

Knowledge Graphs as a backbone for Bilateral AI and Intelligent Data Management

Axel Polleres

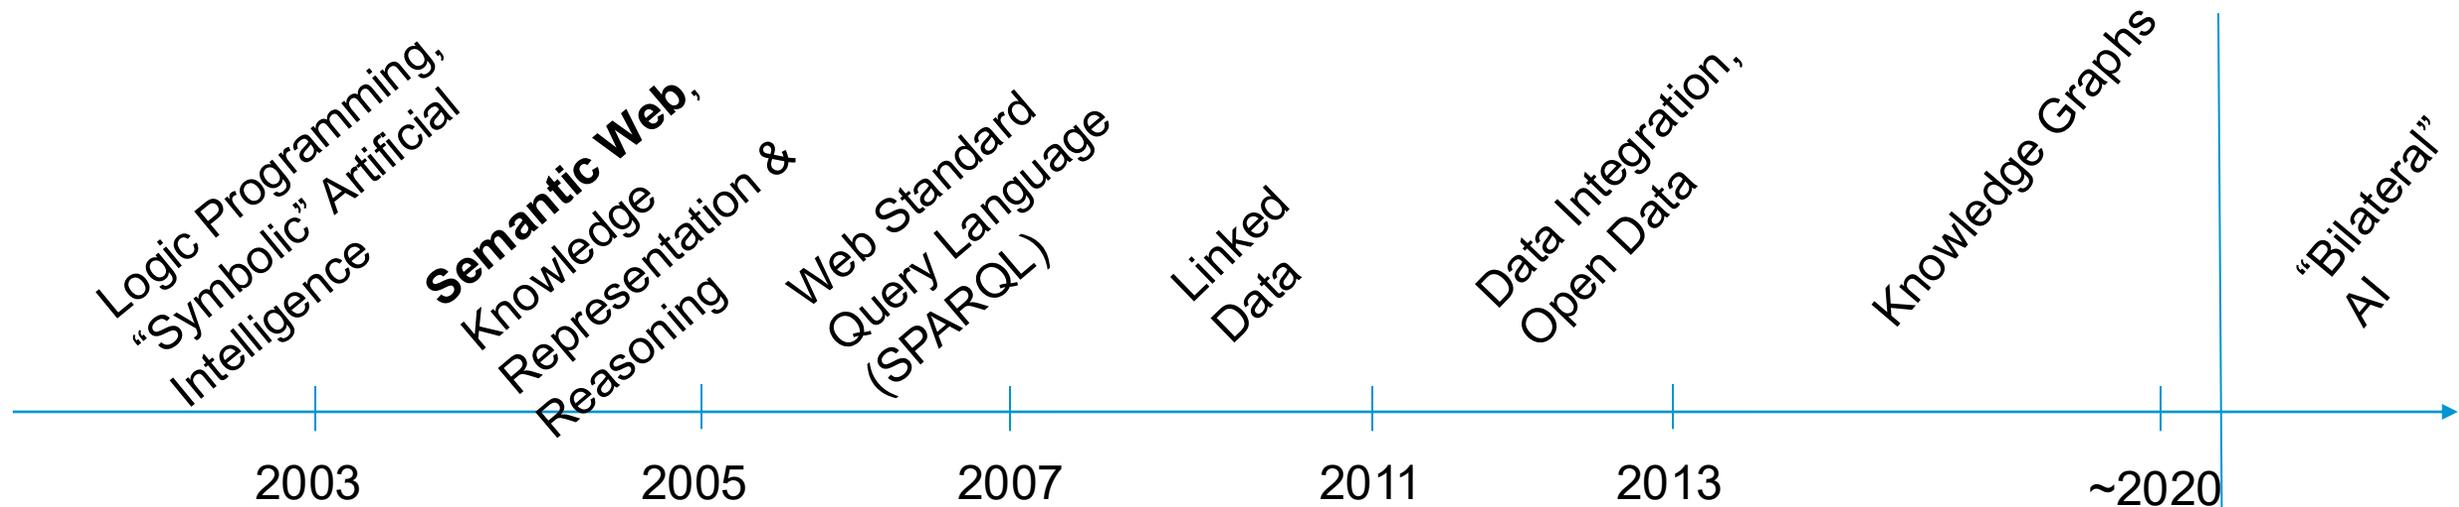
WU
VIENNA

 **EQUIS**
ACCREDITED

 **AACSB**
ACCREDITED

 **ASSOCIATION
AMBA**
ACCREDITED

My background...



TU Vienna



Univ. Innsbruck



Univ. Rey Juan Carlos
Madrid



DERI, NUI Galway,
Ireland



Siemens AG
Österreich



WU Vienna



Instead of its initial focus on **agents** the “Semantic Web” ...



The Semantic Web
A new form of Web content that is meaningful to computers will unleash a revolution of new possibilities
by [TIM BERNERS-LEE, JAMES HENDLER and ORA LASSILA](#)

The entertainment system was belting out the Beatles' "We Can Work It Out" when the phone rang. When Pete answered, his phone turned the sound down by sending a message to all the other *local* devices that had a *volume control*. His sister, Lucy, was on the line from the doctor's office: "Mom needs to see a specialist and then has to have a series of physical therapy sessions. Biweekly or something. I'm going to have my agent set up the appointments." Pete immediately agreed to share the chauffeuring. At the doctor's office, Lucy instructed her Semantic Web agent through her handheld Web browser. The agent promptly retrieved information about Mom's *prescribed treatment* from the doctor's agent, looked up several lists of *providers*, and checked for the ones *in-plan* for Mom's insurance within a *20-mile radius* of her home and with a *rating of excellent or very good* on trusted rating services. It then began trying to find a match between available *appointment times* (supplied by the agents of individual providers through their Web sites) and Pete's and Lucy's busy schedules. (The emphasized keywords indicate terms whose semantics, or meaning, were defined for the agent through the Semantic Web.)

In a few minutes the agent presented them with a plan. Pete didn't like it—University Hospital was all the way across town from Mom's place, and he'd be driving back in the middle of rush hour. He set his own agent to redo the search with stricter preferences about *location* and *time*. Lucy's agent, having *complete trust* in Pete's agent in the context of the present task, automatically assisted by supplying access certificates and shortcuts to the data it had already sorted through.



*“[...] The **agent** promptly retrieved information about Mom's prescribed treatment from the doctor's agent, looked up several lists of providers, and checked for the ones in-plan for Mom's insurance within a 20-mile radius of her home and with a rating of excellent or very good on trusted rating services[...]*”

- *Appointment detection in emails*
- *Semantic Search*
- *Ratings of products/services*

... has then mostly become the basis for the "Web of Data" ...

*"If HTML and the Web made all the online documents look like one huge **book**, RDF, schema and inference languages will make all the data in the world look like **one huge database**"*

Tim Berners-Lee, Weaving the Web, 1999

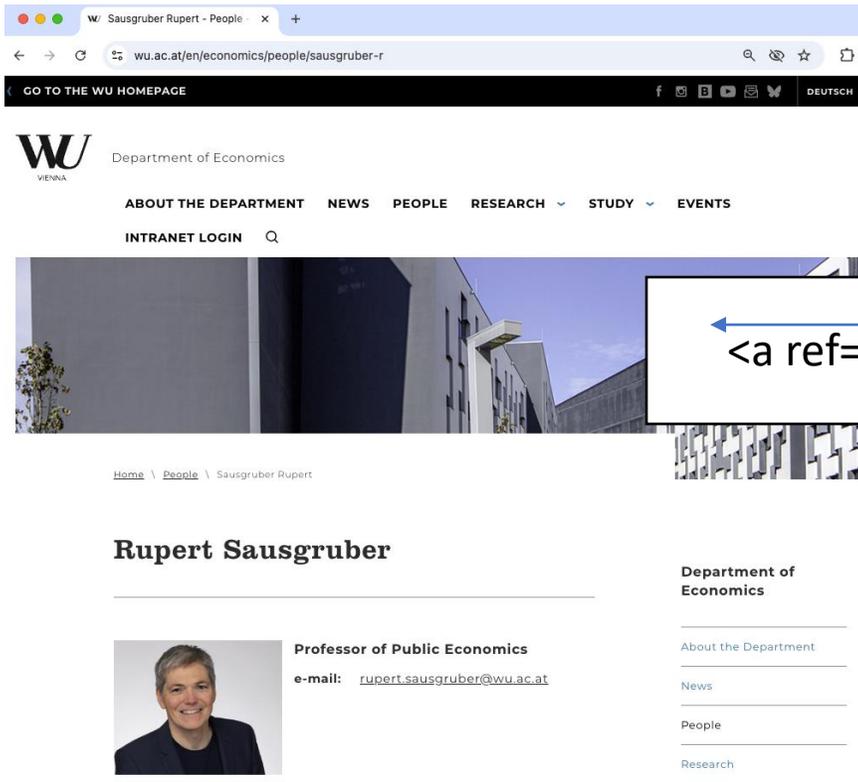
... and its more recent focus on Open Knowledge Graphs...

*" huge knowledge bases, also known as **knowledge graphs**, have been automatically constructed from web data, and have become a key asset for search engines and other use cases.*

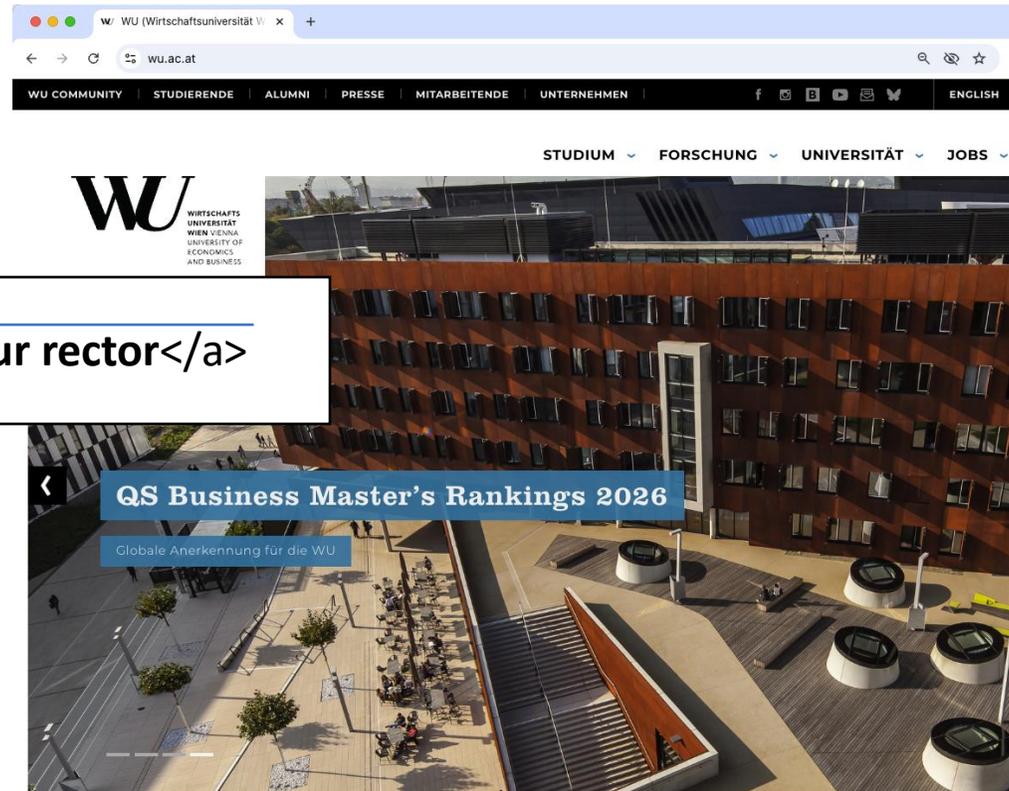
Gerhard Weikum, Knowledge Graphs 2021: A Data Odyssey, VDLB 2021

Phase 1: RDF – Linked “Metadata”

- “Typed links” describing Ressources (such as Web pages)



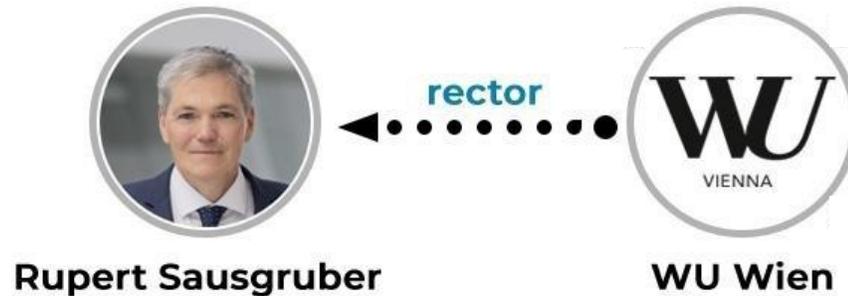
The screenshot shows a browser window with the URL `wu.ac.at/en/economics/people/sausgruber-r`. The page header includes the WU logo and navigation links for 'ABOUT THE DEPARTMENT', 'NEWS', 'PEOPLE', 'RESEARCH', 'STUDY', and 'EVENTS'. The main content area features a profile for **Rupert Sausgruber**, identified as a **Professor of Public Economics** with an email address `rupert.sausgruber@wu.ac.at`. A sidebar on the right lists 'Department of Economics' and links to 'About the Department', 'News', 'People', and 'Research'.



The screenshot shows the WU homepage with the URL `wu.ac.at`. A white box with a blue arrow pointing left is overlaid on the page, containing the HTML code `our rector`. Below this, a blue banner reads 'QS Business Master's Rankings 2026' with the subtext 'Globale Anerkennung für die WU'. The background features an aerial view of the WU building.

Phase 1: RDF (Resource Description framework) – Linked “Metadata”

- Typed links “triples” describing Ressources (such as Web pages)

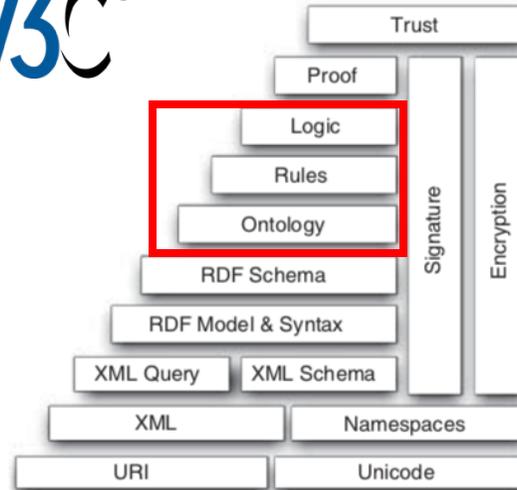
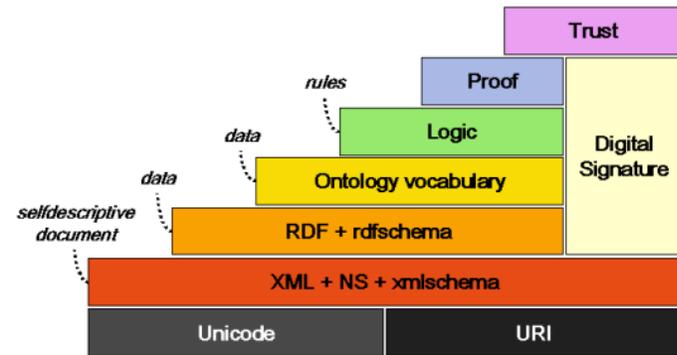


<http://www.wu.ac.at#wuwien> <http://www.wu.ac.at#hasRector> http://www.wu.ac.at/Rupert_Sausgruber

Semantic Web: Standard formats, Reasoning & Logics

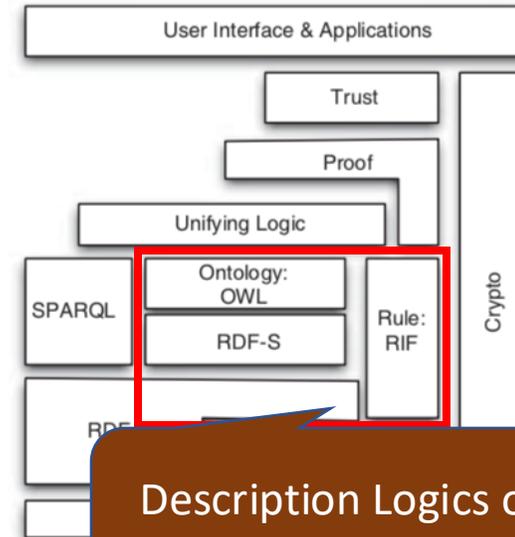


- (2000s - ca. 2009)

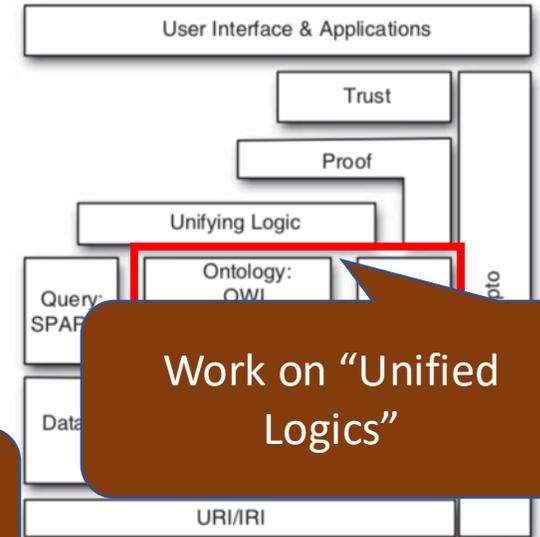


2004^a

^ahttp://www.w3.org/2004/Talks/0319-RDF-WGs/sw_stack...
^b<http://www.w3.org/2007/Talks/0130-sb-W3CTechSemWeb/layerCake-4.png>



Description Logics or Rules?



Work on "Unified Logics"

(Contextualized Reasoning?)

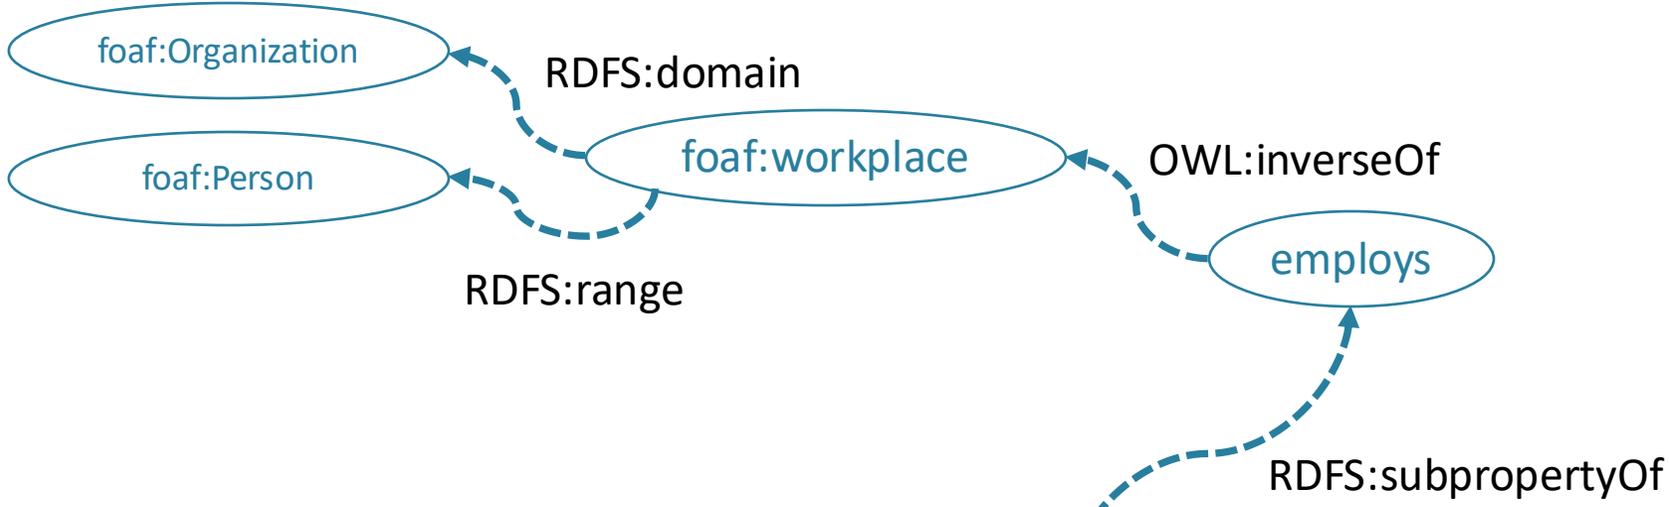
2009^c

Good news! Standards meant a Boost in KR/AI research: We know very well which ontological reasoning approaches are decidable and how they scale → RDF, OWL, SPARQL Ontology-based Data Access but also: constraint checking (SHACL)

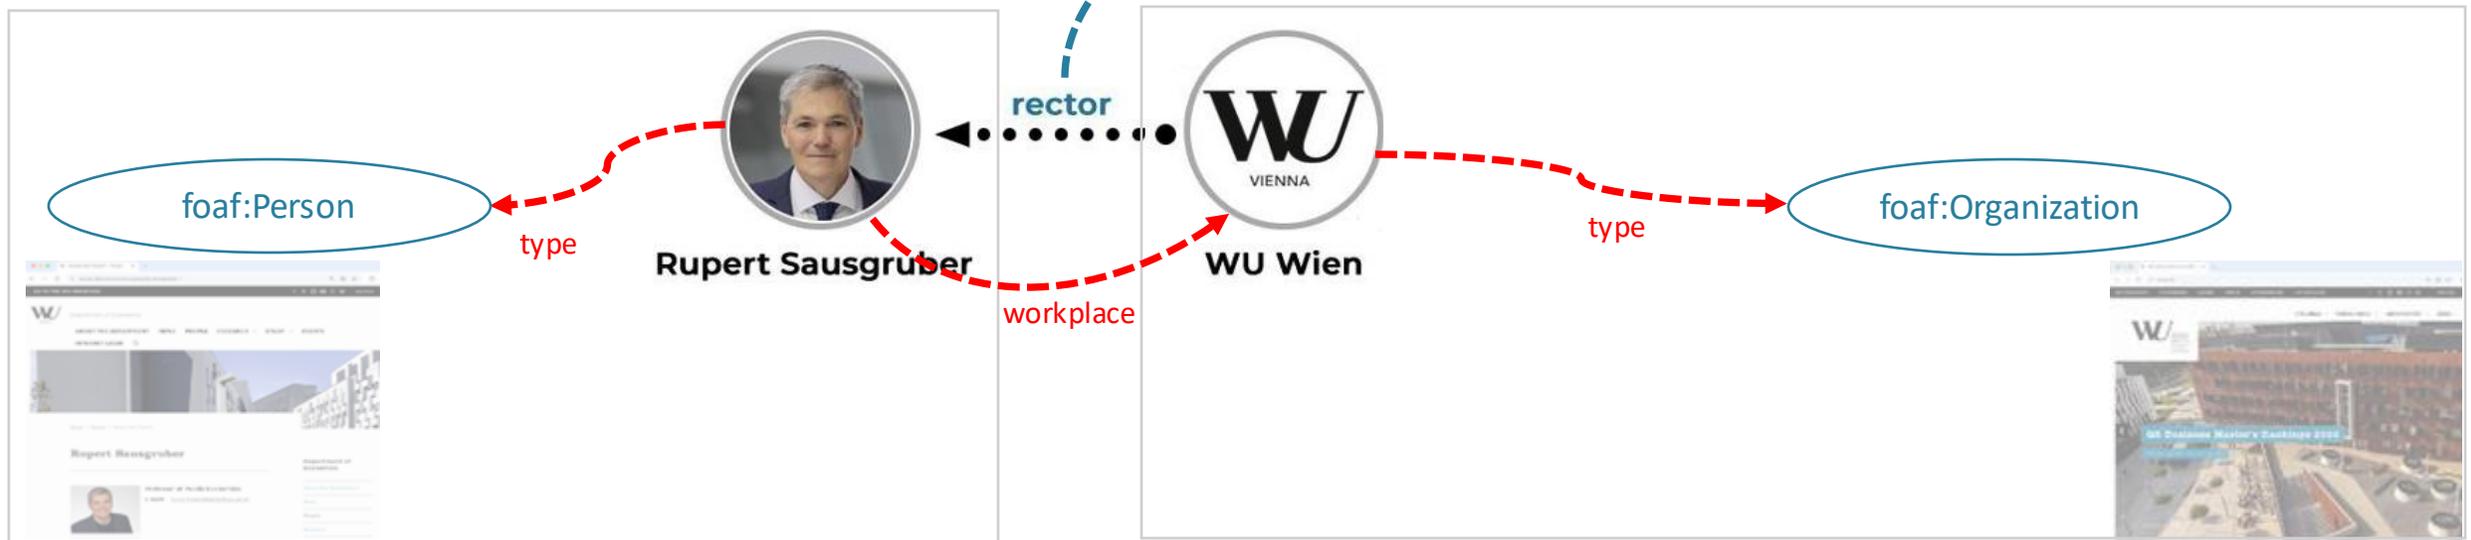
...k, John-Jules Meyer, and Axel Polleres. Norms, Engineering Review, 28(1):107--116, March 2013.

```
SELECT ?P
{ ?P a foaf:Person .
  ?P foaf:workPlace wu:WUWien. }
```

Phase 2: RDF, RDF Schema, OWL Linked “Metadata” with Rules and Ontologies.



FOAF “Ontology”

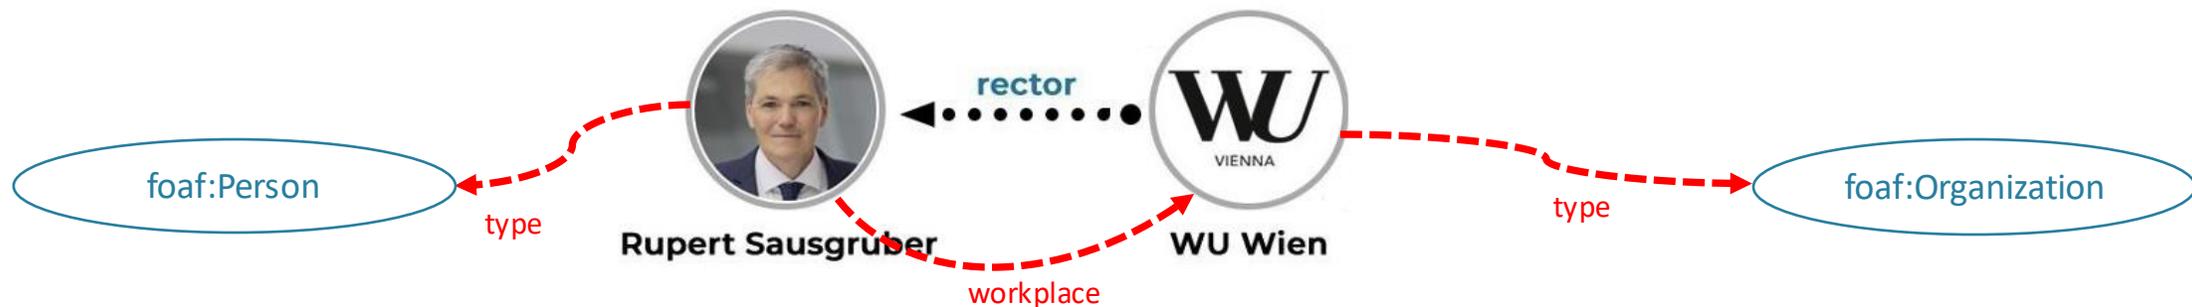


Phase 2: RDF, RDF Schema, OWL, *SPARQL* Linked “Metadata” with Rules and Ontologies.

*"If HTML and the Web made all the online documents look like one huge **book**, RDF, schema and inference languages will make all the data in the world look like **one huge database**"*

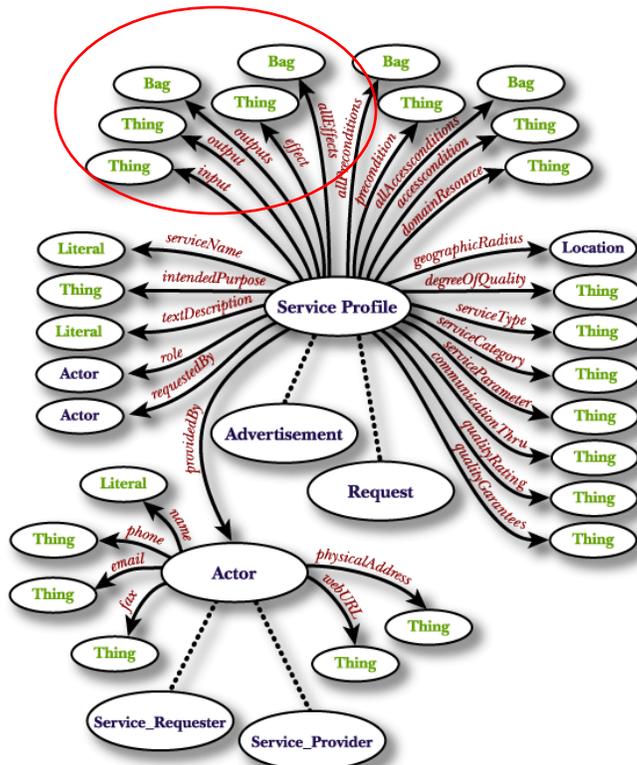
Tim Berners-Lee, Weaving the Web, 1999

```
SELECT ?P
{ ?P a foaf:Person .
  ?P foaf:workPlace wu:WUWien. }
```



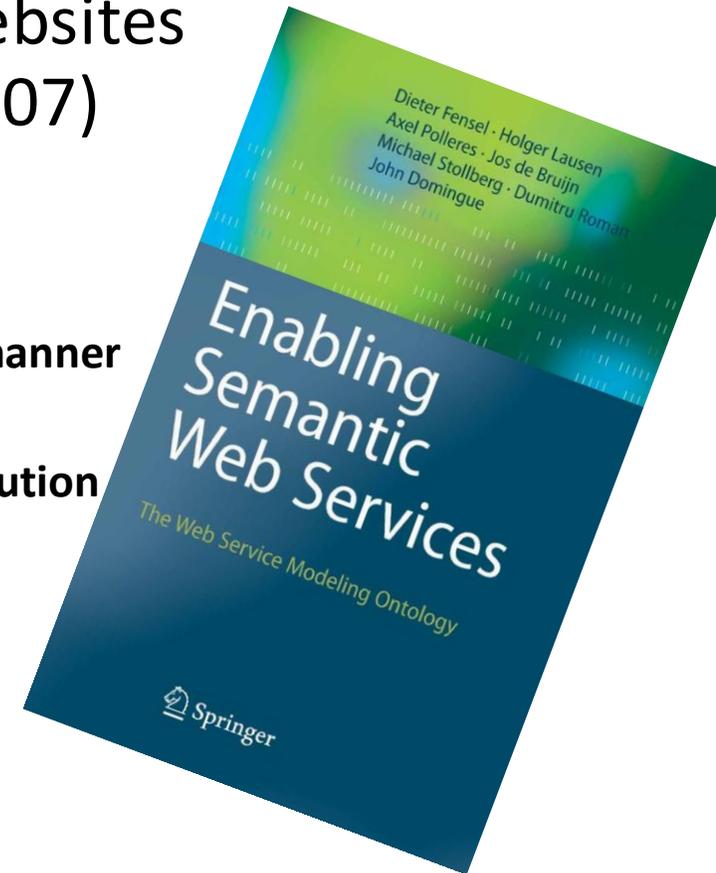
Phase 2 “interlude”

- **Semantic Web Services:** We cannot only describe Websites and Metadata, but also Services and APIs! (~2002-2007)



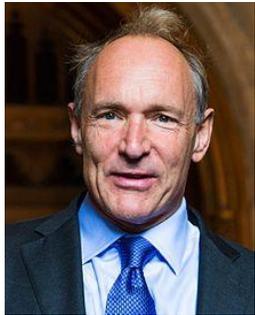
- Main Idea:
 - Describing **Services & Agents** in a **declarative manner**
 - using ontologies...
 - ... should enable **automated composition, execution**

On hindsight: automated composition by purely symbolic inference ... was maybe maybe ahead of its time...



Focus on Data: Linked Data

- (ca. 2006/7 – ca. 2013)
 - Main question: How can I **publish** “Knowledge on the Web” ...



Linked Data Principles

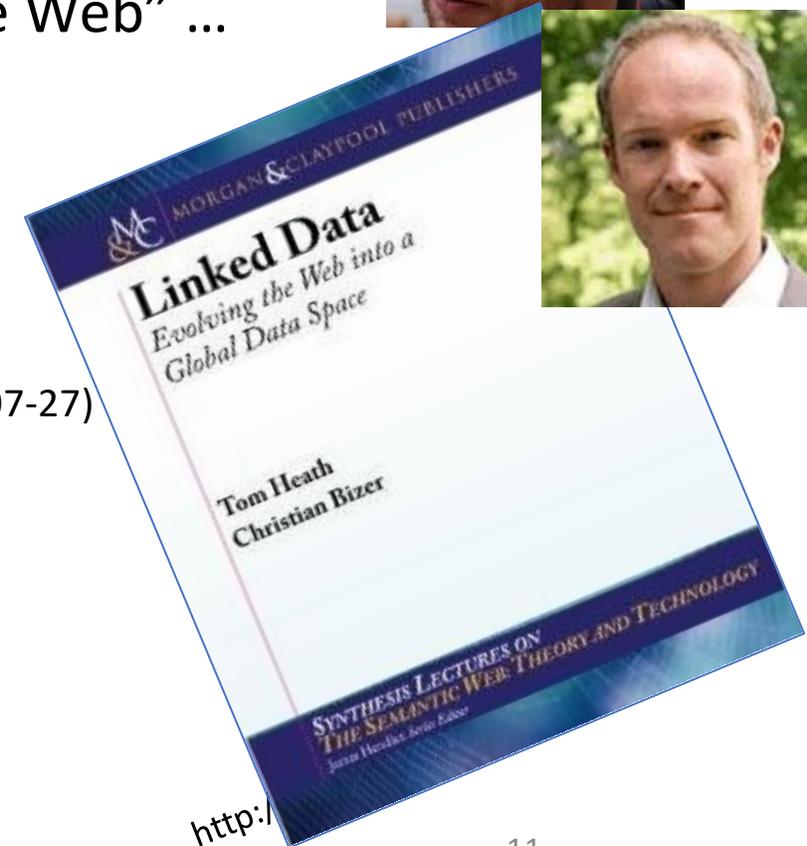
- **LDP1**: use URIs as names for things
- **LDP2**: use HTTP URIs so those names can be dereferenced
- **LDP3**: return useful – RDF? – information upon dereferencing those URIs
- **LDP4**: include links using externally dereferenceable URIs.

<https://www.w3.org/DesignIssues/LinkedData.html> (originally published 2006-07-27)



“A Little Semantics Goes a Long Way” (Jim Hendler)

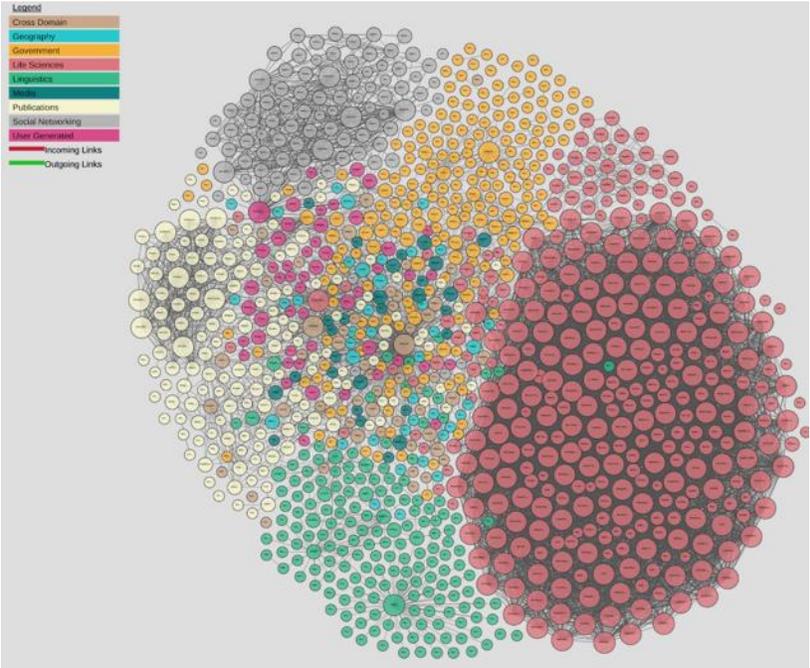
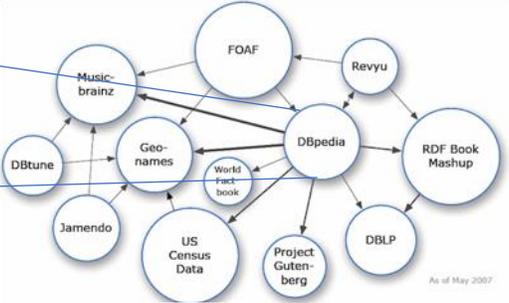
<https://www.cs.rpi.edu/~hendler/LittleSemanticsWeb.html>





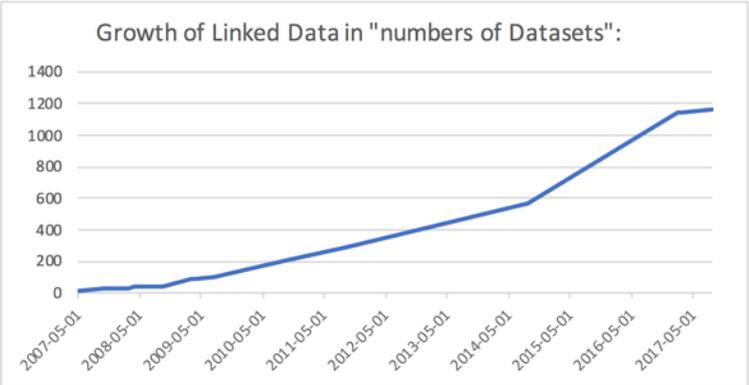
Phase3: From Semantic Web to Linked (Open) Data

- (ca. 2006/7 – 2012)
 - Main question: How can I **publish** “Knowledge on the Web” ...
 - Linked **Open** Data... growth slowed down a bit
 - A lot of active developments to publish and link RDF Data
 - also in Enterprises (“Enterprise Linked Data”)



<http://lod-cloud.net/>

2017-08-22	1163
2017-02-20	1139
2017-01-26	1146
2014-08-30	570
2011-09-19	295
2010-09-22	203
2009-07-14	95
2009-03-27	93
2009-03-05	89
2008-09-18	45
2008-03-31	34
2008-02-28	32
2007-11-10	28
2007-11-07	28
2007-10-08	25
2007-05-01	12



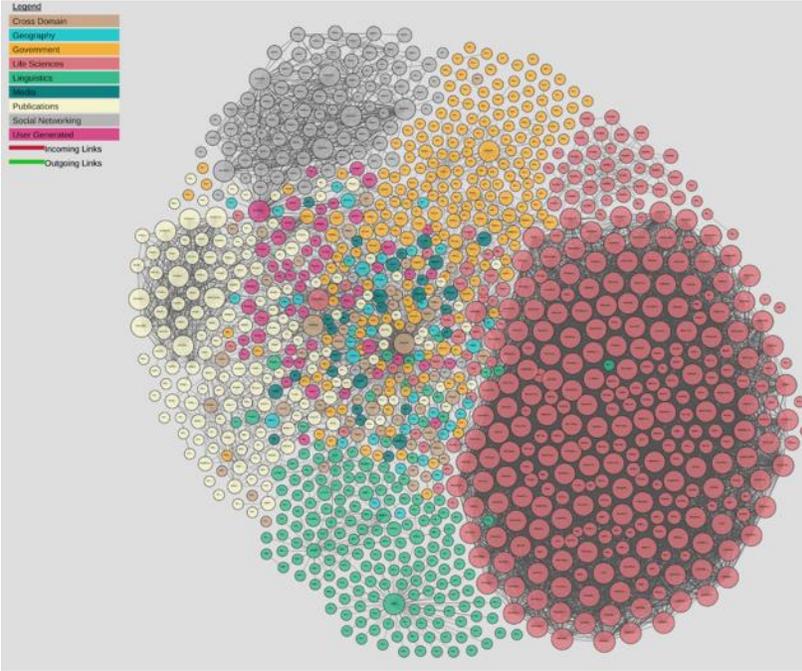
Axel Polleres, Maulik R. Kamdar, Javier D. Fernández, Tania Tudorache, and Mark A. Musen. [A more decentralized vision for linked data](#). In *Decentralizing the Semantic Web (Workshop of ISWC2018)*.



Phase 3: From Semantic Web to Linked (Open) Data

- (ca. 2006/7 – 2013)
 - Main question: How can I **publish** “Knowledge on the Web” ...

• *Side question(s):*
Can deductive symbolic inference and queries scale to a decentralized Web?



International Semantic Web Conference (2014)
Strategies for executing federated queries in SPARQL1.1

Carlos Buil-Aranda^{1*}, Axel Polleres^{2**}, and Jürgen Umbrich²
¹ Department of Computer Science, Pontificia Universidad Católica, Chile
cbuil@ing.puc.cl
² Vienna University of Economy and Business (WU)
{first.last}@wu.ac.at

Reasoning Web. Semantic Technologies for Intelligent Data Access (Reasoning Web 2013)
RDFS & OWL Reasoning for Linked Data

Axel Polleres¹, Aidan Hogan², Renaud Delbru², and Jürgen Umbrich^{2,3}
¹ Siemens AG Österreich, Siemensstraße 90, 1210 Vienna, Austria
² Digital Enterprise Research Institute, National University of Ireland, Galway
³ Fujitsu (Ireland) Limited, Swords, Co. Dublin, Ireland

Abstract. Linked Data promises that a large portion of Web Data will be available as one big interlinked RDF database against which structured queries can be answered. In this lecture we will show how reasoning over RDF Schema (RDFS) and the Web Ontology Language (OWL) can be used to obtain more complete answers for such queries over Linked Data. We first look at the extent to which RDFS and OWL features are supported on the Web. We then introduce two high-level architectures for answering over Linked Data and outline how these can be used to answer (lightweight) RDFS and OWL reasoning queries. We conclude by discussing the challenges faced and discussing theoretical and practical issues.

From Linked Open Data to Knowledge Graphs:



2013: Google adopts Semantic Web ideas under a new name

- Jamie Taylor, Google, Inc., Keynote [ISWC2017](#)

The Power of Knowledge Graph: Interlocking data

Destination suggestions: museum, city

Reservation: airports, airport, airline, flight

Flight Status: terminal, gate, status, leg

Google

Knowledge Graph

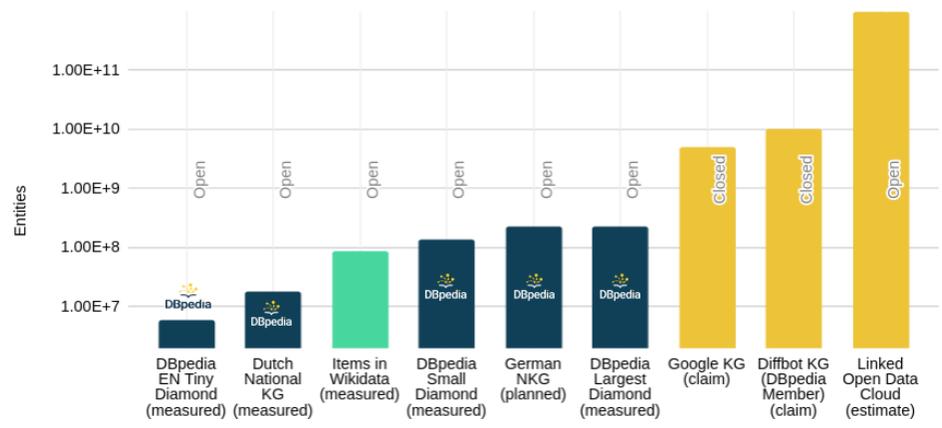
1 Billion Entities
70 Billion Assertions

- Actors, Directors, Movies
- Art Works & Museums
- Cities & Countries
- Islands, Lakes, Lighthouses
- Music Albums & Music Groups
- Planets & Spacecraft
- Roller Coasters & Skyscrapers
- Sports Teams
- [...]

From Linked Open Data to Knowledge Graphs:

Success stories of mainly monolithic (but huge) Knowledge Graphs rather than a network of Linked small KGs:

<https://www.slideshare.net/Frank.van.Harmelen/adoption-of-knowledge-graphs-late-2019>



<https://www.dbpedia.org> 2021



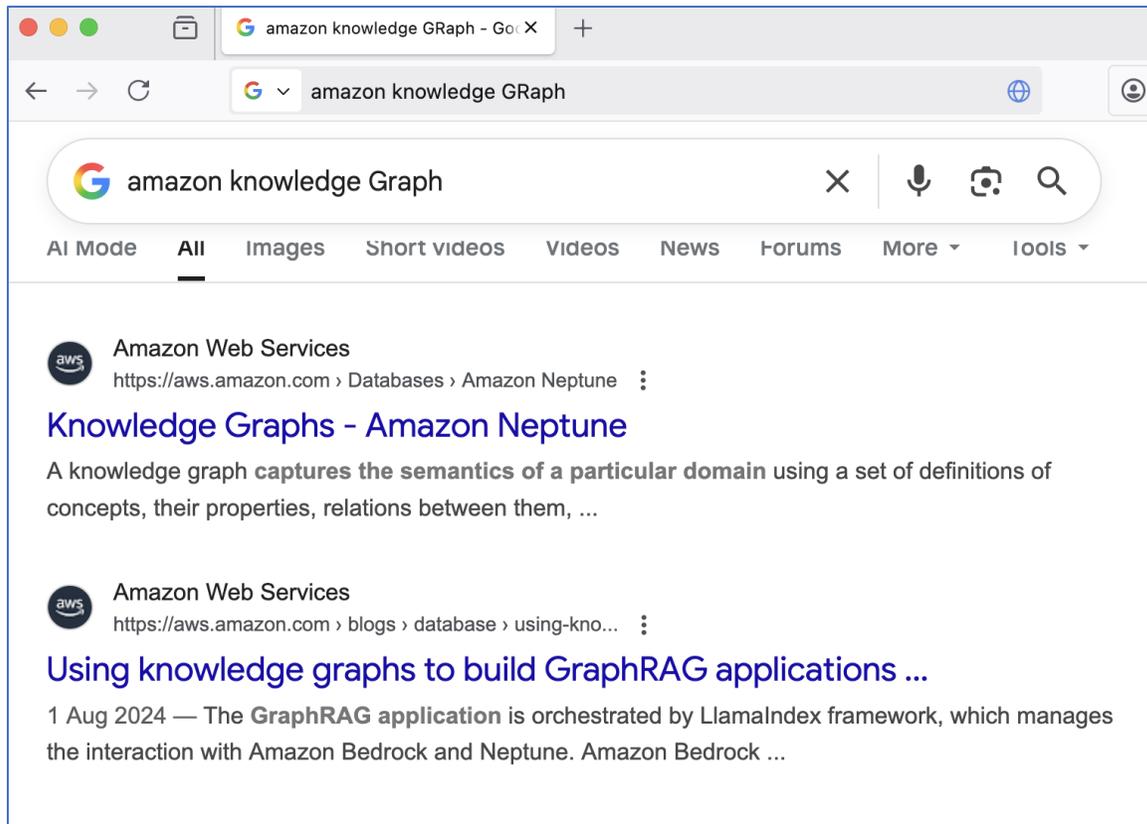
Open KGs (April 2021)	
DBpedia	~4.58m entities
Yago4	~50m entities
Wikidata	~93m entities

N. Noy, Y. Gao, A. Jain, A. Narayanan, A. Hogan et al.: Knowledge Graphs. Co...



It's not only Search – Obviously... User Recommendation, Language Technologies, ...

Since then, industry has become a main driver in scaling and using KG Technologies...



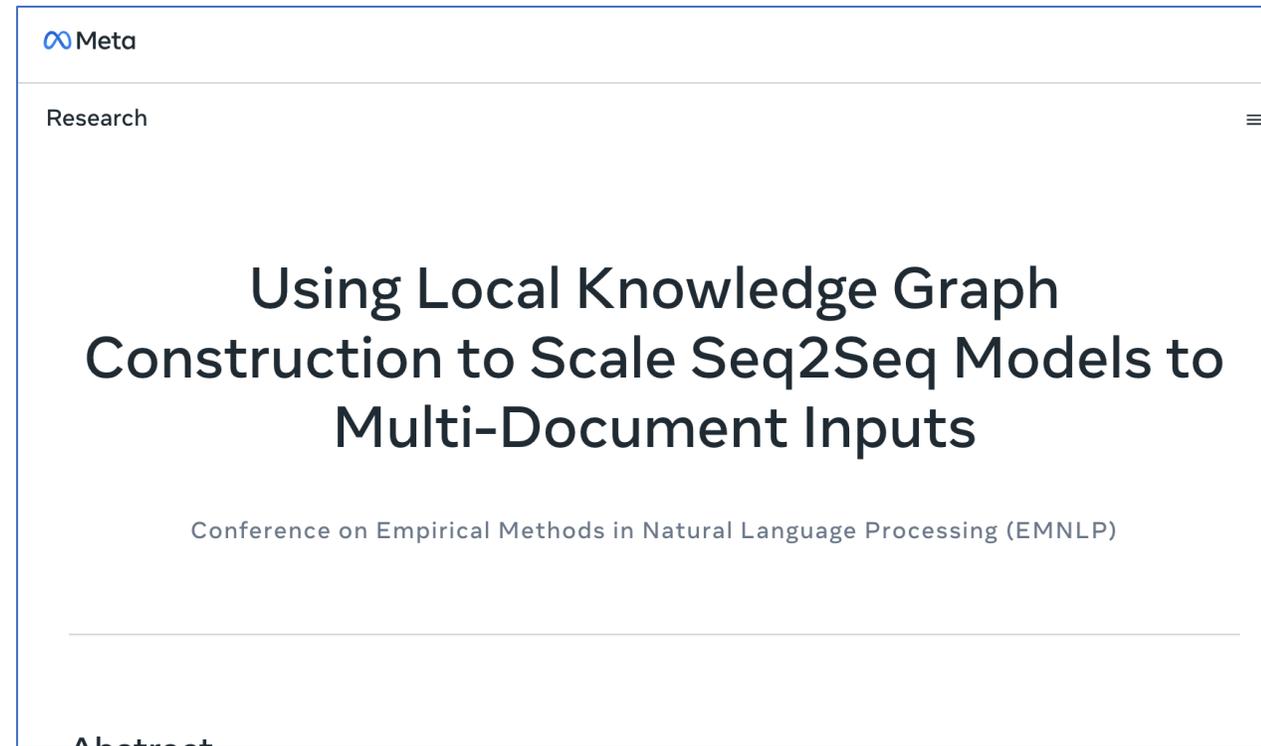
A screenshot of a Google search for "amazon knowledge Graph". The search bar shows the query and the Google logo. Below the search bar, there are tabs for "AI Mode", "All", "Images", "Short videos", "Videos", "News", "Forums", "More", and "Tools". The search results are displayed below, with the first result from Amazon Web Services titled "Knowledge Graphs - Amazon Neptune". The second result is a blog post titled "Using knowledge graphs to build GraphRAG applications ...".

amazon knowledge Graph

AI Mode All Images Short videos Videos News Forums More Tools

Amazon Web Services
https://aws.amazon.com › Databases › Amazon Neptune
Knowledge Graphs - Amazon Neptune
A knowledge graph captures the semantics of a particular domain using a set of definitions of concepts, their properties, relations between them, ...

Amazon Web Services
https://aws.amazon.com › blogs › database › using-kno...
Using knowledge graphs to build GraphRAG applications ...
1 Aug 2024 — The **GraphRAG application** is orchestrated by Llamaindex framework, which manages the interaction with Amazon Bedrock and Neptune. Amazon Bedrock ...



A screenshot of a Meta Research paper titled "Using Local Knowledge Graph Construction to Scale Seq2Seq Models to Multi-Document Inputs". The paper is published in the Conference on Empirical Methods in Natural Language Processing (EMNLP). The word "Abstract" is visible at the bottom of the page.

Meta

Research

Using Local Knowledge Graph Construction to Scale Seq2Seq Models to Multi-Document Inputs

Conference on Empirical Methods in Natural Language Processing (EMNLP)

Abstract

Collaborative, Open Knowledge Graphs:



DBpedia generates a graph from links and facts in Wikipedia's Infoboxes:

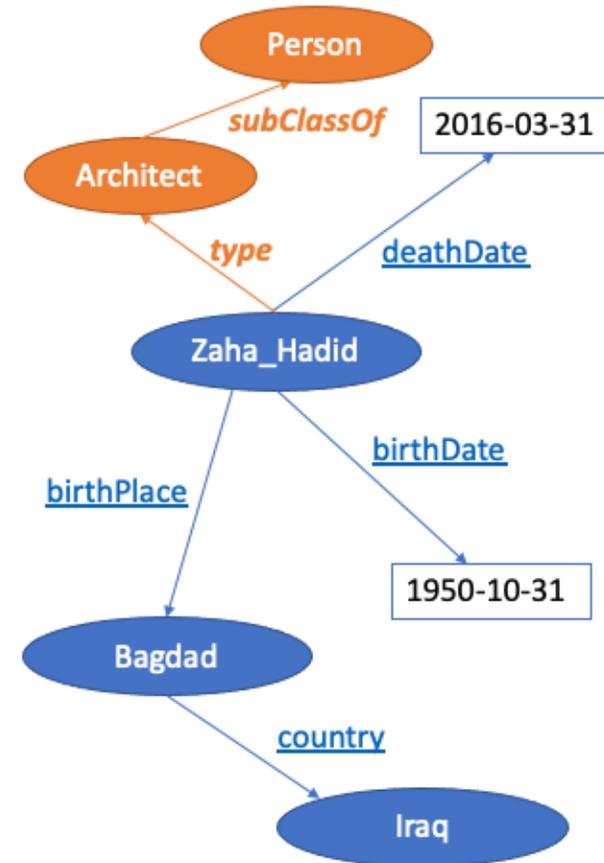
http://wikipedia.org/wiki/Zaha_Hadid

The screenshot shows the Wikipedia article for Zaha Hadid. The infobox contains the following information:

- Born:** Zaha Mohammad Hadid, 31 October 1950, Baghdad, Kingdom of Iraq
- Died:** 31 March 2016 (aged 65), Miami, Florida, U.S.
- Nationality:** Iraq, United Kingdom
- Alma mater:** American University of Beirut, Architectural Association School of Architecture
- Occupation:** Architect
- Parent(s):** Mohammed Hadid, Wajeeha Sabonji
- Practice:** Zaha Hadid Architects
- Buildings:** Vitra Fire Station, MAXXI, Bridge Pavilion, Contemporary Arts Center, Heydar Aliyev Center, Riverside Museum
- Website:** www.zaha-hadid.com

Born	Zaha Mohammad Hadid 31 October 1950 Baghdad, Kingdom of Iraq
Died	31 March 2016 (aged 65) Miami, Florida, U.S.
Nationality	Iraq, United Kingdom
Alma mater	American University of Beirut Architectural Association School of Architecture
Occupation	Architect
Parent(s)	Mohammed Hadid Wajeeha Sabonji
Practice	Zaha Hadid Architects
Buildings	Vitra Fire Station, MAXXI, Bridge Pavilion, Contemporary Arts Center, Heydar Aliyev Center, Riverside Museum
Website	www.zaha-hadid.com

http://dbpedia.org/resource/Zaha_Hadid



Collaborative, Open Knowledge Graphs:



Lionel Messi (Q615)

Argentine association football player

image



occupation

association football player

2 references

FIFA player ID (archived)

229397

1 reference

country of citizenship

Argentina

start time 1987

0 references

Spain

start time 2005

1 reference

Revision history of "Lionel Messi" (Q615)

[View logs for this item](#) ([view abuse log](#))

Filter revisions

Diff selection: Mark the radio buttons of the revisions to compare and hit enter or the button at the bottom

Legend: **(cur)** = difference with latest revision, **(prev)** = difference with preceding revision, **m** = minor edit

(latest | [earliest](#)) View (newer 50 | older 50) (20 | 50 | 100 | 250 | 500)

Compare selected revisions

- (cur | prev) 08:56, 8 December 2024 द्वाद्दग गज्जा (talk | contribs) .. (537,664 bytes) (+92) ..
- (cur | prev) 20:39, 7 December 2024 Sanremofilo (talk | contribs) .. (537,572 bytes) (+363) .. (*messi/5663*) (Tag: Wikidata user interface)
- (cur | prev) 15:21, 2 December 2024 Ytterbyz (talk | contribs) .. (537,209 bytes) (+349) .. (Tag: Wikidata user interface)
- (cur | prev) 23:45, 29 November 2024 Mickey Đại Phát (talk | contribs) .. (536,860 bytes) (-) .. (Wikidata user interface, Mobile termbox)
- (cur | prev) 19:47, 27 November 2024 KrBot (talk | contribs) .. (536,877 bytes) (-14) .. (See *see autofix на / on Property talk:P12924*)

From Linked Open Data to Knowledge Graphs: What's the state of affairs?

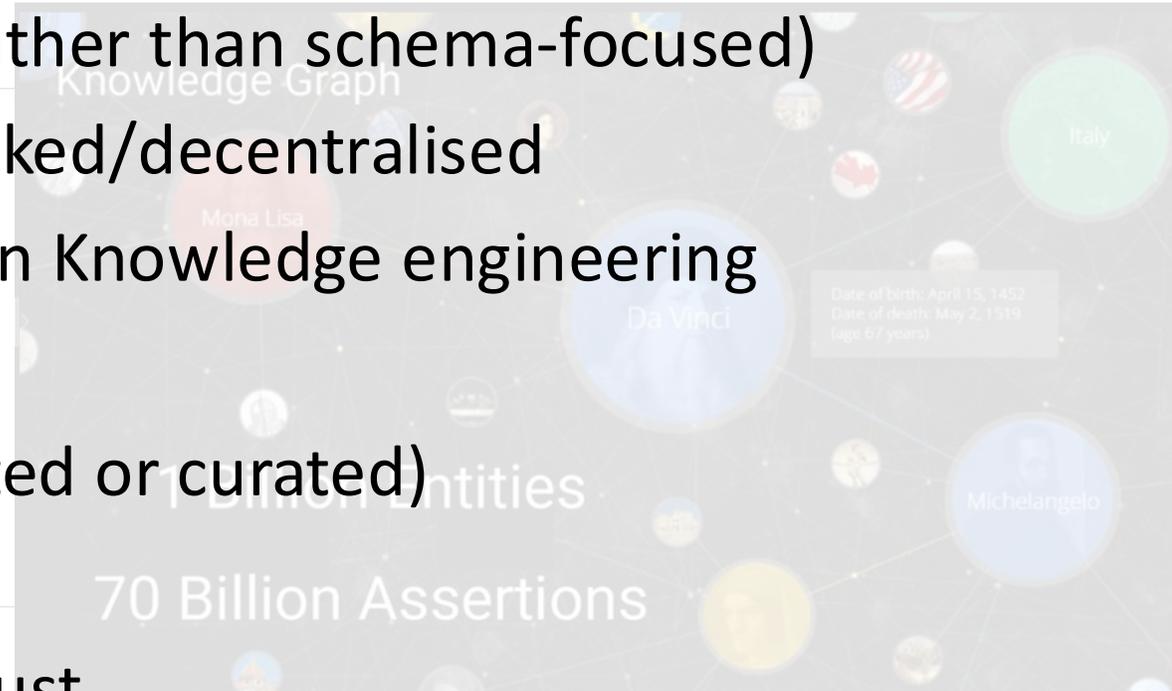
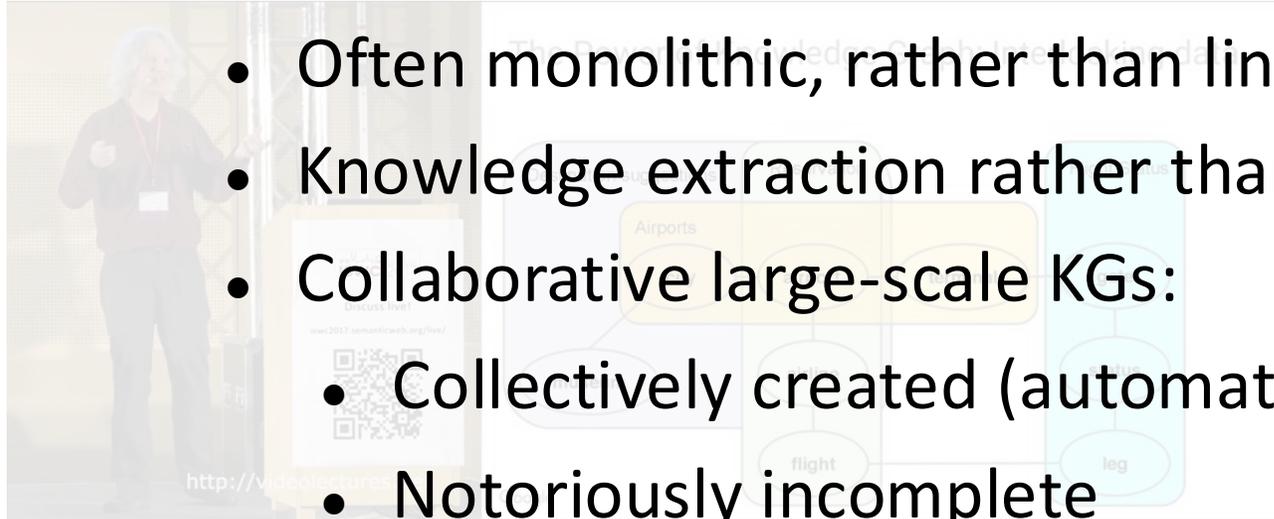
- Jamie Taylor • Large-scale, still data-focused (rather than schema-focused)

- Often monolithic, rather than linked/decentralised
- Knowledge extraction rather than Knowledge engineering
- Collaborative large-scale KGs:
 - Collectively created (automated or curated)
 - Notoriously incomplete
 - (Logical) **consistency** not a must

- Actors, Directors, Movies
- Art Works & Museums
- Cities & Countries
- Islands, Lakes, Lighthouses
- Music Albums & Music Groups
- Planets & Spacecraft
- Roller Coasters & Skyscrapers
- Sports Teams
- [...]
- Enterprise KGs: knowledge necessary to power applications
- *Ontological expressivity* not central – BUT: **Expressing context** is!

For instance:

- Provenance
- Temporal context



Let's have a look at practical examples of such collaboratively curated Knowledge Graphs:

- DBpedia (since 2007)

vs.

- Wikidata (since 2012)

DBpedia



Developer(s)	Leipzig University University of Mannheim
Initial release	10 January 2007 (17 years ago)
Stable release	DBpedia 2016-10 / 4 July 2017
Repository	github.com/dbpedia/
Written in	Scala · Java
Type	Semantic Web · Linked Data
License	GNU General Public License
Website	dbpedia.org

- ✓ • RDF
- ✓ • Standard ontology language (OWL)
- ✓ • SPARQL endpoint
- ✗ • Consistent
- ✗ • Context

Wikidata



- ✓
- ✗
- ✓
- ✗
- ✓

Wikidata

[show]

Type of site Knowledge base · Wiki

Available in Multiple languages

Owner Wikimedia Foundation

Editor Wikimedia community

URL www.wikidata.org/wiki/Wikidata:Main_Page

Commercial No

Registration Optional

Launched 29 October 2012; 12 years ago^[1]

SPARQL: Using KGs to answer questions:

- E.g. from



London

From Wikipedia, the free encyclopedia

This article is about the capital city. For the region of England, see Greater London. For the historic city and financial district within London, see City of London. For other uses, see London (disambiguation).

London (/ˈlɒndən/ [ⓘ]) is the capital and most populous city of England and the United Kingdom.^{[1][2]} Standing on the River Thames in the south east of the island of Great Britain, London has been a major settlement for two millennia. It was founded by the Romans, who named it Londinium.^[6] London's ancient core, the City of London, largely retains its 1.12-square-mile (2.9 km²) medieval boundaries. Since at least the 19th century, "London" has also referred to the metropolis around this core, historically split between Middlesex, Essex, Surrey, Kent, and Hertfordshire,^{[10][11][12]} which today largely makes up Greater London,^{[13][14][15]} governed by the Mayor of London and the London Assembly.^{[16][17][18]}

London is a leading global city^{[19][20]} in the arts, commerce, education, entertainment, fashion, finance, healthcare, media, professional services, research and development, tourism, and transportation.^{[11][21][22]} It is crowned as the world's largest financial centre^{[23][24][25]} and has the fifth- or sixth-largest metropolitan area GDP in the world.^{[11][26][27]} London is a world cultural capital,^{[28][29][30]} it is the world's most-visited city as measured by international arrivals^[31] and has the world's largest city airport system measured by passenger traffic.^[32] London is the world's leading investment destination,^{[33][34][35]} hosting more international retailers^{[36][37]} and ultra high-net-worth individuals^{[38][39]} than any other city. London's universities form the largest concentration of higher education institutes in Europe^[40] and in 2012, London became the first city to have hosted the modern Summer Olympic Games three times.^[41]

Greater London was the largest or any city in the European Union,^[42] and accounting for 13.4% of the UK population.^[43] London's urban area is the second most populous in the EU after Paris, with 8,787,426 inhabitants at the 2011 census.^[44] The city's metropolitan area is the most populous in the EU with 13,879,757 inhabitants,^{[11][45]} while the Greater London Authority states the population of the city-region (covering a large part of the south east) as 22.7 million.^{[46][47]} London was the world's most populous city from around 1831 to 1925.^[48]

London contains four World Heritage Sites: the Tower of London; Kenilworth Castle; the site comprising the Palace of Westminster, Westminster Abbey, and St Margaret's Church; and the historic settlement of Greenwich (in which the Royal Observatory, Greenwich marks the Prime Meridian, 0° longitude, and GMT).^[49] Other famous landmarks include Buckingham Palace, the London Eye, Piccadilly Circus, St Paul's Cathedral, Tower Bridge, Trafalgar Square, and The Shard. London is home to numerous museums, galleries, libraries, sporting events, and other cultural institutions, including the British Museum, National Gallery, Natural History Museum, Tate Modern, British Library, and West End theatres.^[50] The London Underground is the oldest underground railway network in the world.

<https://en.wikipedia.org/wiki/London>

Automatic Extractors

- One of the central datasets of the Linked Open Data-Cloud
- RDF extracted from Wikipedia-Infoboxes
- You can use a language called SPARQL endpoint (roughly: SQL for RDF) to do **structured queries** over RDF:
 - „Cities in the UK with more than 1M population“:

DBpedia

Browse using - Formats - Faceted Browser - Sparql Endpoint

About: London

An Entity of Type: populated place, from Named Graph: <http://dbpedia.org/>, within Data Space: <http://dbpedia.org/>

<http://dbpedia.org/resource/London>

London (/ˈlɒndən/ [ⓘ]) is the capital and most populous city of England and the United Kingdom.^{[1][2]} Standing on the River Thames in the south east of the island of Great Britain, London has been a major settlement for two millennia. It was founded by the Romans, who named it Londinium. London's ancient core, the City of London, largely retains its 1.12-square-mile (2.9 km²) medieval boundaries. Since at least the 19th century, "London" has also referred to the metropolis around this core, historically split between Middlesex, Essex, Surrey, Kent, and Hertfordshire, which today largely makes up Greater London, governed by the Mayor of London and the London Assembly.

Property	Value
dbpedia:PopulatedPlace/areaTotal	1572.0
dbpedia:PopulatedPlace/populationDensity	5518.0
dbpedia:abstract	London (/ˈlɒndən/ [ⓘ]) is the capital and most populous city of England and the United Kingdom. ^{[1][2]} Standing on the River Thames in the south east of the island of Great Britain, London has been a major settlement for two millennia. It was founded by the Romans, who named it Londinium.

Structured queries (SPARQL):

<https://api.triplydb.com/s/gZZskqRpQ>

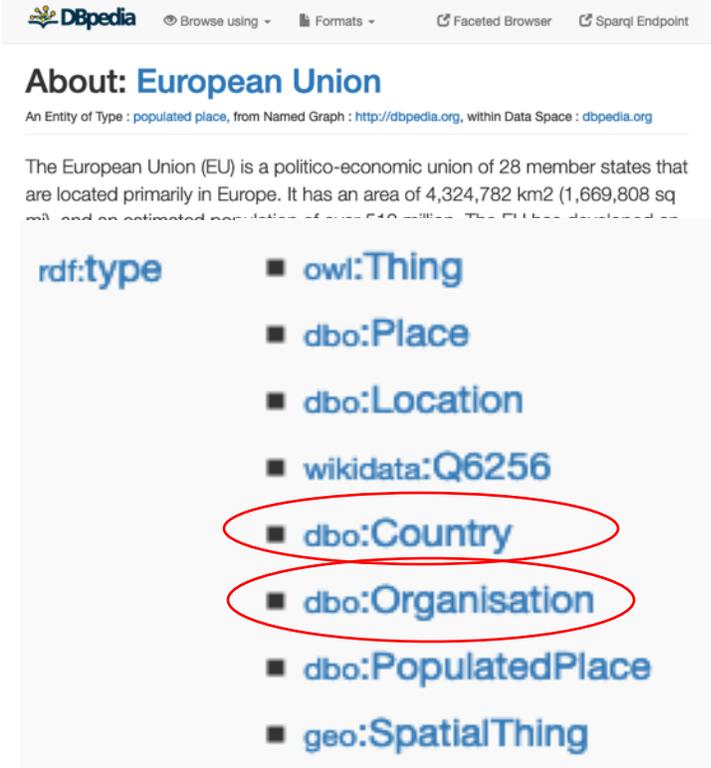
```
PREFIX : <http://dbpedia.org/resource/>
PREFIX dbo: <http://dbpedia.org/ontology/>
PREFIX yago: <http://dbpedia.org/class/yago/>

SELECT DISTINCT ?city ?pop WHERE {
    ?city a schema:City .
    ?city dbo:country :United_Kingdom.
    ?city dbo:populationTotal ?pop

    FILTER ( ?pop > 1000000 )
}
```

Knowledge Graphs like Dbpedia are not logically consistent! ☹️ [1]

- E.g. 



DBpedia

Browse using - Formats - Faceted Browser - Sparql Endpoint

About: European Union

An Entity of Type : populated place, from Named Graph : <http://dbpedia.org>, within Data Space : dbpedia.org

The European Union (EU) is a politico-economic union of 28 member states that are located primarily in Europe. It has an area of 4,324,782 km² (1,669,808 sq mi) and an estimated population of over 540 million. The EU has developed an

rdf:type

- owl:Thing
- dbo:Place
- dbo:Location
- wikidata:Q6256
- **dbo:Country**
- **dbo:Organisation**
- dbo:PopulatedPlace
- geo:SpatialThing

Dbpedia Ontology:

`dbo:Agent owl:disjointWith dbo:Place.`

`dbo:Country rdfs:subClassOf dbo:Place.`

`dbo:Organisation rdfs:subClassOf dbo:Agent.`



1. Stefan Bischof, Markus Krötzsch, Axel Polleres, and Sebastian Rudolph. Schema-agnostic query rewriting in SPARQL 1.1. In *Proceedings of the 13th International Semantic Web Conference (ISWC 2014)*, Lecture Notes in Computer Science (LNCS). Springer, October 2014. [[.pdf](#)]

Wikidata is also not “consistent”, but doesn’t use OWL

image	
occupation	<p>association football player</p> <p>▶ 2 references</p>
FIFA player ID (archived)	<p>229397</p> <p>▶ 1 reference</p>
country of citizenship	<p>Argentina</p> <p>start time 1987</p> <p>▼ 0 references</p>
	<p>Spain</p> <p>start time 2005</p> <p>▶ 1 reference</p>

The same question as before in Wikidata:

Note: Wikidata does not even use standard OWL

- “Simple” surface [query](#):

Which cities in the UK have more than 1M people?

```
SELECT DISTINCT ?city WHERE {  
  ?city wdt:P31/wdt:P279* wd:Q515.  
  ?city wdt:P1082 ?population .  
  ?city wdt:P17 wd:Q38 .  
  FILTER (?population > 1000000) }
```

Note: Wikidata uses numeric IDs

city (Q515)
large and permanent human settlement

population (P1082)
number of people inhabiting the place; number of people of subject

country (P17)
sovereign state of this item
United Kingdom (Q145)
country in Europe

instance of (P31)
that class of which this subject is a particular example and member. (Subject typically an individual member with Proper Name label.) Different from P279 (subclass of).

subclass of (P279)
all instances of these items are instances of those items; this item is a class (subset) of that item. Not to be confused with Property:P31 (instance of).

- What’s this?

The same question as before in Wikidata:

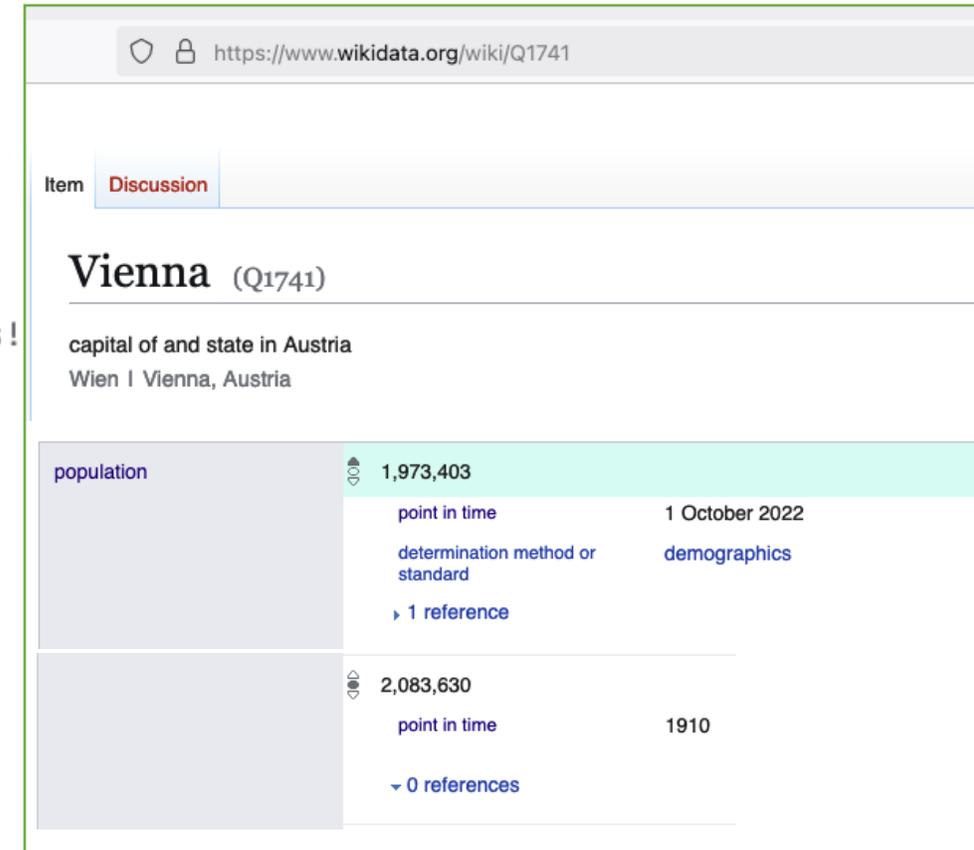
<https://w.wiki/BqRX>

Which cities in the Austria have more than 1M/2M people?

```
SELECT DISTINCT ?City ?Pop
{
  ?City wdt:P17 wd:Q40;
        wdt:P31/wdt:P279* wd:Q515;
        wdt:P1082 ?Pop.
  FILTER (?Pop > 1000000)
  # note: Vienna historically had more than 2M inhabitants!
  # FILTER (?Pop > 2000000)
}
```

Note: Wikidata also has such contextual information!!!!

So, WHEN did Vienna have 2M inhabitants?



The screenshot shows the Wikidata page for Vienna (Q1741). The page title is "Vienna (Q1741)" and it is identified as the "capital of and state in Austria". The page displays two population entries:

property	value	point in time	determination method or standard
population	1,973,403	1 October 2022	demographics
population	2,083,630	1910	

The same question as before in Wikidata:

<https://w.wiki/BqRj>

Which cities in the Austria have more than 1M/2M people?

```
SELECT DISTINCT ?City ?Pop ?Timepoint
{
  ?City wdt:P17 wd:Q40;
        wdt:P31/wdt:P279* wd:Q515;
        p:P1082 ?Stmnt.
  ?Stmnt ps:P1082 ?Pop;
         pq:P585 ?Timepoint.
  # FILTER (?Pop > 1000000)
  # note: Vienna historically had more than 2M inhabitants!
  FILTER (?Pop > 2000000)
}
```

So, WHEN did Vienna have 2M inhabitants? Works!!!!

But needs an understanding of **Wikidata's proprietary RDF reification model** to model context!

See our recent ISWC2024 tutorial: <https://ww101.ai.wu.ac.at/>

Phase 4: Knowledge Graphs: “(Less Linked) Graph data with Metadata”

- Aka Labelled property graphs



Messi
(wd:Q615)



country of citizenship (wdt:P27)

{startTime: 1987}



Spain
(wd:Q29)



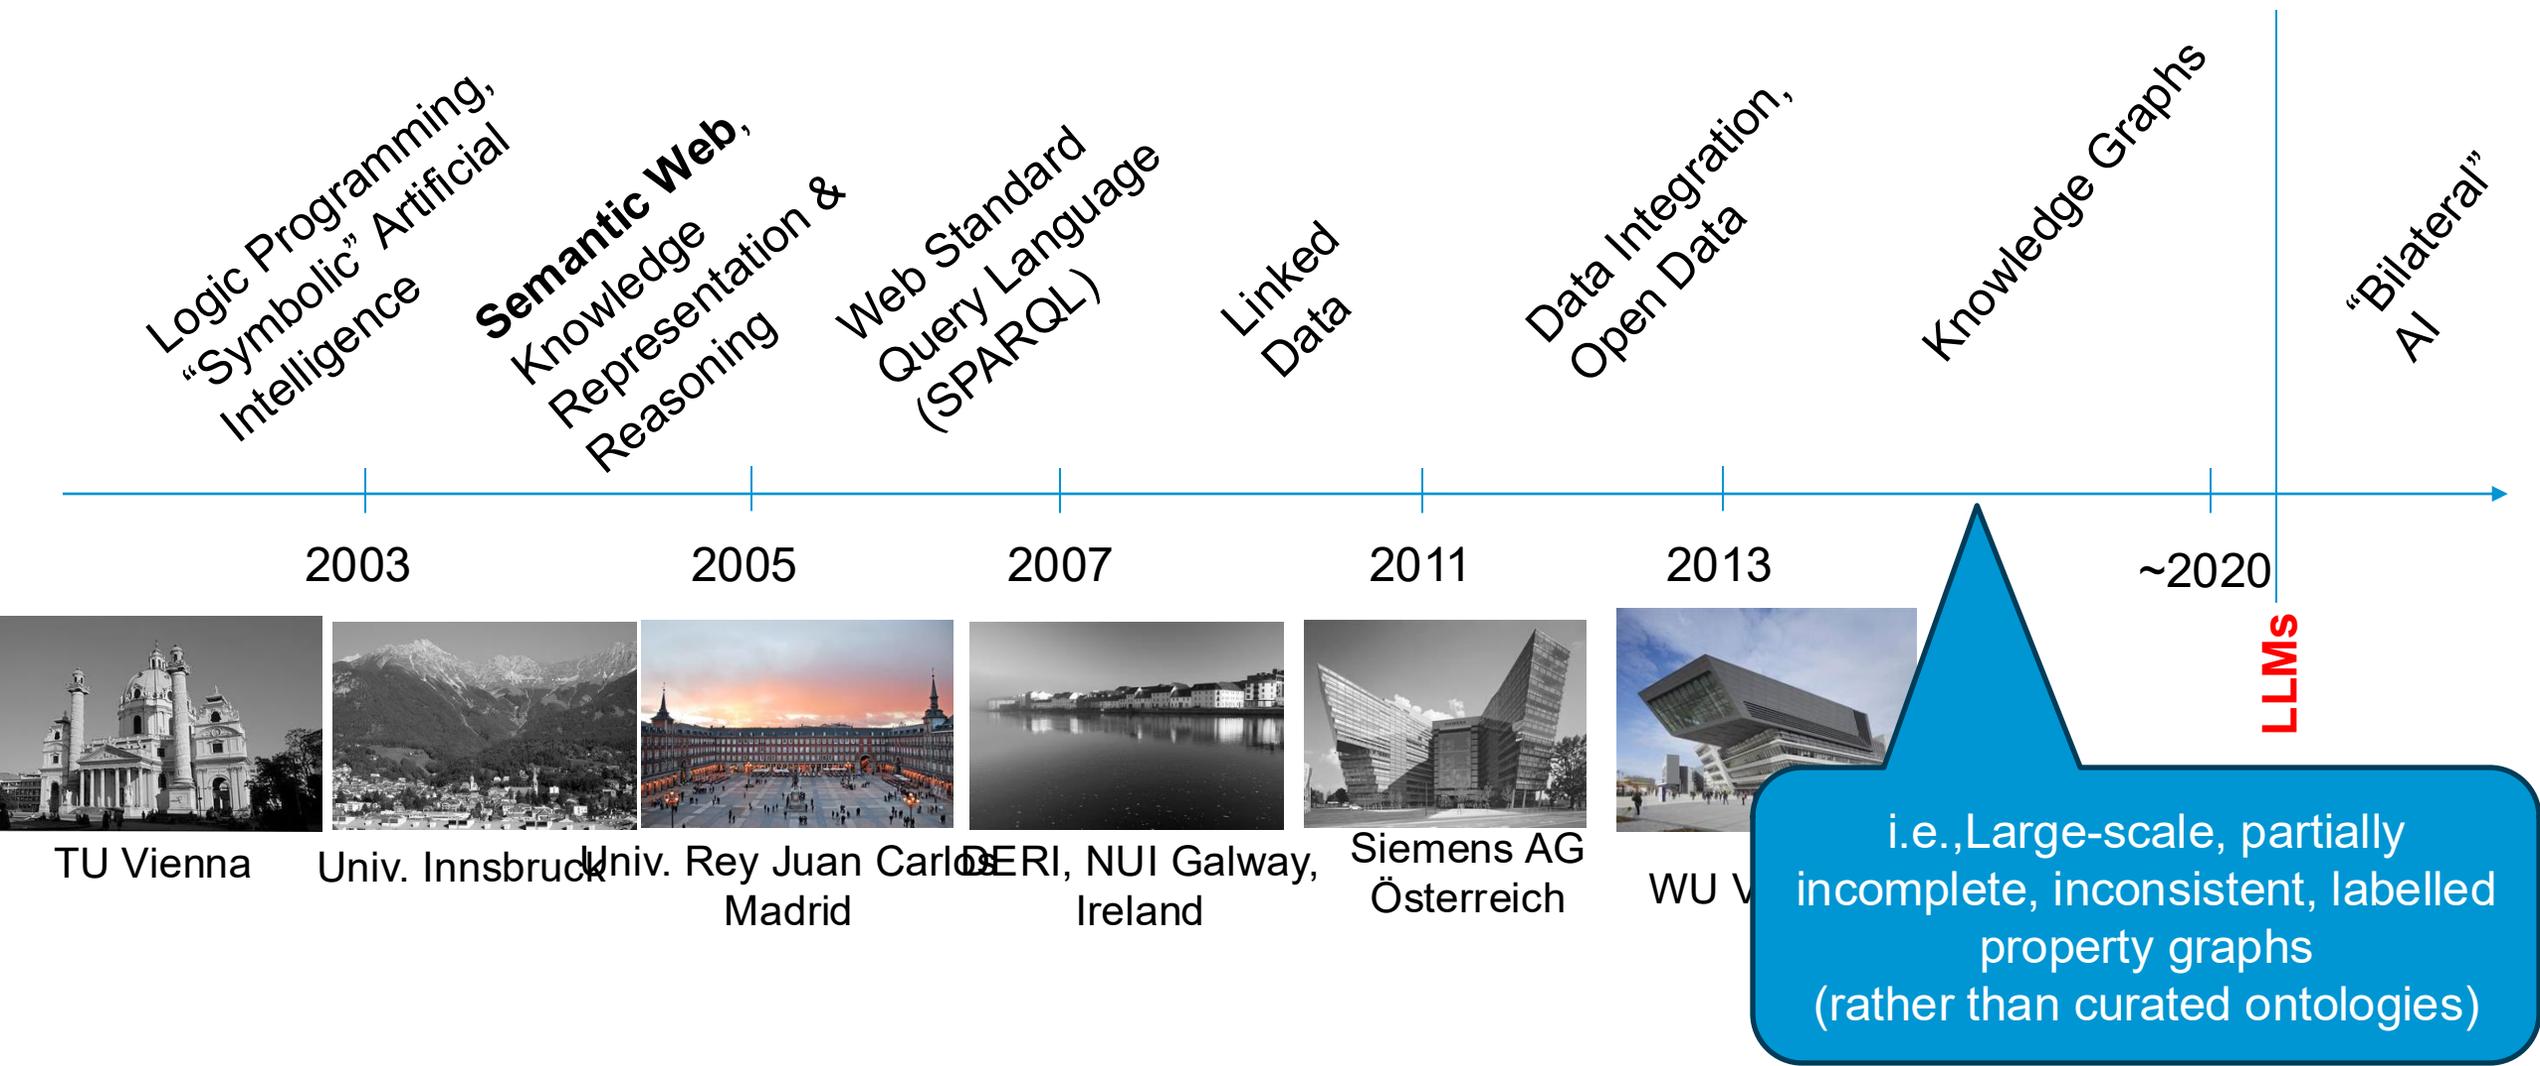
Argentina
(wd:Q414)

{startTime: 2005,
reference: https://cadenaser.com/ser/2005/09/26/deportes/1127690892_850215.html}

country of citizenship (wdt:P27)

country of citizenship	Argentina	
	start time	1987
	0 references	
Spain		
	start time	2005
	1 reference	

So, what are these KGs actually good for in the age of LLMs



So, for what are these KGs actually good for in the age of LLMs and AI?

- Interesting Note - IBM Watson - Jeopardy! (2011)

“Super-human” Question Answering was achieved by Knowledge Graphs before the LLM hype!

A screenshot of the IBM Watson Developer Cloud website. The page is titled "Natural Language Understanding" and shows an "Introduction" section. The API endpoint is listed as `https://gateway.watsonplatform.net/natural-language-understanding/api/v1`. The page also includes a "Curl" tab and a "Python" tab. The API Reference section lists various endpoints and features like "Analyze", "Categories", "Concepts", "Emotion", "Entities", "Keywords", "Metadata", "Relations", "Semantic Roles", and "Sentiment".

<https://youtu.be/P0Obm0DBvwl?t=951>

Some of our own research in this area:



- How good or bad can KGs deal with Question answering?
 - *By essentially translating Questions to SPARQL queries (Watson)?*
 - *By other bespoke techniques such as Message passing:*

Svitlana Vakulenko, Javier Fernández, Axel Polleres, Maarten de Rijke, and Michael Cochez. Message passing for complex question answering over knowledge graphs. In Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM2019, pages 1431--1440, Beijing, China, November 2019. ACM.

Some of our own research in this area:

- How good or bad can KGs deal with Question answering?
 - *By essentially translating Questions to SPARQL queries (Watson)?*
 - *By other bespoke techniques such as Message passing:*

Svitlana Vakulenko, Javier Fernández, Axel Polleres, Maarten de Rijke, and Michael Cochez. Message passing for complex question answering over knowledge graphs. In Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM2019, pages 1431--1440, Beijing, China, November 2019. ACM.

- How good or bad are LLMs with Question answering compared with KGs?

Gerhard Georg Klager and Axel Polleres. Is GPT fit for KGQA? -- preliminary results. In Proceedings of the International Workshop on Knowledge Graph Generation from Text (Text2KG2023), co-located with Extended Semantic Web Conference 2023 (ESWC 2023), May 2023.

Admittedly things have improved dramatically, since then!

So, for what are these KGs actually good for in the age of LLMs and AI?

For instance:

- Explainable answers, fact-checking against hallucinating
- Potentially less resource consumption!

Large generative model	Knowledge graph lookup
• 6 tokens input	• Find item out of 100m
• 60 tokens output (2 tokens)	• Find key out of 10k
• 96 layers	• Logarithmic operations
• 175 billion parameters	

- KGC23 Keynote: “The Future of Knowledge Graphs in a V LLMs — Denny Vrandečić, Wikimedia”

<https://www.youtube.com/watch?v=ww99npDh4cg>



ChatGPT ▾ Share

Which cities in Austria have more than 2M inhabitants?

No city in Austria has more than 2 million inhabitants. Austria's capital and largest city, **Vienna (Wien)**, has a population of approximately **1.9 million people** as of recent estimates. It is the only city in Austria with a population exceeding 1 million. Other cities, such as Graz, Linz, and Salzburg, have much smaller populations, typically below 300,000.

If you're interested in any specific details about Austrian cities, feel free to ask!

Did Vienna ever have more than 2M inhabitants?

Yes, Vienna once had a population exceeding 2 million. During the early 20th century, particularly around the year **1910**, Vienna's population peaked at approximately **2.1 million people**. This was during the Austro-Hungarian Empire, when Vienna was a major imperial capital and one of the largest cities in Europe.

After World War I and the dissolution of the Austro-Hungarian Empire in 1918, Vienna's population began to decline. This was due to several factors, including changes in national borders, economic challenges, and shifts in urbanization trends. The city's population has been growing again in recent decades but has not yet surpassed 2 million.

What's good for what? LLMs, Search Engines, KGs



Large Language Models, Knowledge Graphs and Search Engines: A Crossroads for Answering Users' Questions

Aidan Hogan
ahogan@dcc.uchile.cl
IMFD; DCC, Universidad de Chile
Santiago, Chile

Xin Luna Dong
lunadong@meta.com
Meta Reality Labs
Redmond, Washington, USA



Denny Vrandečić
denny@wikimedia.org
Wikimedia Foundation
San Francisco, California, USA

Gerhard Weikum
weikum@mpi-inf.mpg.de
Max Planck Institute for Informatics
Saarbrücken, Germany



SEARCH ENGINES ASSUME BOTH HUMANS AND MACHINES ARE STUPID

∴ WIDELY USED AND WIDELY USEFUL



KNOWLEDGE GRAPHS ASSUME MACHINES ARE STUPID AND HUMANS SMART

∴ BARELY USED AND BARELY USEFUL



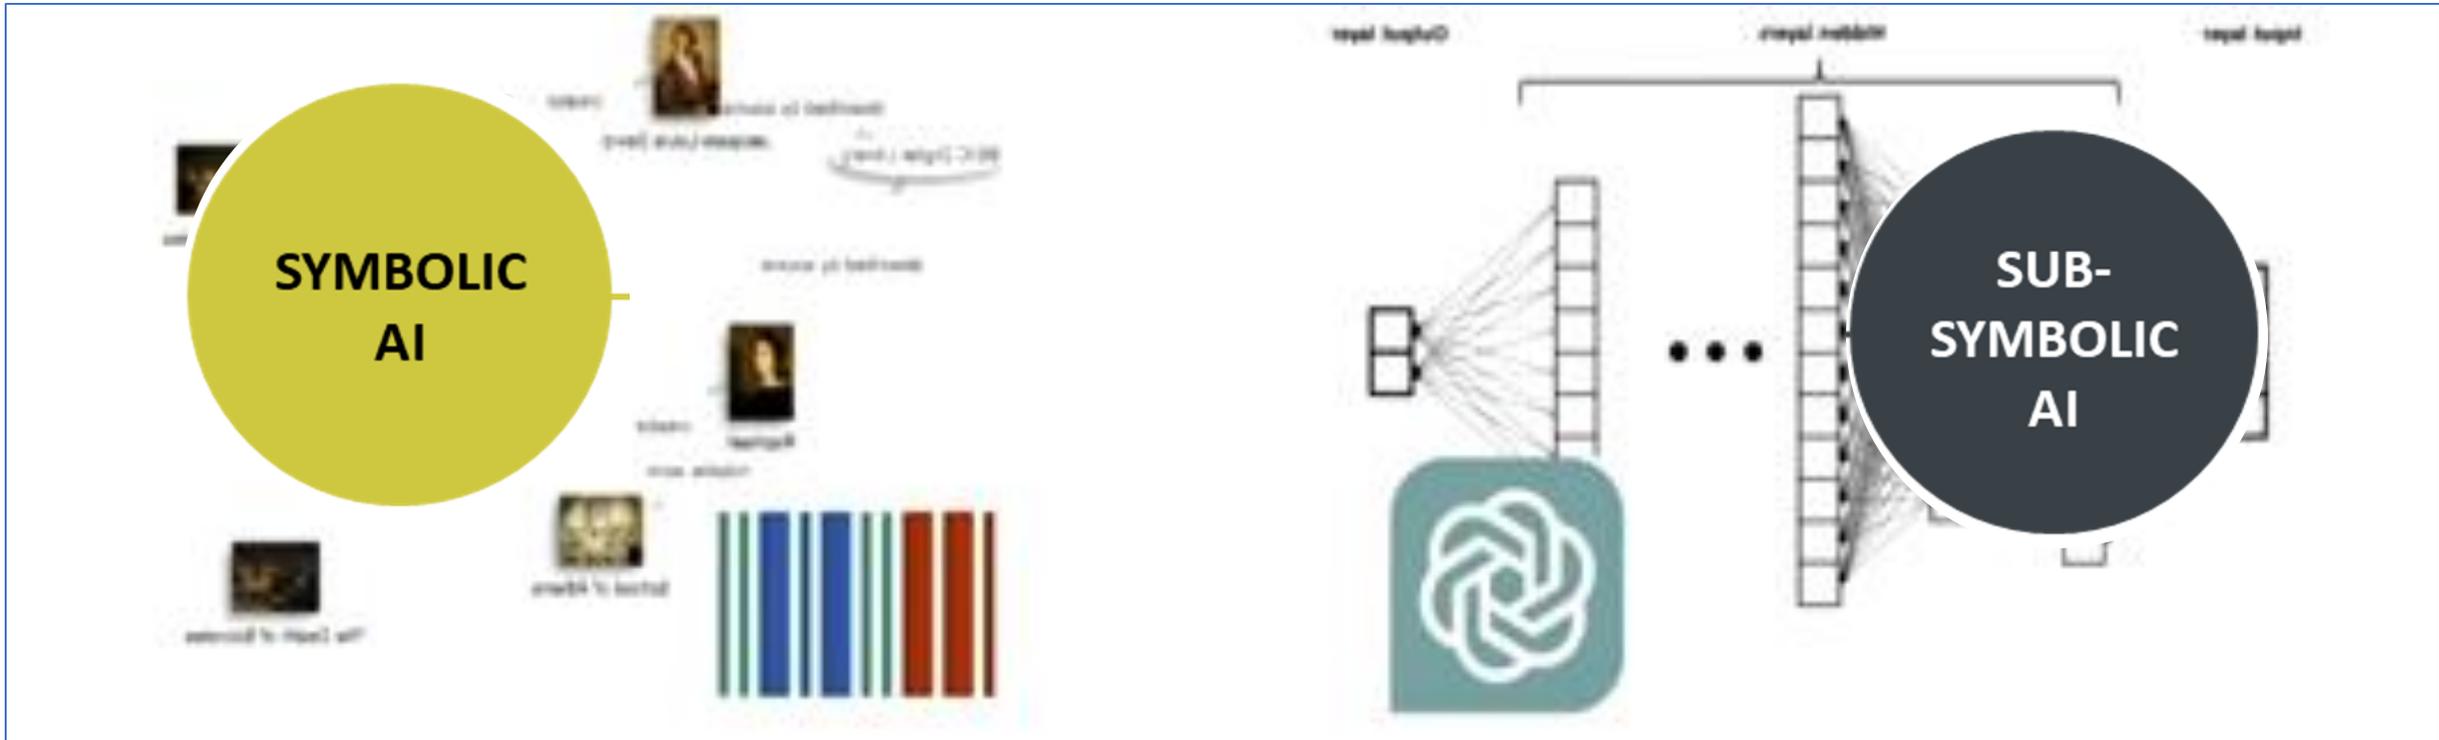
LARGE LANGUAGE MODELS ASSUME MACHINES ARE SMART AND HUMANS STUPID

∴ WIDELY USED AND BARELY USEFUL

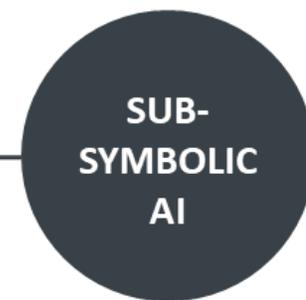
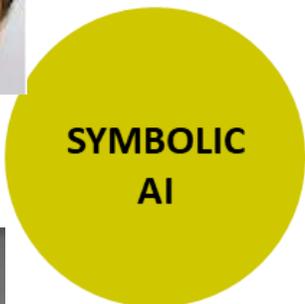


Dimension	SE	KG	LM
Precision	+ stores corpus - noisy content	+ stores corpus + precise operators	- abstracts corpus - hallucinations
Coverage	+ broad coverage	+ domain specific - patchy coverage	+ broad coverage - poor long tail
Freshness	+ quick updates + news often text	+ deprecation - structure lag	- slow updates - cold start
Generation	- no generation	+ ontologies/rule + graph learning	+ text generation
Synthesis	- no integration - no synthesis	+ data integration + synthesis	+ text integration + synthesis
Transparency	+ provenance - opaque ranks	+ algorithmic + provenance	- black box - no provenance
Determinism	+ deterministic	+ deterministic	- unstable results - randomness
Curation	+ curatable - opaque ranks	+ curatable	- indirect curation - unpredictable
Fairness	+ content as-is - bias in ranks	+ content as-is - bias in coverage	- generative biases - needs safeguards
Usability	+ natural language + simple queries	- structured - complex queries	+ natural language + conversational
Expressivity	- simple queries - ambiguity	+ complex queries - lacks nuance	+ complex queries - ambiguity
Efficiency	+ retrieval-based + simple queries	+ retrieval-based - complex queries	- inference-based - costly training
Multilingual	- lang. dependent	+ lang. agnostic - manual labels	+ multilinguality - variable results
Context	- limited context - not interactive	- limited context - not personalizable	+ in-context learn. + personalizable

So, let us better combine KGs and LLMs!



Bilateral AI



RM7: Ethical AI systems

- RM1: Graph-based structures
- RM2: Context
- RM3: Knowledge & learning
- RM4: Reasoning
- RM5: Causality
- RM6: Explainable AI

RM8: Demonstration & benchmarking



Bilateral AI



Christoph Lampert
Machine Learning
Trustworthy Learning

- Machine Learning and Computer Vision group
- ELLIS Unit ISTA



TECHNISCHE UNIVERSITÄT WIEN



Agata Ciabattoni
Logic Reasoning



Thomas Eiter
Symbolic AI
Knowledge representation

- Institute for Logic and Computation



Axel Polleres
Knowledge Graphs

- Institute for Data Process and Knowledge Management



Gerhard Friedrich
Symbolic AI
Model-based reasoning

- Institute for Artificial Intelligence and Cybersecurity



Robert Legenstein
Machine Learning
Computational Neuroscience

- Institute of Machine Learning and Neural Computation

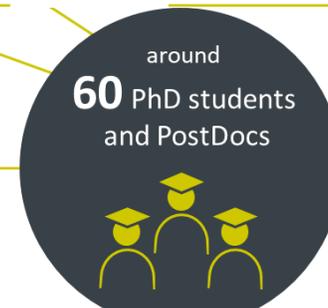
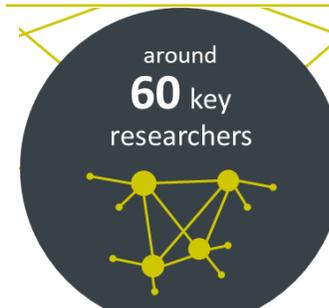
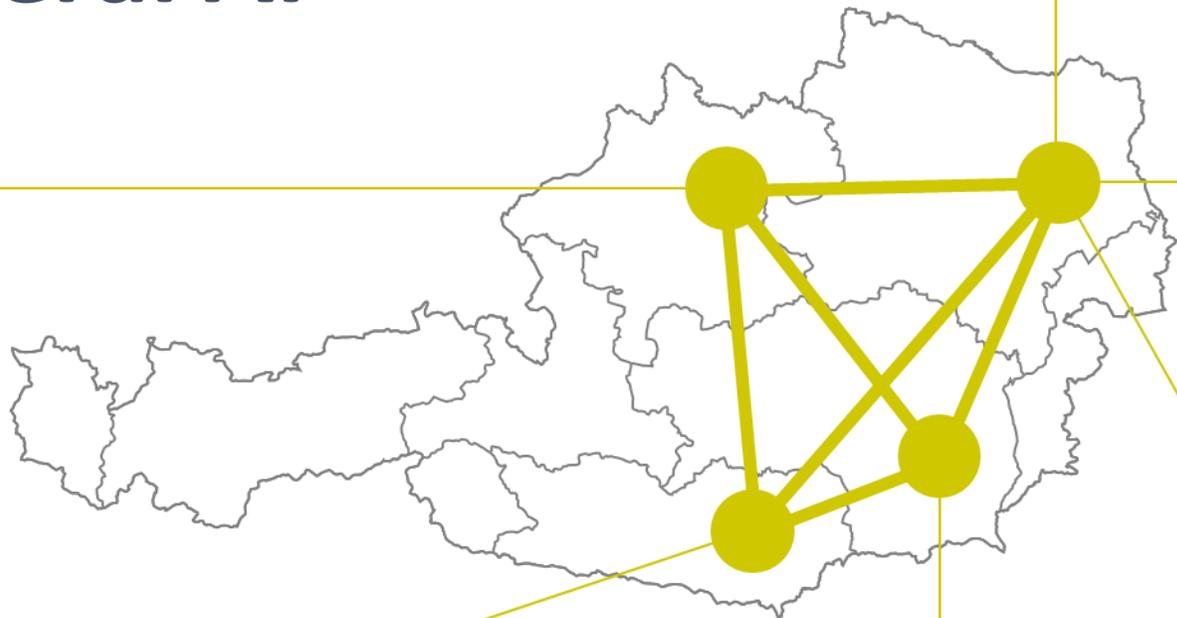


Martina Seidl
Symbolic AI
SAT Solving
Formal methods



Sepp Hochreiter
Machine Learning
LSTM
Vanishing gradient

- Institute for Machine Learning
- ELLIS Unit Linz
- LIT AI Lab
- Institute for Symbolic Artificial Intelligence



BILATERAL AI – First results and plans from our Institute and other BILAI partners...



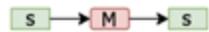
What's next?



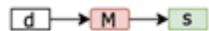
Ingredient 1: We have a pretty good understanding of Engineering **Neuro-Symbolic Systems...**

Atomic-Patterns

A1: s-M-s



A2: d-M-s

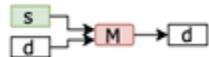


A3: s-M-d

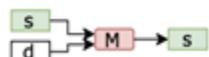


Fusion-Patterns

F1: d/s-M-d



F2: d/s-M-s

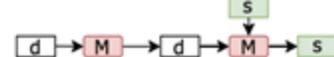


F3: d/s-M-d/s

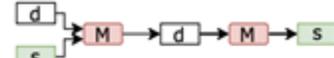


T-Patterns

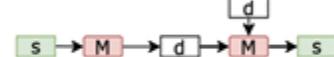
T1: {d-M-d/s}-M-s



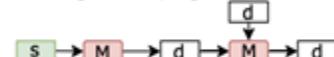
T2: {d/s}-M-d-M-s



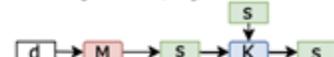
T3: {s-M-d/d}-M-s



T4: {s-M-d/d}-M-d



T5: {d-M-s/s}-K-s

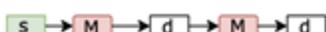


I-Patterns

I1: s-M-d-M-s

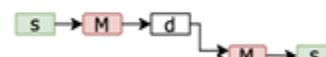


I2: s-M-d-M-d

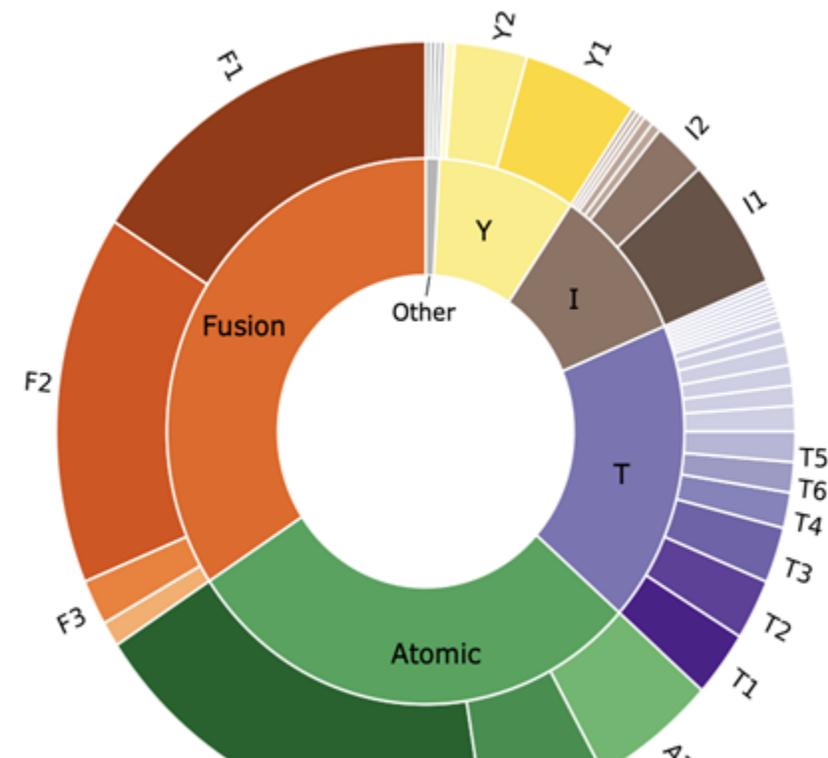
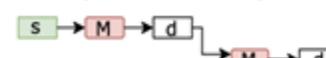


Y-Patterns

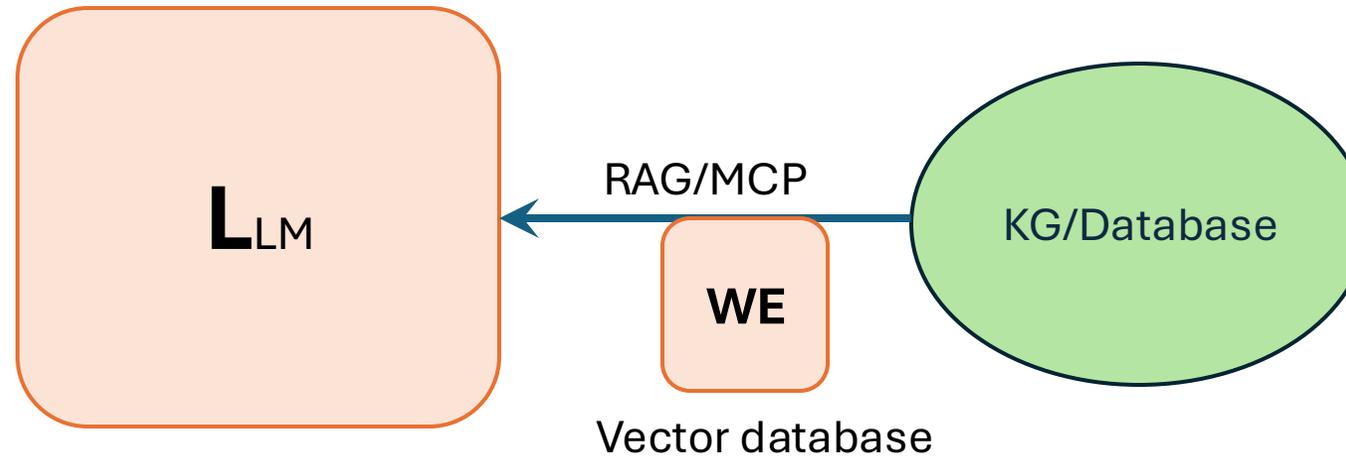
Y1: {s-M-d/d-M-d}-M-s



Y2: {s-M-d/d-M-d}-M-d



One such architecture... For instance:



LLMs get better ... but still (plausibly) hallucinate

“are supposed to be PhD level”



AP Please give me the top 10 publications by Axel Polleres

I'll search for information about Axel Polleres' top publications.

Axel Polleres top publications 10 results

Let me fetch his Google Scholar page to see his most cited publications.

Failed to fetch https://scholar.google.com/citations?user=R-SmVOYAAAAJ&hl=en

Axel Polleres most cited papers SPARQL RDF 10 results

Let me fetch his DBLP page which should have a comprehensive list of his publications.

Fetches Axel Polleres dblp.org

Based on Axel Polleres' extensive publication record with over 18,778 citations (Axel Polleres), here are his top 10 most influential publications:

Top 10 Publications by Axel Polleres

1. **"Knowledge Graphs"** (2021/2022) Co-authored with Aidan Hogan, Eva Blomqvist, and others, published in ACM Computing Surveys and as a book in the Synthesis Lectures series [dblp: Axel Polleres](#) ok
2. **"A survey of RDF stores & SPARQL engines for querying knowledge graphs"** (2022) Published in The VLDB Journal [A survey of RDF stores & S...](#) !
3. **"From SPARQL to rules (and back)"** (2007) A highly cited paper on the relationship between SPARQL and rule systems [ResearchGate](#) [ACM Digital Library](#) ok
4. **"Evaluating query and storage strategies for RDF archives"** (2019) Co-authored with Javier D. Fernández, Jürgen Umbrich, and Magnus Knuth [dblp: Axel Polleres](#) ok
5. **"Binary RDF representation for publication and exchange (HDT)"** (2013) Published in Journal of Web Semantics, introducing the HDT format for RDF compression [dblp: Axel Polleres](#) ok
6. **"A fine-grained evaluation of SPARQL endpoint federation systems"** (2016) Co-authored with Muhammad Saleem and others, evaluating federated SPARQL query systems [A fine-grained evaluation o...](#) !
7. **"Federating queries in SPARQL 1.1: Syntax, semantics and evaluation"** (2013) Co-authored with Carlos Buil-Aranda, Marcelo Arenas, and Óscar Corcho [dblp: Axel Polleres](#) ok
8. **"RDFS and OWL Reasoning for Linked Data"** (2013) Published in Reasoning Web, co-authored with Aidan Hogan, Renaud Delbru, and Jürgen Umbrich [dblp: Axel Polleres](#) ok
9. **"Mapping between RDF and XML with XSPARQL"** (2012) Co-authored with Stefan Bischof, Stefan Decker, Thomas Krennwallner, and Nuno Lopes [dblp: Axel Polleres](#) ok
10. **"Everything you always wanted to know about blank nodes"** (2014) Published in Journal of Web Semantics with Aidan Hogan, Marcelo Arenas, and Alejandro Mallea ok

Here are **10 influential publications** by Axel Polleres, including papers, book-chapters, and working documents. These are among his most cited or otherwise important contributions. If you want, I can also fetch full download links or more recent ones.

Top 10 Publications by Axel Polleres

#	Title	Co-authors / Where / When	
1	<i>Web Service Modeling Ontology (WSMO)</i>	Jos de Bruijn, Dieter Fensel, Uwe Keller, Holger Lausen, Rubén Lara, Michael Kifer, Martin Hepp, Birgitta König-Ries, Jacek Kopecky, Dumitru Roman, John Domingue, Eyal Oren + Axel Polleres etc. (W3C Member Submission, 2005) w3c +3	ok
2	<i>Binary RDF representation for publication and exchange (HDT)</i>	Javier D. Fernández, Miguel A. Martínez-Prieto, Claudio Gutiérrez, Axel Polleres, Mario Arias (Journal of Web Semantics, 2013) repositorio.latin... +4	ok
3	<i>From SPARQL to Rules (and back)</i>	Axel Polleres (WWW 2007) dl.acm.org +2	ok
4	<i>SPARQL1.1: New Features and Friends (OWL2, RIF)</i>	Axel Polleres (in Web Reasoning and Rule Systems (RR 2010), Lecture Notes in Computer Science) link.springer.com +2	ok
5	<i>A General Framework for Representing, Reasoning and Querying with Annotated Semantic Web Data</i>	Antoine Zimmermann, Nuno Lopes, Axel Polleres, Umberto Straccia (2011) arxiv.org +1	ok
6	<i>Are We Better Off With Just One Ontology on the Web?</i>	Armin Haller, Axel Polleres (Semantic Web Journal, 2020) WU Wirtschafts... +2	ok
7	<i>Evaluating Query and Storage Strategies for RDF Archives</i>	Javier David Fernández García, Jürgen Umbrich, Axel Polleres, Magnus Knuth (Date ~ 2015) WU Wirtschafts... +1	ok
8	<i>Serializing RDF in Compressed Space (HDT++)</i>	Antonio Hernández-Ilerra, Miguel A. Martínez-Prieto, Javier David Fernández, Axel Polleres WU Wirtschafts... +1	!
9	<i>How (Well) Do Datalog, SPARQL and RIF Interplay?</i>	Axel Polleres (in Datalog 2.0 2012, LNCS) link.springer.com +1	ok
10	<i>Semantic Web Languages and Semantic Web Services as Application Areas for Answer Set Programming</i>	Axel Polleres (Dagstuhl Seminar Proceedings, 2005) drops.dagstuhl.de +1	ok



Inspired by Georg Gottlob's SEMANTICS2025 keynote

Heng Ji's SEMANTICS2025 keynote

“Larger models are not necessarily better!”

General Data management issue,
not KG-specific!

But are KGs (& Databases) actually better?

Google asks it's
users to manually
correct/merge
results

DBLP or scopus
attempt to curate
ambiguities by
constraints (not always
successful..)



The screenshot shows a research profile for Axel Polleres at the Vienna University of Economics and Business. A 'Merged citations' window is open, displaying a list of articles with their citation counts. A 'Search dblp' window is also visible, showing author search results for 'Thomas Müller' with various disambiguation options.

Profile: Axel Polleres, Vienna University of Economics and Business (WU - Wirtschaftsuniversität Wien). Research interests: Data & Knowledge Management, Artificial Intelligence, Logic Programming, Semantic Web, Knowledge Graphs.

Merged citations:

Article Title	Citation Count
Knowledge graphs A Hogan, E Blomqvist, M Cochez, C d'Amato, GD Melo, C Gutierrez, ... ACM Computing Surveys (Csur) 54 (4), 1-37, 2021	2591
Anisa Rula, Lukas Schmelzeisen, Juan F. Sequeda, Steffen Staab, and Antoine Zimmermann A Hogan, E Blomqvist, M Cochez, C d'Amato, G De Melo, C Gutierrez, ... Knowledge graphs, 2020	344 *
Knowledge graphs H Aidan, E Blomqvist, M Cochez, C d'Amato, G de Melo, C Gutiérrez, ... Mines Saint-Etienne, 2020	21 *
Knowledge graphs. CoRR abs/2003.02320 (2020) A Hogan, E Blomqvist, M Cochez, C d'Amato, G de Melo, C Gutiérrez, ... URL: https://arxiv.org/abs/2003.02320, 2003	13 *

Search dblp results:

- Thomas Müller — *disambiguation page*
- Thomas Müller⁰⁰⁰¹
TU Munich, Faculty of Computer Science, Germany
- Thomas Müller⁰⁰⁰²
Eastphalia University of Applied Sciences, Braunschweig/Wolfsbüttel, Germany
- Thomas Müller⁰⁰⁰³
Friedrich Schiller University Jena, Institute of Computer Science, Germany
- Thomas Müller⁰⁰⁰⁴
Fraunhofer Institute for Optronics, System Technologies and Image Exploitation (IOSB), Karlsruhe, Germany
- show all*

Likely matches:

- Thomas Müller⁰⁰¹³ 
NVIDIA Research, Zürich, Switzerland
- Thomas Müller-Gronbach
University of Passau, Germany
- Thomas Müller⁰⁰⁰⁹ 
Google Research
- Thomas Müller⁰⁰⁰⁷ 
Department of Philosophy, University of Konstanz, Germany



But are KGs actually better?

Example from Wikidata...



Tilikum (Q127236)

[Item](#) [Discussion](#)

[Read](#) [View history](#) [☆](#)

dugout canoe

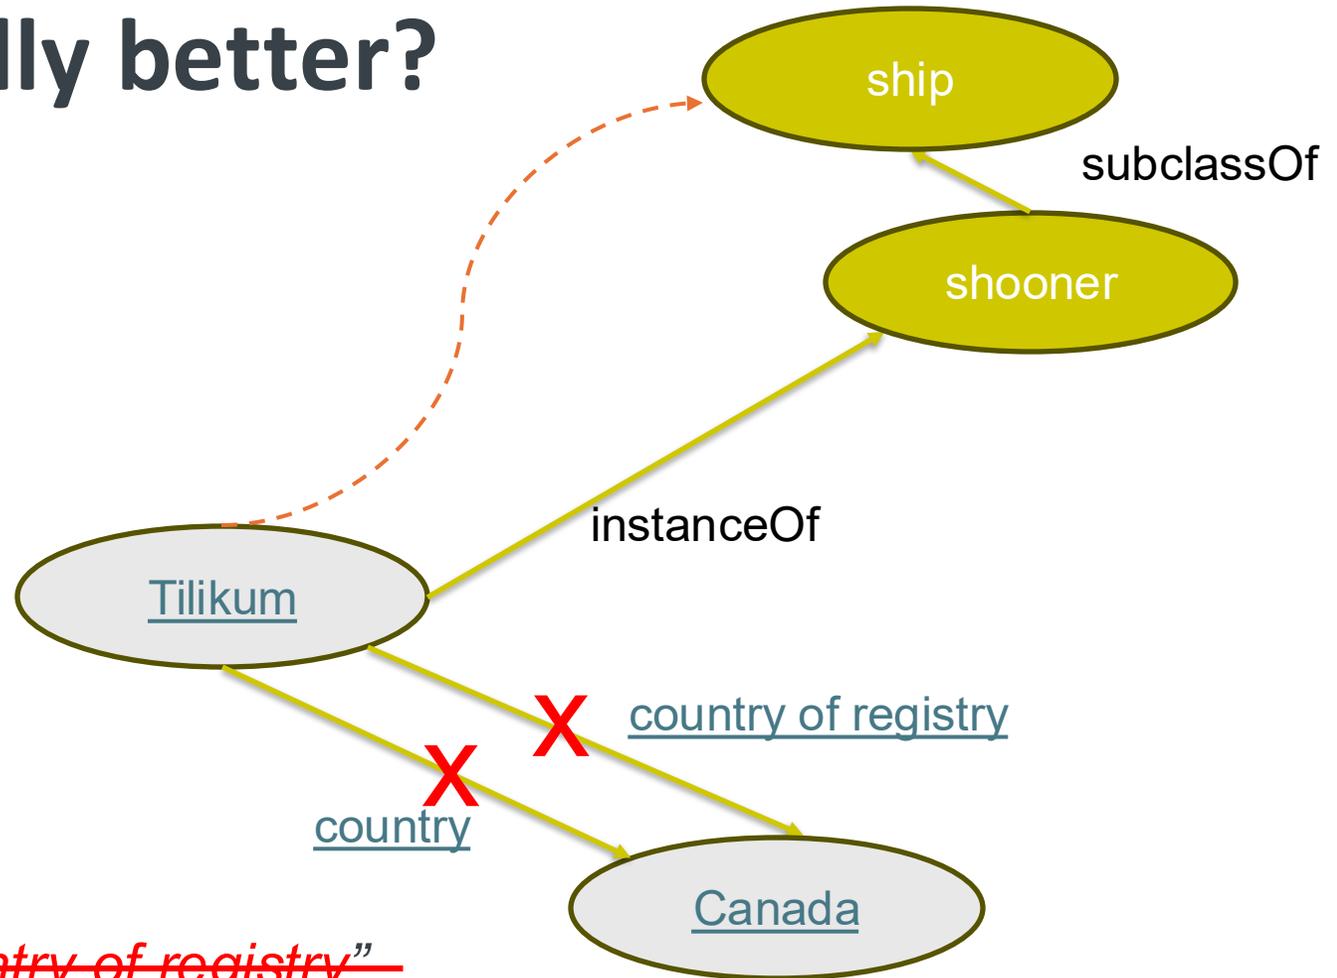
[edit](#)

[In more languages](#)

[Configure](#)

Language	Label	Description	Also known as
default for all languages	No label defined	–	
English	Tilikum	dugout canoe	
British English	Tilikum	No description defined	
American English	No label defined	No description defined	
German	Tilikum	Segelboot	

But are KGs actually better?



Ontologies + Constraints,
such as:

~~“country conflicts with country of registry”~~

- “i.e., an entity should not have both a “country” and a “country of registry”

Why KGs can't (scalably) be correct:

Collaborative Editing

Multi-sourced Automatic Extraction

Knowledge is distributed and evolving!

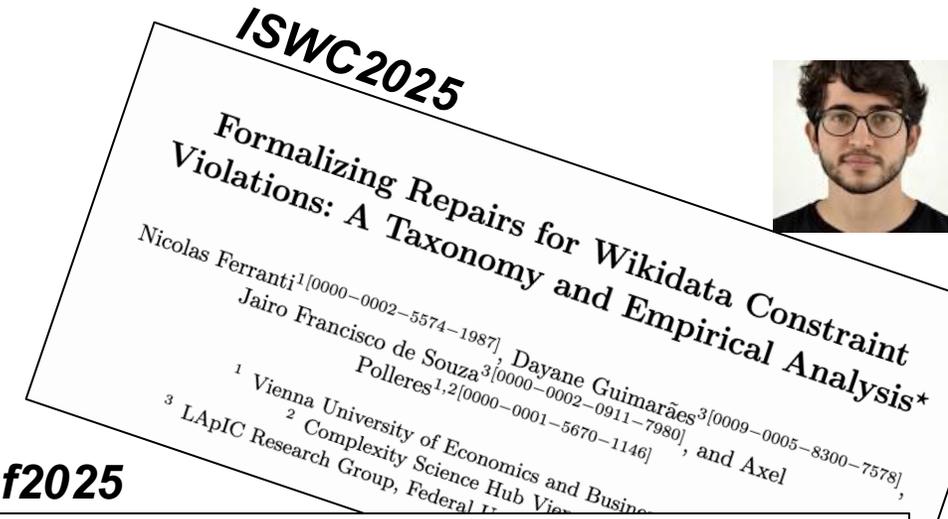
- *Remote changes affect consistency*
- *Ontologies & Constraints evolve as well*

What can we do about it?

- Best of both worlds needs **Bilateral AI** approaches!
- **Symbolic** methods to repair inconsistencies

Forthcoming...

- **Subsymbolic** methods to resolve inconsistencies
 - **LLMs & KG embeddings** for resolving inconsistencies
 - **GNNs** for *learning* repairs (*from historical data, from users' repairs*)
 - Iterative **RAG** and **Agent-based** pipelines to fix and construct KGs from text



Webconf2025



Further starting points in BILAI:

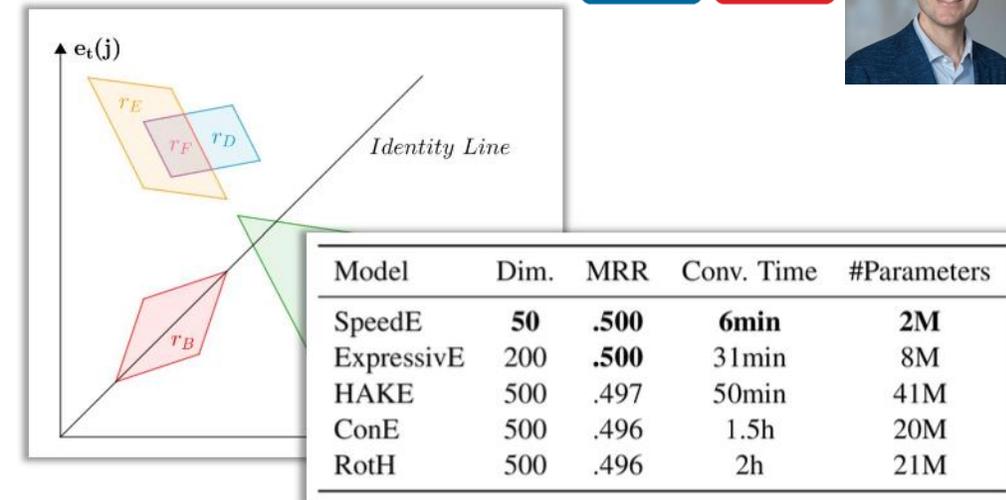
Notably, in BILAI, colleagues from TU Vienna (Sallinger, Pavlovic) work on KG Embeddings that can capture (certain) rules and constraints already:

- Similar to Word embeddings and LLMs, **Knowledge Graph Embeddings** allow to
 - predict missing edges in incomplete KGs
 - predict inconsistencies & possible repairs

Open Problems:

- Scaling KG Embeddings to full, decentralized KGs ...
- ... but (1) **modularization** might help here, relation to the
- (2) corresponding trend to LLMs-based “multi-agent frameworks”

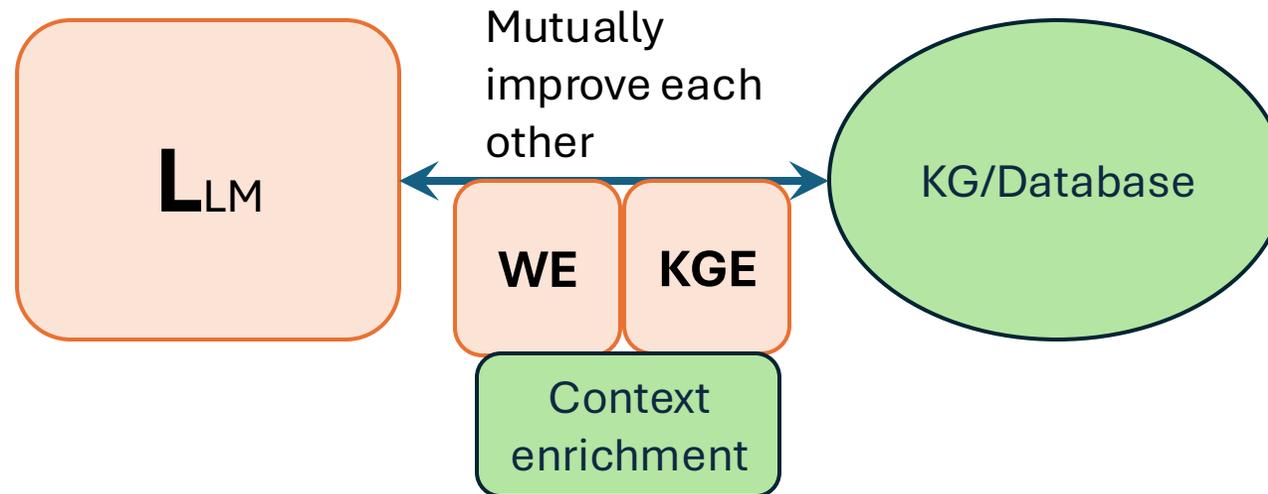
Slide: Emanuel Sallinger



Logical Rule	ExpressivE	RotatE	TransE	BoxE	ComplEx	DistMult
Symmetry: $r_1(X, Y) \Rightarrow r_1(Y, X)$	✓	✓	✗	✓	✓	✓
Anti-symmetry: $r_1(X, Y) \Rightarrow \neg r_1(Y, X)$	✓	✓	✓	✓	✓	✗
Inversion: $r_1(X, Y) \Leftrightarrow r_2(Y, X)$	✓	✓	✓	✓	✓	✗
Comp. def.: $r_1(X, Y) \wedge r_2(Y, Z) \Leftrightarrow r_3(X, Z)$	✓	✓	✓	✗	✗	✗
Gen. comp.: $r_1(X, Y) \wedge r_2(Y, Z) \Rightarrow r_3(X, Z)$	✓	✗	✗	✗	✗	✗
Hierarchy: $r_1(X, Y) \Rightarrow r_2(X, Y)$	✓	✗	✗	✓	✓	✓
Intersection: $r_1(X, Y) \wedge r_2(X, Y) \Rightarrow r_3(X, Y)$	✓	✓	✓	✓	✗	✗
Mutual exclusion: $r_1(X, Y) \wedge r_2(X, Y) \Rightarrow \perp$	✓	✓	✓	✓	✓	✓



So we actually need this?



...

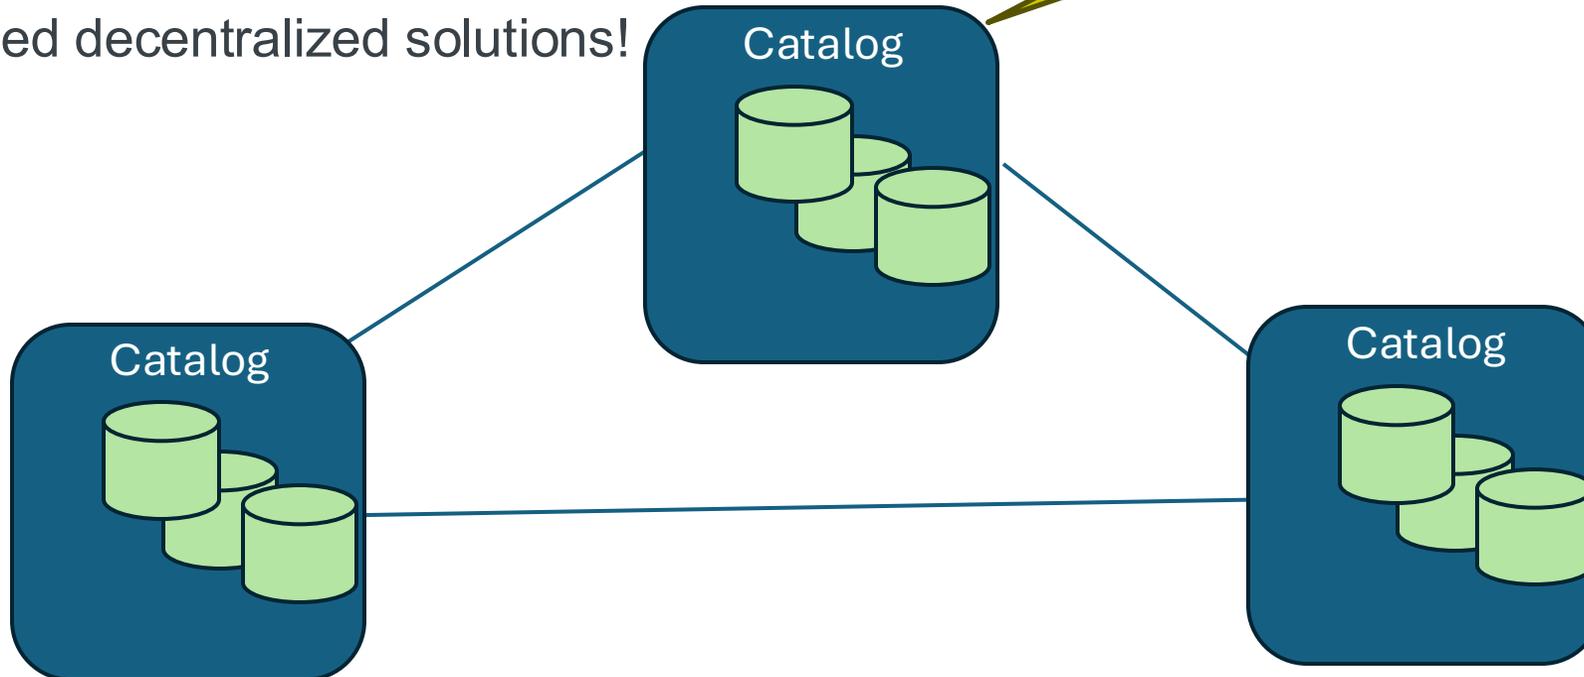
Quite some work at current ISWC2025, e.g. RAGE-KG workshop is looking into such extensions...

What's next? Where should this go?

Caveat: Our data and knowledge is typically:

- **Distributed/decentralized** and
- **not always sharable** with a central (API) model
- → we need decentralized solutions!

“Our knowledge will never go into a central LLM” – BILAI industry network keynote Stefan Rohringer (Infineon)



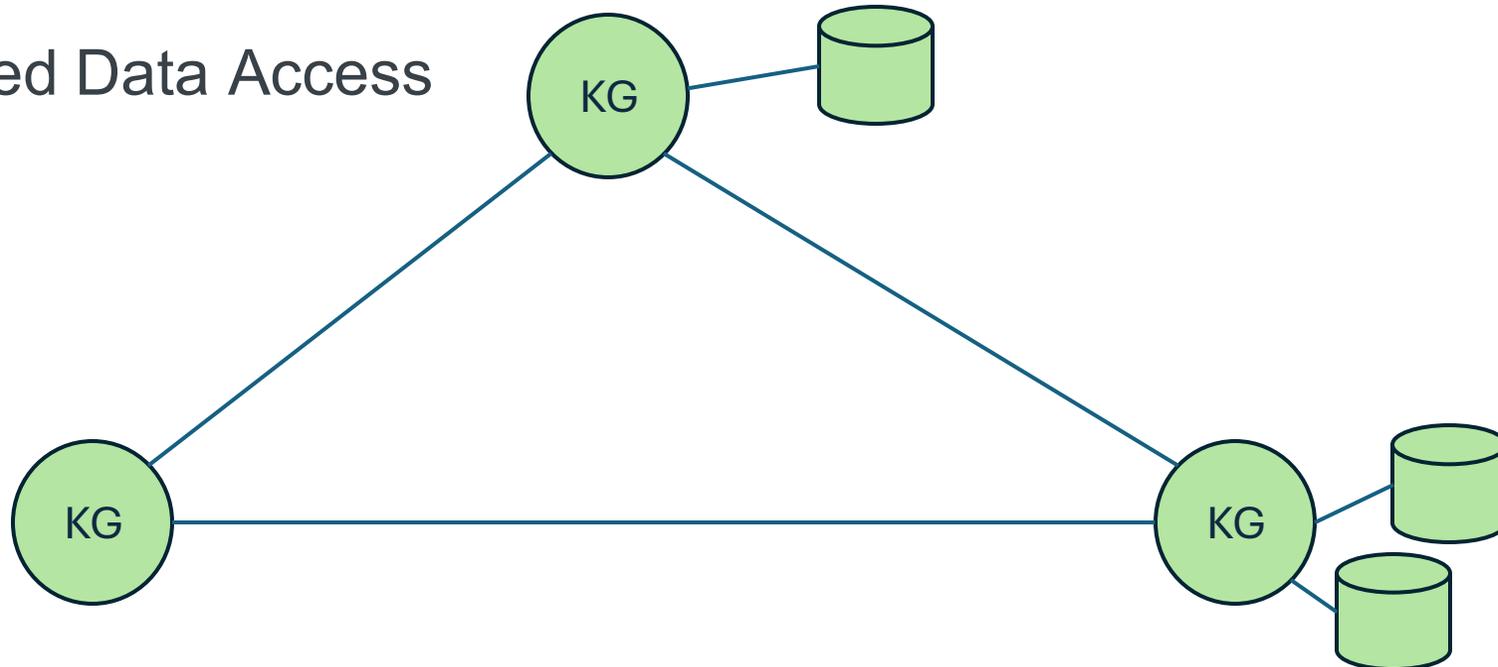
Ingredient 4: “Data Spaces”



What's next? Where should this go?

Ingredient 2: Expertise on (Symbolic) Data Integration and Repair. Before the LLM hype, we long worked on *purely symbolic* solutions to integrating *decentralised* Knowledge

- Linked Data
- Ontology-based Data Access



Scalable data access to decentralized KGs:



Starting Points:

- graph partitioning and splitting processing between clients and SPARQL endpoints the performance, the problems of central query endpoints can be significantly alleviated
- Investigating how the *asynchronous evolution* of KGs affects consistency

Interesting Open Problems:

- E.g. Synchronisation/Updates/Repairs in Polyglott Databases



SWJ (2024) <https://doi.org/10.3233/SW-243571>

Wikidata Query Service

```

1 # Which settlements had ever over 1M population?
2
3 SELECT DISTINCT ?City ?Country ?Pop (Min(?Timepoint) AS ?FirstTime)
4 {
5   ?City wdt:P17 ?Country;
6         wdt:P31/wdt:P279* wd:Q486972;
7         p:P1082 ?Stmnt.
8   ?Stmnt ps:P1082 ?Pop;
9         pq:P585 ?Timepoint.
10  FILTER (?Pop > 1000000)
11
12 } GROUP BY ?City ?Country ?Pop
  
```

Query timeout limit reached

SHORT-PAPER | OPEN ACCESS

The Massive Problem of Remote Changes in Ontology Reuse

Authors: Romania Pernisch, Daniil Dobryi, Axel Polleres | [Authors Info & Claims](#)

WWW '25: Companion Proceedings of the ACM on Web Conference 2025 • Pages 1254 - 1258 • <https://doi.org/10.1145/3701716.3715478>



Isn't there something missing?



Ora Lassila (Keynote ISWC2024)

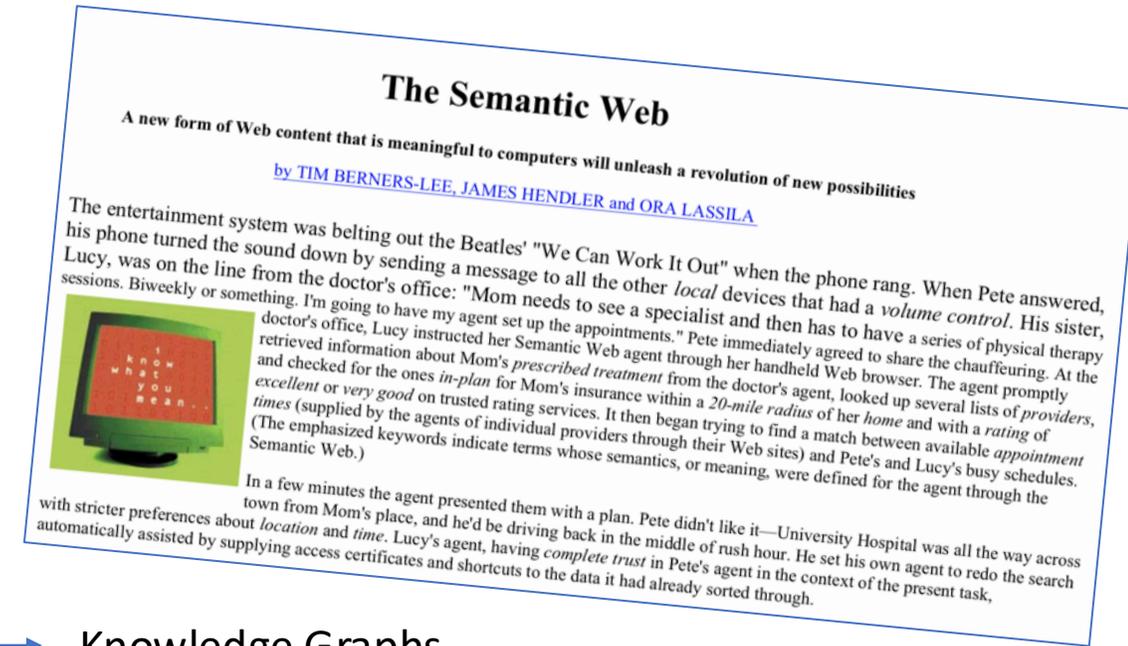
<https://www.lassila.org/publications/2024/lassila-iswc2024-keynote.pdf>

The realization of autonomous agents minimally requires these:

1. KR & reasoning
2. Planning
3. Ability to converse with the agents

LLMs will give you #3, but not #1 or #2

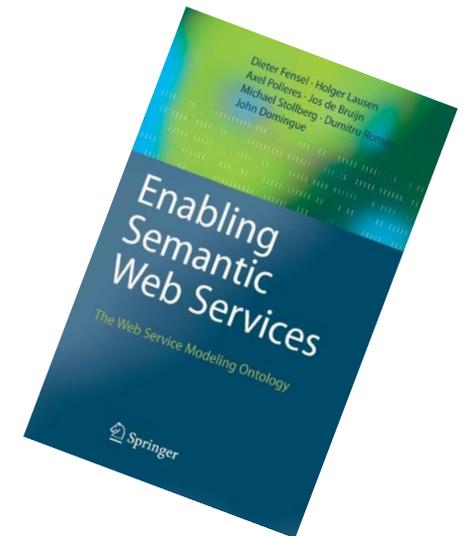
- (despite what you hear)
- "agentic", "agentive", ... huh?
- good news: we already have #1 and #2



Knowledge Graphs

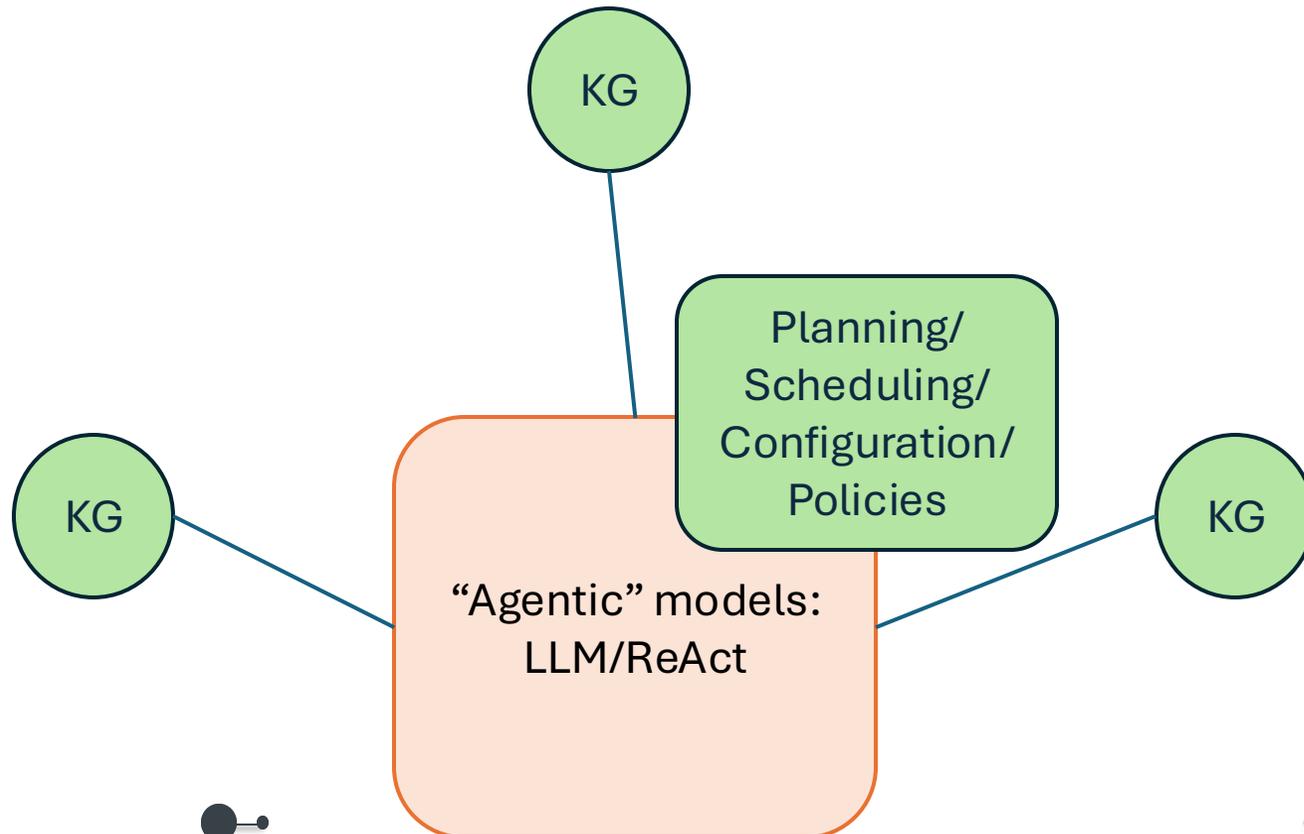
"Agentic" Models ... and maybe a revival?

LLMs



What's next? Where should this go?

Ingredient3: *Combinations* of LLMs and decentral KGs via **RAG** and **Agent frameworks** are one example of such Neuro-symbolic architectures:



Example: Federated querying of KGs via agents (work in progress):

Prompt: “Write a federated SPARQL query that queries all of Tim Berners-Lee’s publications from DBLP the DOIs of which also appear on Wikidata”

```

PREFIX dblp: < http://dblp.org/rdf/schema#>
PREFIX bibo: <http://purl.org/ontology/bibo/>
PREFIX wdt: <http://www.wikidata.org/prop/direct/>
SELECT DISTINCT ?dblpPub ?title ?doi ?wikidataItem
WHERE {
  SERVICE < https://dblp.org/sparql > {
    ?dblpPub a bibo:Article ;
    dblp: author ?author ;
    dblp:title ?title ;
    bibo: doi ?doi .
    ?author dblp: name "Tim Berners-Lee" .}
  SERVICE <https://query.wikidata.org/sparql> {
    ?wikidataItem wdt:P356 ?doi . } }
  
```

What's next? Where should this go?

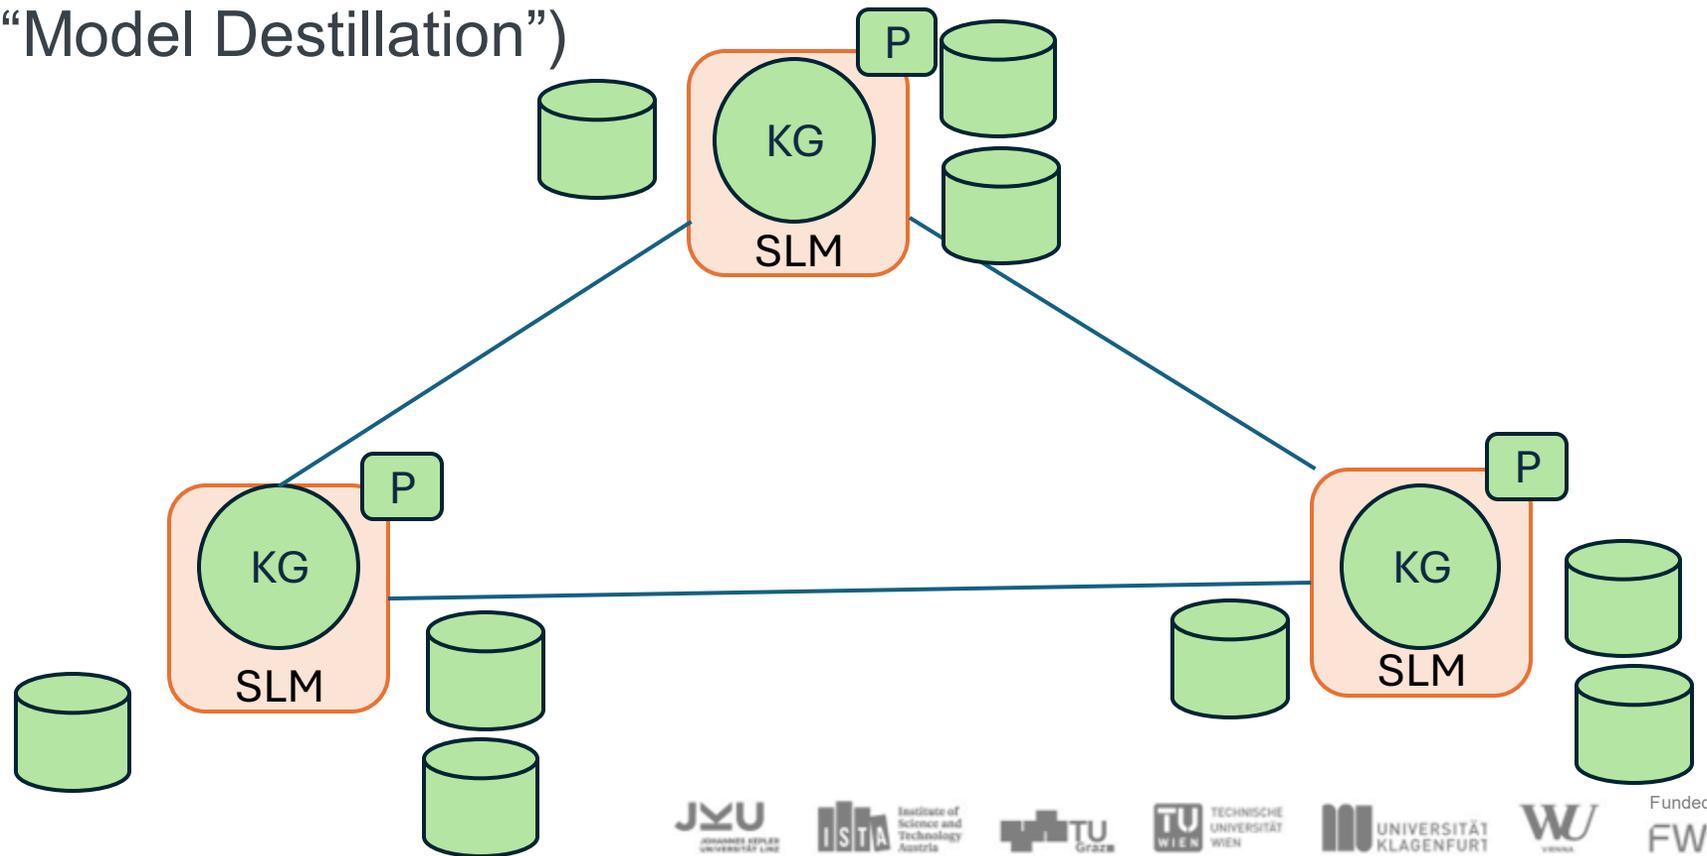
What is a solution we should jointly work towards?

Neuro-symbolic agents

Trustable Decentralized Agents

Small models (e.g. "Model Distillation")

"Agent Spaces"



Thanks!

Austrian National “Cluster of Excellence” BILAI:

- Broad AI
- (Knowledge) Graph-Based AI plays a key role in BILAI

Knowledge Graphs are symbolic AI with “features” of LLMs:

- Large-scale
- Not always consistent

(**Dezentralized**) Agents, KGs and LLM combined, could realize BILAI’s vision of a broad, **robust** AI, and change the way we access and manage Data

- **Ingredient1:**
Neurosymbolic Architectures
- **Ingredient2:**
scalable distributed, federated data symbolic processing



Currently running!