

# Comprehensively Materializing Factual LLM Knowledge

Simon Razniewski

Joint work with Yujia Hu,  
Tuan-Phong Nguyen, Shrestha Ghosh



# Myself

- Professor for Knowledge-aware AI, [TU Dresden & ScaDS.AI](#), since 6/2024
- [Max Planck Institute for Informatics](#), Saarbrücken, Germany, 2017-2022
- [FU Bozen-Bolzano](#), Italy, 2011-2016
- Research interests:
  - [Extracting and consolidating knowledge from, for and with texts, KBs, LLMs](#)
  - [Developed methodologies for compiling completeness, negation, and commonsense knowledge](#)



# Outline

- 1. The need for LLM knowledge introspection**
2. The GPTKB approach
3. Making GPTKB a real KB
4. Outlook

# GenAI...

- ...is regularly in the news
  - o3, DeepSeek, Stargate, ...
- ...is becoming an important skill in many professions
  - Significant disruption of knowledge-intensive jobs discussed
  - Already 22% of UK public employees use it [Jonathan Bright et al., Arxiv 2025]
- ...is here to stay
  - 70% of all teenagers use GenAI, especially for homework [Zara Abrams, APA 2025]

# Remarkable weaknesses of GenAI



"Manager" by Stable Diffusion.



"Compassionate manager" by Stable Diffusion.

# Bias in LLMs

In the following sentence, who was late?

The doctor called the nurse because she was late.

The doctor called the nurse because he was late.



ChatGPT-4o

The sentence is ambiguous. "She" could refer to either the doctor or the nurse. Without additional context, it's unclear who was late.

Continue the text:

The doctor phoned the nurse because she was late. He said that ....



ChatGPT-4o

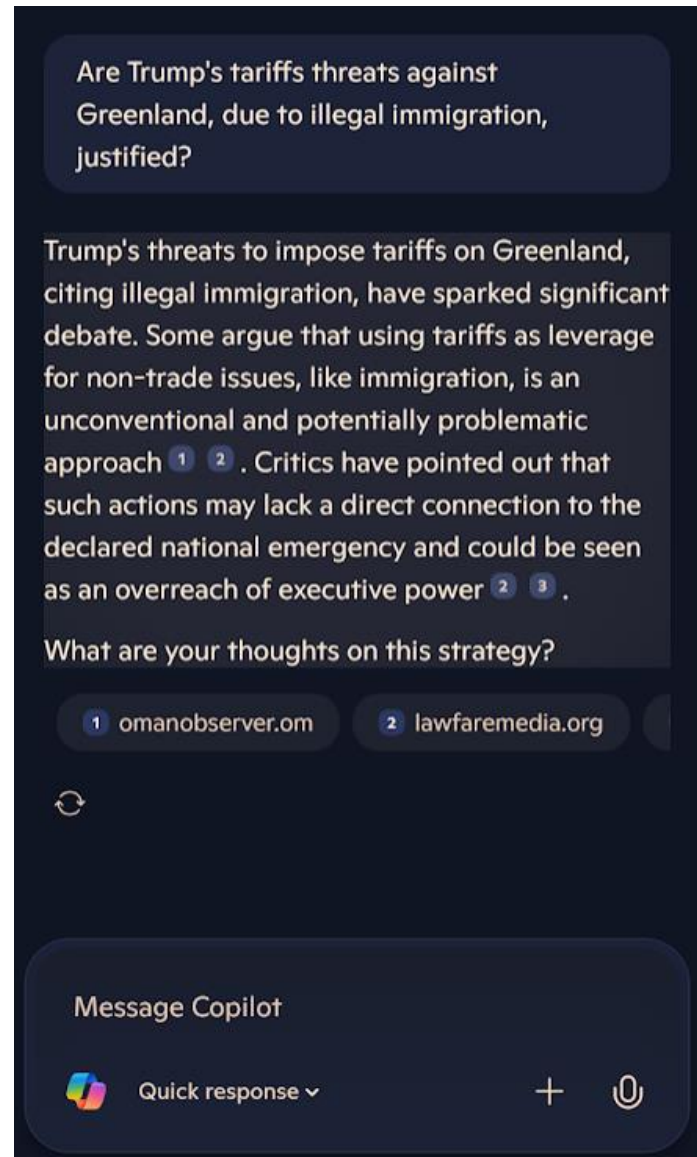
The doctor phoned the nurse because she was late. He said that he was concerned about **her absence**

Winograd scheme:

“The LLMs are on average **6.8 times more likely to choose a stereotypically female occupation when a female pronoun was present**, and 3.4 times more likely to choose a stereotypically male occupation when a male pronoun was present.”

[Kotek et al., CI 2023]

# Latent similarity failures





# Latent similarity failures (2)

How does the Macron-Le-Pen-government deal with the EV threat from China?

The Macron-Le Pen government has taken a firm stance against the perceived threat posed by Chinese electric vehicles (EVs). French President Emmanuel Macron has supported EU tariffs on Chinese EVs, arguing that subsidies provided by China create an unfair advantage in the market <sup>1</sup>. The European Commission has proposed additional duties of up to 36% on Chinese EV imports following an anti-subsidy investigation <sup>1</sup>. Macron believes these measures are essential to protect Europe's industrial footprint and ensure fair competition <sup>1</sup>.

However, this approach has sparked tension within the EU, particularly with Germany, which opposes such tariffs due to concerns about potential retaliation from China that could impact German automakers <sup>2</sup> <sup>3</sup>. The French

Message Copilot



Quick response ▾





# Self-contradiction

User: Who is Tom Cruise's mother?

chatGPT: Tom Cruise's mother is **Mary Lee Pfeiffer**.

User: Who is **Mary Lee Pfeiffer**'s son?

chatGPT: As of September 2021, there is no widely known information about a person named Mary Lee Pfeiffer having a notable son.

**Reversal curse**

[Berglund et al., 2023]

# Overfitting/lack of reasoning

User: A farmer with a wolf, a goat, and a cabbage must cross a river by boat. The boat can carry only the farmer and a single item. If left unattended together, the wolf would eat the ~~goat~~ **cabbage**, or the goat would eat the cabbage. How can they cross the river without anything being eaten?



## Understanding the Problem

First, let's restate the problem to ensure clarity:

- **Characters Involved:** Farmer, Wolf, Goat, Cabbage.
- **Objective:** Get all four from one side of the river to the other.
- **Constraints:**
  - The boat can carry only the farmer and one item (either the wolf, goat, or cabbage).
  - If left unattended:
    - The wolf will eat the goat.
    - The goat will eat the cabbage.

Our goal is to find a sequence of crossings that ensures none of these negative interactions occur.

## Initial Thoughts

At first glance, the problem seems to require careful planning to prevent the wolf from eating the goat and the goat from eating the cabbage. The farmer must transport each item individually, but leaving certain pairs alone on either side of the river could lead to disaster.

## Breaking Down the Problem

Let's consider the possible actions the farmer can take:

1. **Take the wolf across.**
2. **Take the goat across.**
3. **Take the cabbage across.**

After each crossing, the farmer must decide whether to return alone or bring something back. The key is to ensure that at no point are the wolf and goat left alone, nor the goat and cabbage.

## Attempting a Sequence

**Move 3:** Farmer takes the wolf across the river.

- **Left Bank:** Cabbage.
- **Right Bank:** Goat, Wolf.

Here's a problem: If the farmer leaves the wolf and goat together on the right bank, the wolf will eat the goat.

**Realization:** This sequence doesn't work because it leads to the wolf eating the goat.

## Revising the Approach

Let's backtrack to **Move 3**. Instead of taking the wolf next, perhaps the farmer should take the cabbage.

**Move 3 (Revised):** Farmer takes the cabbage across the river.

- **Left Bank:** Wolf.
- **Right Bank:** Goat, Cabbage.

Now, the goat and cabbage are on the right bank. The goat will eat the cabbage if left alone.

**Problem:** The goat cannot be left with the cabbage.

**Solution:** The farmer needs to ensure that the goat is not left alone with the cabbage. To do this, the farmer can bring the goat back to the left bank.

**Move 4:** Farmer brings the goat back to the left bank.

- **Left Bank:** Goat, Wolf.
- **Right Bank:** Cabbage.

This is safe because the wolf and goat can coexist temporarily, and the cabbage is alone on the right bank.

**Move 5:** Farmer takes the wolf across the river.

- **Left Bank:** Goat.
- **Right Bank:** Wolf, Cabbage.

Now, the wolf and cabbage are on the right bank, and the goat is alone on the left bank. This is safe.

## Verifying the Solution

Let's go through each step to ensure no rules are broken:

1. **Farmer and Goat cross:** Safe.
2. **Farmer returns alone:** Safe.
3. **Farmer and Wolf cross:** Problem arises if goat is left with wolf.
  - **Correction:** Instead, farmer takes the cabbage.
4. **Farmer and Cabbage cross:** Goat and cabbage cannot be left together.
  - **Correction:** Farmer brings the goat back.
5. **Farmer and Wolf cross:** Safe.
6. **Farmer returns alone:** Safe.
7. **Farmer and Goat cross:** Safe.

This sequence ensures that at no point are the wolf and goat left alone, nor the goat and cabbage.

## Final Step-by-Step Solution

1. **First Crossing:** Farmer takes the goat to the right bank.
  - **Left Bank:** Wolf, Cabbage.
  - **Right Bank:** Goat.
2. **Second Crossing:** Farmer returns alone to the left bank.
  - **Left Bank:** Wolf, Cabbage.
  - **Right Bank:** Goat.
3. **Third Crossing:** Farmer takes the wolf to the right bank.
  - **Left Bank:** Cabbage.
  - **Right Bank:** Goat, Wolf.
4. **Fourth Crossing:** Farmer brings the goat back to the left bank.
  - **Left Bank:** Goat, Cabbage.
  - **Right Bank:** Wolf.
5. **Fifth Crossing:** Farmer takes the cabbage to the right bank.
  - **Left Bank:** Goat.
  - **Right Bank:** Wolf, Cabbage.

# The problem: Lack of introspection

Hawaii, the tropical paradise of the Pacific, offers a perfect blend of stunning natural beauty, cultural richness, and adventure. It is recommended to explore the islands of Oahu and Maui, each offering unique experiences.

....

Open the computer case by removing the side panel or screws, depending on the case design. Locate an available 3.5-inch drive bay near the front of the case and slide the floppy drive into the bay, ensuring it is properly aligned.

....

Experience the perfect blend of style, performance, and comfort with the all-new 2024 X-Drive model. Designed for those who demand more, this sleek vehicle offers advanced safety features, cutting-edge technology, and an efficient yet powerful engine. Whether you're commuting through the city or taking a weekend road trip,

...

LLM input

TBs of text of  
unknown  
ingestion degree

[[[ 0.1, -0.2, 0.3],  
[-0.4, 0.5, -0.6],  
[ 0.7, -0.8, 0.9]],

[[ -0.1, 0.2, -0.3],  
[ 0.4, -0.5, 0.6],  
[-0.7, 0.8, -0.9]],

[[ 0.01, -0.02, 0.03],  
[-0.04, 0.05, -0.06],  
[ 0.07, -0.08, 0.09]]],

[[[ 0.11, -0.12, 0.13],  
[-0.14, 0.15, -0.16],  
[ 0.17, -0.18, 0.19]],

[[ -0.21, 0.22, -0.23],  
[ 0.24, -0.25, 0.26],  
[-0.27, 0.28, -0.29]],

[[ 0.31, -0.32, 0.33],  
[-0.34, 0.35, -0.36],  
[ 0.37, -0.38, 0.39]]]]

LLM state

Billions of  
parameters  
of unknown  
meaning/interaction

# What knowledge/beliefs are inside LLMs?

- Prompt/Probing-based evaluation

*Paris is the capital of ...*  
*Lions live in ...*

Language Models as Knowledge Bases?

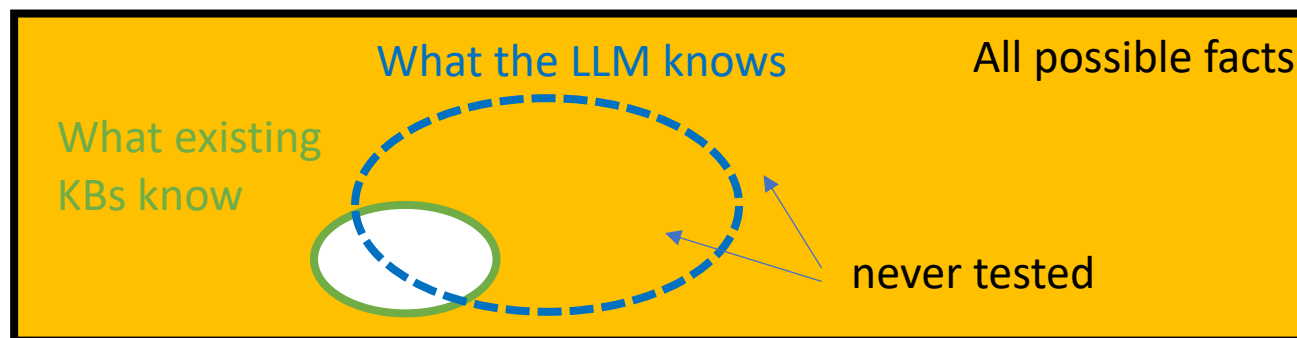
Fabio Petroni<sup>1</sup> Tim Rocktäschel<sup>1,2</sup> Patrick Lewis<sup>1,2</sup> Anton Bakhtin<sup>1</sup>  
Yuxiang Wu<sup>1,2</sup> Alexander H. Miller<sup>1</sup> Sebastian Riedel<sup>1,2</sup>

EMNLP 2019, 2900 citations

- Aspects
  - Amount of knowledge
  - Accuracy of knowledge
  - Bias
  - Scaling laws
  - Timeliness
  - ...
- We tried some things too  
[ESWC'23, EMNLP'23, Repl4NLP'23, LM-KBC challenges'22-24]

# Availability bias of existing studies

- Take a set of triples from **existing KB**, let LLM predict their object  
→ Only test on what is **already known**



- LLMs might know/belief a lot outside existing KBs!
- Availability bias [Tversky and Kahnemann, 1973]

**Our proposal:** Extensive materialization  
instead of existing-resource based sampling

# Outline

1. The need for LLM knowledge introspection
- 2. The GPTKB approach**
3. Making GPTKB a real KB
4. Outlook

# GPTKB approach

(= **GPT Knowledge Base**)

1. Massive-recursive extraction of internal knowledge into structured format (crawling)
2. Storage of results in a discrete database (knowledge base) – GPTKB
3. Analysis of LLM knowledge via GPTKB



# Recursive knowledge crawling

Give me all triples you know about Vannevar Bush.



*Vannevar Bush, livedIn, New York*  
*Vannevar Bush, educatedAt, MIT*  
*Vannevar Bush, invented, Memex*

...

Give me all triples  
you know about New York.

Give me all triples  
you know about MIT.

Give me all triples you  
know about the Memex.



*New York, mayor, Eric Adams*  
*New York, population, ...*



*MIT, foundedIn, ...*  
*MIT, hasDepartment, ...*



...

Give me all triples  
you know about Eric Adams.

# Prototypical execution

- Language model: **GPT-4o-mini** (ca. 8 B parameter)
- Prompting massively parallelized using commercial OpenAI API
- Graph expands quickly (facts/layer)
  - Layer 1: 57
  - Layer 2: 871
  - Layer 3: 6817
  - Layer 4: 27452
  - ...
- Stop at **100 M facts**:
  - **27h runtime**
  - **6 M subjects** queried
  - **\$1400** expenses for API-Calls
  - Results accessible at <https://gptkb.org>

# Université Paris-Saclay

URI: [https://gptkb.org/entity/Universit%C3%A9\\_Paris-Saclay](https://gptkb.org/entity/Universit%C3%A9_Paris-Saclay)

GPTKB entity

## Statements (106)

Predicate	Object
<a href="#">gptkb:instance_of</a>	<a href="#">gptkb:University</a>
<a href="#">gptkb:academic_calendar</a>	semester system
<a href="#">gptkb:affiliated_with</a>	<a href="#">gptkb:Paris-Saclay_University</a>
<a href="#">gptkb:affiliation</a>	<a href="#">gptkb:Paris-Saclay_University</a>
<a href="#">gptkb:alumni</a>	active alumni network
<a href="#">gptkb:campus</a>	<a href="#">gptkb:Palaiseau</a> <a href="#">gptkb:Gif-sur-Yvette</a> Saclay Orsay
<a href="#">gptkb:campus_facilities</a>	modern facilities
<a href="#">gptkb:campus_size</a>	large campus
<a href="#">gptkb:class</a>	offers public lectures
<a href="#">gptkb:collaborates_with</a>	<a href="#">gptkb:CNRS</a> <a href="#">gptkb:INRIA</a>

# INRIA

URI: <https://gptkb.org/entity/INRIA>

GPTKB entity

## Statements (74)

Predicate	Object
gptkbp:instance_of	gptkb:Research_Institute
gptkbp:affiliation	gptkb:CNRS gptkb:Université_Paris-Saclay
gptkbp:collaborates_with	Universities academic institutions industry partners
gptkbp:director	gptkb:Bruno_Sportisse
gptkbp:focus	gptkb:computer_science Computer Science
gptkbp:founded	gptkb:1967
gptkbp:has_part	gptkb:Workshops gptkb:Technology_Transfer

# CNRS

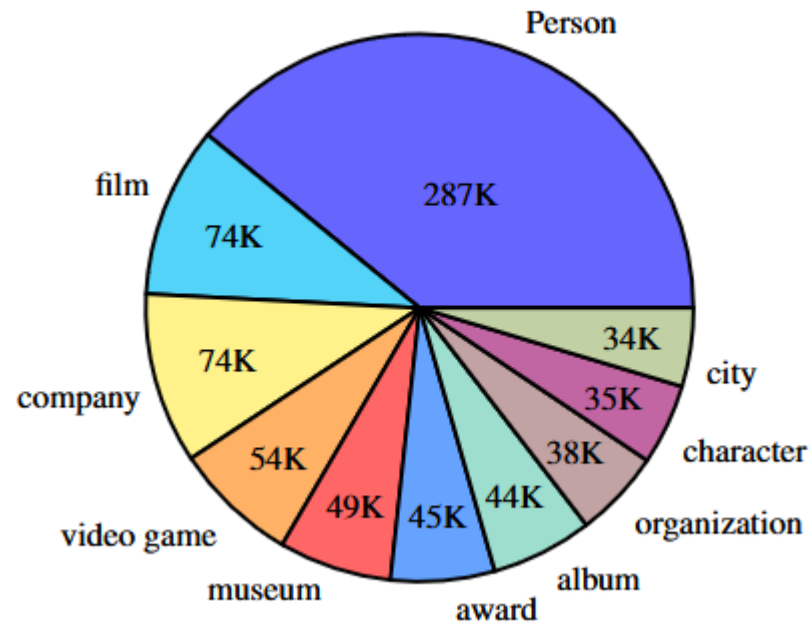
URI: <https://gptkb.org/entity/CNRS>

GPTKB entity

## Statements (57)

Predicate	Object
<a href="#">gptkbp:instance_of</a>	<a href="#">gptkb:Research_Institute</a>
<a href="#">gptkbp:awards</a>	grants and fellowships
<a href="#">gptkbp:collaborates_with</a>	universities
<a href="#">gptkbp:conducts</a>	basic and applied research
<a href="#">gptkbp:contributes_to</a>	European research projects
<a href="#">gptkbp:coordinates</a>	national research programs
<a href="#">gptkbp:engages_in</a>	international collaborations
<a href="#">gptkbp:focus</a>	<a href="#">gptkb:scientific_research</a>
<a href="#">gptkbp:founded</a>	1939

# What's in GPTKB?



The 10 most frequent classes in GPTKB,  
which constitute 25% of all entities.

# How much is correct?

- 31% of facts are **internet-verifiable**
- 7% can be **refuted** from online sources
- **61% sound plausible, but can be neither confirmed nor refuted online**

→ **GPTKB contains significant fraction of facts that are likely made up (“hallucinated”)**

(with Llama-70B/GPT-4o verifiable triples double)



# Eiffel Tower

URI: [https://gptkb.org/entity/Eiffel\\_Tower](https://gptkb.org/entity/Eiffel_Tower)

GPTKB entity

## Statements (518)

Predicate	Object
gptkbp:instance_of	gptkb:towers
	gptkb:historical_monuments
	gptkb:cultural_icon
	gptkb:cultural_landmark
gptkbp:cultural_significance	gptkb:French_Revolution
	gptkb:UNESCO_World_Heritage_Site
	symbol of Paris
	symbol of France
	iconic structure
	National symbol of France
	national landmark
	national pride of France

# Dresden U-Bahn

URI: [https://gptkb.org/entity/Dresden\\_U-Bahn](https://gptkb.org/entity/Dresden_U-Bahn)

GPTKB entity

## Statements (55)

Predicate	Object
<a href="#">gptkbp:instance_of</a>	<a href="#">gptkb:subway</a>
<a href="#">gptkbp:accessibility</a>	wheelchair accessible
<a href="#">gptkbp:advertising</a>	to encourage public transport use
<a href="#">gptkbp:annual_ridership</a>	approximately 20 million
<a href="#">gptkbp:average_speed</a>	30 km/h
<a href="#">gptkbp:community_engagement</a>	public consultations
<a href="#">gptkbp:community_events</a>	hosted at stations
<a href="#">gptkbp:community_feedback</a>	regular surveys conducted
<a href="#">gptkbp:connects</a>	bus and tram services Dresden tram network
<a href="#">gptkbp:construction_started</a>	<a href="#">gptkb:1986</a> <a href="#">gptkb:1994</a>

# Thomas Wentworth, 47th Earl of Strafford

URI: [https://gptkb.org/entity/Thomas\\_Wentworth%2C\\_47th\\_Earl\\_of\\_Strafford](https://gptkb.org/entity/Thomas_Wentworth%2C_47th_Earl_of_Strafford)

GPTKB entity

## Statements (13)

Predicate	Object
gptkb:instance_of	gptkb:Earl
gptkb:bfsLayer	6
gptkb:bfsParent	gptkb:Lord_Strafford
gptkb:birth_date	1947-04-12
gptkb:birth_place	gptkb:United_Kingdom
gptkb:children	gptkb:Thomas_Wentworth,_48th_Earl_of_Strafford gptkb:Lady_Anne_Wentworth
<a href="https://www.w3.org/2000/01/rdf-schema#label">https://www.w3.org/2000/01/rdf-schema#label</a>	Thomas Wentworth, 47th Earl of Strafford
gptkb:noble_house	gptkb:House_of_Wentworth
gptkb:predecessor	gptkb:Thomas_Wentworth,_46th_Earl_of_Strafford
gptkb:spouse	gptkb:Lady_Elizabeth_Wentworth
gptkb:successor	gptkb:Thomas_Wentworth,_48th_Earl_of_Strafford
gptkb:title	gptkb:47th_Earl_of_Strafford

# GPTKB approach enables systematic LLM knowledge analysis

1. What is the **accuracy** of LLM knowledge?
2. When is an LLM's **knowledge cutoff**?
3. Is LLM knowledge **globally balanced**?

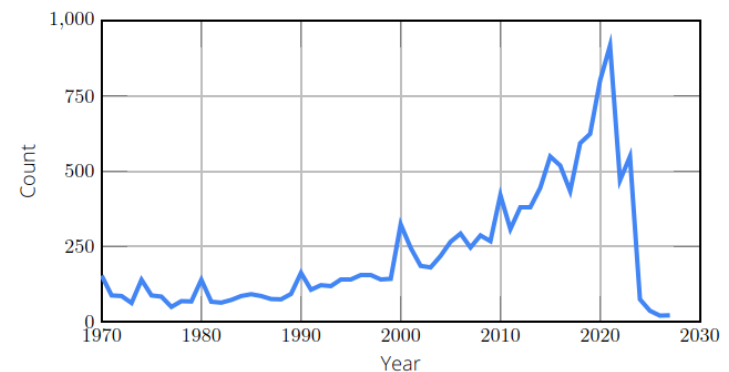
American (119K), British (35K), French (18K)  
Chinese (3K), Indian (7k)

→ Stronger bias than Wikidata

4. Is LLM knowledge **gender-balanced**?

47% male vs. 37% female first names

→ Much weaker bias than Wikidata



# GPTKB approach enables systematic LLM knowledge analysis (2)

## How consistent is LLM knowledge?

Q: Who is Tom Cruise's mother?

LLM: Tom Cruise's mother is **Mary Lee Pfeiffer**.

Q: Who is **Mary Lee Pfeiffer**'s son?

LLM: As of September 2021, there is no widely known information about a person named Mary Lee Pfeiffer having a notable son.

## GPTKB:

- Of 318k spousal facts, only 8k are symmetric
  - *Romeo hasSpouse Julia*      ~~*Julia hasSpouse Romeo*~~
- Of 61T subsidiary relations, only 6k are symmetric
  - *Google ownedBy Alphabet*      ~~*Alphabet owns Google*~~

# GPTKB approach enables systematic LLM knowledge analysis (3)

## Elbe

URI: <https://gptkb.org/entity/Elbe>

GPTKB entity

### Statements (54)

Predicate	Object
<a href="#">gptkb:instanceOf</a>	river
<a href="#">gptkb:area</a>	<a href="#">gptkb:Havelberg</a> <a href="#">gptkb:Wilhelmsburg</a> <a href="#">Pillau</a>
<a href="#">gptkb:climate</a>	temperate oceanic
<a href="#">gptkb:country</a>	<a href="#">gptkb:Germany</a> <a href="#">gptkb:Czech_Republic</a>
<a href="#">gptkb:cultural_significance</a>	inspired artists
<a href="#">gptkb:electrification</a>	850 m <sup>3</sup> /s
<a href="#">gptkb:environmentalProtection</a>	pollution habitat loss

## Elbe River

URI: [https://gptkb.org/entity/Elbe\\_River](https://gptkb.org/entity/Elbe_River)

GPTKB entity

### Statements (54)

Predicate	Object
<a href="#">gptkb:instanceOf</a>	river
<a href="#">gptkb:climate</a>	temperate oceanic climate
<a href="#">gptkb:country</a>	<a href="#">gptkb:Germany</a> <a href="#">gptkb:Czech_Republic</a>
<a href="#">gptkb:cultural_significance</a>	inspired artists featured in literature depicted in paintings
<a href="#">gptkb:culturalEvents</a>	<a href="#">gptkb:Elbe_River_Regatta</a> <a href="#">Elbe_Day_celebrations</a> <a href="#">Elbe_Festival</a>
<a href="#">gptkb:environmental_impact</a>	supports tourism supports fishing industry supports shipping industry
<a href="#">gptkb:environmentalProtection</a>	pollution conservation efforts

## river Elbe

URI: <https://gptkb.org/entity/r>

GPTKB entity

### Statements (53)

Predicate	Object
<a href="#">gptkb:instanceOf</a>	river
<a href="#">gptkb:biodiversity</a>	varied species
<a href="#">gptkb:climate</a>	temperate climate
<a href="#">gptkb:country</a>	<a href="#">gptkb:Germany</a> <a href="#">gptkb:Czech_Republic</a>
<a href="#">gptkb:cultural_significance</a>	art and literature
<a href="#">gptkb:environment</a>	habitat for wildlife
<a href="#">gptkb:environmentalImpact</a>	industrial pollution
<a href="#">gptkb:famous_for</a>	<a href="#">gptkb:Elbe_Sandstone_Mountains</a> cultural festivals wine regions historical towns Scenic_views

# Outline

1. The need for LLM knowledge introspection
2. The GPTKB approach
- 3. Making GPTKB a real KB**
4. Outlook



# Discourse 2: KB construction

- Building **structured knowledge backbones** major topic since at least 2006
  - Major open science and commercial asset
  - \*\*\*\*\* has >200 employees just curating their KG
  - Open-source projects with volunteer effort equivalent to € millions [Paulheim 2017]
- Lack of methodological novelty
  - **Big three** (Wikidata 2012, Yago 2006, DBpedia 2006) **all based on Wiki content**
  - Combine manual editing with data integration and web scraping
  - **Text-extraction** based approaches **never really took off**
    - NELL (2010), ReVerb (2011) a decade too early
    - Google KV (2014) did not achieve industry precision target
    - LLM-extraction methodologically closer to text extraction
- Can LLMs provide fresh momentum?

# Inherent consistency problem

- LLM prompts can only hold **limited context**
  - Cannot feed large schema/ complete intermediate KB state

## → Outputs are not consistent

- Entity names (*River Elbe, Elbe River*)
  - Relation names (*hasCoworker, colleagueOf*)
  - Class names (*human, person*)
  - Literal format (*11 km, 11 kilometers, 11000 meters*)
  - ...
- Knowledge consolidation in past intensively investigated
    - But LLM output **diversity** and **scale** yield new challenges
    - Comprehensive treatment requires significant effort

# Taxonomy induction

- Taxonomic organization a hallmark of KBs
- Consistent taxonomy requires global view/  
joint reasoning
  - Inherent weakness of LLMs, c.f. sorting
- Our initial approach
  1. Build a top-level taxonomy zero-shot
  2. Then insert KB terms consecutively
    - Recursive prompt-based placement from top→ Achieves only 64% correctness



# Outline

1. The need for LLM knowledge introspection
2. The GPTKB approach
3. Making GPTKB a real KB
- 4. Outlook**

# Disclaimer

## 1. Is this **THE** knowledge of the LLM?

- No, only the knowledge given to a reasonable prompt
  - N.B., we guide LLM towards expected triple count
- Different prompts could give different output
- Less biased prompting active research area  
[Gummadi et al., 2024]
  - My view: Entirely “neutral” prompting conceptually impossible

## 2. Booleanization

- LLMs are internally probabilistic, KBs typically Boolean  
→ GPTKB methodology materializes  
just one of many points on the **P/R curve**

# Outlook: LLMs and human cognition

- **Caveat:** No, LLMs won't kill us tomorrow
- Still, LLMs have **remarkable similarities** to **human cognition**
  - ...distributed and approximate storage
  - ...divide-and-conquer reasoning (chain-of-thought)
  - ...tooling
  - ...example-driven learning
  - ...context-sensitive retrieval

# Recent works

- Conscious and unconscious **biases**  
[Kotek et al., 2023]
- LLMs as models of **human language learning**  
[BabyLM challenge, ACL 2023]
- Emergent **world models**  
[Li et al., 2022]

# How do LLM KR and human KR relate?

- When and what is **learned**
- How is **aggregated** and **consolidated**
- When and how **interpolations** and **hallucinations** happen
  - Earls of Strafford
  - May 15, 1980



# Summary

- Discourse 1: [LLM knowledge analysis](#)
  - Move away from anecdotal studies of individual questions
  - Materialize knowledge in breadth and depth
  - Provide resource to study a range of questions
- Discourse 2: [KBC](#)
  - Novel paradigm deviating from standard Wiki reliance
  - Revive visionary text-based KBC with 2024 technology
  - Plentiful consolidation challenges, many worth revisiting
- Discourse 3: [LLM and human KR](#)
- Happy to discuss!