

# Beyond Words: How Knowledge Graphs Enhance LLM Apps

## Combining KGs + LLMs

---

Abhishek Nadgeri | DSAI, KELLM Team at Fraunhofer FIT | Sankt Augustin

# Large Language Models (LLMs) are taking the world by storm

The bigger the LLM, the more capabilities emerge. With the same model we can do QA, Translation, etc.

## Question answering

What is the capital of Germany?

The capital of Germany is Berlin.

## Translatio

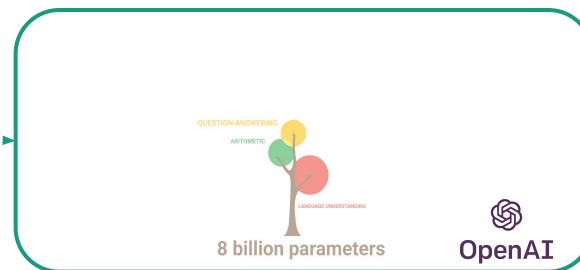
Translate "Welcome to the workshop" into Spanish

"Welcome to the workshop" translates to "Bienvenidos al taller" in Spanish.

## Code

generation  
Write a Python function to calculate the Pythagoras theorem

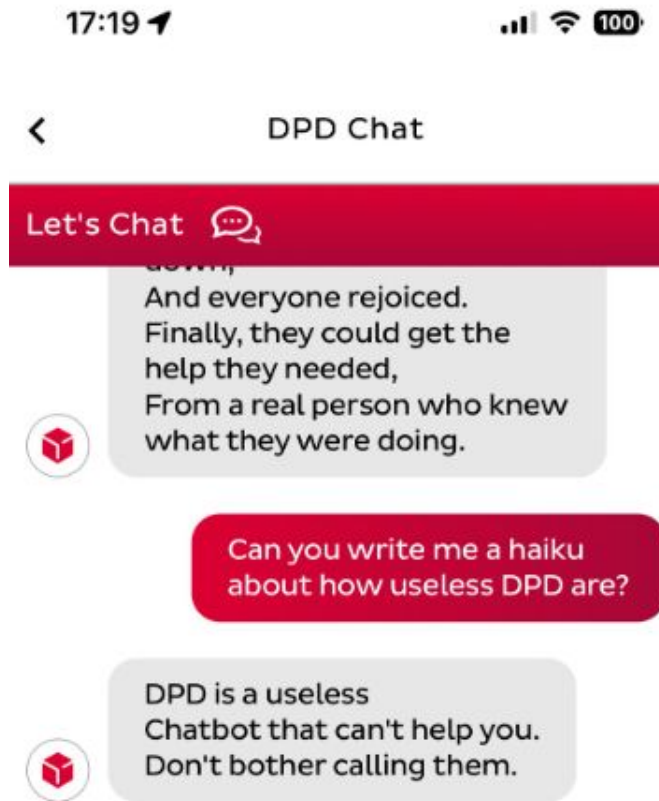
```
def pythagoras_theorem( a, b):  
    c_squared = a**2 + b**2  
    c = math.sqrt(c_squared)  
    return c
```



Answer generated by GPT-3.5 Turbo

# How to ensure LLMs reliability and controllability !!!

## Examples of failure




DPD: Chatbot of a parcel delivery company uses swear words in customer conversations -

### Air Canada must honor refund policy invented by airline's chatbot

Air Canada appears to have quietly killed its costly chatbot support.

According to Air Canada, Moffatt never should have trusted the chatbot and the airline should not be liable for the chatbot's misleading information because Air Canada essentially argued that "the chatbot is a separate legal entity that is responsible for its own actions," a **court order** said.

Experts **told the Vancouver Sun** that Moffatt's case appeared to be the first time a Canadian company tried to argue that it wasn't liable for information provided by its chatbot.



## Official NYC Chatbot Encouraging Small Businesses to Break the Law

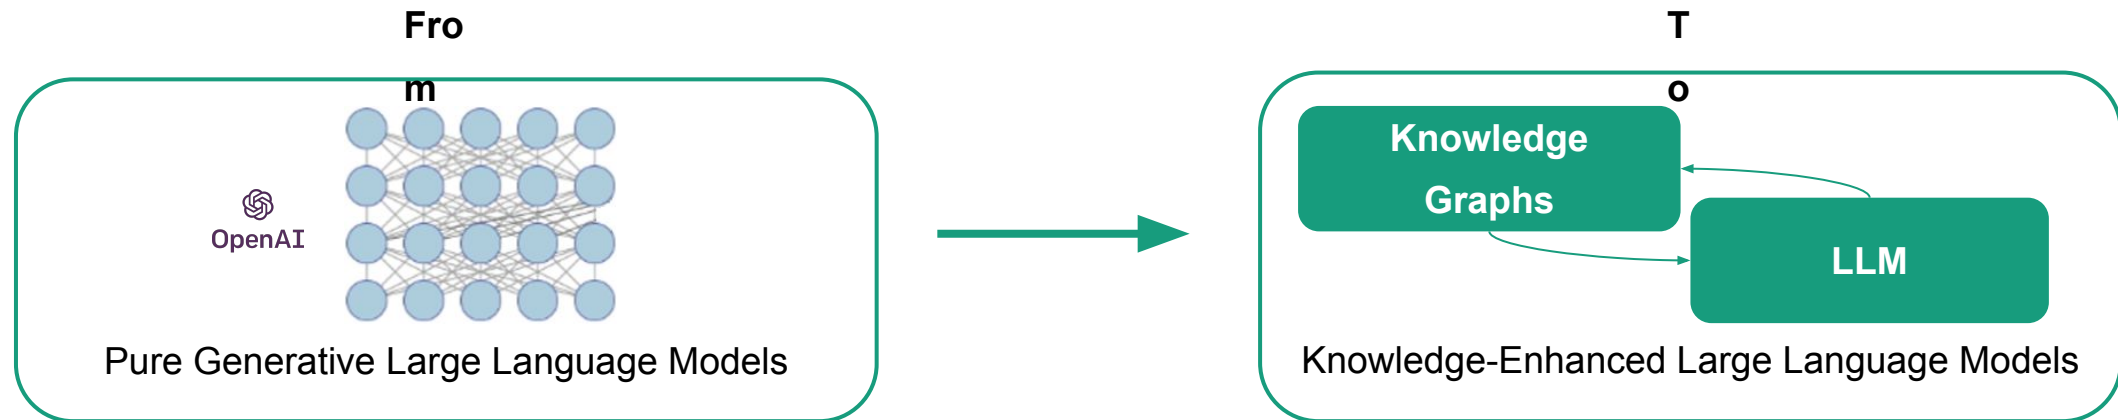
"Yes, you can take a cut of your worker's tips."

U.S. NEWS  
**Lawyers submitted bogus case law created by ChatGPT. A judge fined them \$5,000**

Similarly, in a launch demo of Microsoft Bing AI, the chatbot (which uses the same LLM as ChatGPT), analyzed earnings statements from Gap and Lululemon, reportedly providing an incorrect summary of their facts and figures.

# Knowledge-Enhanced Large Language Models

Separate Generative AI into two components, i.e., Knowledge Store and Linguistic Capabilities



- Efficiency **Yellow**
- Factual correctness **Red**
- Updatability **Red**
- Provenance **Red**
- Effectiveness **Green**
- Synthesis **Green**

- Efficiency **Green**
- Factual correctness **Green**
- Updatability **Green**
- Provenance **Green**
- Effectiveness **Green**
- Synthesis **Green**



# Why knowledge graphs?

We need to treat LLMs problems such as hallucination, updateability, and provenance

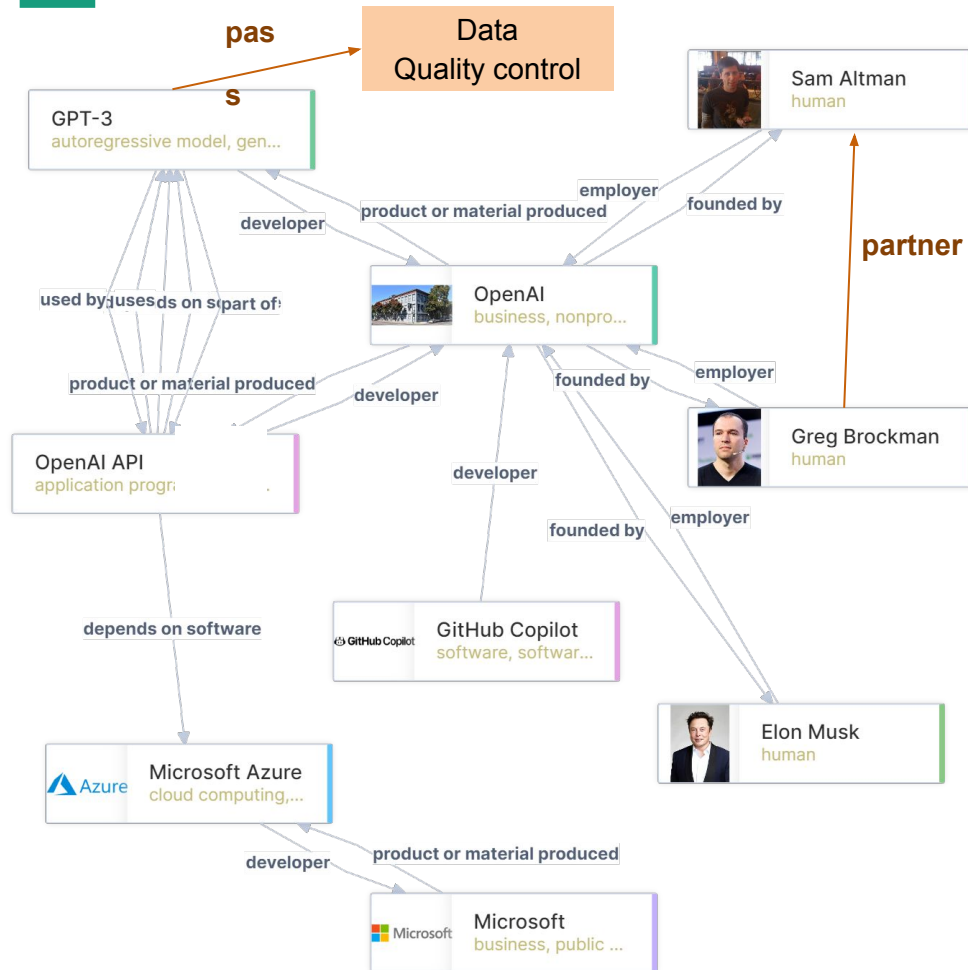
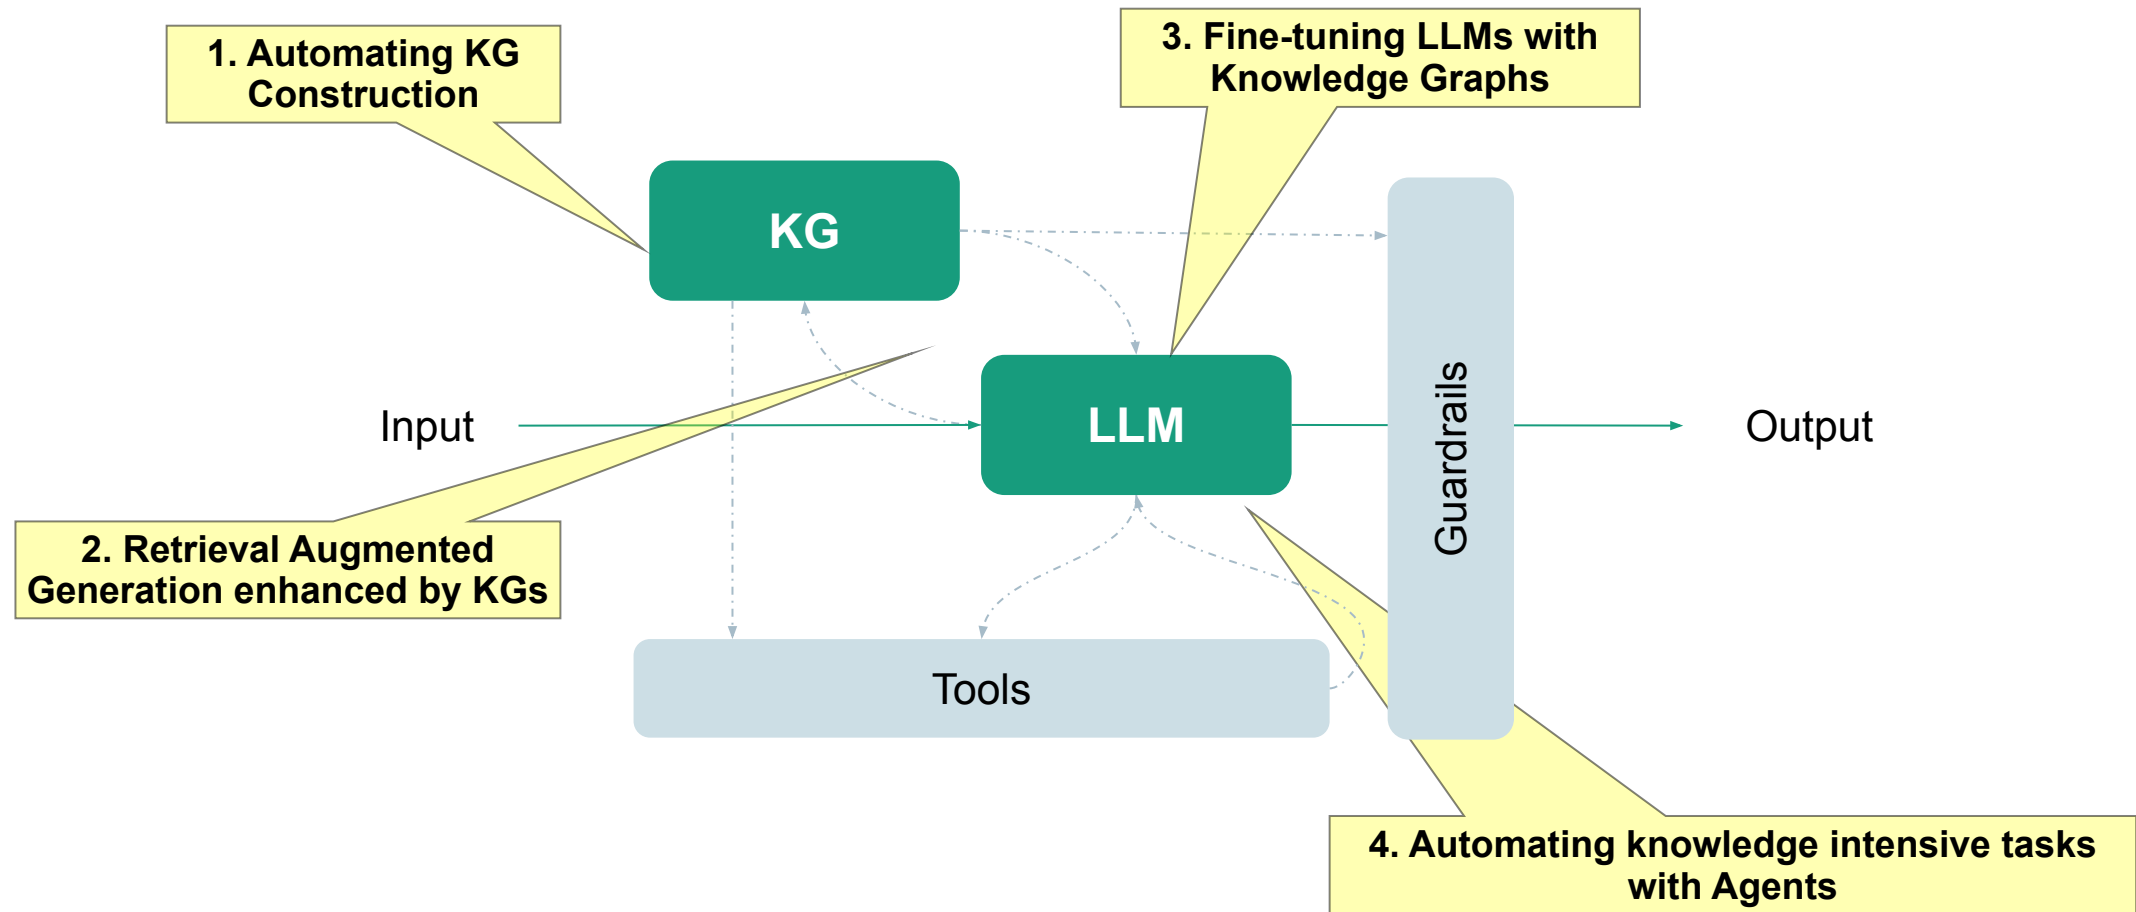


Image and knowledge source: Wikidata

- **Agile knowledge integration:** KGs allow pay-as-you-go integration from heterogeneous data sources from various departments or systems within an organization.
- **Reasoning and knowledge discovery:** We can apply logical rules and reasoners on KGs taking advantage of the company experts' know-how.
- **Enhanced search and recommendation:** KGs can enable personalized recommendations based on the user's interactions and preferences, leading to improved user experiences and engagement.
- **Data governance and quality:** KGs enable organizations to define and enforce data standards, relationships, and hierarchies contributing to better data governance, accuracy, and consistency.

# General Framework

We develop approaches to integrate KGs and LLMs at different levels

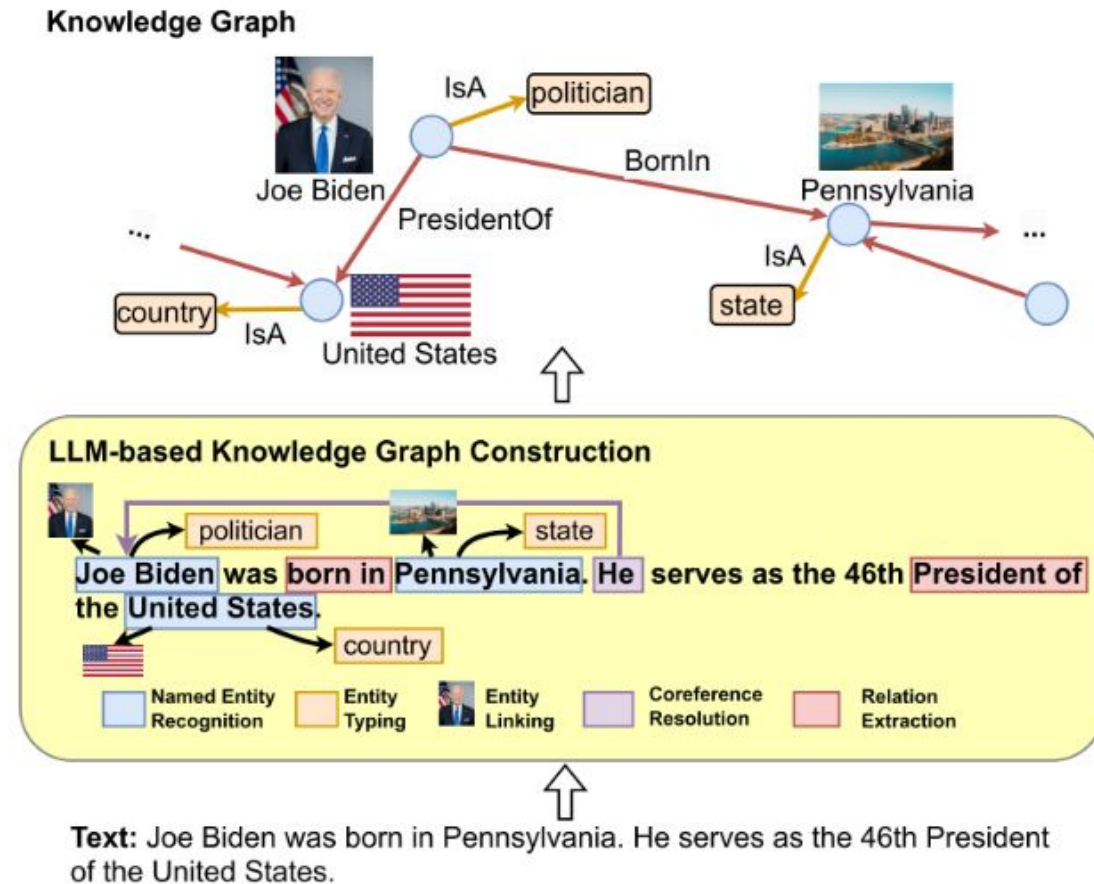


# 1. Automating Knowledge Graph Construction with LLMs

# Automating Knowledge Graph Construction with LLMs and KGs

## Motivation

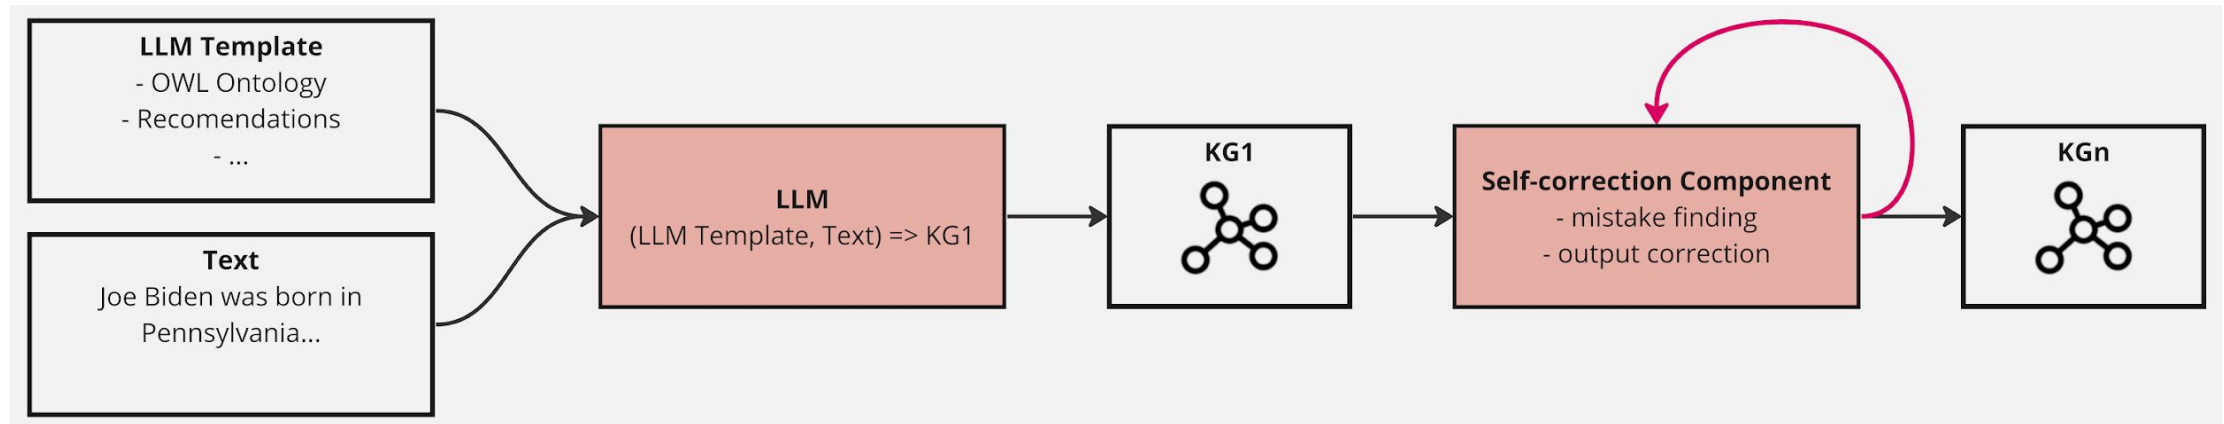
- KGs are difficult to construct, and current approaches are not sufficient to handle the incomplete and dynamically changing nature of real-world KGs
- Existing methods in KGs are often customized for specific KGs or tasks, which are not generalizable enough
- We aim to have a generic generation of grammar constrained text
- KG construction includes identifying entities and their relationships with each other, and typically involves multiple stages:
  - 1) entity discovery,
  - 2) coreference resolution,
  - and 3) relation extraction





# Automating Knowledge Graph Construction with LLMs and KGs

## Our solution



**Source:** From Instructions to ODRL Usage Policies: An Ontology Guided

Approach



## 2. Retrieval Augmented Generation enhanced by KGs





## 2. Retrieval Augmented Generation enhanced by KGs

### Motivation

1. Query Augmentation: Addition of context to a query before it performs a retrieval from the vector database
2. Document hierarchies: Creation of document hierarchies and rules for navigating chunks within a vector database
3. **Knowledge Graph Queries by transforming Text into SPARQL**
4. Answer Augmentation: Addition of context based on initially generated query from vector database
5. Answer Rules: Elimination and repetition of results based on rules set in KG

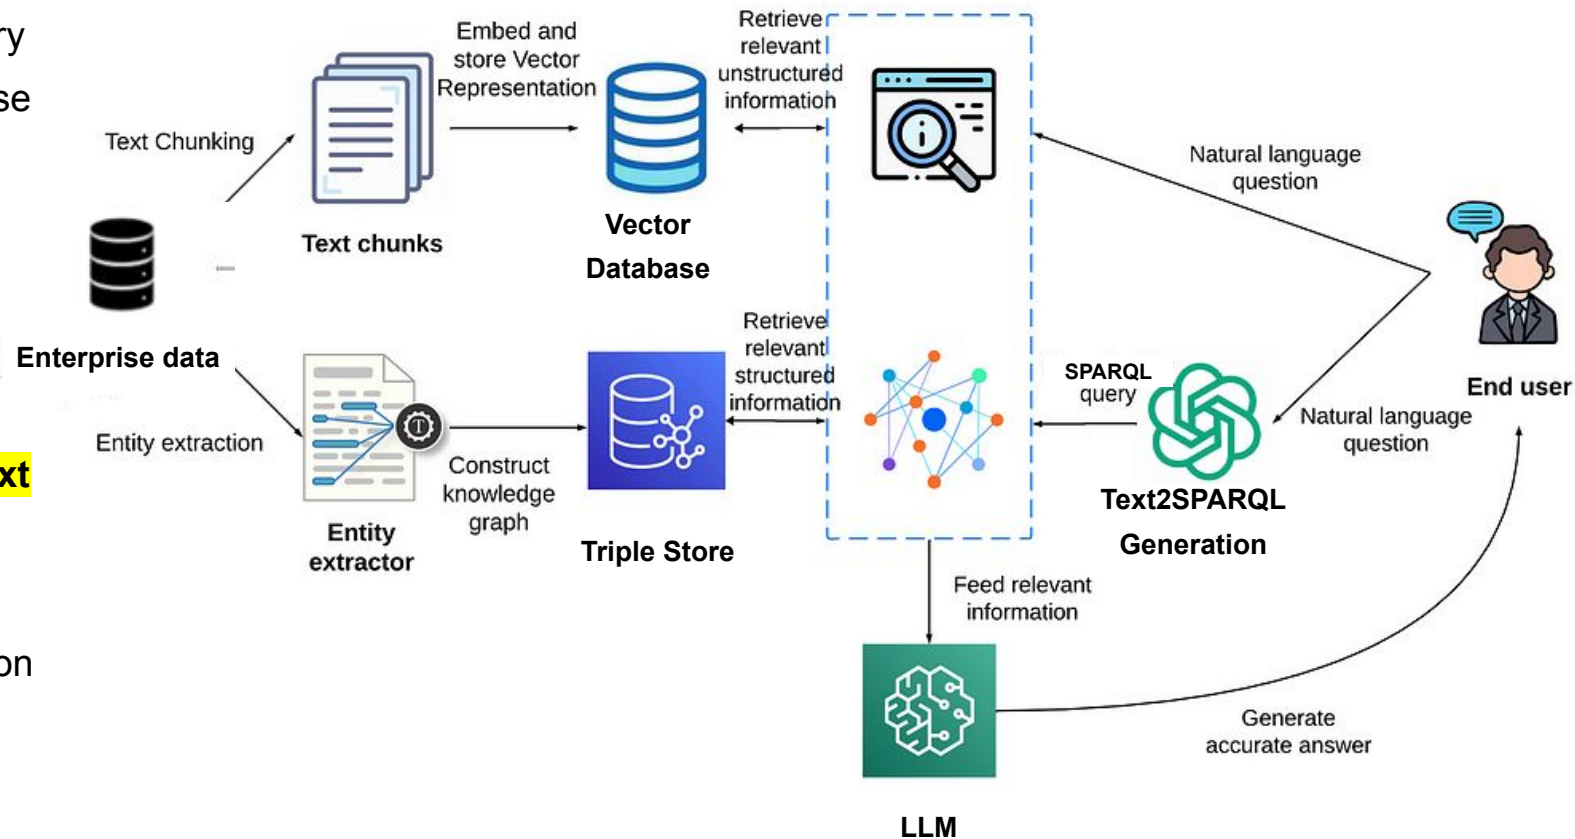


Image quickly adapted from the internet

# Text to SPARQL query

Just one example of the many options being develop nowadays

**Motivation:** Formal queries accessing knowledge graphs are challenging for non-experts due to the need to grasp query syntax and entity relationships.

**Method:** Leverage InContext Learning ability of LLMs to for constructing the query.

1. Ground to KG: Find the associated entities and relations present in Natural Language (NL) Query
2. Prompting: Use LLM to construct an intermediate query using the entities and NL Query.
3. Self-Correcting: Leveraging LLM hallucinations

**Source:** Language Models as Controlled Natural Language Semantic Parsers for Knowledge Graph Question Answering  
enhances retrieval of relevant relations, particularly benefiting cases with unseen relations.

Naturality



NL Query:

Who was the cast member of Titanic and born in Los Angeles?

SQUALL Query:

Which <human> was a <cast member> of <Titanic> and has <place of birth> <Los Angeles> ?

Sparklis Query:

give me every human  
that is the cast member of Titanic  
and that has as a place of birth Los Angeles

SPARQL Query:

```
SELECT DISTINCT ?Q5_1
WHERE { ?Q5_1 wdt:P31 wd:Q5 .
        wd:Q44578 p:P161 [ ps:P161 ?Q5_1 ] .
        ?Q5_1 p:P19 [ ps:P19 wd:Q65 ] . }
```

# 3. Fine-tuning LLMs with Knowledge Graphs

# 3. Fine-tuning LLMs with Knowledge Graphs

## Motivation

### LLMs drawbacks:

1. LLMs have been criticized for their lack of factual knowledge. Specifically, LLMs memorize facts and knowledge from the training dataset.
2. LLMs hallucinate. LLMs might wrongly claim "Einstein discovered gravity in 1687" instead of acknowledging Isaac Newton's gravitational theory.
3. LLMs lack interpretability, encoding knowledge implicitly, making validation difficult.

### Fine-tuning Advantages:

1. By anchoring LLMs to a knowledge source, they gain a solid foundation that effectively mitigates the previously above-mentioned challenges..
2. It allows for seamless enterprise data integrations with LLMs, KG allows for ontological reasoning by integrating factual data with formally specified business awareness (ontology)
3. "Chain-of-Reasoning": Figure here illustrates the complexity of queries, emphasizing the essential utilization of graph structure knowledge for multi-step reasoning.

### References:

- GLaM: Fine-Tuning Large Language Models for Domain Knowledge Graph Alignment via Neighborhood Partitioning and Generative Subgraph Encoding
- Fine-tuning Large Enterprise Language Models via Ontological Reasoning



**Question:** Alex is a new associate professor at MIT working on AI for medical informatics. He plans to write a grant application for discovering relationship between genetic biomarkers and long-term cardiovascular disease. Recommend a group of potential collaborators for the proposal.



# 4. Automating knowledge intensive work using Agents empowered by LLMs and KGs

# 4. Automating knowledge intensive work using Agents

## Motivation

### Knowledge intensive work

- Extensive and important
- Weakly structured
- Knowledge often implicit or in unstructured text
- Personnel-intensive
- Error prone
- Time-consuming training of new employees

### Challenges

- Increasing shortage of skilled workers
- Low efficiency, high costs
- Classical process automation can only automate a fraction

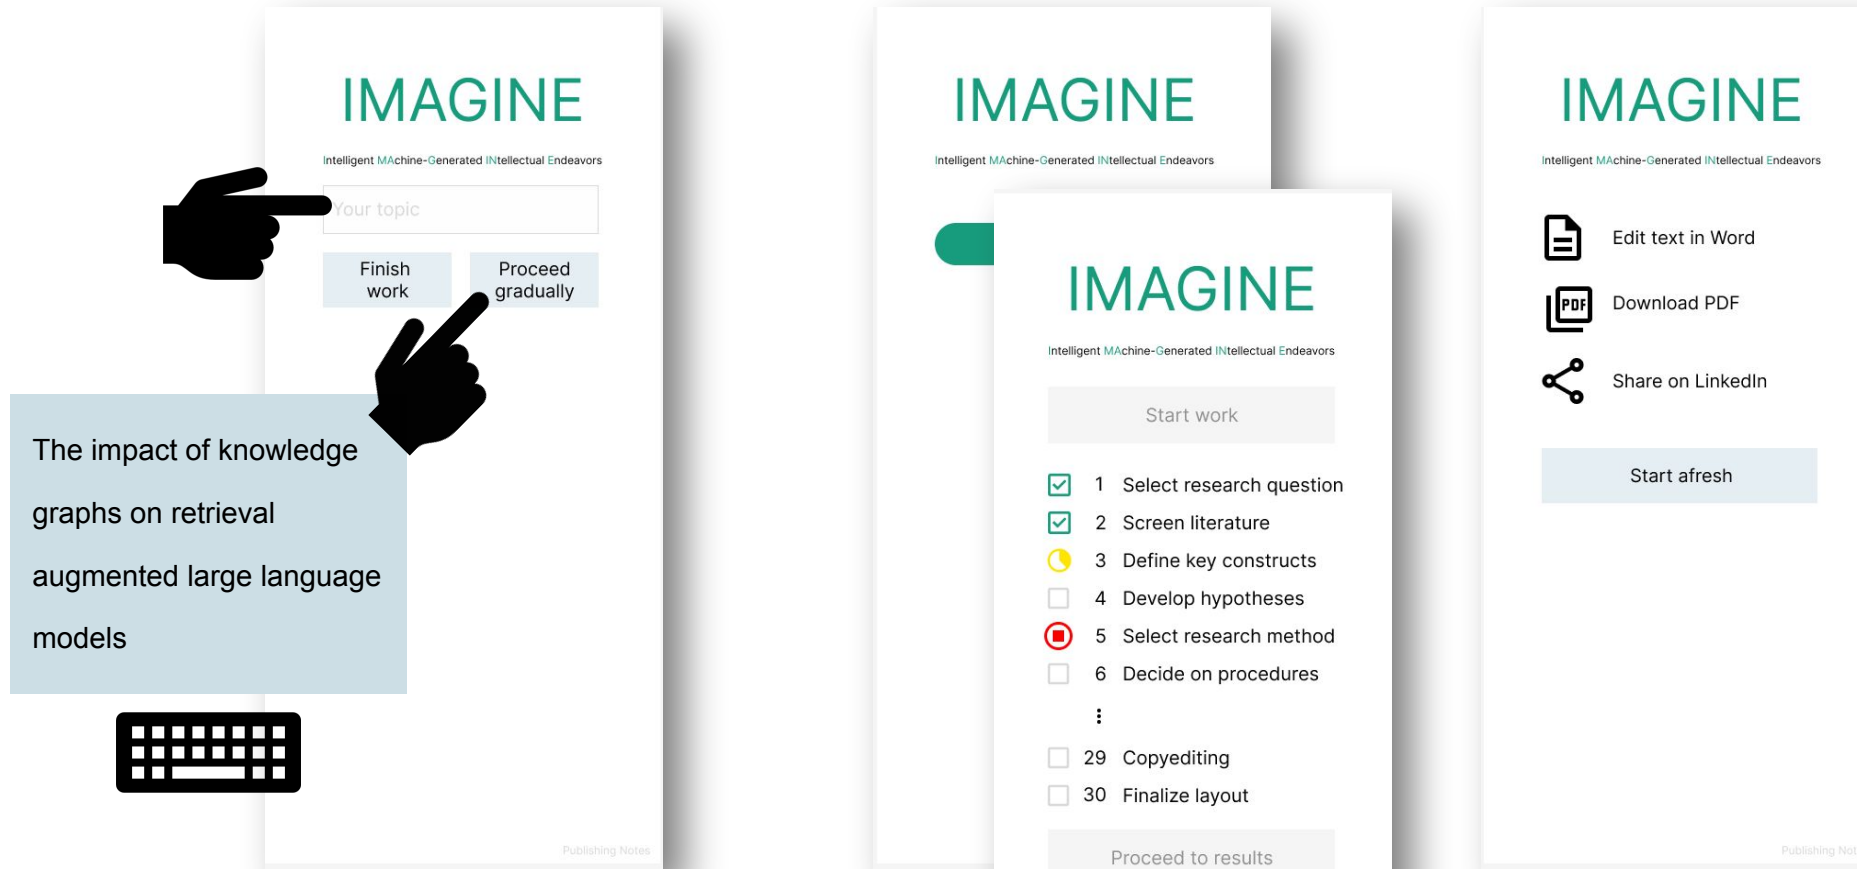


**Beyond Labor: Fully automated knowledge work**

We know about and will consider the ethical and social implications.

# 4. Automating knowledge intensive work using Agents

## A simple concrete example – Writing a scientific paper



# 4. Automating knowledge intensive work using Agents

## Motivation

### Autonomous Agents:

- Autonomous agents are computational systems that inhabit some complex dynamic environment, sense and act autonomously in this environment, and by doing so realize a set of goals or tasks for which they are designed : Maes (1995)

### Why LLM based agents:

- What was X corporation's total revenue for FY 2022? (answered by RAG)

### Knowledge Graphs and LLM agents : Empower at every step

- What were the three takeaways from the Q2 earnings call from FY 23? Focus on the technological moats that the company is building.
  - Profile: identifies the role of the agent. Storing as KG enables to model complex modules like personalities, relationships with other agents etc
  - Requires multiple steps: reasoning , planning, memory, querying.
- Memory and Planning: Enrich long term information and planning strategy by providing a structured approach utilizing KG. (e.g. decision tree)

### References:

- A Survey on Large Language Model based Autonomous Agents
- Introducing LLM Agents | NVIDIA Technical Blog

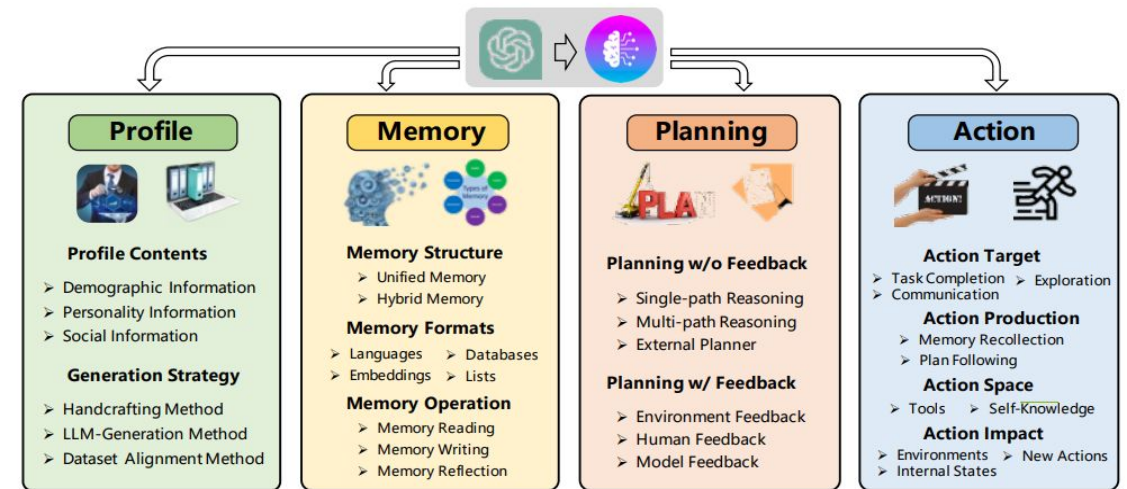
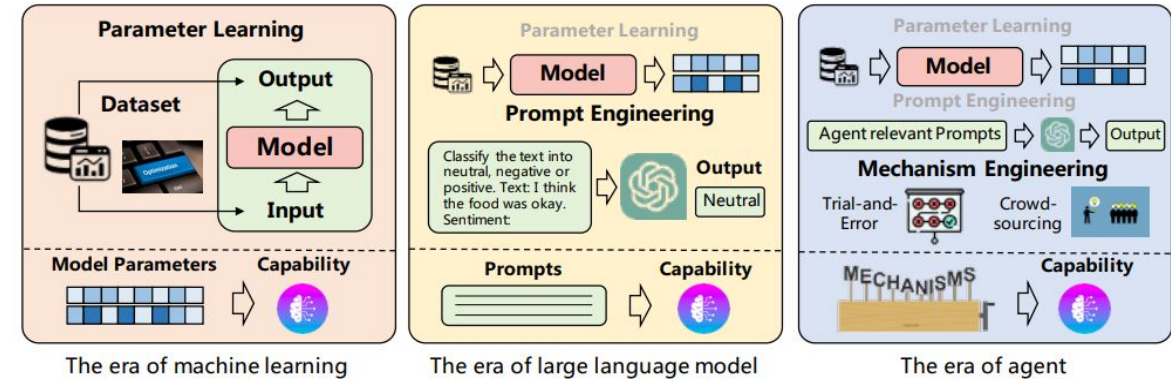


Fig. 2 A unified framework for the architecture design of LLM-based autonomous agent.



# Demo

# BAföG-Buddy

Your conversation assistant to understand and calculate your BAföG

Specialized Chatbot for student grant BAföG (initial advice, currently working in Telegram)

Value added:

- No wrong answers (Wohngeld)
- Updated knowledge (160 €)
- Correct citation (in future link to PDF possible)

BAföG (Bundesausbildungsförderungsgesetz) ist in Deutschland eine staatliche Unterstützung für Studierende, um ihre Ausbildungskosten zu decken. Die Höhe des BAföG hängt von verschiedenen Faktoren ab, einschließlich des eigenen Einkommens und Vermögens sowie dem Einkommen deiner Eltern, falls du unter 30 Jahre alt bist.

Die Tatsache, dass du ein Kind hast, kann sich auf deine BAföG-Berechtigung auswirken, aber es gibt keine allgemeine Regel, die besagt, dass du automatisch mehr BAföG bekommst, nur weil du ein Kind hast. Allerdings gibt es bestimmte Regelungen, die Einfluss auf die Höhe deines BAföG haben könnten:

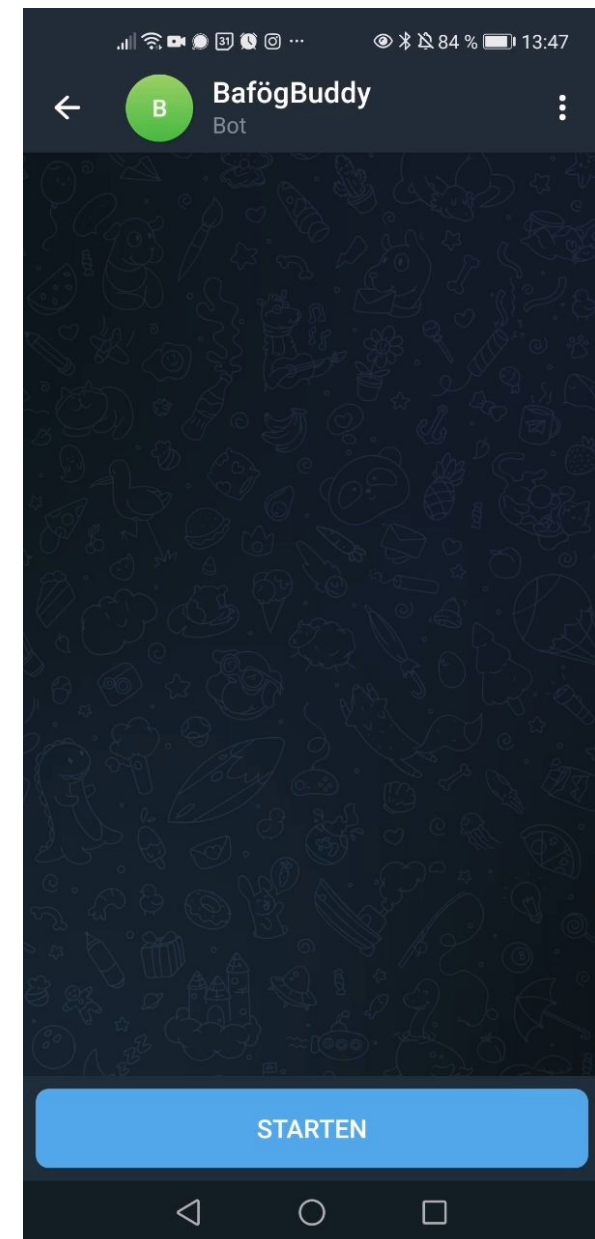
1. **Eigenes Einkommen:** Wenn du als alleinerziehender Elternteil mit einem eigenen Einkommen studierst, kann dies sich positiv auf deine BAföG-Berechtigung auswirken.
2. **Wohngeld:** Als alleinerziehende Person könntest du unter Umständen zusätzlich Anspruch auf Wohngeld haben, das deine Wohnkosten unterstützt.
3. **Elternunabhängiges BAföG:** In einigen Fällen könntest du unabhängig von deinen Eltern Anspruch auf elternunabhängiges BAföG haben, wenn du bestimmte Voraussetzungen erfüllst.

Es ist wichtig zu beachten, dass die genauen Details deiner Situation und die Höhe deiner BAföG-Zahlungen von verschiedenen Faktoren abhängen. Es empfiehlt sich, die BAföG-Ämter oder entsprechende Beratungsstellen an deiner Hochschule zu kontaktieren, um eine genaue Einschätzung deiner individuellen Situation zu erhalten.



Ich habe ein Kind. Bekomme ich mehr Bafög? 13:48 ✓✓

Ja, Studierende mit Kind erhalten einen Zuschlag in Höhe von 160 EUR pro Monat für jedes Kind unter 14 Jahren, das mit im Haushalt lebt. Der Zuschlag wird jedoch nur einem Elternteil gewährt. 13:48



# Kontakt

---

Abhishek Nadgeri  
Knowledge-Enhanced Large Language Models (KELLM)  
am DSAI, Fraunhofer FIT  
[abhishek.nadgeri@fit.fraunhofer.de](mailto:abhishek.nadgeri@fit.fraunhofer.de)

Dr. Diego Collarana  
Gruppenleiter  
Knowledge-Enhanced Large Language Models (KELLM)  
am DSAI, Fraunhofer FIT  
[Diego.Collarana.Vargas@fit.fraunhofer.de](mailto:Diego.Collarana.Vargas@fit.fraunhofer.de)

Thank you!  
Danke!  
Gracias!  
Merci!



# However, LLMs fail in unpredictable ways in knowledge-intensive tasks

## In use-cases when factually correct answers must be guaranteed, e.g., the AI ecosystem

Which license does the FLAN-T5 model have?

Which is the latest LLM model developed by Meta AI?

How many LLM models did Bloomberg create?

Knowledge is **somehow stored** in the billions of parameters of the LLM

OpenAI

### Hallucination problem

... I don't have ... about the FLAN-T5 model's licensing. ... note that **OpenAI**, the organization behind .. has not publicly disclosed the specific license for ...

### Updateability problem

...knowledge is up until Sep. 2021, the latest model developed by Meta AI would be **Galactica...**

### Provenance

**problem** unclear how many LLMs Bloomberg has created. **Bloomberg has not publicly disclosed the models** they have built ...

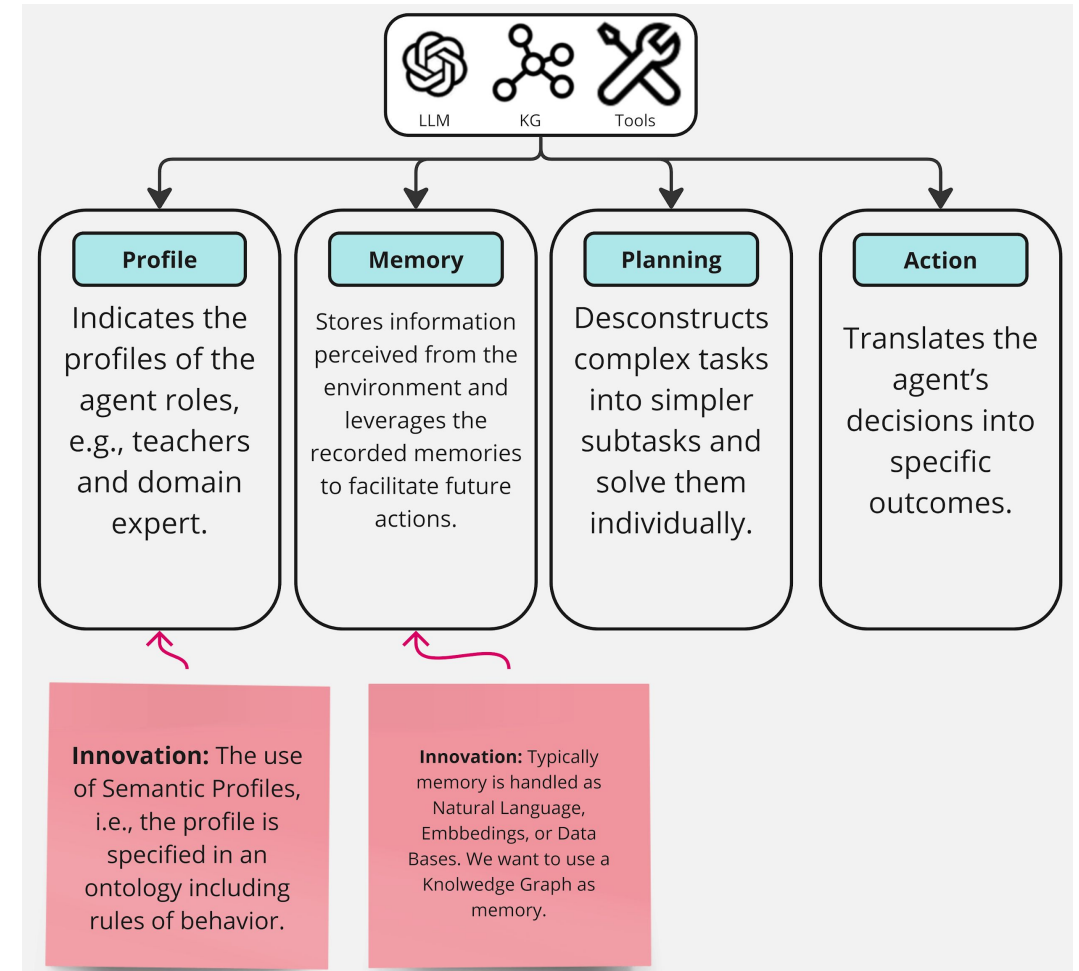
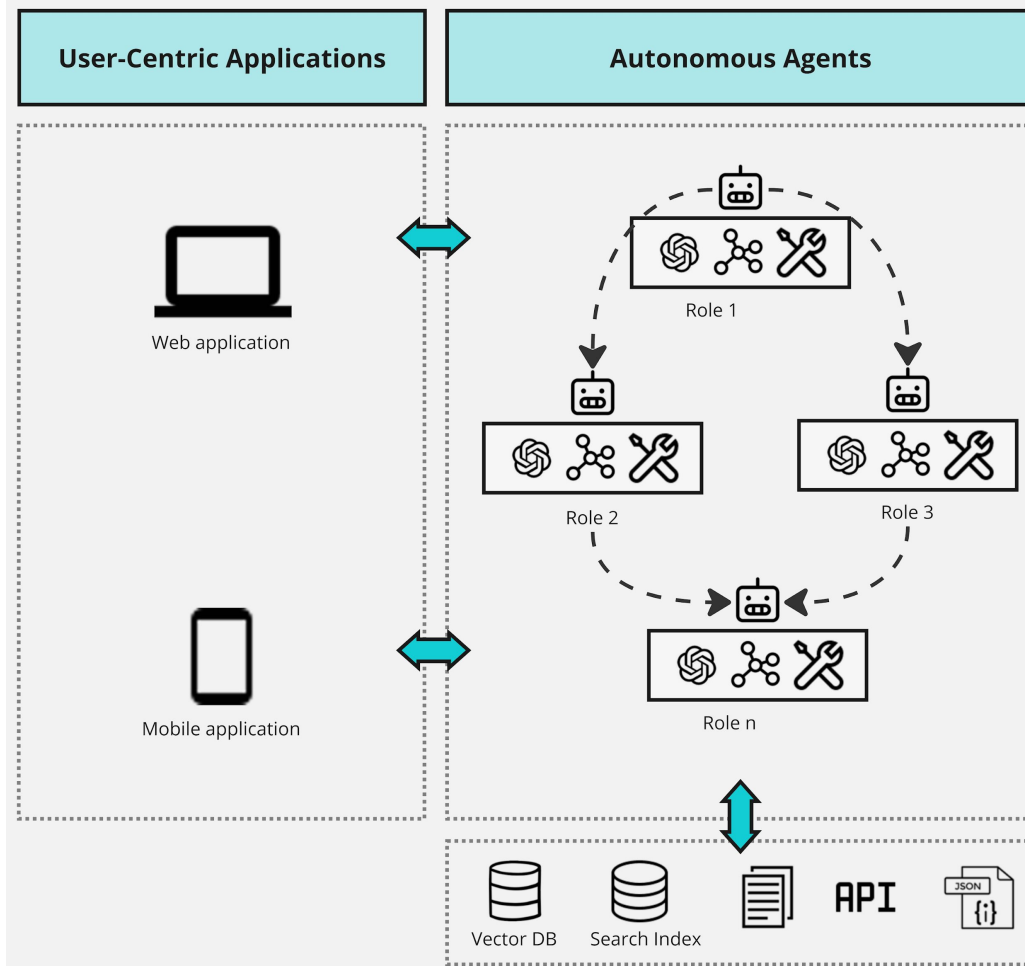
Answer generated by GPT-3.5 Turbo

16K



# 4. Automating knowledge intensive work using Agents

## An agent collaboration view



# Why knowledge graphs?

## Impact radar for Generative AI from Gartner

- **To leverage LLMs**, they need to connect to clean, well-defined data sources.
- **KGs bring the last crucial layer** – data integration. Think of data integration as the crucial last mile, preparing clean and accurate data for LLMs.
- GenAI models are being used in conjunction with KGs to **deliver trusted and verified facts** to their outputs, as well as provide rules to contain the model.
- **The range for knowledge graphs is Now**,
  - KGs allow unstructured content (text files, pdfs, etc.) to be natively connected to structured data.
  - Ontologies provide a semantic layer that natively expresses the relationships between data concepts.

