returns a 'begin' data type, it is converted to a 'double'. When the 'SHORTEST_PATH' function is called with a custom weight function as an argument, the closure is called for each edge in the graph. The closure should return the weight distance value for the edge or the edge's parent edge. The closure is executed in the context of the 'WEIGHTEDPATH' object and has access to the edge's properties and the current state of the SHORTEST_PATH search.

Prompt Tokens: 72 Response Tokens: 110

Thank you!



www.blueboot.com

