

Augmenting Knowledge-Based Conversational Search Systems with Large Language Models

Manuel Klettner

12.06.2023, Master's Thesis Kick-Off Presentation

Chair of Software Engineering for Business Information Systems (sebis)
Department of Computer Science
School of Computation, Information and Technology (CIT)
Technical University of Munich (TUM)
www.matthes.in.tum.de

1. Background
2. Motivation
3. Scope
4. Goals
5. Research Questions
6. Initial Results
7. Next Steps & Timeline
8. Questions

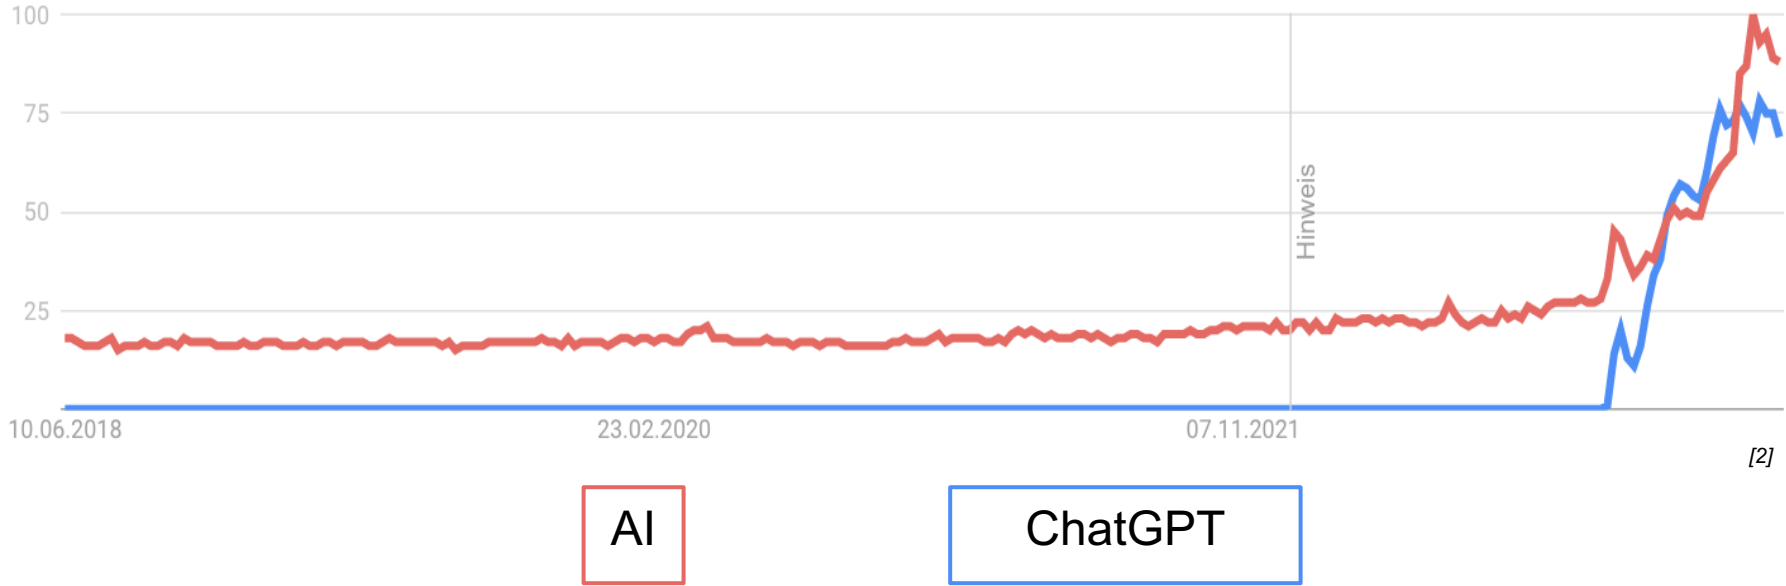
 [https://www.reuters.com > technology > chatgpt-sets-record-fastest-growing-user-base-analyst-n...](https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note)

ChatGPT sets record for fastest-growing user base - analyst note

Feb 2, 2023 · Feb 1 (Reuters) - **ChatGPT**, the popular chatbot from OpenAI, is estimated to have reached **100 million** monthly active users in January, just two months after launch, making it the...

[1]


Google Trends



AI


ChatGPT


Hallucinations

 CNN.com
<https://www.cnn.com> › tech › google-ai-bard-demo-error ⋮
Bard: Google's AI chatbot had an inaccurate response ...
9 Feb 2023 — Shares in Google's parent company Alphabet fell 7.7% Wednesday, wiping \$100 billion off its market value, after the inaccurate response from ...

[3]

Outdated information

 Summarise the most important news of today

 I apologize, but as an AI language model, I don't have real-time data or access to the internet to provide you with the most up-to-date news. My training only goes up until September 2021, so any news or events that have occurred since then are beyond my knowledge. I recommend checking reliable news sources or conducting an internet search to find the latest news updates.

Reliability

MA

How many countries start with the letter v?



There are currently two countries whose names start with the letter "V":

1. Vanuatu
2. Vatican City

Please note that this information is accurate as of my knowledge cutoff in September 2021. There may have been changes or new countries established since then.

MA

What about Vietnam and Venezuela?

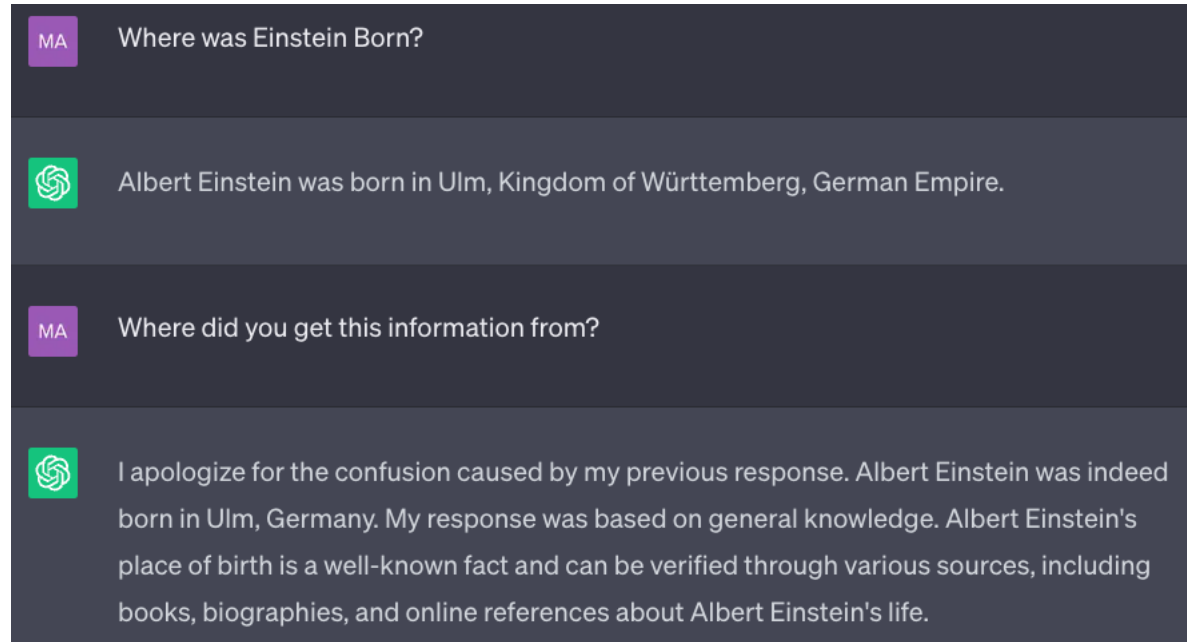


Apologies for the oversight. You are correct, and I apologize for the omission in my previous response. In addition to Vanuatu and Vatican City, there are indeed two more countries whose names start with the letter "V":

3. Vietnam
4. Venezuela

Thank you for pointing out the error, and I appreciate your understanding.

Data Lineage




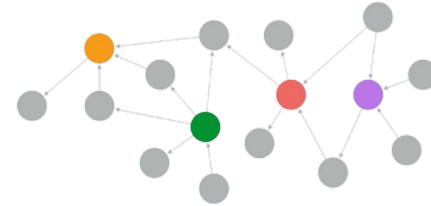
MA Where was Einstein Born?

Albert Einstein was born in Ulm, Kingdom of Württemberg, German Empire.

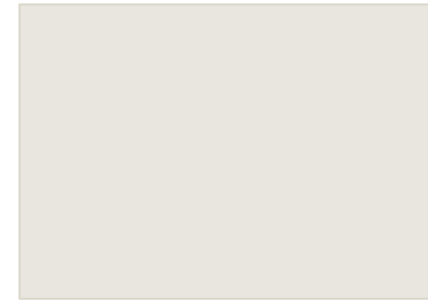
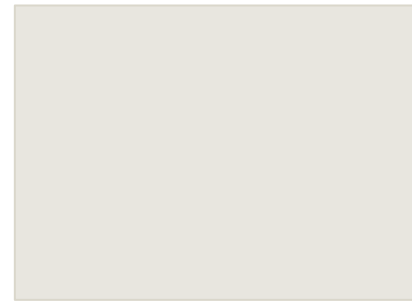
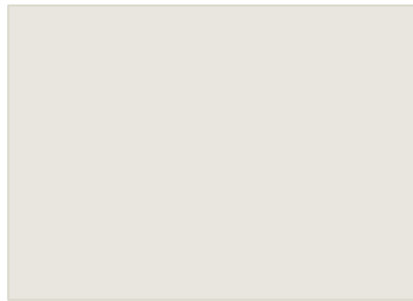
MA Where did you get this information from?

I apologize for the confusion caused by my previous response. Albert Einstein was indeed born in Ulm, Germany. My response was based on general knowledge. Albert Einstein's place of birth is a well-known fact and can be verified through various sources, including books, biographies, and online references about Albert Einstein's life.

 Knowledge grounding mechanisms are required for truthful responses



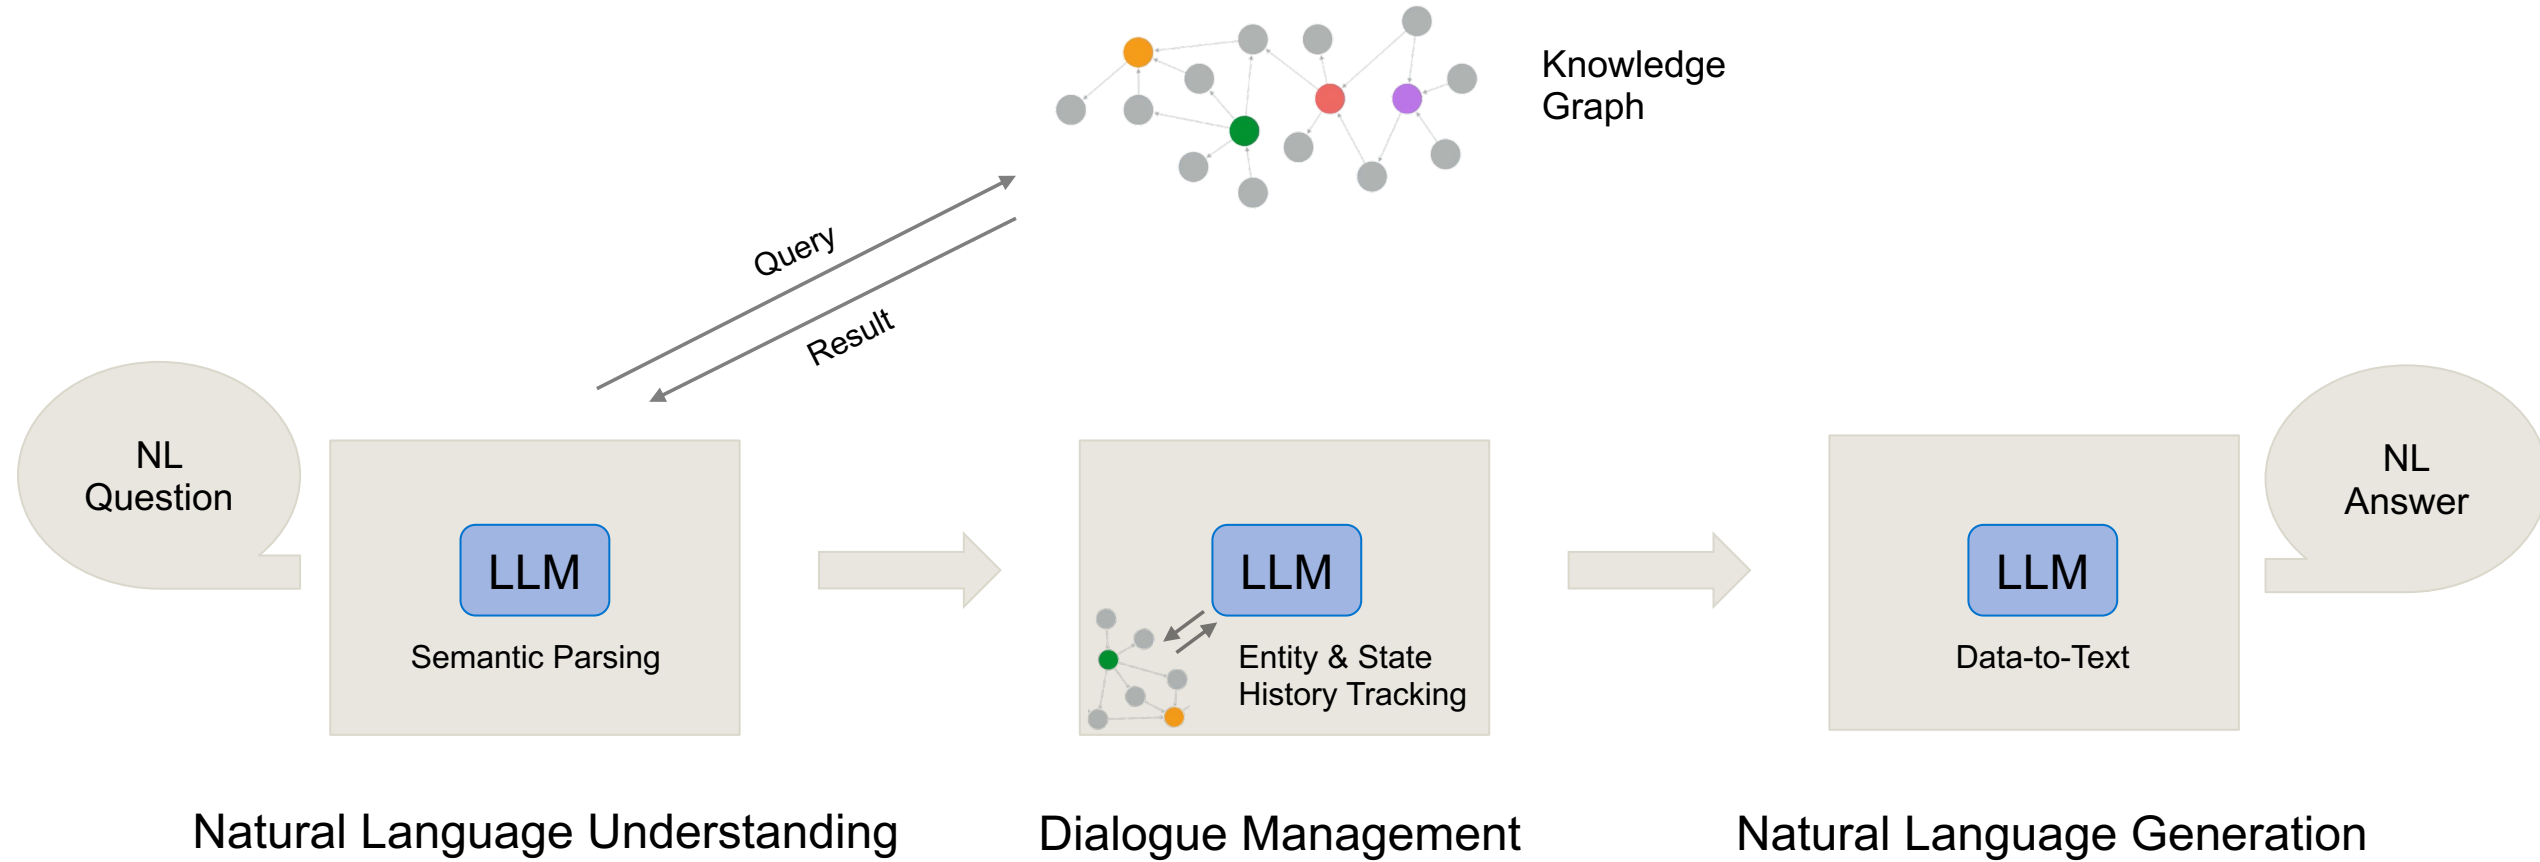
Knowledge Graph

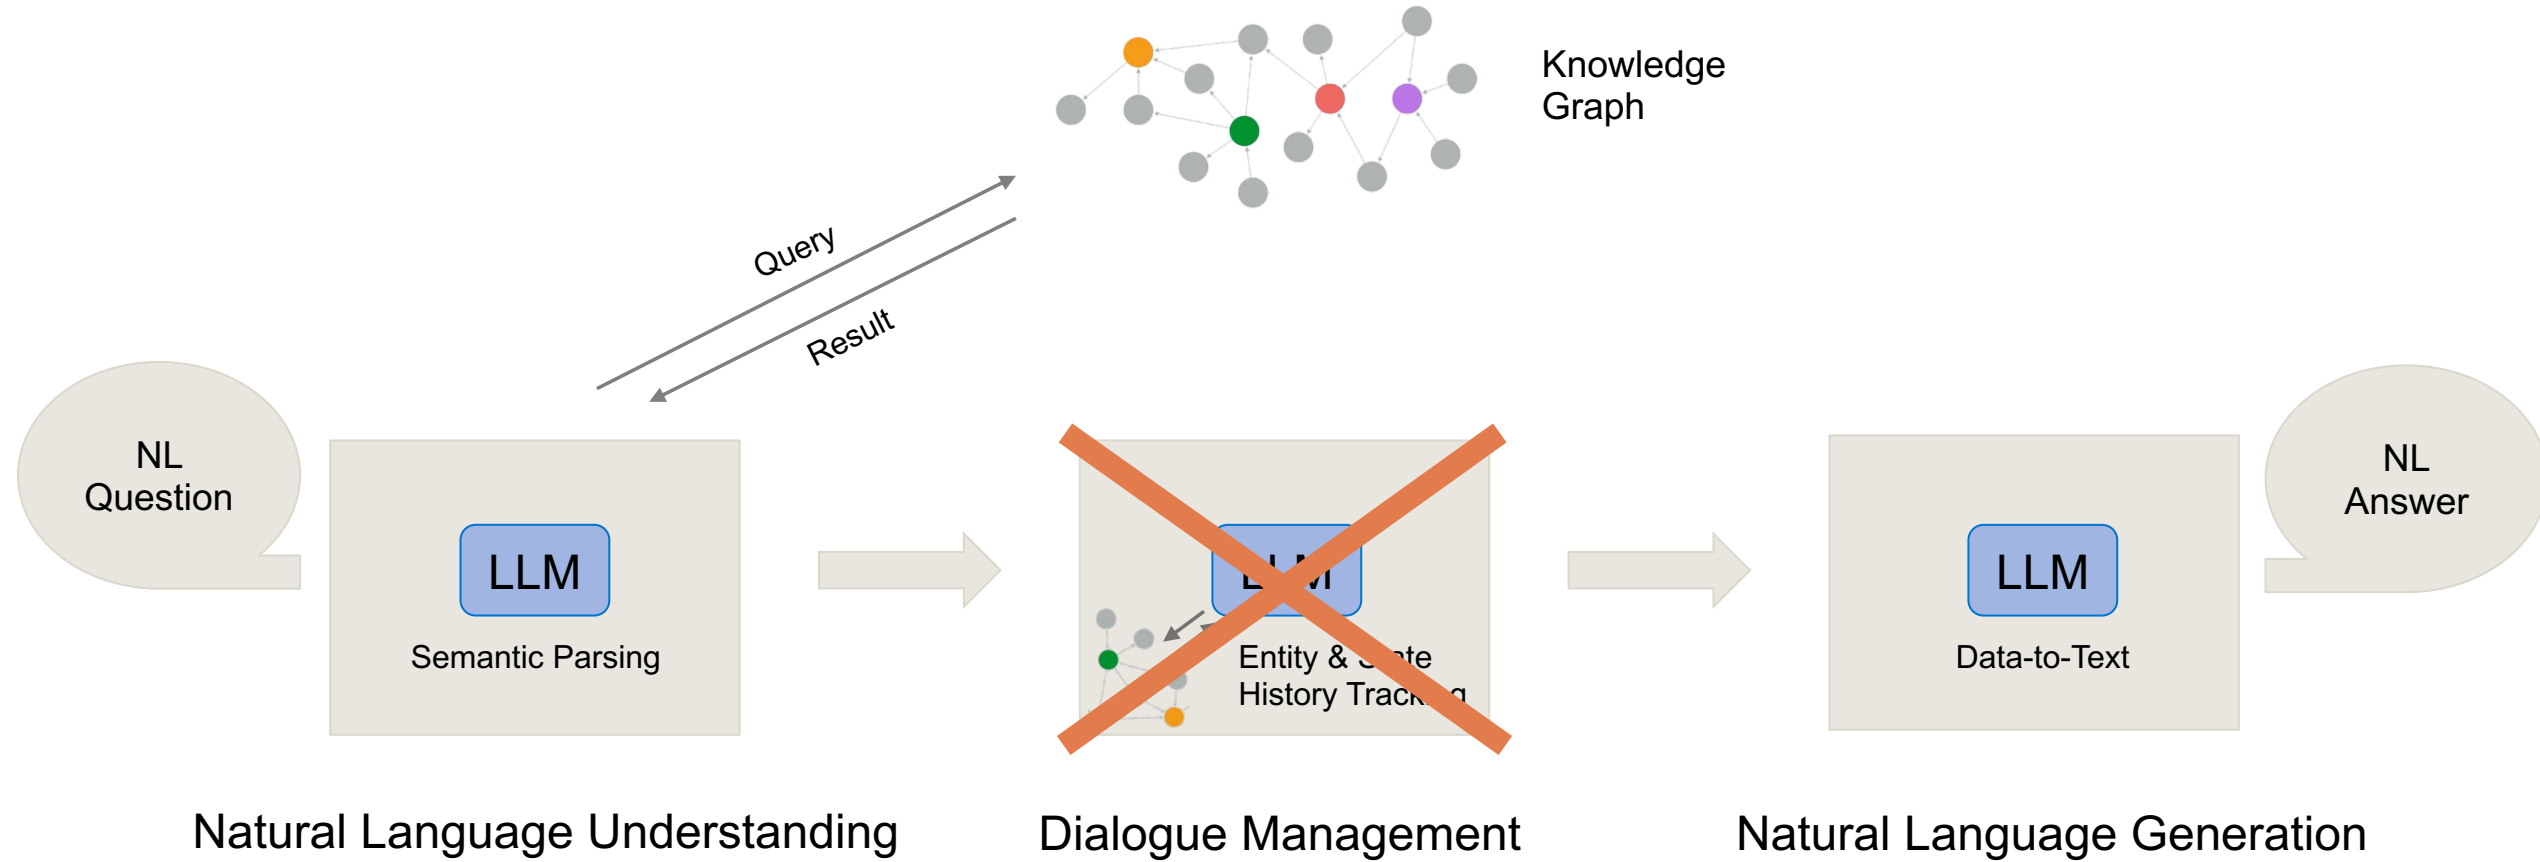


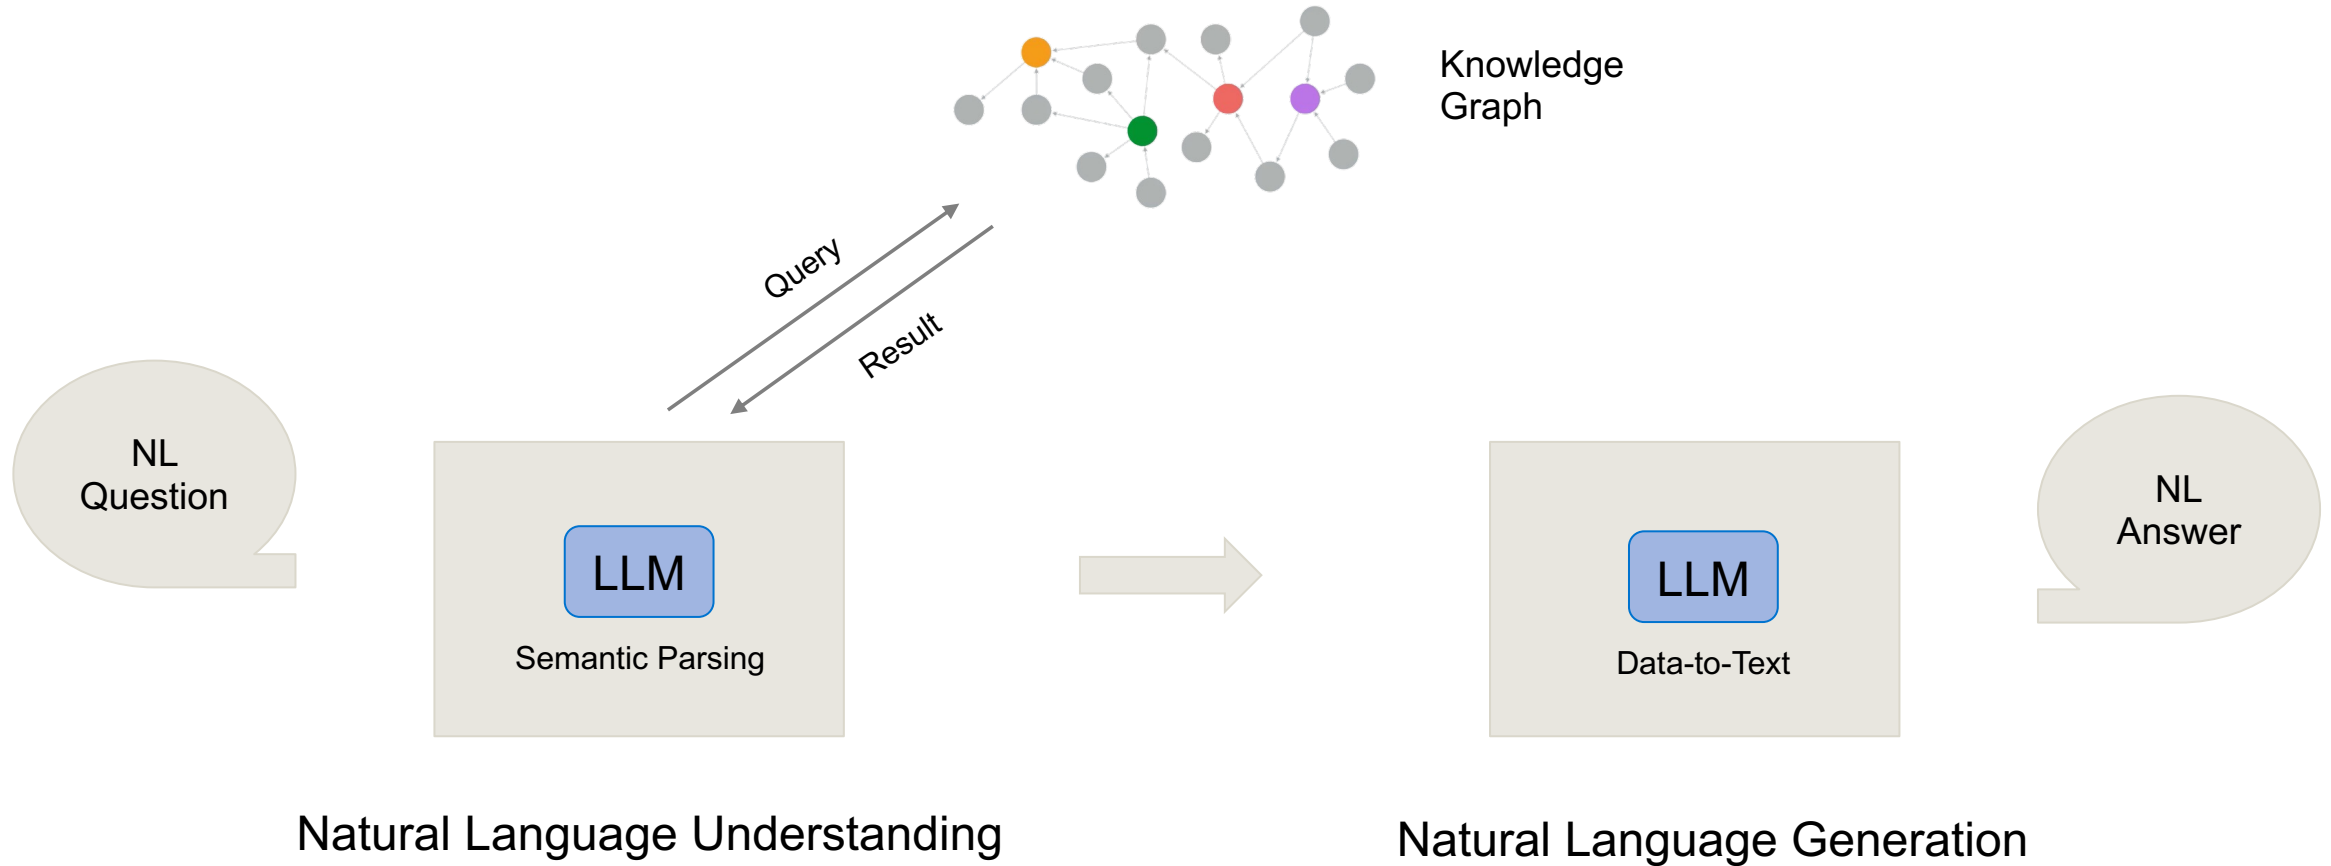
Natural Language Understanding

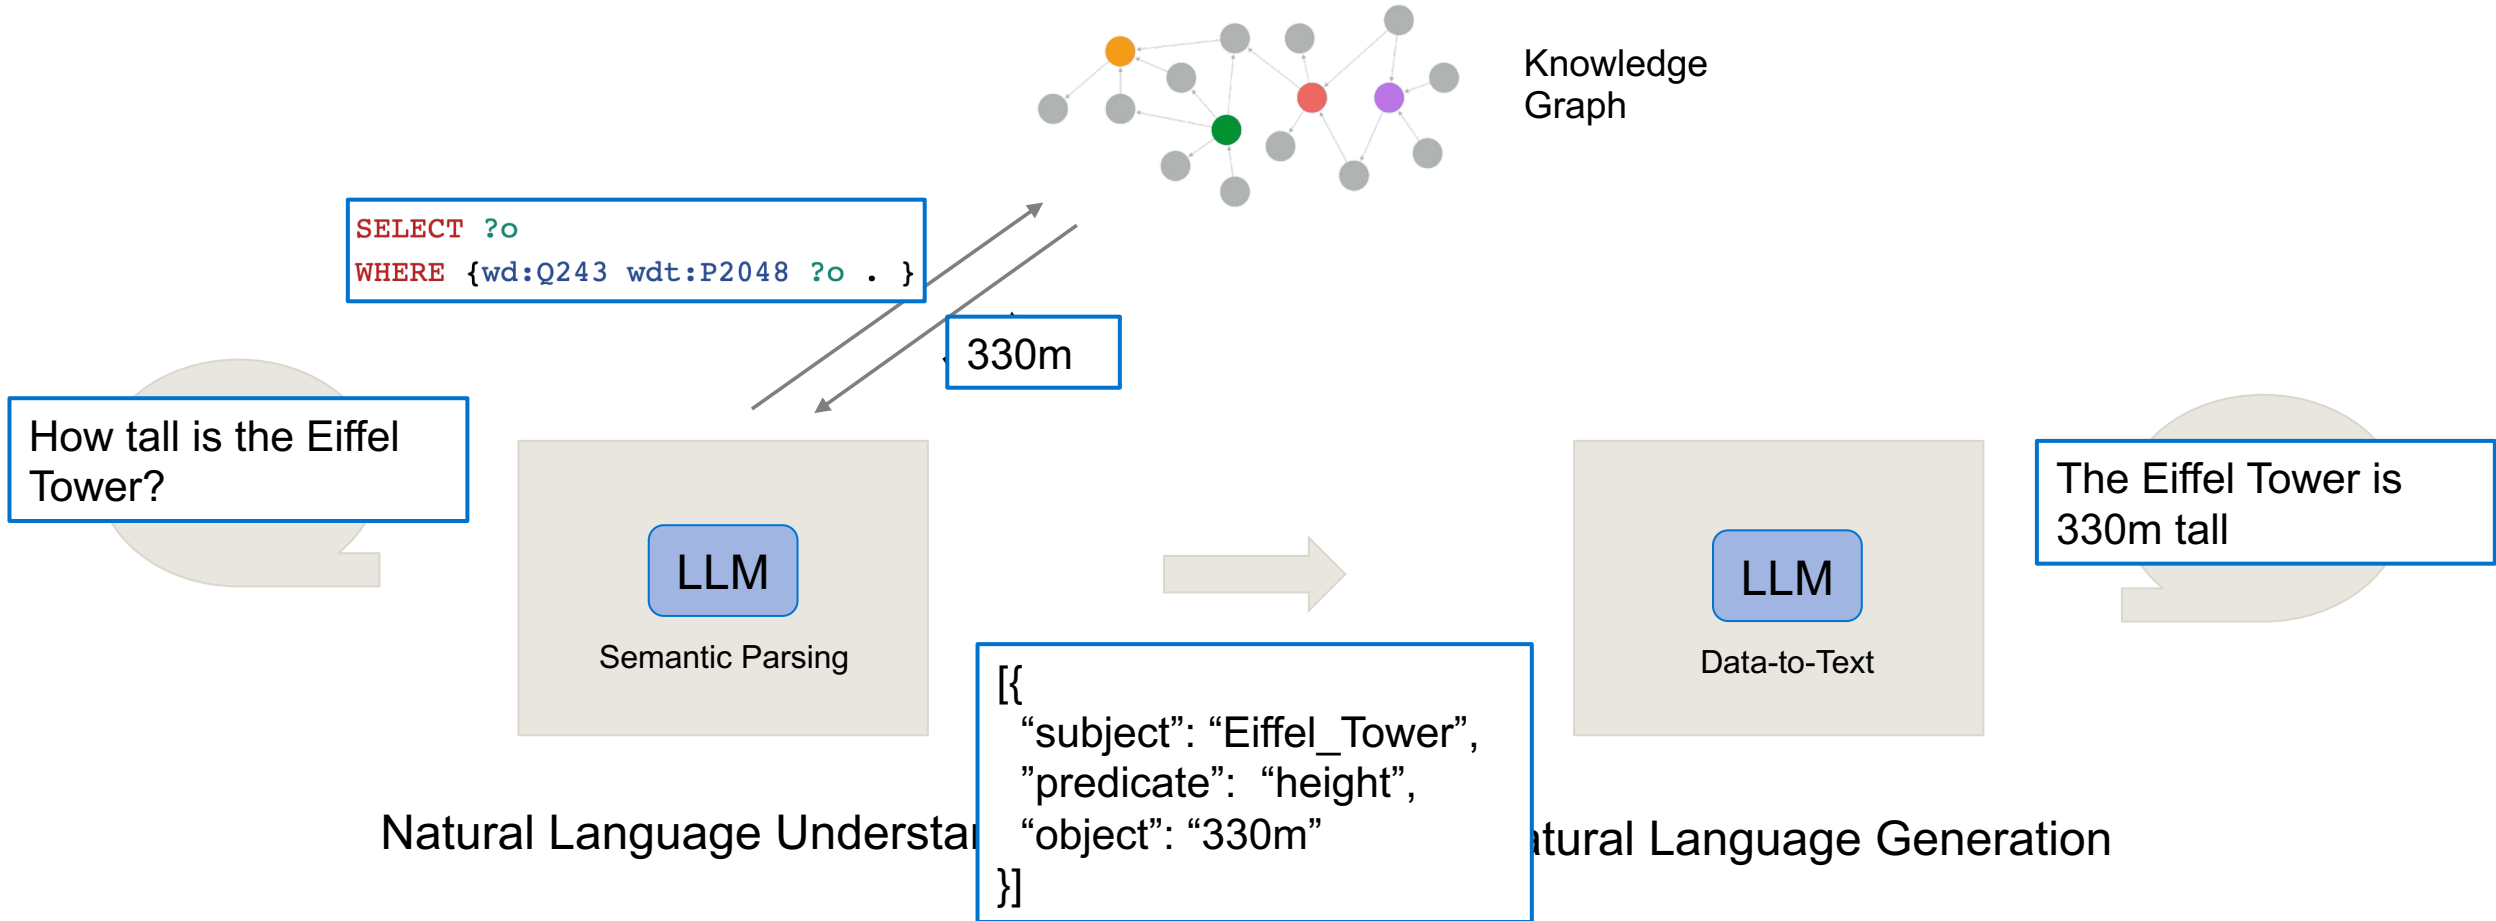
Dialogue Management

Natural Language Generation









- Use LLM to generate SPARQL queries based on user question (Semantic Parsing)
- Transform resulting data into natural language answer (Text Generation)
- Gain insights in performance of LLMs with different amounts of parameters and pre/fine-tuning procedures
- Acquire understanding of the influence different prompting techniques have on the before mentioned tasks

1.

Which previous studies have investigated using Large Language Models for the tasks of semantic parsing and text generation?

2.

What selection of Large Language Models and Prompting techniques are suitable for a comparative analysis of the considered tasks?

3.

How can we systematically evaluate the performance of the tested Large Language Models?

Zero-Shot

MA Classify the following text into "positive" or "negative" and only return the category.
Text: "Coding is fun."

 Positive

Few-Shot

MA Classify the following text into "positive" or "negative" and only return the category.

Text: "Coding is fun."
Category: positive


Text: "Housing prices in Munich are too expensive"
Category: negative

Text: "TUM is an excellent university"
Category:

 positive

Chain-of-Thought (CoT)

MA Lucy has 5 apples. She receives two more from Bob, eats one of them and gives half of the apples that are left to Anna. How many apples are left for Lucy?
Let's think step by step,

 Step 1: Lucy has 5 apples.
Step 2: She receives two more from Bob. So now Lucy has $5 + 2 = 7$ apples.

Other

- System instruction
- Repeating key instructions
- Least-to-Most Prompting
- Mock-Acknowledgement
- Naming the Agent
- ...

Zero-Shot

MA Classify the following text into "positive" or "negative" and only return the category.
Text: "Coding is fun."

 Positive

Few-Shot

MA Classify the following text into "positive" or "negative" and only return the category.

Text: "Coding is fun."
Category: positive


Text: "Housing prices in Munich are too expensive"
Category: negative

Text: "TUM is an excellent university"
Category:

 positive

Chain-of-Thought (CoT)

MA Lucy has 5 apples. She receives two more from Bob, eats one of them and gives half of the apples that are left to Anna. How many apples are left for Lucy?
Let's think step by step,

 Step 1: Lucy has 5 apples.
Step 2: She receives two more from Bob. So now Lucy has $5 + 2 = 7$ apples.

Other

- System instruction
- Repeating key instructions
- Least-to-Most Prompting
- Mock-Acknowledgement
- Naming the Agent
- ...

LLMs	LLaMA	Vicuna	LLaMA-LoRA	GPT-3.5-Turbo	Text-Davinci-003
Access	Open Source*	Open Source*	Open Source*	Closed Source	Closed Source
Parameters	7B	7B	7B	~150B**	175B
Training	Pre-Trained	Fine-Tuned (Chat)	Fine-Tuned (Our Task)	Reinforcement Learning from Human Feedback (RLHF)	Fine-Tuned (Instruction-Following)

*Non-commercial license

**Not officially acknowledged

[4], [5], [6], [7]

WebNLG

- Category: Athlete
- Tripletset:

```
[{'object': 'FC_Karpaty_Lviv', 'property': 'club', 'subject': 'Aleksandre_Guruli'},  
{'object': 'FC_Dinamo_Batumi', 'property': 'club', 'subject': 'Aleksandre_Guruli'}]
```
- Lexicalization: "Aleksandre Guruli played for FC Karpaty Lviv and FC Dinamo Batumi."

[8]

Vicuna

Prompt:

Generate a concise **textual description** for the given set of triples. Ensure that the generated output only includes the provided information from the triples.

Result:

The given set of triples describes two objects, FC_Karpaty_Lviv and FC_Dinamo_Batumi, both of which have the property "club" and the subject "Aleksandre_Guruli".

LLaMA

```
def get_few_shot_prompt() -> str:
    return """
    Tripletset: [{"object": "Community_of_Madrid", "property": "isPartOf", "subject": "Paracuellos_de_Jarama"}]
    Output: Paracuellos de Jarama is part of the community of Madrid.

    Tripletset: [{"object": "Ska_punk", "property": "genre", "subject": "Aaron_Bertram"}]
    Output: Aaron Bertram performs ska punk music.

    Tripletset: [{'object': 'Soul_music', 'property': 'genre', 'subject': 'Aleksandra_Kovač'}]
    Output:
    """
```

```
def get_few_shot_prompt_no_newline() -> str:
    return """
    Tripletset: [{"object": "Community_of_Madrid", "property": "isPartOf", "subject": "Paracuellos_de_Jarama"}]
    Output: Paracuellos de Jarama is part of the community of Madrid.

    Tripletset: [{"object": "Ska_punk", "property": "genre", "subject": "Aaron_Bertram"}]
    Output: Aaron Bertram performs ska punk music.

    Tripletset: [{'object': 'Soul_music', 'property': 'genre', 'subject': 'Aleksandra_Kovač'}]
    Output: """
```

LLaMA

```
def get_few_shot_prompt() -> str:
    return """
    Tripletset: [{"object": "Community_of_Madrid", "property": "isPartOf", "subject": "Paracuellos_de_Jarama"}]
    Output: Paracuellos de Jarama is part of the community of Madrid.

    Tripletset: [{"object": "Ska_punk", "property": "genre", "subject": "Aaron_Bertram"}]
    Output: Aaron Bertram performs ska punk music.

    Tripletset: [{'object': 'Soul_music', 'property': 'genre', 'subject': 'Aleksandra_Kovač'}]
    Output:
    """
```

```
def get_few_shot_prompt_no_newline() -> str:
    return """
    Tripletset: [{"object": "Community_of_Madrid", "property": "isPartOf", "subject": "Paracuellos_de_Jarama"}]
    Output: Paracuellos de Jarama is part of the community of Madrid.

    Tripletset: [{"object": "Ska_punk", "property": "genre", "subject": "Aaron_Bertram"}]
    Output: Aaron Bertram performs ska punk music.

    Tripletset: [{'object': 'Soul_music', 'property': 'genre', 'subject': 'Aleksandra_Kovač'}]
    Output: """
```

Initial Results - Insights

LLaMA

```
def get_few_shot_prompt() -> str:
    return """
    Tripletset: [{"object": "Ska_punk", "property": "genre", "subject": "Aaron_Bertram"}]
    Output: Paracuellos de Jarama

    Tripletset: [{"object": "Ska_punk", "property": "genre", "subject": "Aaron_Bertram"}]
    Output: Aaron Bertram

    Tripletset: [{"object": "Ska_punk", "property": "genre", "subject": "Aleksandra_Kovač"}]
    Output: """
```

```
"""
def __init__(self, subject, predicate, object):
    self.subject = subject
    self.predicate = predicate
    self.object = object

def __repr__(self):
    return "Tripletset: {0} {1} {"
```

```
def get_few_shot_prompt_no_newline() -> str:
    return """
    Tripletset: [{"object": "Ska_punk", "property": "genre", "subject": "Aaron_Bertram"}]
    Output: Paracuellos de Jarama

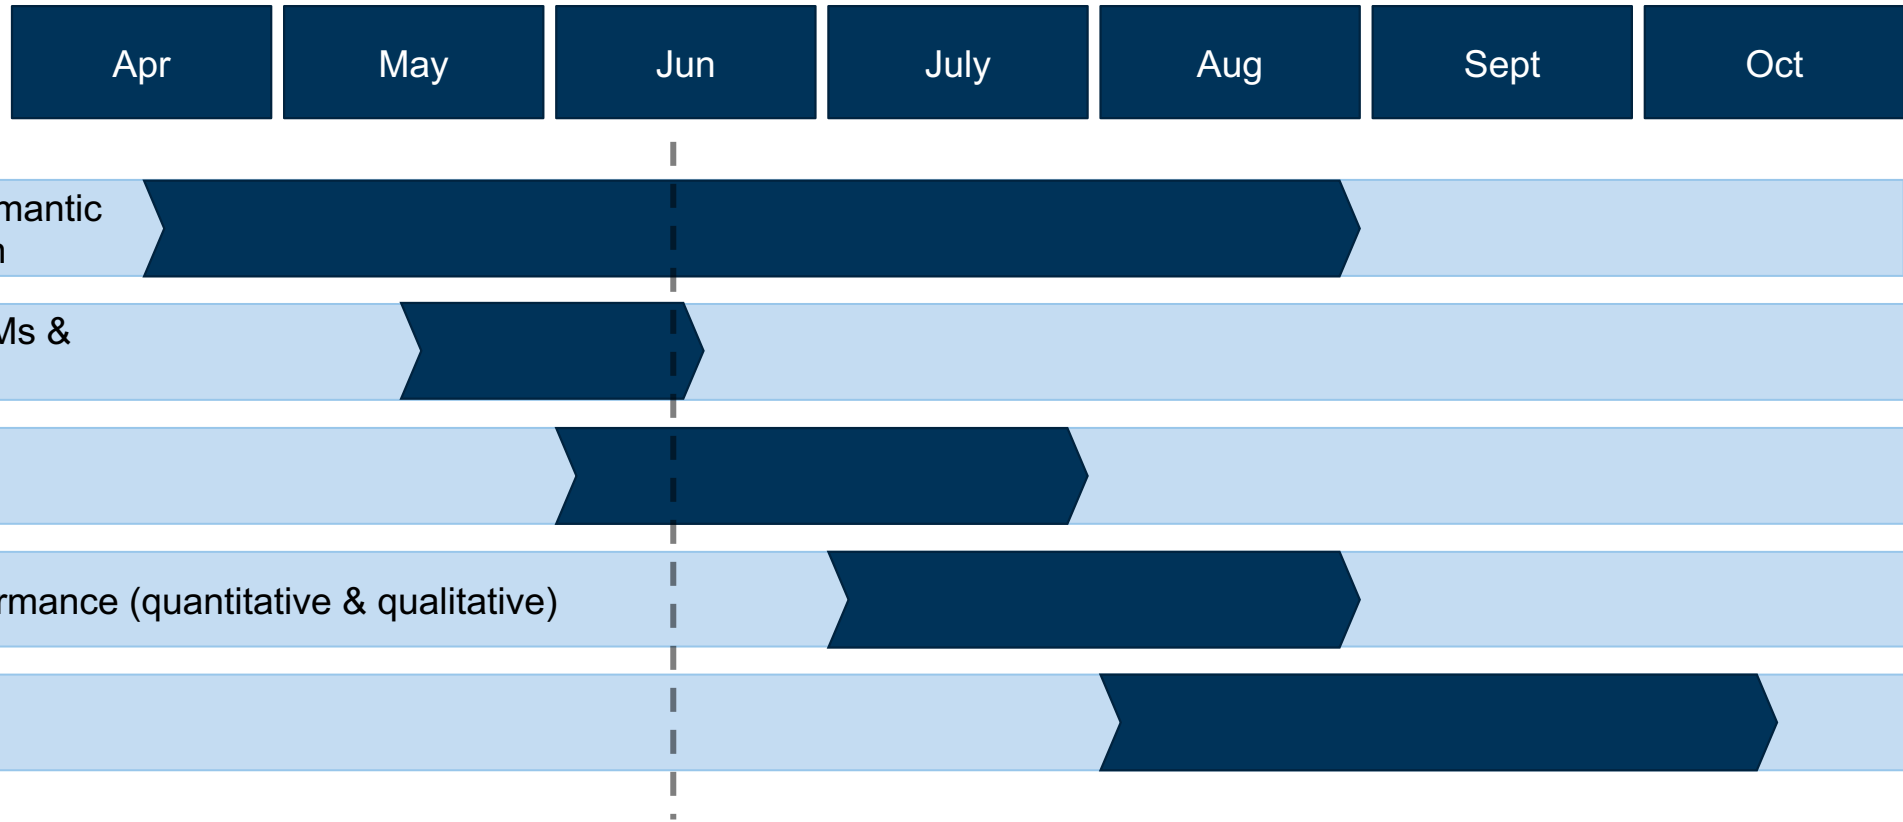
    Tripletset: [{"object": "Ska_punk", "property": "genre", "subject": "Aaron_Bertram"}]
    Output: Aaron Bertram

    Tripletset: [{"object": "Ska_punk", "property": "genre", "subject": "Aleksandra_Kovač"}]
    Output: """
```

```
Aleksandra Kovač performs soul music.

Tripletset: [{"object": "Ska_punk", "property": "genre", "subject": "Aaron_Bertram"}, {"object": "Ska_punk", "property": "genre", "subject": "Aleksandra_Kovač"}]
```

Next Steps & Timeline



- [1] Hu, Krystal: ChatGPT sets record for fastest-growing user base - analyst note
- [2] Google Trends: ChatGPT & AI
- [3] Thorbecke, Catherine: Google shares lose \$100 billion after company's AI chatbot makes an error during demo
- [4] Touvron, H. et al.: LLaMA: Open and Efficient Foundation Language Models
- [5] Chiang, W.-L. et al.: Vicuna: An Open-Source Chatbot Impressing GPT-4 with 90%* ChatGPT Quality
- [6] Hu, E. J.: LoRA: Low-Rank Adaptation of Large Language Models
- [7] OpenAI: Models-Overview
- [8] Castro Ferreira, T. et al.: The 2020 Bilingual, Bi-Directional WebNLG+ Shared Task: Overview and Evaluation Results (WebNLG+ 2020)



Prof. Dr.

Florian Matthes

Technical University of Munich (TUM)
TUM School of CIT
Department of Computer Science (CS)
Chair of Software Engineering for Business
Information Systems (sebis)

Boltzmannstraße 3
85748 Garching bei München

+49.89.289.17132
matthes@in.tum.de
www.matthes.in.tum.de

