

Knowledge Graphs in the age of Large Language Models

EGC Keynote

Aidan Hogan

DCC, Universidad de Chile

Instituto Milenio Fundamentos de los Datos



UNIVERSIDAD
DE CHILE



Instituto Milenio
Fundamentos
de los datos



The problem with the Web



Example 1.1

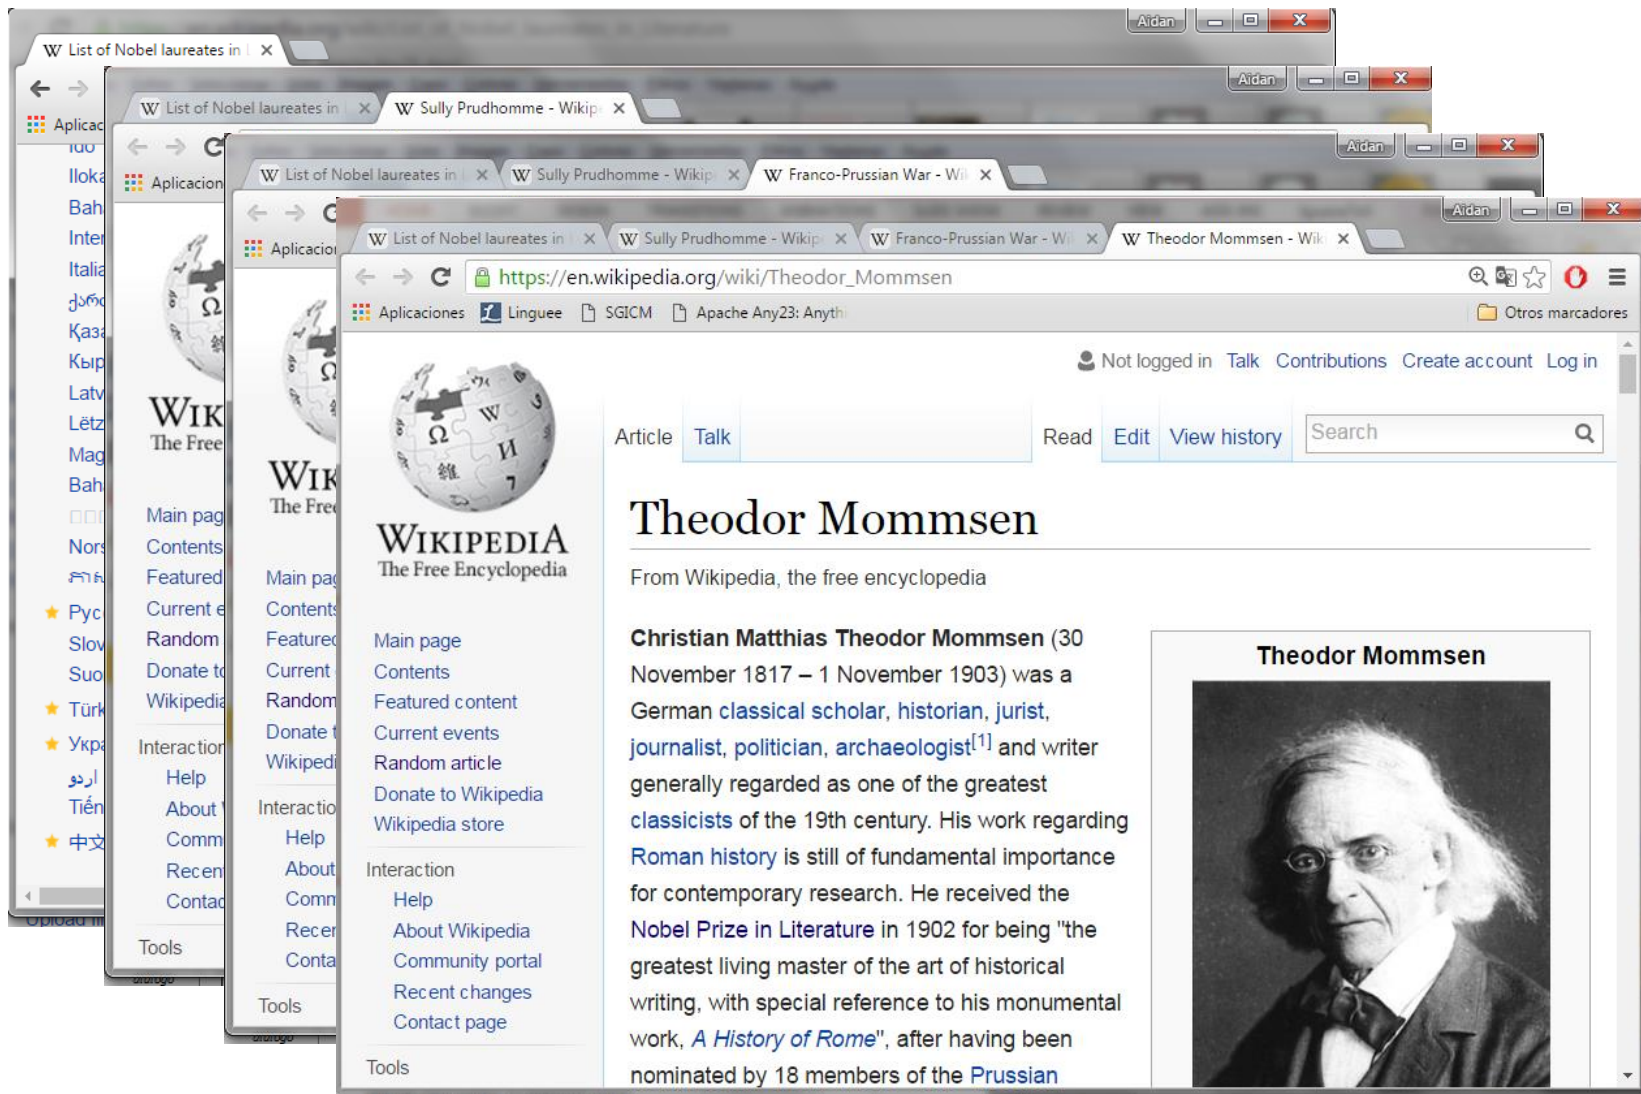


Julie is a literature student. For her thesis, she needs to find a list of Nobel Laureates in Literature who fought in a war, the year they were awarded the Nobel prize, and the name of the war(s) they fought in.



Aidan Hogan

The Web of Data



Web
Data

Nobel Laureates

We list Nobel laureates ordered by category then year.

Literature

Year	Awardee
1901	Sully Prudhomme
...	...
1953	Winston Churchill
1954	Ernest Hemingway
...	...
1964	Jean-Paul Sartre [†]
...	...
1983	William Golding
...	...
1985	Claude Simon
...	...

[†]Sartre declined.

Peace

Year	Awardee
1901	Henry Dunant
1901	Frédéric Passy
...	...

William G. Golding

Sir William Gerald Golding CBE (1911–1993) was a British novelist, playwright and poet. His most celebrated works include the novels "[Lord of the Flies](#)", and "[Rites of Passage](#)" for which he was awarded the [Booker Prize](#).

...

He joined the [Royal Navy](#) in 1940, commanding a landing ship at the [Normandy Invasion](#) on [D-Day](#).

Invasion of Normandy

On June 6th 1944 ([D-Day](#)), the [Western Allies](#) of [World War II](#) assaulted [Nazi](#) occupation forces from the sea at [Normandy](#), as part of [Operation Overlord](#).

What Nobel Laureates fought in WWII?

William G. Golding, ...

Aidan Hogan

The Web of Data

Lucys Sedgyefyz

Ny sozf Lucys sedgyefyz ugwygyw ci befjyugi fhyl iyeg.

Sofygefdgy

Iyeg	Enegwyw
1901	Zdssi Rgdwhummy
...	...
1953	Nolzful Bhdgbhoss
1954	Yglyzf Hymoljnei
...	...
1964	Vyel-Reds Zegfgy [†]
...	...
1983	Nossoem Juswolj
...	...
1985	Bsedwy Zomul
...	...

[†] Zegfgy wybsolyw.

Ryeby

Iyeg	Enegwyw
1901	Hylgi Wdlelf
1901	Pgýwýgob Rezzi
...	...

...

Nossoem J. Juswolj

Zog Nossoem Jygesw Juswolj BCY (1911–1993) nez e Cgofozh lutysozf, rseigojhf elw ruyf. Hoz muzf bysycgefyw nugqz olbsdwy fhy lutysz "[Sugw up fhy Psoyz](#)", elw "[Gofyz up Rezzejy](#)" pug nhobh hy nez enegwyw fhy "[Cuuqyg Rgoay](#)".

...

Hy vuolyw fhy [Guies Leti](#) ol 1940, bummelwolj e selwolj zhor ef fhy [Lugmelwi Oltezoul](#) ul [W-Wei](#).

...

Oltezoul up Lugmelwi

Ul Vdly 6th 1944 ([W-Wei](#)), fhy [Nyzfygl Essoyz](#) up [Nugsw Neg OO](#) ezzedsfyw [Leao](#) ubbdrefoul pugbyz pgum fhy zye ef [Lugmelwi](#), ez regf up [Urygefoul Utygsugw](#).

...

Ol nhobh negz hety Lucys Sedgyefyz pudjhf?

Nugsw Neg OO, ...

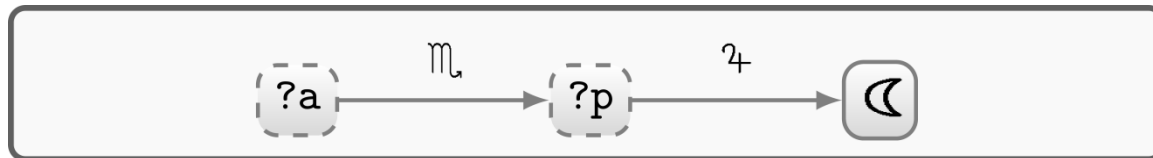
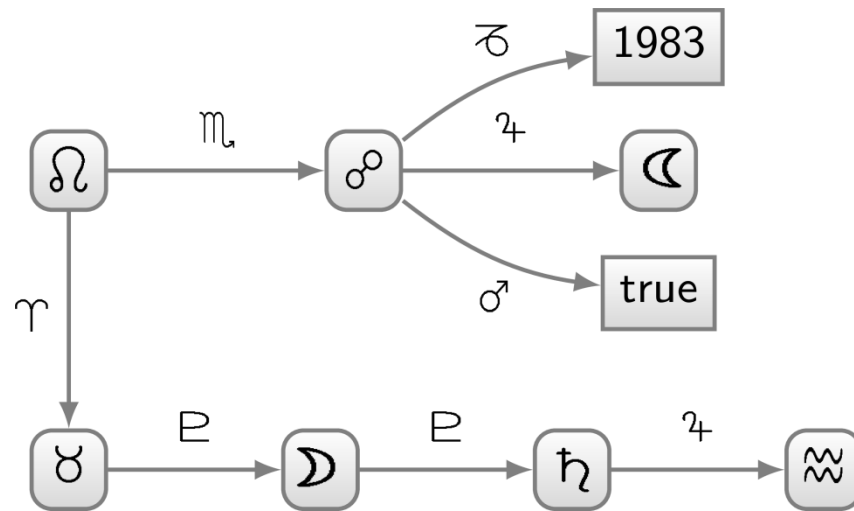
Aidan Hogan

The Web of Data



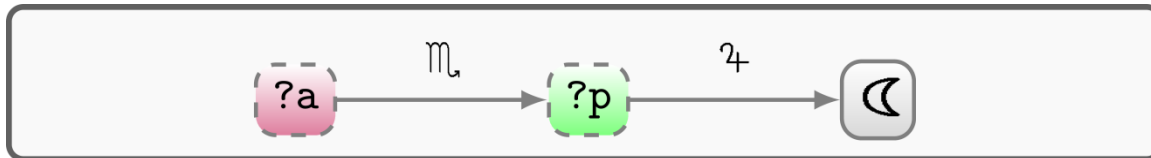
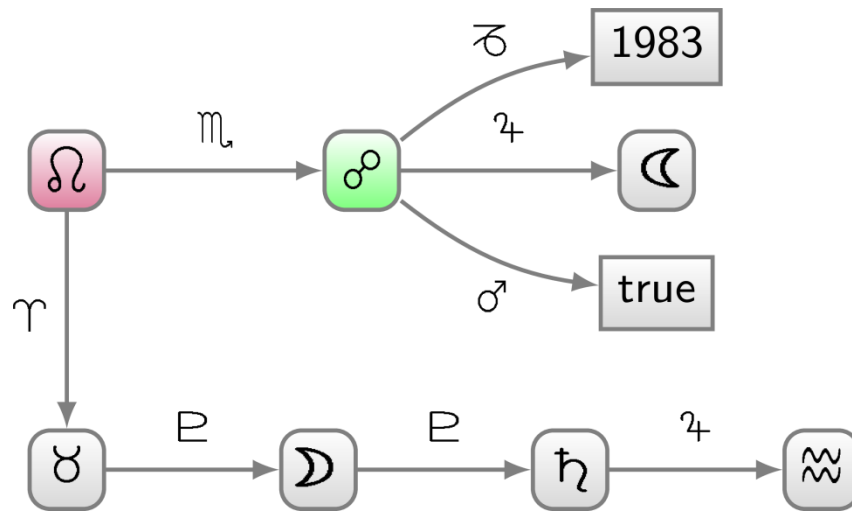
Structure and symbols





Aidan Hogan

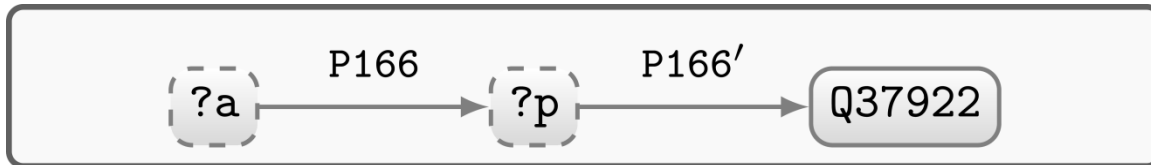
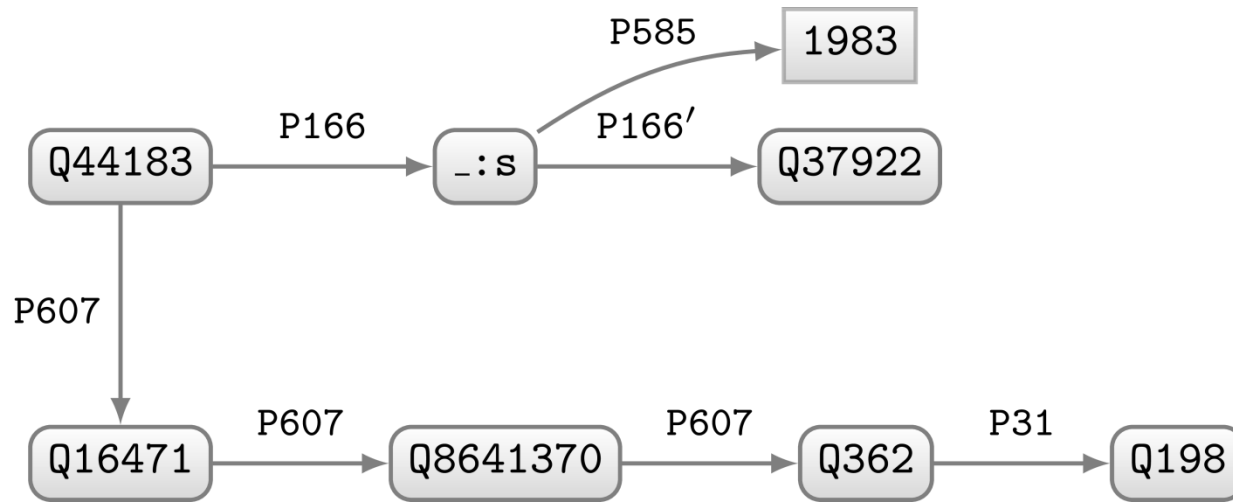
The Web of Data



?a	?p
Ω	♂

Aidan Hogan

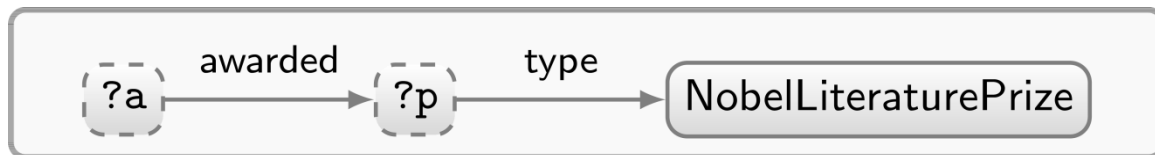
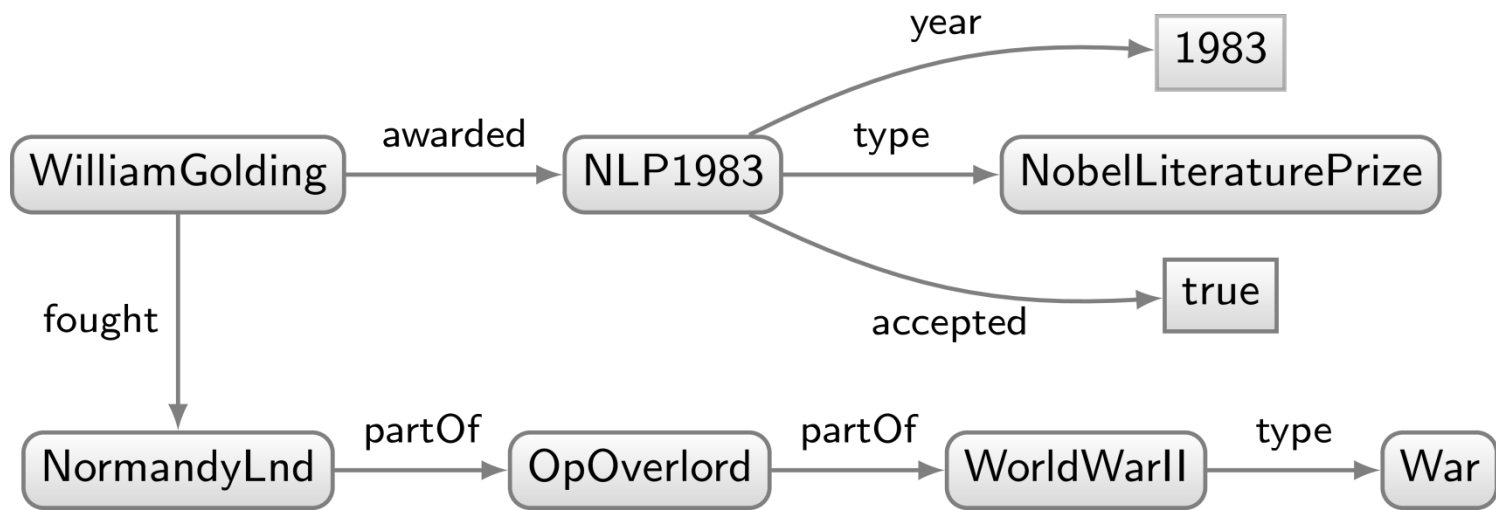
The Web of Data



?a	?p
Q44183	_:s

Aidan Hogan

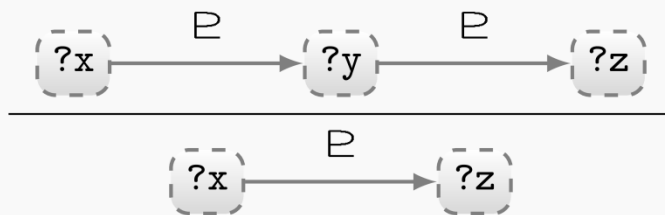
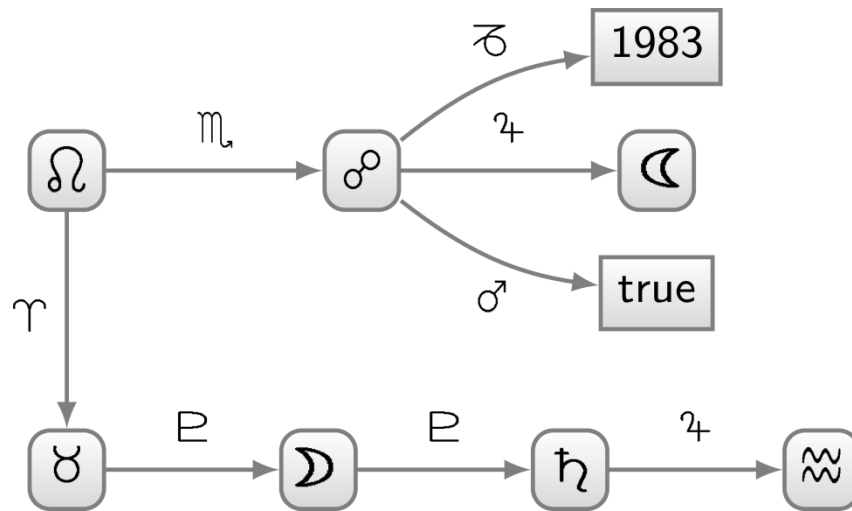
The Web of Data



?a	?p
WilliamGolding	NLP1983

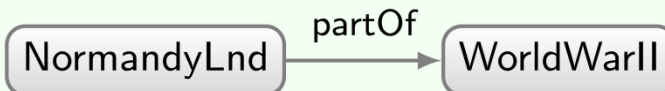
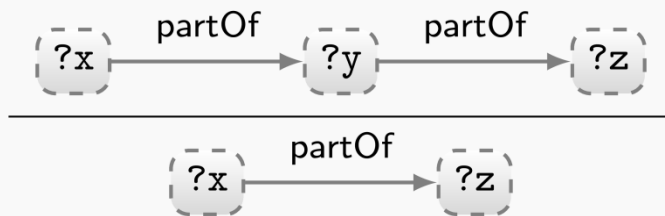
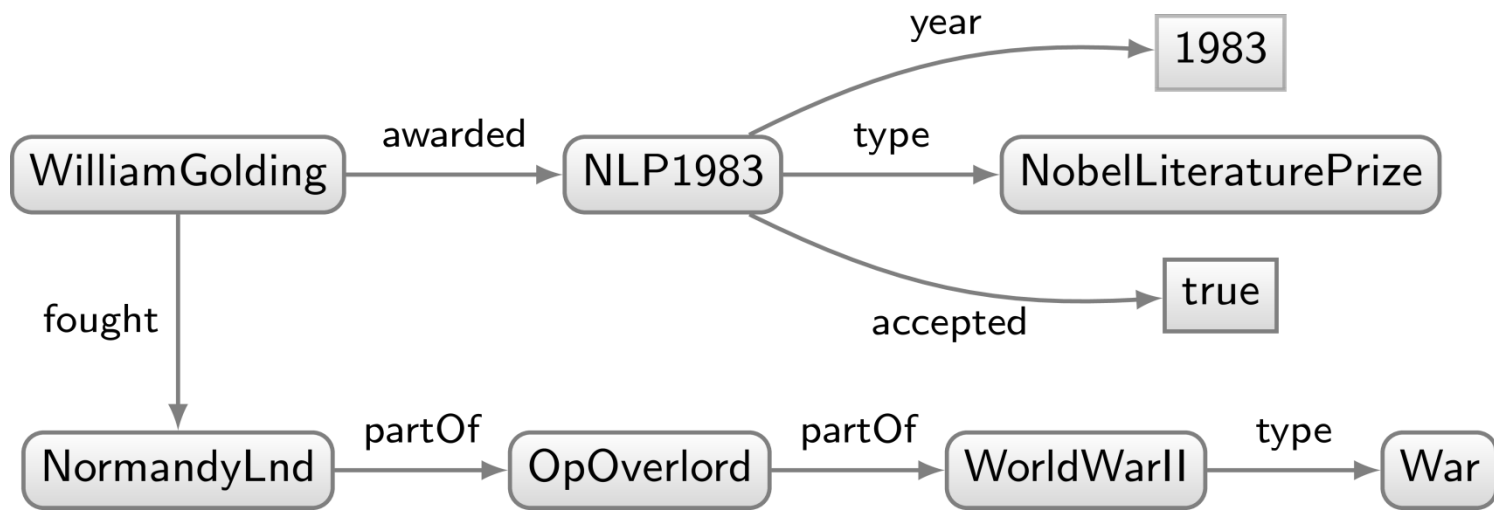
Aidan Hogan

The Web of Data



Aidan Hogan

The Web of Data



Aidan Hogan

The Web of Data



The problem with the Web?



Lucys Sedgyefyz

Ny sozf Lucys sedgyefyz ugwygyw ci befjyugi fhyl iyeg.

Sofygefdgy

Iyeg	Enegwyw
1901	Zdssi Rgdwhummy
...	...
1953	Nolzful Bhdgbhoss
1954	Yglyzf Hymoljnei
...	...
1964	Vyel-Reds Zegfgy [†]
...	...
1983	Nossoem Juswolj
...	...
1985	Bsedwy Zomul
...	...

[†] Zegfgy wybsolyw.

Ryeby

Iyeg	Enegwyw
1901	Hylgi Wdlelf
1901	Pgýwýgob Rezzi
...	...

...

Nossoem J. Juswolj

Zog Nossoem Jygesw Juswolj BCY (1911–1993) nez e Cgofozh lutysozf, rseigojhf elw ruyf. Hoz muzf bysycgefyw nugqz olbsdwy fhy lutysz "[Sugw up fhy Psoyz](#)", elw "[Gofyz up Rezzejy](#)" pug nhobh hy nez enegwyw fhy "[Cuuqyg Rgoay](#)".

...

Hy vuolyw fhy [Guies Leti](#) ol 1940, bummelwolj e selwolj zhor ef fhy [Lugmelwi Oltezoul](#) ul [W-Wei](#).

...

Oltezoul up Lugmelwi

Ul Vdly 6th 1944 ([W-Wei](#)), fhy [Nyzfygl Essoyz](#) up [Nugsw Neg OO](#) ezzedsfyw [Leao](#) ubbdrefoul pugbyz pgum fhy zye ef [Lugmelwi](#), ez regf up [Urygefoul Utygsugw](#).

...

Ol nhobh negz hety Lucys Sedgyefyz pudjhf?

Nugsw Neg OO, ...

Aidan Hogan

The Web of Data

Information may be split over multiple clauses that use references such as pronouns that may be difficult to resolve; for example, compare the “*she*” in both:

*“Jane tried to follow her late grandmother’s recipe, but **she** had not included details of the cooking times.”*

and

*“Jane tried to follow her late grandmother’s recipe, but **she** had forgotten to buy lemons”.*

Aidan Hogan

The Web of Data



A

Jane tried to follow her late grandmother's recipe, but she had not included details of the cooking times. Who forgot to include the cooking times: Jane or her grandmother?



Jane's late grandmother forgot to include the cooking times in the recipe.

A

Jane tried to follow her late grandmother's recipe, but she had forgotten to buy lemons. Who forgot to buy lemons: Jane or her grandmother?



Jane forgot to buy lemons.

A

Why was Jane's grandmother late?



The information provided in the previous questions does not mention why Jane's grandmother was late. If you could provide more context or details, I'd be happy to help you with a specific answer.

A



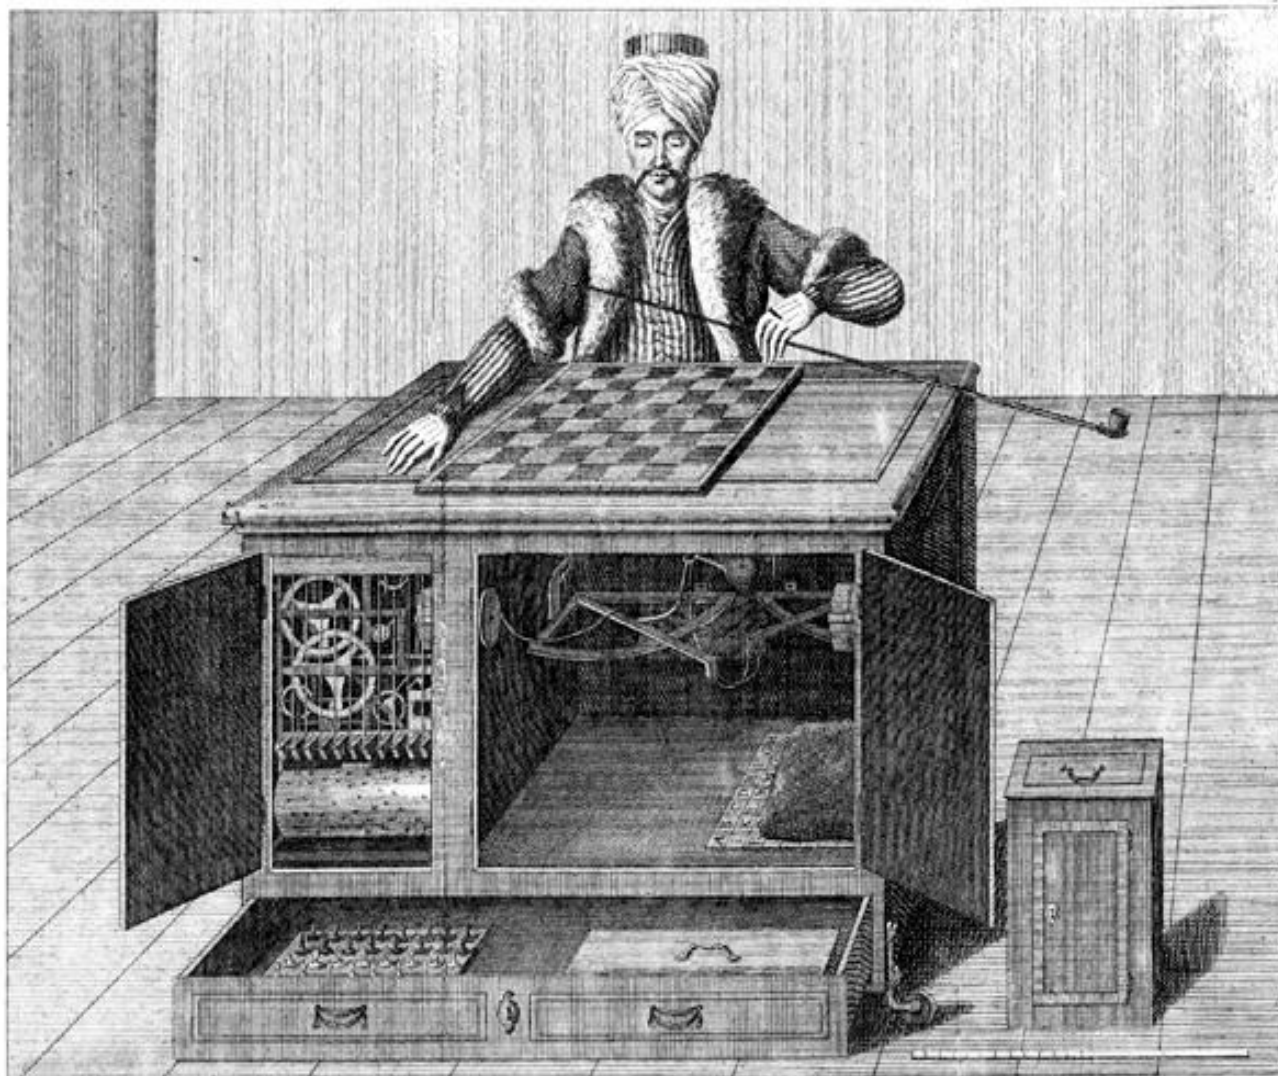
Why was Jane's grandmother late?

The phrase "late grandmother" typically means that Jane's grandmother has passed away, rather than being late in the sense of tardiness. In this context, "late" is used to indicate that the person is deceased. Therefore, the sentence doesn't provide information about Jane's grandmother being late for an event; it indicates that she is no longer alive.



Back to the future



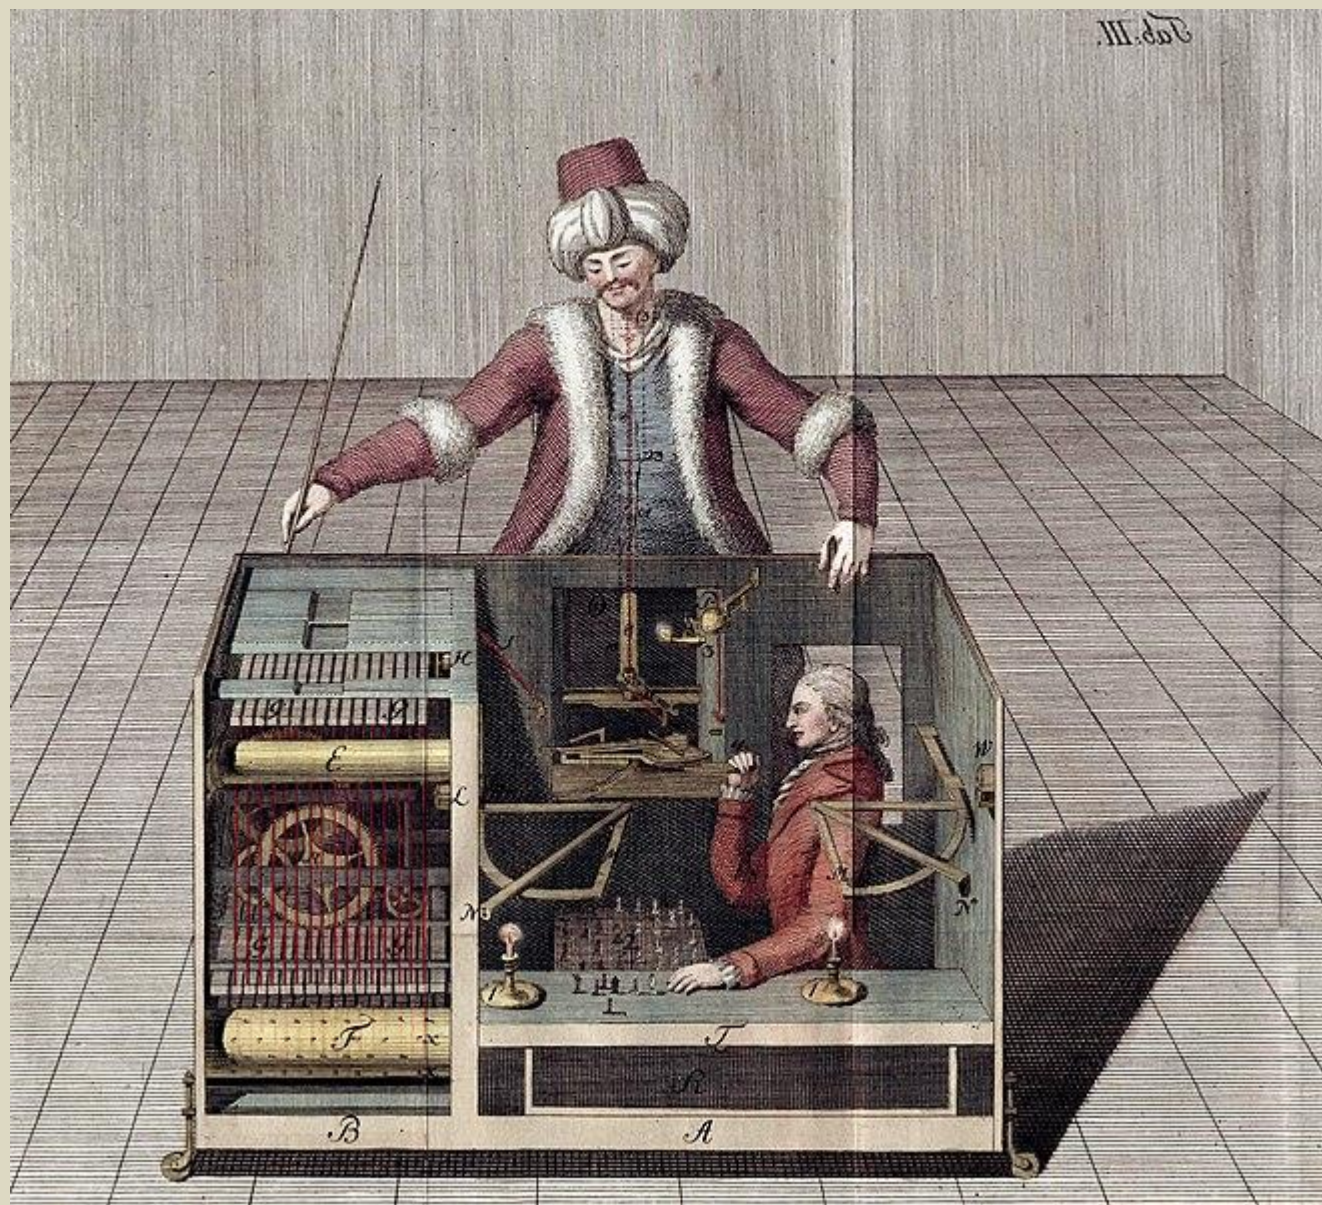


W. de Kempelen del.

Che a Mechoh exend. Basilea.

P.G. Rintz sc.

Der Schach-Spieler, wie er vor dem Spiele gezeiget wird von Herrn L. Souverel Chees, tel qu'on le montre avant le jeu par devant.



A PROPOSAL FOR THE
DARTMOUTH SUMMER RESEARCH PROJECT
ON ARTIFICIAL INTELLIGENCE

J. McCarthy, Dartmouth College
M. L. Minsky, Harvard University
N. Rochester, I. B. M. Corporation
C. E. Shannon, Bell Telephone Laboratories

August 31, 1955

A Proposal for the
DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

August 31, 1955

A PROPOSAL FOR THE

1) Automatic Computers

If a machine can do a job, then an automatic calculator can be programmed to simulate the machine. The speeds and memory capacities of present computers may be insufficient to simulate many of the higher functions of the human brain, but the major obstacle is not lack of machine capacity, but our inability to write programs taking full advantage of what we have.

August 31, 1955

2) How Can a Computer be Programmed to Use a Language

It may be speculated that a large part of human thought consists of manipulating words according to rules of reasoning and rules of conjecture. From this point of view, forming a generalization consists of admitting a new word and some rules whereby sentences containing it imply and are implied by others. This idea has never been very precisely formulated nor have examples been worked out.

I could car
write a lot
this.

3. Neuron Nets

How can a set of (hypothetical) neurons be arranged so as to form concepts. Considerable theoretical and experimental work has been done on this problem by Uttley, Rashevsky and his group, Farley and Clark, Pitts and McCulloch, Minsky, Rochester and Holland, and others. Partial results have been obtained but the problem needs more theoretical work.

5) Self-Improvement

Probably a truly intelligent machine will carry out activities which may best be described as self-improvement. Some schemes for doing this have been proposed and are worth further study. It seems likely that this question can be studied abstractly as well.

7) Randomness and Creativity

Randomness is an artifact appearing in the H.C. method of solution.

A fairly attractive and yet clearly incomplete conjecture is that the difference between creative thinking and unimaginative competent thinking lies in the injection of a some randomness. The randomness must be guided by intuition to be efficient. In other words, the educated guess or the hunch include controlled randomness in otherwise orderly thinking.

problem
the H.C.
method
that a
"hunch"
includes
some

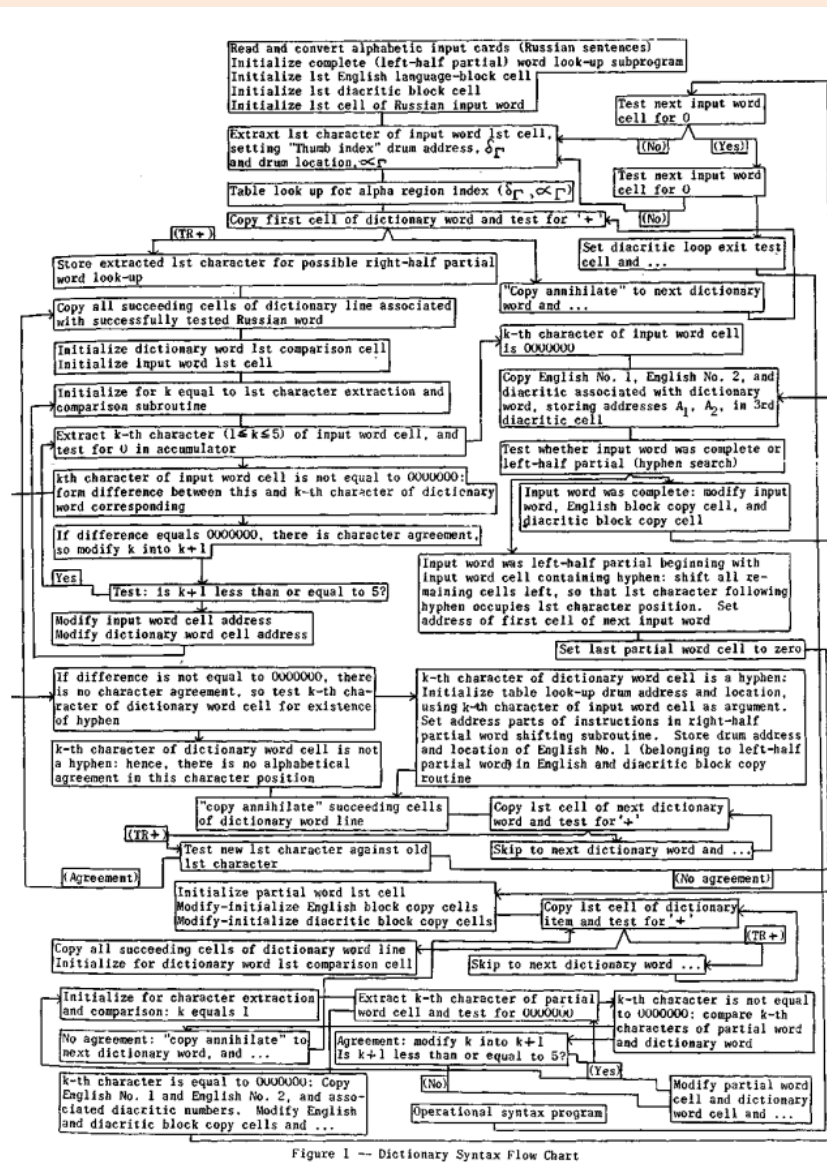


Figure 1 -- Dictionary Syntax Flow Chart

- Jacob Ornstein: Mechanical Translation: New Challenge to Communication. Science vol. 122 (1955).

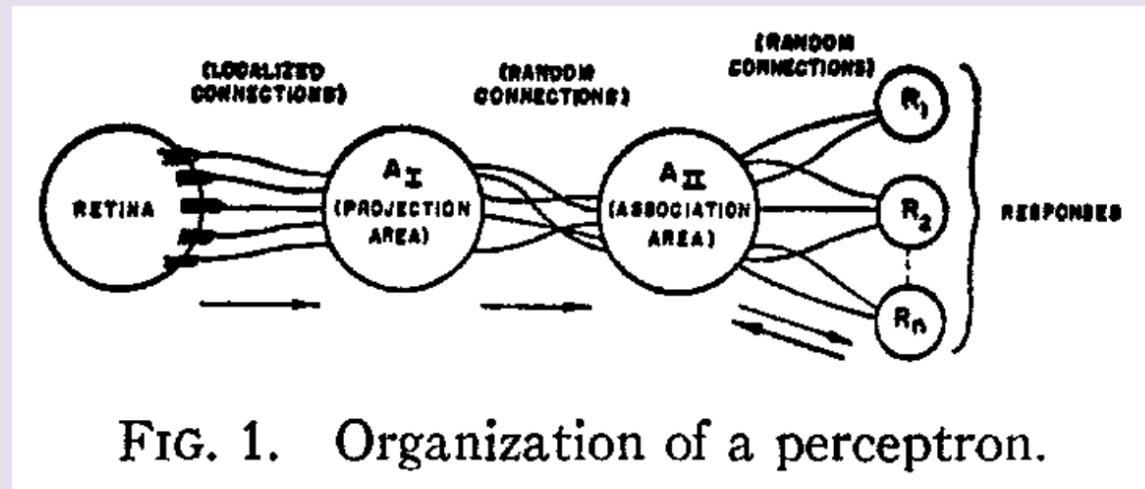


FIG. 1. Organization of a perceptron.

ory (2, 3, 9, 14). The alternative approach, which stems from the tradition of British empiricism, hazards the guess that the images of stimuli may never really be recorded at all, and that the central nervous system simply acts as an intricate switching network, where retention takes the form of new connections, or pathways, between centers of activity. In many

- Frank Rosenblatt: The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain. Psychological Review. 65 (6): 386–408 (1958)

THE NEW YORK TIMES, TUESDAY, JULY 8, 1958.

NEW NAVY DEVICE LEARNS BY DOING

Psychologist Shows Embryo
of Computer Designed to
Read and Grow Wiser

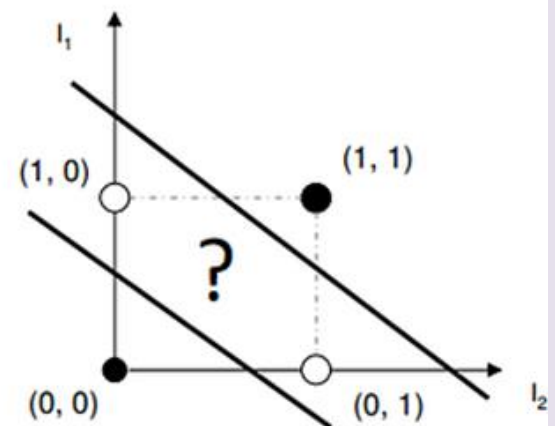
WASHINGTON, July 7 (UPI)—The Navy revealed the embryo of an electronic computer today that it expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence.

The embryo—the Weather Bureau's \$2,000,000 "704" computer—learned to differentiate between right and left after fifty attempts in the Navy's demonstration for newsmen.

The service said it would use this principle to build the first of its Perceptron thinking machines that will be able to read and write. It is expected to be finished in about a year at a cost of \$100,000.

Later Perceptrons will be able to recognize people and call out their names and instantly translate speech in one language to speech or writing in another language, it was predicted.

XOR		
I_1	I_2	out
0	0	0
0	1	1
1	0	1
1	1	0



Source: Solving XOR with a single Perceptron
Lucas Araújo.

- Frank Rosenblatt: The Perceptron: A Mathematical Model for Emulating the Storage and Organization of Information in the Human Brain (1958), pp. 386–408 (1958)

Welcome to

EEEEEE	LL	IIII	ZZZZZZZ	AAAAA
EE	LL	II	ZZ	AA AA
EEEEE	LL	II	ZZZ	AAAAAAA
EE	LL	II	ZZ	AA AA
EEEEEE	LLLLL	IIII	ZZZZZZZ	AA AA

Eliza is a mock Rogerian psychotherapist.

The original program was described by Joseph Weizenbaum in 1966.

This implementation by Norbert Landsteiner 2005.

ELIZA: Is something troubling you ?

YOU: Men are all alike.

ELIZA: What is the connection, do you suppose ?

YOU: They're always bugging us about something or other.

ELIZA: Can you think of a specific example ?

YOU: Well, my boyfriend made me come here.

ELIZA: Is it important to you that your boyfriend made you come here ?

YOU: He says I'm depressed much of the time.

ELIZA: I am sorry to hear that you are depressed.

YOU: It's true. I am unhappy.

ELIZA: Can you explain what made you unhappy ?

YOU: ■

CULTURE DESK

THE SAMANTHA TEST


By Brian Christian

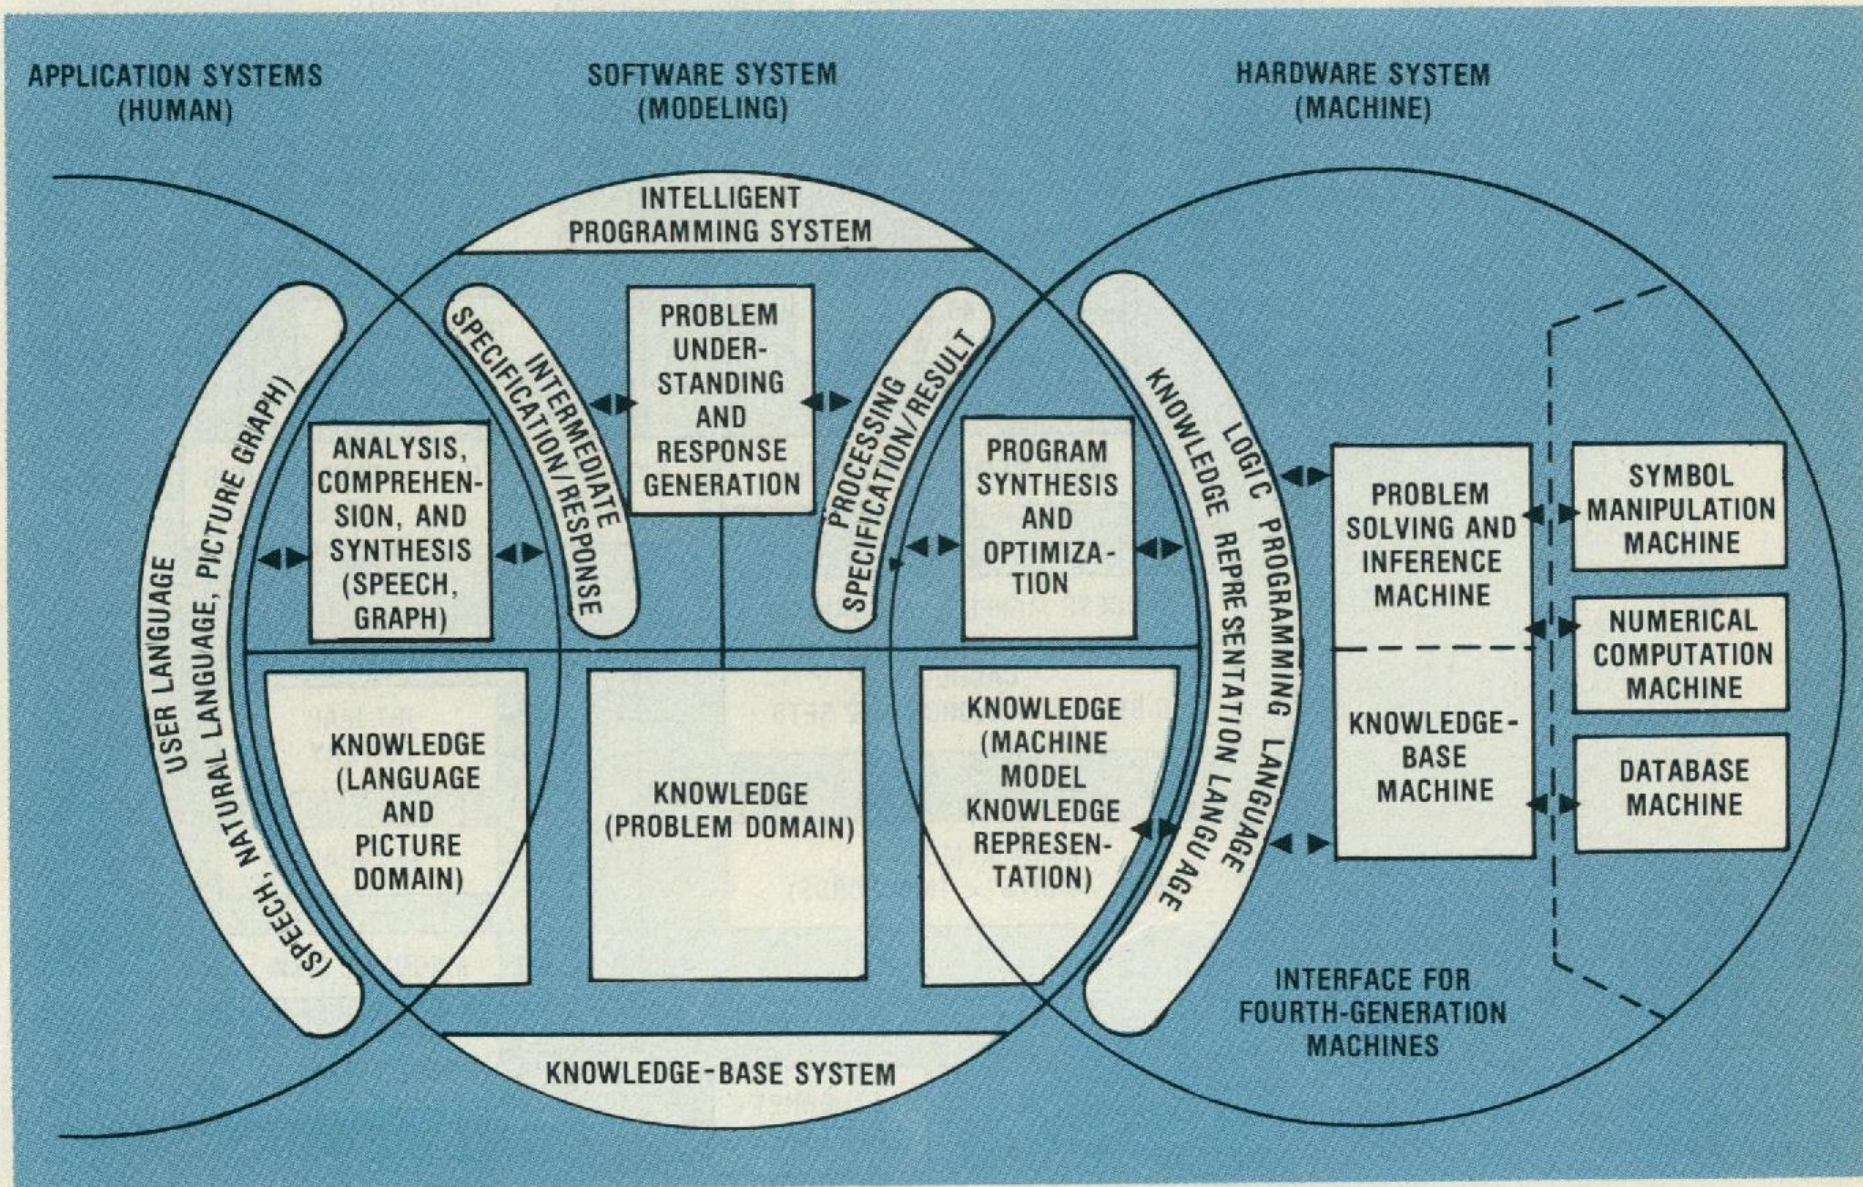
December 30, 2013

The original program was described by Joseph Weizenbaum in 1966.

This implementa

Eliza was a milestone in computer understanding of natural language. Yet Weizenbaum was more concerned with how users seemed to form an emotional relationship with the program, which consisted of nothing more than a few hundred lines of code. “I was startled to see how quickly and how very deeply people conversing with DOCTOR became emotionally involved with the computer and how unequivocally they anthropomorphized it,” he wrote. “Once my secretary, who had watched me work on the program for many months and therefore surely knew it to be merely a computer program, started conversing with it. After only a few interchanges with it, she asked me to leave the room.” He continued, “What I had not realized is that extremely short exposures to a relatively simple computer program could induce powerful delusional thinking in quite normal people.”

ELIZA: Is somethi
YOU: Men are al
ELIZA: What is th
YOU: They're al
ELIZA: Can you th
YOU: Well, my b
ELIZA: Is it impo
YOU: He says I'
ELIZA: I am sorry
YOU: It's true.
ELIZA: Can you ex
YOU: 



- Tohru Moto-Oka, Harold S. Stone: Fifth-Generation Computer Systems: A Japanese Project. Computer 17(3): 6-13 (1984)



Pandora's box?



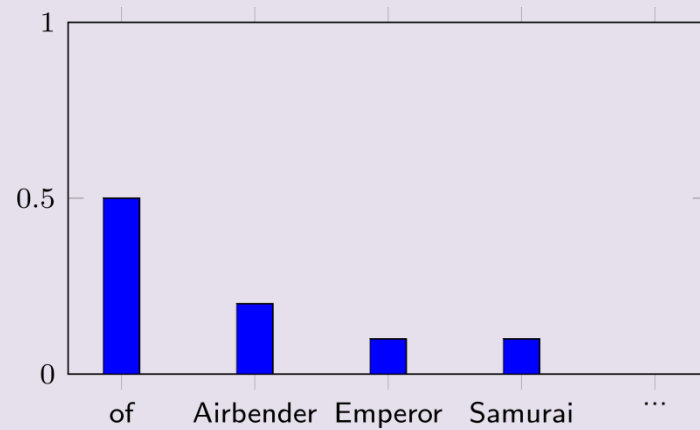




COMMON
CRAWL



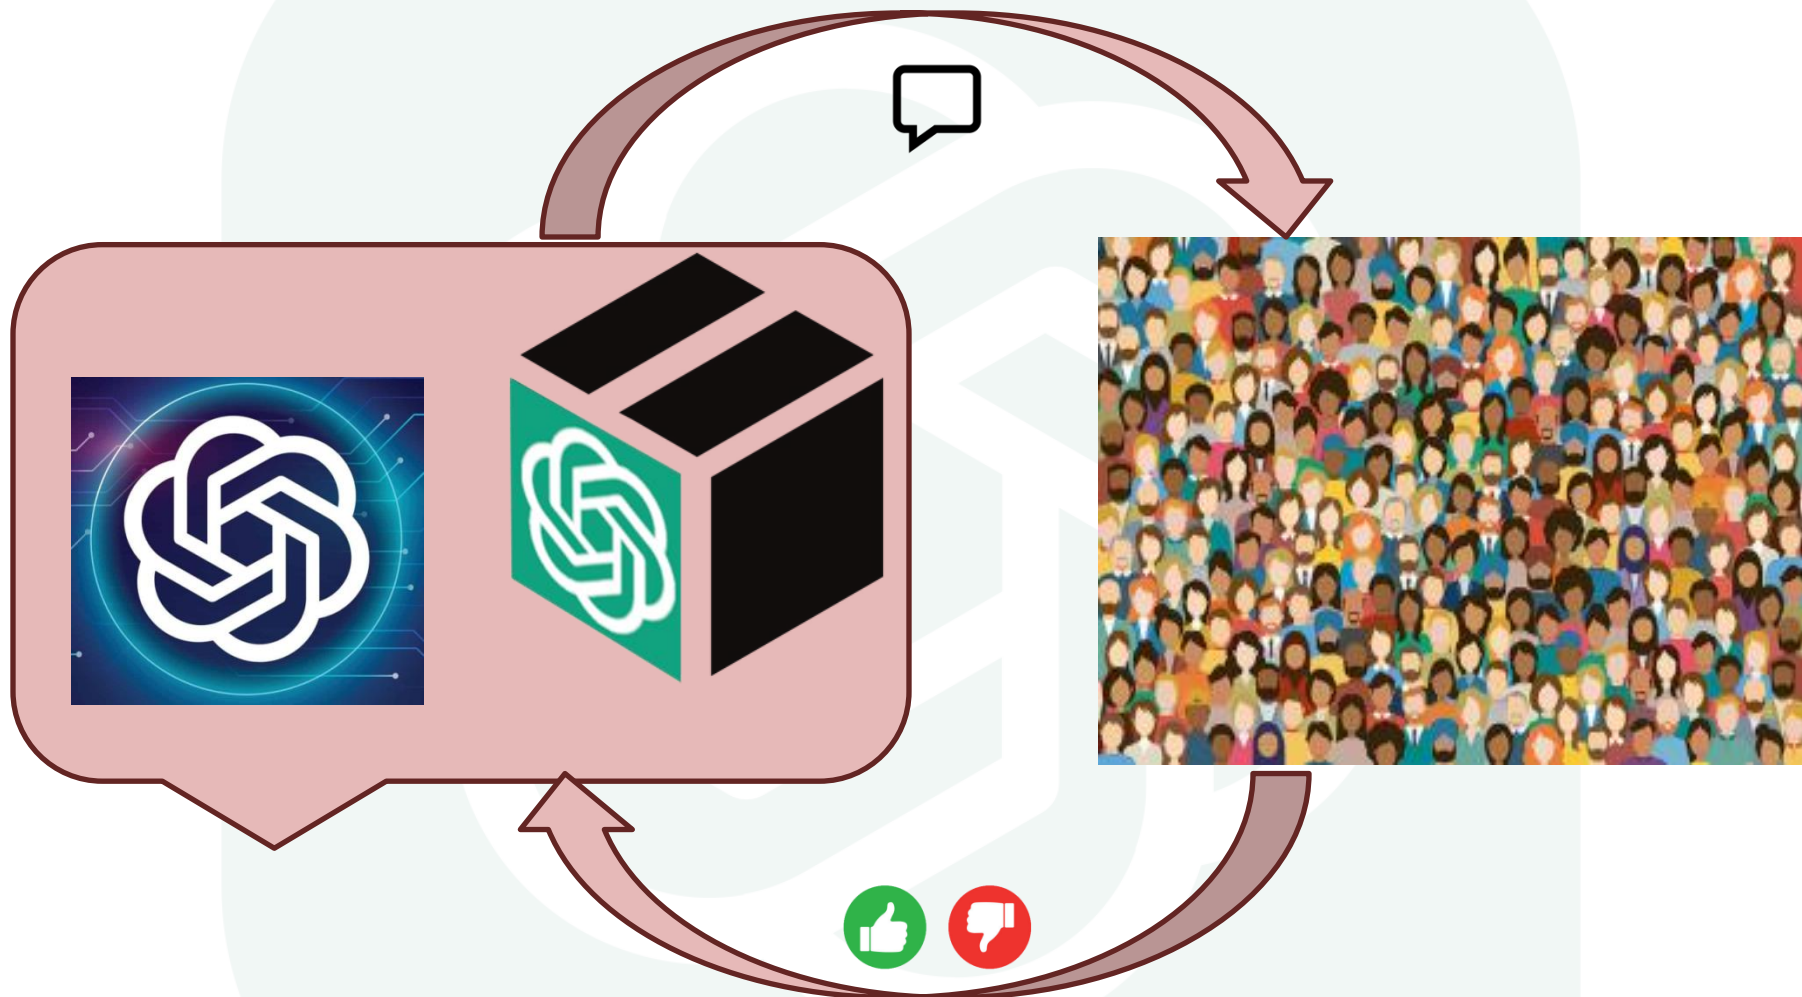
My favourite movie is The Last ____






LLAMA 2

BL  **M**





The KG is dead,
long live the LLM?



Language Models sounds the Death Knell of Knowledge Graphs

Kunal Suri
Optum, India
kunal_suri@optum.com

Atul Singh
Optum, India
atul_singh18@optum.com

Prakhar Mishra
Optum, India
prakhar_mishra29@optum.com

Swapna Sourav Rout
Optum, India
rout.swapnasourav@optum.com

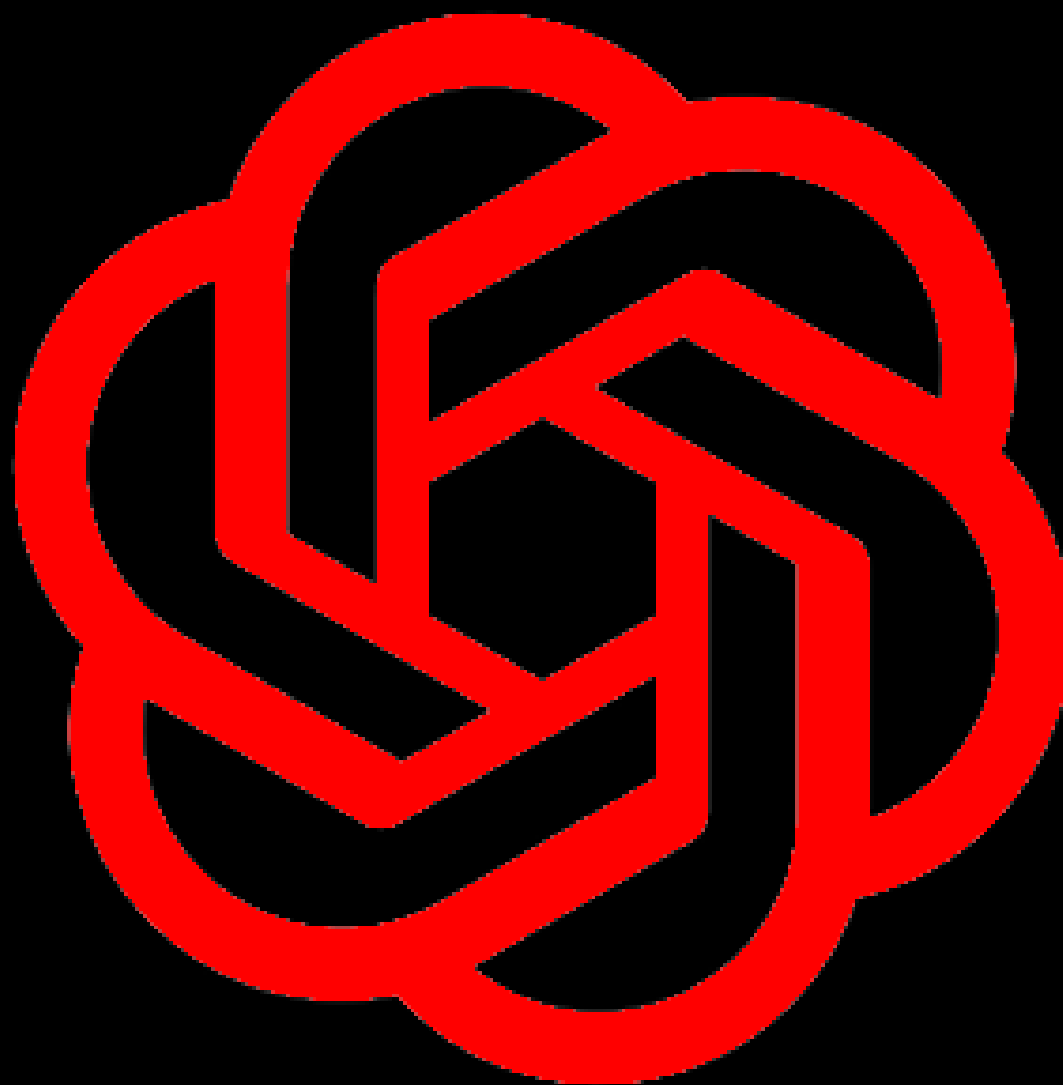
Rajesh Sabapathy
Optum, India
rajesh_sabapathy@uhc.com

Abstract—Healthcare domain generates a lot of unstructured and semi-structured text. Natural Language processing (NLP) has been used extensively to process this data. Deep Learning based NLP especially Large Language Models (LLMs) such as BERT have found broad acceptance and are used extensively for many applications. A Language Model is a probability distribution over a word sequence. Self-supervised Learning on a large corpus of data automatically generates deep learning-based language models. BioBERT and Med-BERT are language models pre-trained for the healthcare domain. Healthcare uses typical NLP tasks such as question answering, information extraction, named entity recognition, and search to simplify and improve processes. However, to ensure robust application of the results, NLP practitioners need to normalize and standardize them. One of the main ways of achieving normalization and standardization is the use of Knowledge Graphs. A Knowledge Graph captures concepts and their relationships for a specific domain, but their creation is time-consuming and requires manual intervention from domain experts, which can prove expensive. SNOMED CT (Systematized Nomenclature of Medicine - Clinical Terms), Unified Medical Language System (UMLS), and Gene Ontology (GO) are popular ontologies from the healthcare domain. SNOMED CT and UMLS capture concepts such as disease,

relationships among different entities makes knowledge graphs redundant for many applications.

This paper uses similar terms from SNOMED-CT KG and passes them through a language model for the healthcare domain BioRedditBERT to get a 768-dimensional dense vector representation. The paper presents the results for analyzing these embeddings. The experiments presented in the paper validate that similar terms cluster together. The paper uses simple heuristics to assign names to clusters. The results show that the cluster names match the names in the KG. Finally, the experiments demonstrate that the cosine similarity of vector representation of similar terms is high and vice versa.

Our contributions include: (i) We propose a study to demonstrate the value and application of Large Language Models (LLMs) in comparison to Knowledge Graph-based approaches for the task of synonym extraction. (ii) We extensively evaluate our approach on a standard, widely accepted dataset, and the results are encouraging.



A



Which Turing Award winners were born in Latin America?

As of my last update in September 2021, Adolfo Guzmán-Arenas is a Turing Award winner born in Latin America. He was born in Guatemala. He received the Turing Award in 2002 along with his colleague Manuel Blum for their work on computational complexity theory and the development of the concept of a "proof system."

Please note that my information might not be current, so I recommend verifying with the latest sources to confirm the most up-to-date information about Turing Award winners from Latin America.

A



Which Turing Award winners were born in Latin America?

As of my last update in September 2021, there are no Turing Award winners who were born in Latin America. The Turing Award is often considered the "Nobel Prize of Computer Science" and is awarded annually by the Association for Computing Machinery (ACM). The recipients have primarily been from the United States, Europe, and Asia, reflecting in part the historical investment and focus on computer science research in these regions.

However, it's important to note that the field of computer science is increasingly global, and outstanding contributions are being made by researchers from all around the world, including Latin America. While there have not been Turing Award winners born in Latin America as of my last update, this could change in the future as the field continues to evolve and diversify.

Please verify this information for the most current data, as it might have changed after my last training data in September 2021.

A



Can you give me a SPARQL query for Wikidata to find Turing Award winners who were born in Latin America?

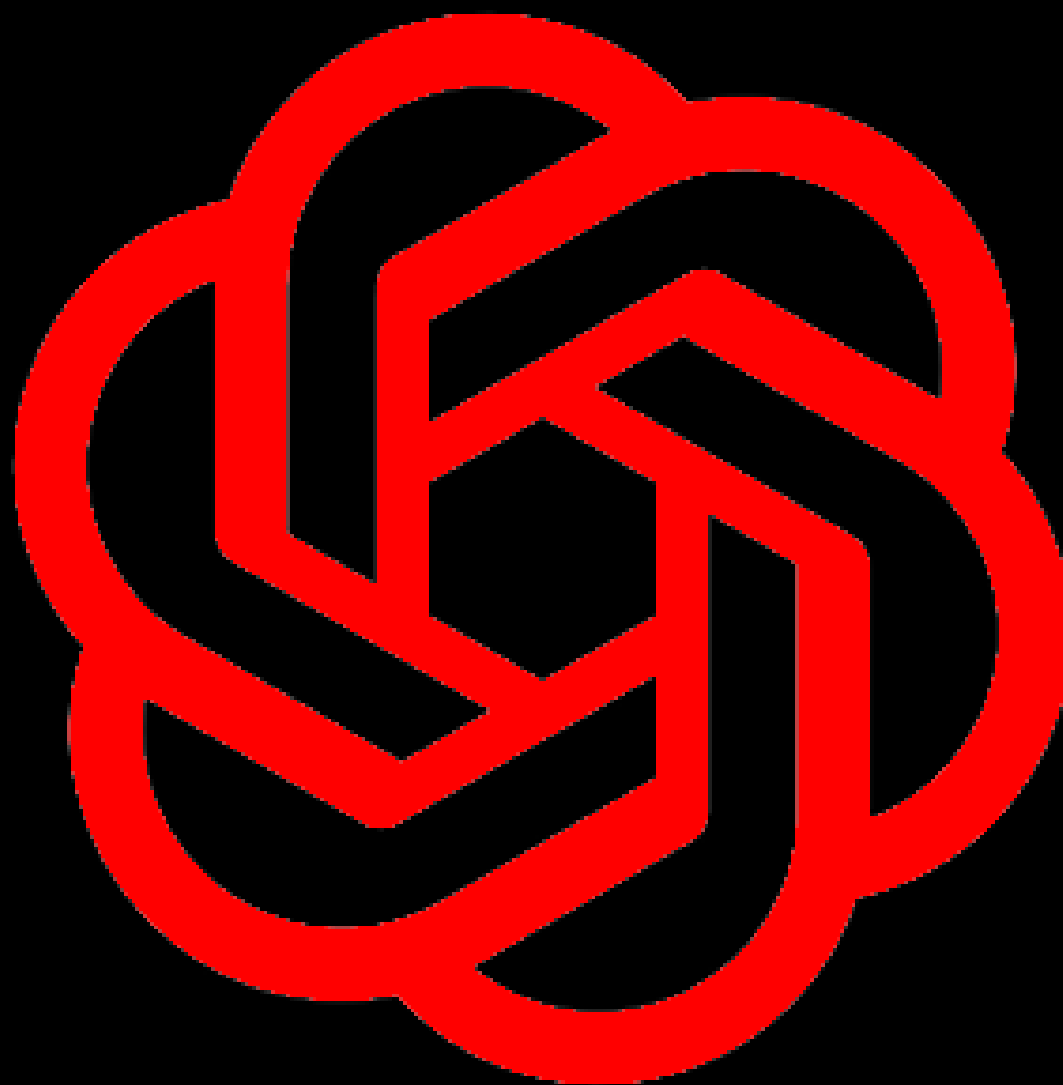
Certainly!

```
# Turing Award winners born in Latin America
SELECT ?person ?personLabel ?birthPlace ?birthPlaceLabel ?award ?awardLabel WHERE {
  ?person p:P166 ?statement.
  ?statement ps:P166 wd:Q7251; pq:P582 ?awardDate. # Turing Award
  ?person wdt:P19 ?birthPlace. # Place of birth
  ?birthPlace wdt:P30 wd:Q12521. # Latin America
  SERVICE wikibase:label { bd:serviceParam wikibase:language "[AUTO_LANGUAGE]". }
}
ORDER BY ?awardDate
```



LLMs + KGs: A dream team?







LLM

KG

Hallucinations



Opaqueness:



Intransigence:



Staleness:



Bias:



Costs:



Short tail:



Sanitised:



Non-determinism:




































Usability:



Factual only:

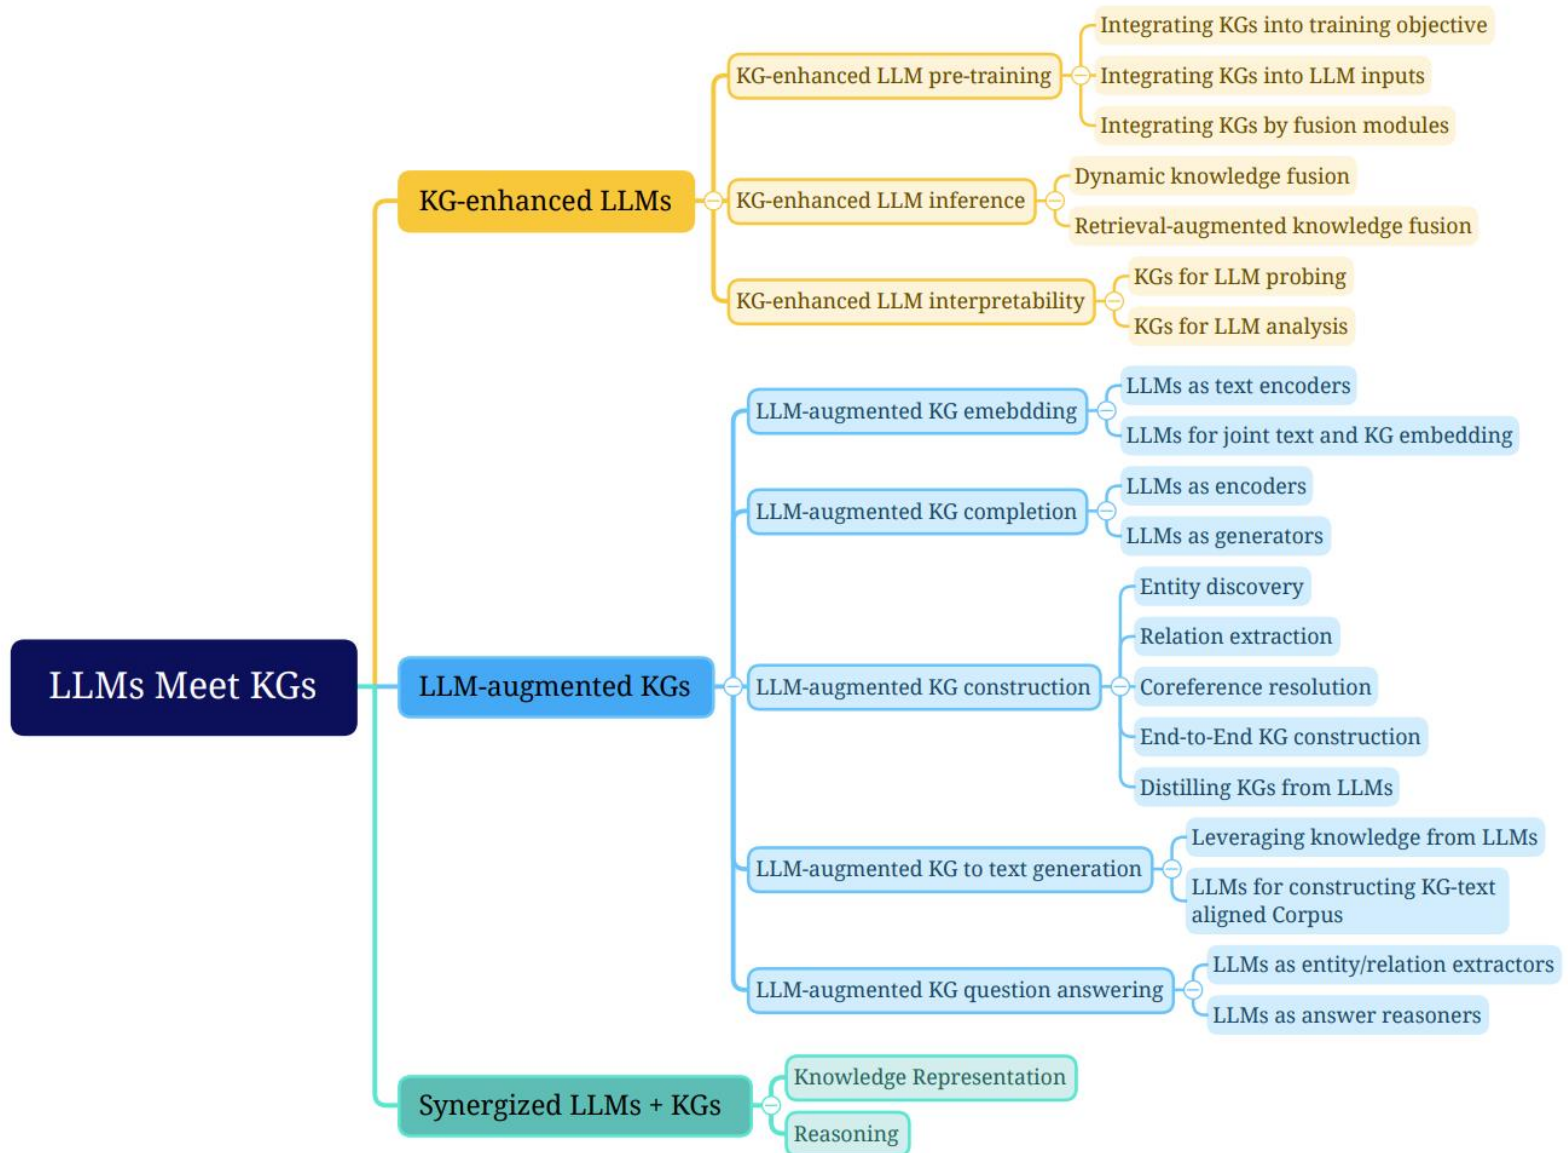


	LLM	KG	KG + LLM
Hallucinations			
Opaqueness:			
Intransigence:			
Staleness:			
Bias:			
Costs:			
Short tail:			
Sanitised:			
Non-determinism:			
Usability:			
Factual only:			

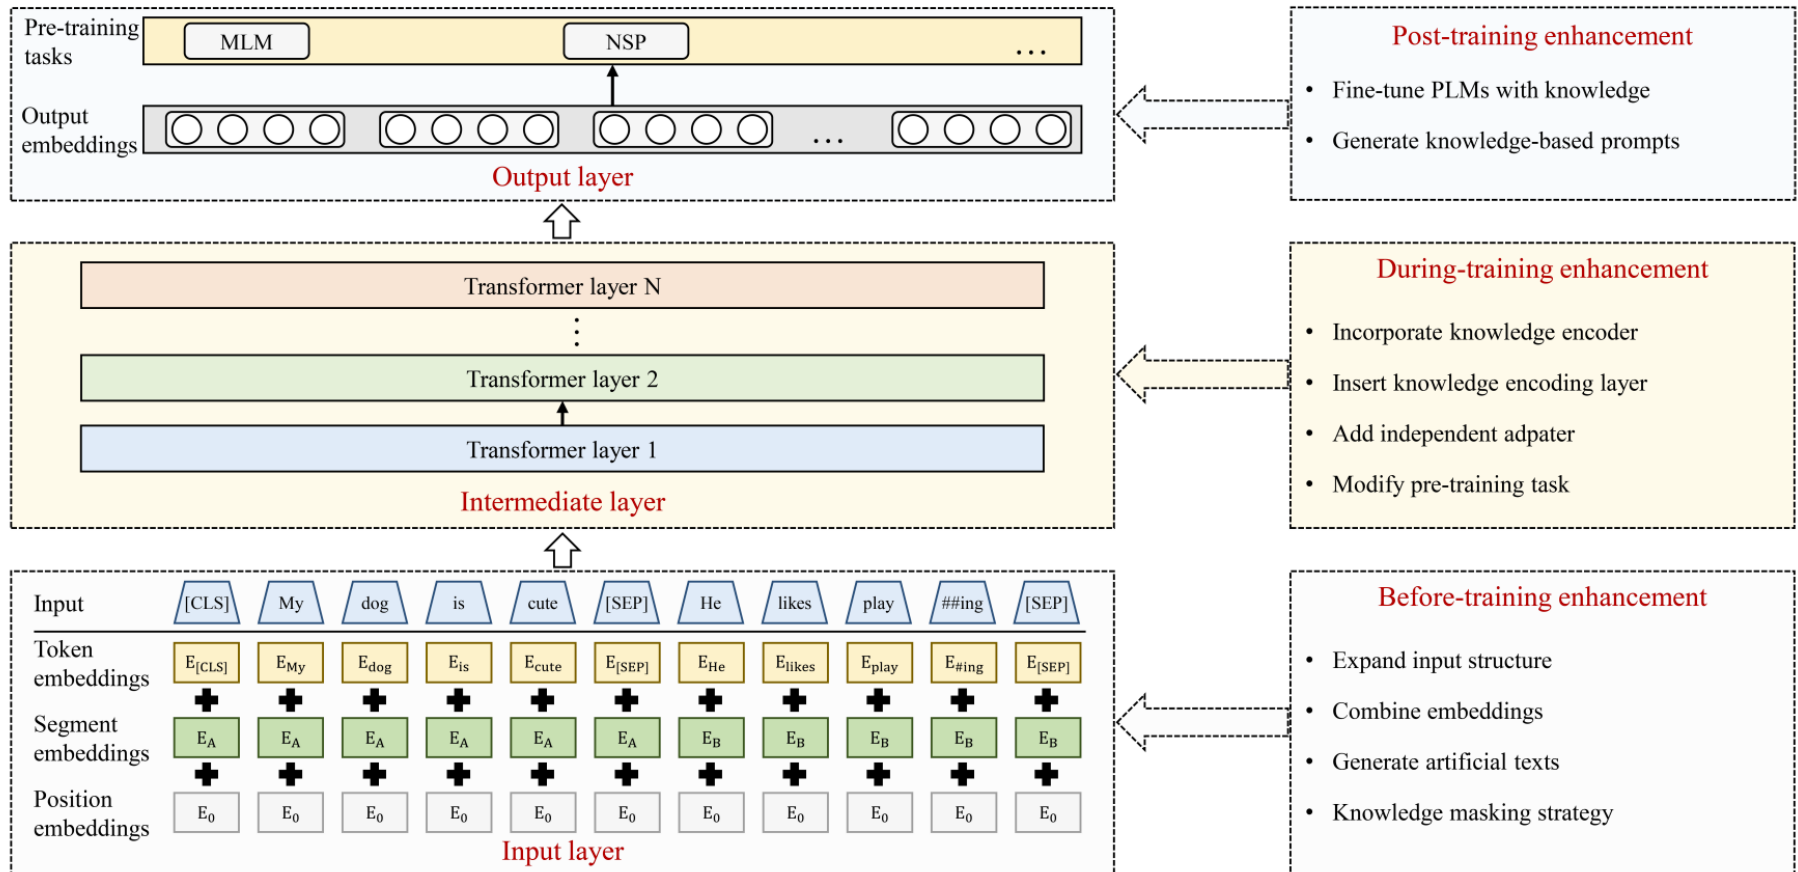


KGs + LLMS: Ongoing Work





- Shirui Pan, Linhao Luo, Yufei Wang, Chen Chen, Jiapu Wang, Xindong Wu: Unifying Large Language Models and Knowledge Graphs: A Roadmap. CoRR abs/2306.08302 (2023).



- Linyao Yang, Hongyang Chen, Zhao Li, Xiao Ding, Xindong Wu: ChatGPT is not Enough: Enhancing Large Language Models with Knowledge Graphs for Fact-aware Language Modeling. CoRR abs/2306.11489 (2023).

A decorative graphic in the top-right corner consisting of a grid of small black dots. Overlaid on this grid are several colorful, stylized, and somewhat abstract shapes in shades of blue, green, and purple, resembling a corner bracket or a stylized 'L' shape.

LLMs can make KGs more usable

A decorative graphic in the bottom-left corner, similar to the one in the top-right, featuring a grid of small black dots with colorful, stylized shapes in shades of blue, green, and purple overlaid.

A



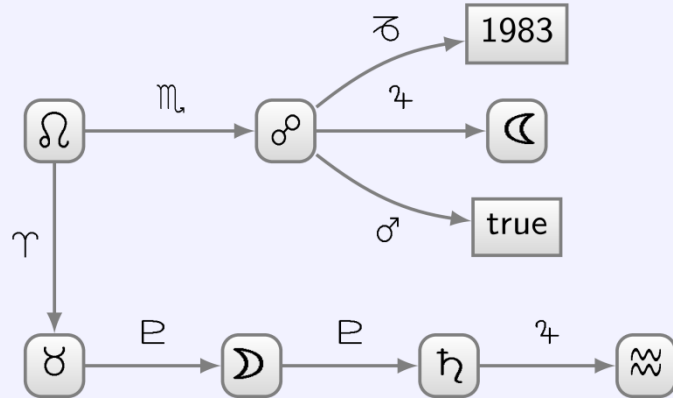
Can you give me a SPARQL query for Wikidata to find Turing Award winners who were born in Latin America?

Certainly!

```
# Turing Award winners born in Latin America
SELECT ?person ?personLabel ?birthPlace ?birthPlaceLabel ?award ?awardLabel WHERE {
  ?person p:P166 ?statement.
  ?statement ps:P166 wd:Q7251; pq:P582 ?awardDate. # Turing Award
  ?person wdt:P19 ?birthPlace. # Place of birth
  ?birthPlace wdt:P30 wd:Q12521. # Latin America
  SERVICE wikibase:label { bd:serviceParam wikibase:language "[AUTO_LANGUAGE]". }
}
ORDER BY ?awardDate
```



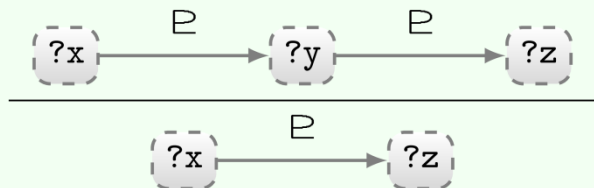
DATA



QUERY



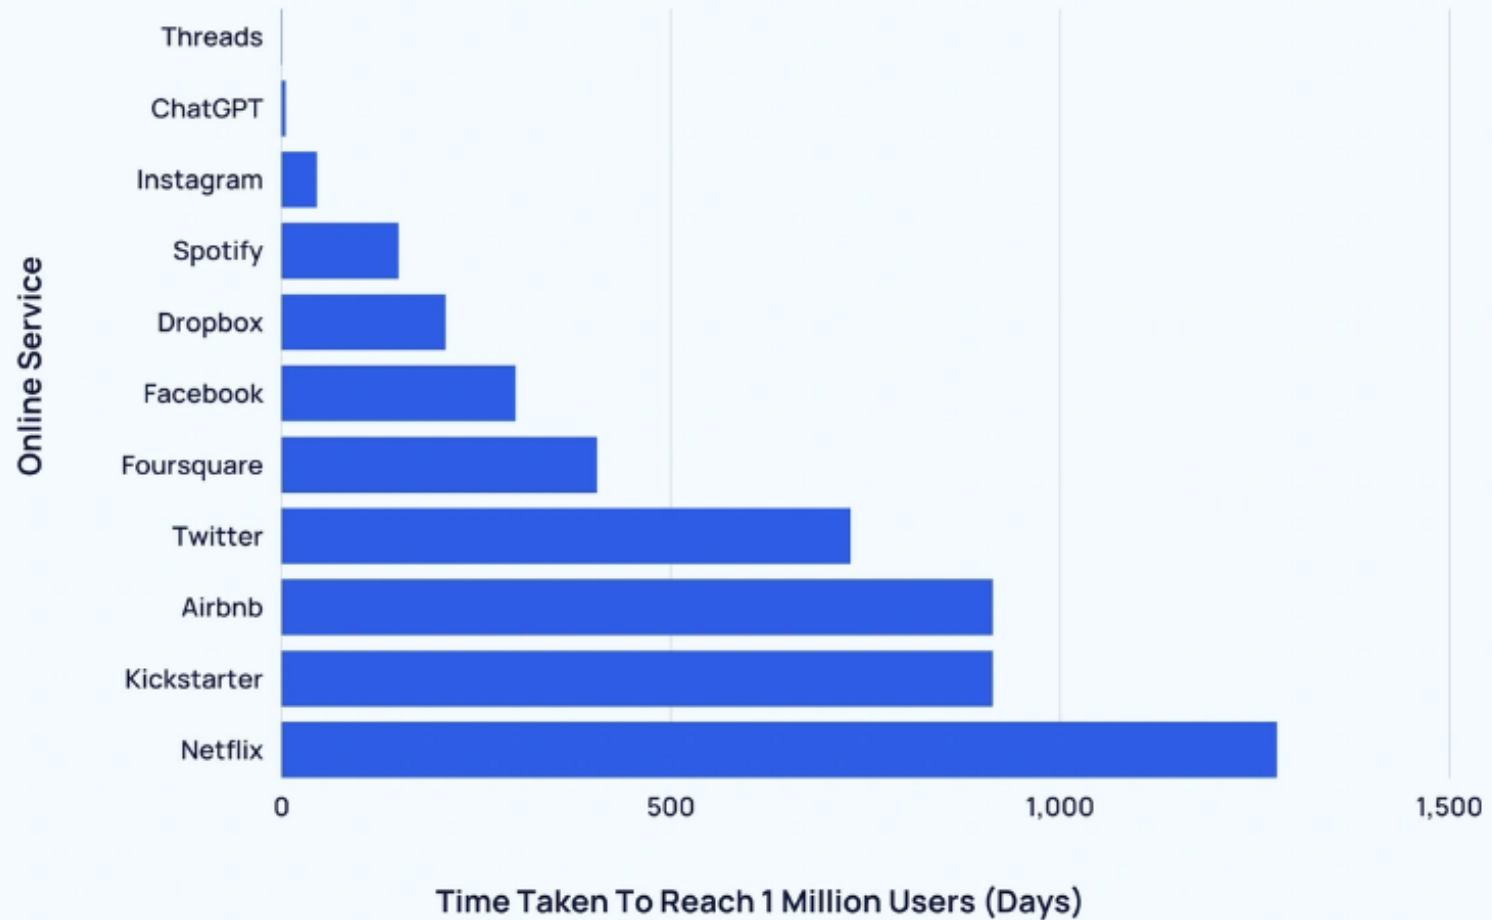
LOGIC



LANGUAGE



Time taken to reach 1 million users



Conclusions

- Neural vs. symbolic an old debate, but why not both?
- LLMs solve some tasks that KGs do (and do not)
- KGs solve some tasks that LLMs do (and do not)
- LLMs unlikely to replace KGs (and vice versa)
- LLMs + KGs = A fertile research topic

LLMs can make KGs more usable
and bring them to the masses
(LLMs will likely take the credit, but that's okay)



Shameless plugs



K N O W L E D G E G R A P H S

Aidan Hogan

IMFD, DCC, Universidad de Chile

Eva Blomqvist

Linköping University

Michael Cochez

Vrije Universiteit Amsterdam and Discovery Lab, Elsevier

Claudia d'Amato

University of Bari

<https://kgbook.org/>

Give us feedback!



<https://dagstuhl.de/tgdk>

Transactions on Graph Data and Knowledge

ISSN 2942-7517

Publications

[General](#)


[Editorial Board](#)

[Selection/Review](#)

[Author Instructions](#)

[Editor Instructions](#)

Volumes

 Volume

Transactions on Graph Data and Knowledge, Volume 1

[TGDK, Volume 1](#)

- [Issue 1](#) - Special Issue on Trends in Graph Data and Knowledge (2023)

News

 2023-11-30 [TGDK - Call for Speci...](#)

 2023-11-22 [TGDK - Call For Papers](#)

 2023-11-21 [New Journal at Dagst...](#)

References

- Agarwal+21:** Oshin Agarwal, Heming Ge, Siamak Shakeri, Rami Al-Rfou: Knowledge Graph Based Synthetic Corpus Generation for Knowledge-Enhanced Language Model Pre-training. NAACL-HLT 2021: 3554-3565
- Ahmed+23** Abdullah Fathi Ahmed, Asep Fajar Firmansyah, Mohamed Ahmed Sherif, Diego Moussallem, and Axel-Cyrille Ngonga Ngomo Explainable Integration of Knowledge Graphs using Large Language Models. NLDB 2023: 124-139
- AxelssonS2023*** Agnes Axelsson, Gabriel Skantze Using Large Language Models for Zero-Shot Natural Language Generation from Knowledge Graphs. CoRR abs/2307.07312 (2023).
- Cenikj+23** Gjorgjina Cenikj, Lidija Strojnik, Risto Angelski, Nives Ogrinc, Barbara Koroušić Seljak, Tome Eftimov: From language models to large-scale food and biomedical knowledge graphs. Scientific Reports, 2023.
- ChoudharyR2023*** Nurendra Choudhary, Chandan K. Reddy: Complex Logical Reasoning over Knowledge Graphs using Large Language Models. CoRR abs/2305.01157 (2023).
- Colon+21*** Pedro Colon-Hernandez, Catherine Havasi, Jason B. Alonso, Matthew Huggins, Cynthia Breazeal: Combining pre-trained language models and structured knowledge. CoRR abs/2101.12294 (2021).
- Fei+20** Hao Fei, Yafeng Ren, Yue Zhang, Donghong Ji and Xiaohui Liang: Enriching contextualized language model from knowledge graph for biomedical information extraction. Briefings Bioinform. 22(3) (2021)
- Hao+23** Shibo Hao, Bowen Tan, Kaiwen Tang, Bin Ni, Xiyan Shao, Hengzhe Zhang, Eric P. Xing, Zhiting Hu: BertNet: Harvesting Knowledge Graphs with Arbitrary Relations from Pretrained Language Models. ACL (Findings) 2023: 5000-5015

- Hu+22** Ziniu Hu, Yichong Xu, Wenhao Yu, Shuohang Wang, Ziyi Yang, Chenguang Zhu, Kai-Wei Chang, Yizhou Sun Empowering Language Models with Knowledge Graph Reasoning for Question Answering. EMNLP 2022: 9562-9581
- JSun+23*** Jiashuo Sun, Chengjin Xu, Lumingyuan Tang, Saizhuo Wang, Chen Lin, Yeyun Gong, Heung-Yeung Shum, Jian Guo Think-on-Graph: Deep and Responsible Reasoning of Large Language Model with Knowledge Graph. CoRR abs/2307.07697 (2023).
- Kin+20** Bosung Kim, Taesuk Hong, Youngjoong Ko, Jungyun Seo: Multi-Task Learning for Knowledge Graph Completion with Pre-trained Language Models. COLING 2020: 1737-1743
- KSun+23*** Kai Sun, Yifan Ethan Xu, Hanwen Zha, Yue Liu, Xin Luna Dong: Head-to-Tail: How Knowledgeable are Large Language Models (LLM)? A.K.A. Will LLMs Replace Knowledge Graphs? CoRR abs/2308.10168 (2023).
- Lehmann+23*** Jens Lehmann, Preetam Gattogi, Dhananjay Bhandiwad, Sébastien Ferré, Sahar Vahdati: Language Models as Controlled Natural Language Semantic Parsers for Knowledge Graph Question Answering. Amazon TR, 2023.
- Li+21** Junyi Li, Tianyi Tang, Wayne Xin Zhao, Zhicheng Wei, Nicholas Jing Yuan, Ji-Rong Wen: Few-shot Knowledge Graph-to-Text Generation with Pretrained Language Models. ACL/IJCNLP (Findings) 2021: 1558-1568
- LiuW2023*** Chang Liu, Bo Wu: Evaluating large language models on graphs: performance insights and comparative analysis. CoRR abs/2308.11224 (2023).
- Logan+19** Robert L. Logan IV, Nelson F. Liu, Matthew E. Peters, Matt Gardner, Sameer Singh: Barack's Wife Hillary: Using Knowledge Graphs for Fact-Aware Language Modeling. ACL (1) 2019: 5962-5971.

Meyer+23* Lars-Peter Meyer, Claus Stadler, Johannes Frey, Norman Radtke, Kurt Junghanns, Roy Meissner, Gordian Dziwis, Kirill Bulert, and Michael Martin LLM-assisted Knowledge Graph Engineering: Experiments with Chat-GPT. CoRR abs/2307.06917 (2023)

Mruthyunjaya+23* Vishwas Mruthyunjaya, Pouya Pezeshkpour, Estevam Hruschka, Nikita Bhutani: Rethinking Language Models as Symbolic Knowledge Graph. CoRR abs/2308.13676 (2023).

Omar+23* Reham Omar, Omij Mangukiya, Panos Kalnis, Essam Mansour: Chat-GPT versus Traditional Question Answering for Knowledge Graphs: Current Status and Future Directions Towards Knowledge Graph Chatbots. CoRR abs/2302.06466 (2023).

Pan+23 Jeff Z. Pan, Simon Razniewski, Jan-Christoph Kalo, Sneha Singhania, Jiaoyan Chen, Stefan Dietze, Hajira Jabeen, Janna Omeliyanenko, Wen Zhang, Matteo Lissandrini, Russa Biswas, Gerard de Melo, Angela Bonifati, Edlira Vakaj, Mauro Dragoni, Damien Graux Large Language Models and Knowledge Graphs: Opportunities and Challenges. TGDK 2023.

Pan+23* Shirui Pan, Linhao Luo, Yufei Wang, Chen Chen, Jiapu Wang, Xindong Wu: Unifying Large Language Models and Knowledge Graphs: A Roadmap. CoRR abs/2306.08302 (2023).

Trajanoska+23* Milena Trajanoska, Riste Stojanov, Dimitar Trajanov: Enhancing Knowledge Graph Construction Using Large Language Models. CoRR abs/2305.04676 (2023).

Vassiliou+23 Giannis Vassiliou, Nikolaos Papadakis, Haridimos Kondylakis: SummaryGPT: Leveraging ChatGPT for Summarizing Knowledge Graphs, ESWC Demo 2023.

- Verga+21** Pat Verga, Haitian Sun, Livio Baldini Soares, William W. Cohen: Adaptable and Interpretable Neural Memory Over Symbolic Knowledge. NAACL-HLT 2021: 3678-3691
- Wang+20*** Chenguang Wang, Xiao Liu, Dawn Song: Language models are open knowledge graphs. CoRR abs/2010.11967 (2020).
- YangaT20** Xinran Yanga, Ilaria Tiddi Creative Storytelling with Language Models and Knowledge Graphs, CIKM (Workshops) 2020.
- Yang+23*** Linyao Yang, Hongyang Chen, Zhao Li, Xiao Ding, Xindong Wu: ChatGPT is not Enough: Enhancing Large Language Models with Knowledge Graphs for Fact-aware Language Modeling. CoRR abs/2306.11489 (2023).
- Yao+23*** Liang Yao, Jiazhen Peng, Chengsheng Mao, Yuan Luo: Exploring Large Language Models for Knowledge Graph Completion. CoRR abs/2308.13916 (2023).
- Yasunaga+21** Michihiro Yasunaga, Hongyu Ren, Antoine Bosselut, Percy Liang, Jure Leskovec: QA-GNN: Reasoning with Language Models and Knowledge Graphs for Question Answering. NAACL-HLT 2021: 535-546
- Yasunaga+22** Michihiro Yasunaga, Antoine Bosselut, Hongyu Ren, Xikun Zhang, Christopher D Manning, Percy Liang, Jure Leskovec: Deep Bidirectional Language-Knowledge Graph Pretraining, NeurIPS 2022.
- Yu+22** Donghan Yu, Chenguang Zhu, Yiming Yang, Michael Zeng: JAKET: Joint Pre-training of Knowledge Graph and Language Understanding. AAAI 2022: 11630-11638.
- Zhang+23*** Bohui Zhang, Ioannis Reklos, Nitisha Jain, Albert Meroño Peñuela, Elena Simperl: Using Large Language Models for Knowledge Engineering (LLMKE): A Case Study on Wikidata. CoRR abs/2309.08491 (2023).

Zhu+23* Yuqi Zhu, Xiaohan Wang, Jing Chen, Shuofei Qiao, Yixin Ou, Yunzhao Yao, Shumin Deng, Huajun Chen, Ningyu Zhang: LLMs for Knowledge Graph Construction and Reasoning: Recent Capabilities and Future Opportunities. CoRR abs/2305.13168 (2023).