

Knowledge Graphs – a key component in Bilateral AI



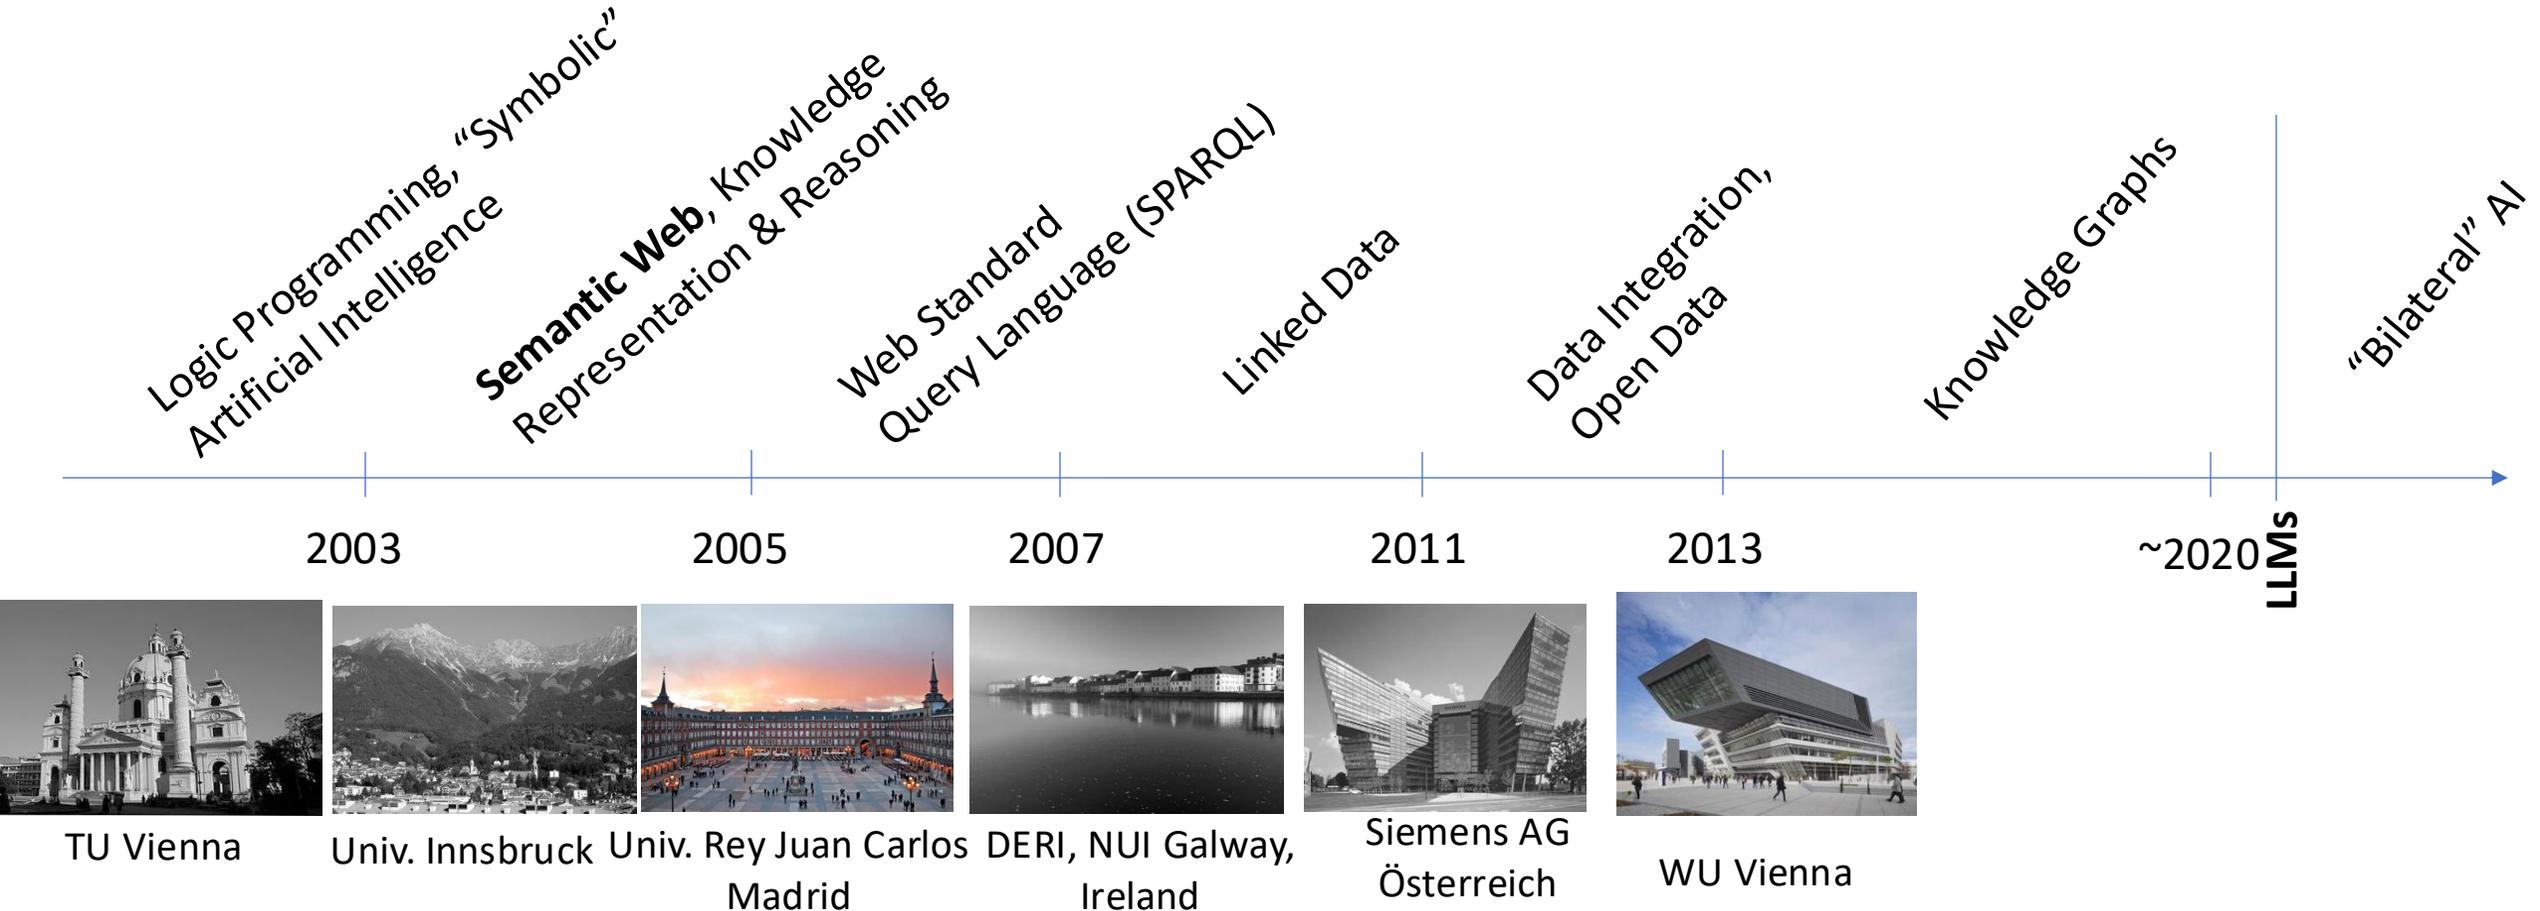
Axel Polleres



Institute for Data , Process and Knowledge Management



Great to be back!



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...

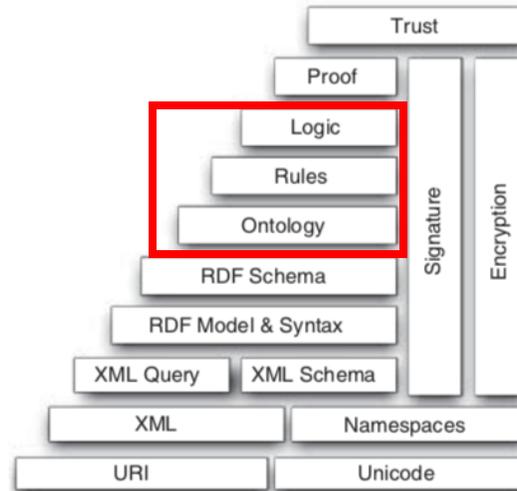
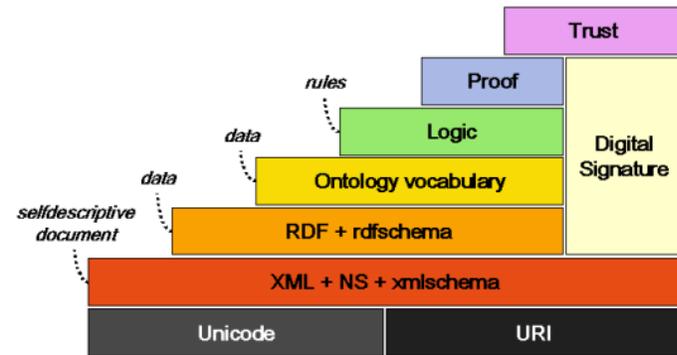
*" huge knowledge bases, also known as **knowledge graphs**, have been automatically constructed from web data, and have become a key asset for search engines and other use cases.*

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

Semantic Web: Standard formats, Reasoning & Logics

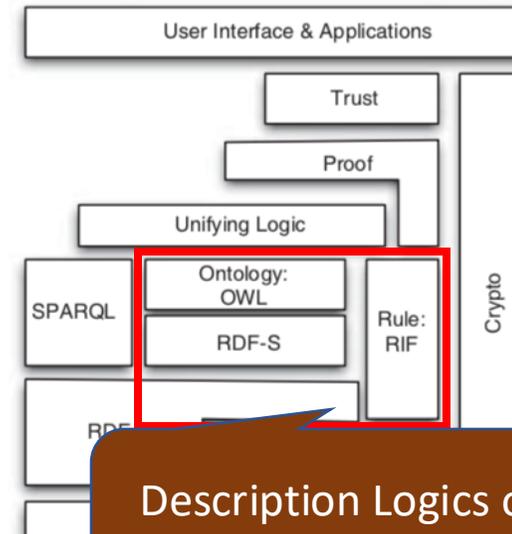


- (2000s - ca. 2009)

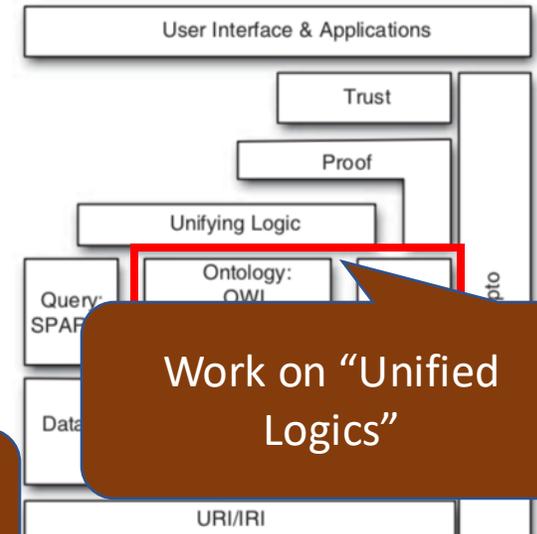


2004^a

^ahttp://www.w3.org/2004/Talks/0319-RDF-WGs/sw_stack...
^b<http://www.w3.org/2007/Talks/0130-sb-W3CTechSemWeb/layerCake-4.png>
^c<http://www.w3c.it/talks/2009/athena/images/layerCake.png>



Description Logics or Rules?



Work on "Unified Logics"

(Contextualized Reasoning?)

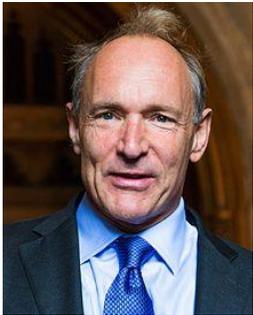
2009^c

Good news! Boost in KR/AI research:
We know very well which ontological reasoning approaches are decidable and how they scale
→ OWL, OBDA, but also: constraint checking (SHACL)



Focus on Data: Linked Data

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



Linked Data Principles

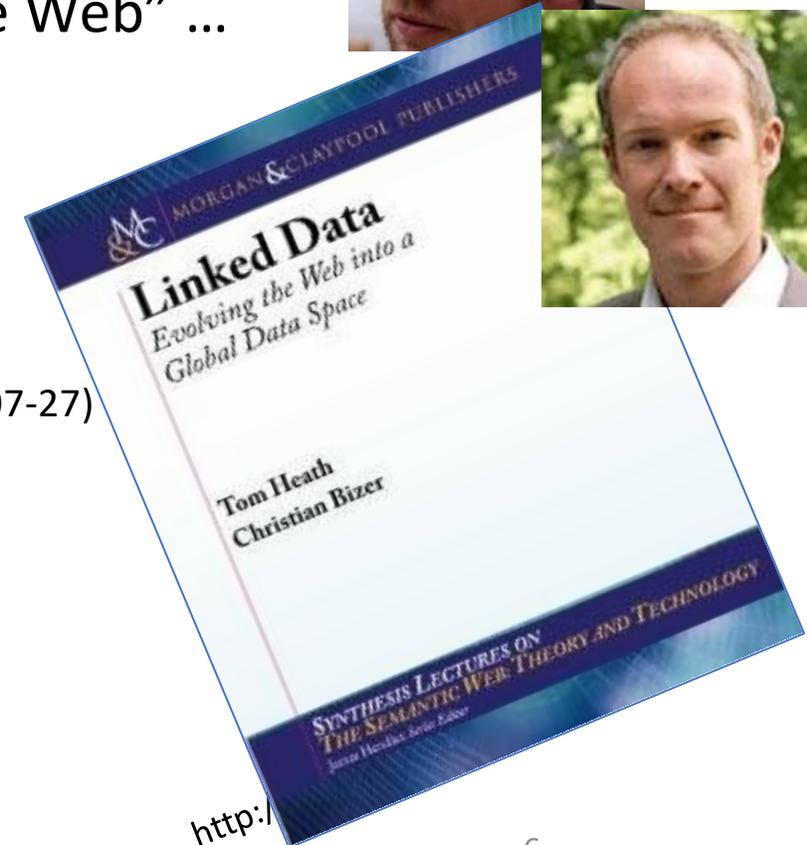
- **LDP1**: use URIs as names for things
- **LDP2**: use HTTP URIs so those names can be dereferenced
- **LDP3**: return useful – RDF? – information upon dereferencing those URIs
- **LDP4**: include links using externally dereferenceable URIs.

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



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

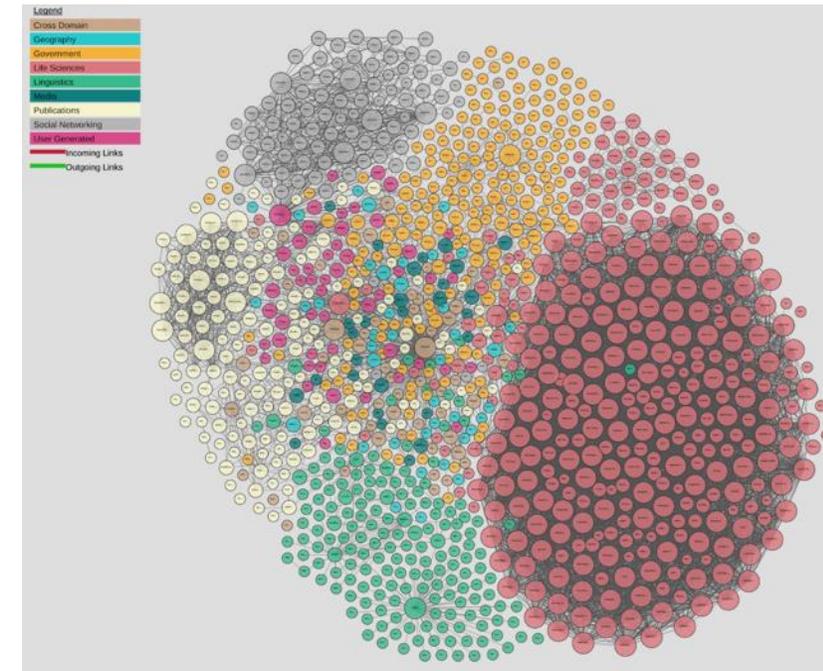
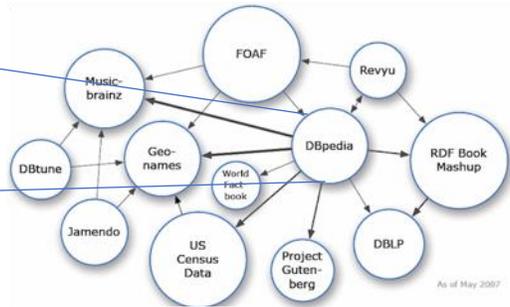
<https://www.cs.rpi.edu/~hendler/LittleSemanticsWeb.html>





From Semantic Web to Linked (Open) Data

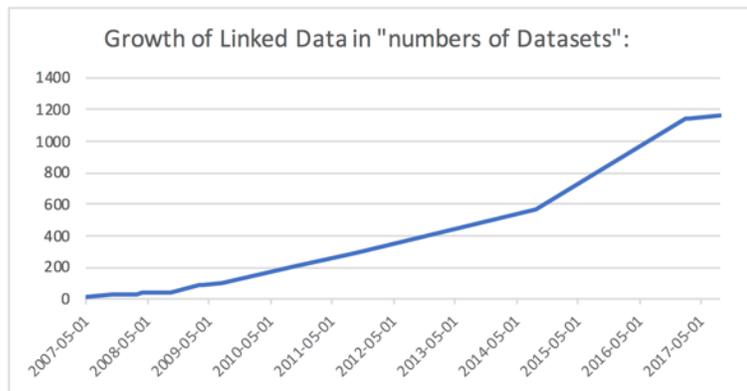
- (ca. 2006/7 – 2012)
 - Main question: How can I **publish** “Knowledge on the Web” ...
 - Linked **Open** Data... growth slowed down a bit
 - A lot of active developments to publish and link RDF Data
 - also in Enterprises (“Enterprise Linked Data”)



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

2017-08-22
2017-02-20
2017-01-26
2014-08-30
2011-09-19
2010-09-22
2009-07-14
2009-03-27
2009-03-05
2008-09-18
2008-03-31
2008-02-28
2007-11-10
2007-11-07
2007-10-08
2007-05-01

1163
1139
1146
570
295
203
95
93
89
45
34
32
28
28
25
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)*.

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

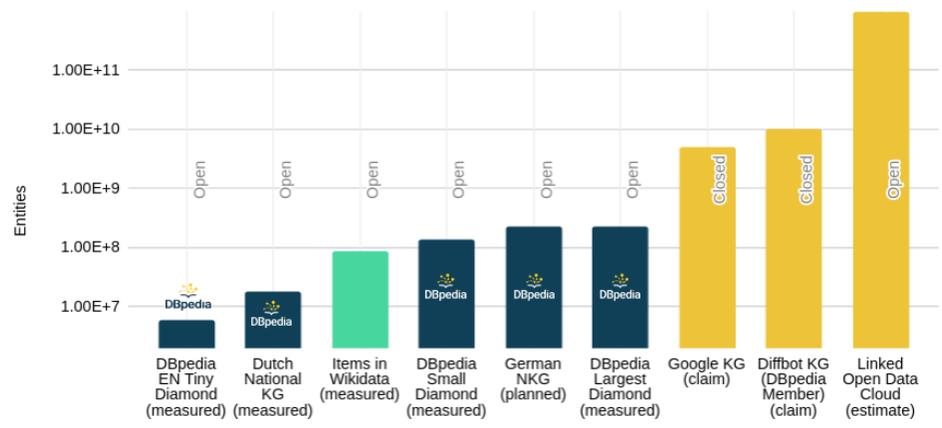
1 Billion Entities
70 Billion Assertions

- 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...



Collaborative, Open Knowledge Graphs:

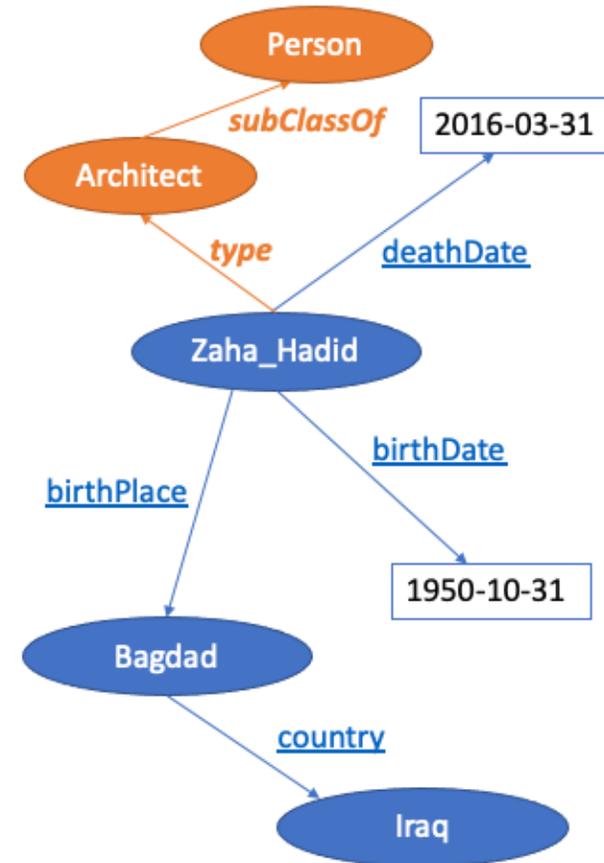


DBpedia generates a graph from links and facts in Wikipedia's Infoboxes:

http://wikipedia.org/wiki/Zaha_Hadid

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



occupation: association football player
2 references

FIFA player ID (archived): 229397
1 reference

country of citizenship: Argentina
start time: 1987
0 references

country of citizenship: Spain
start time: 2005

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 द्वाद्गन 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 [autofix](#) на / on [Property talk:P12924](#))

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

- Jamie Taylor, <http://www.jamietaylor.org.uk/>
- 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

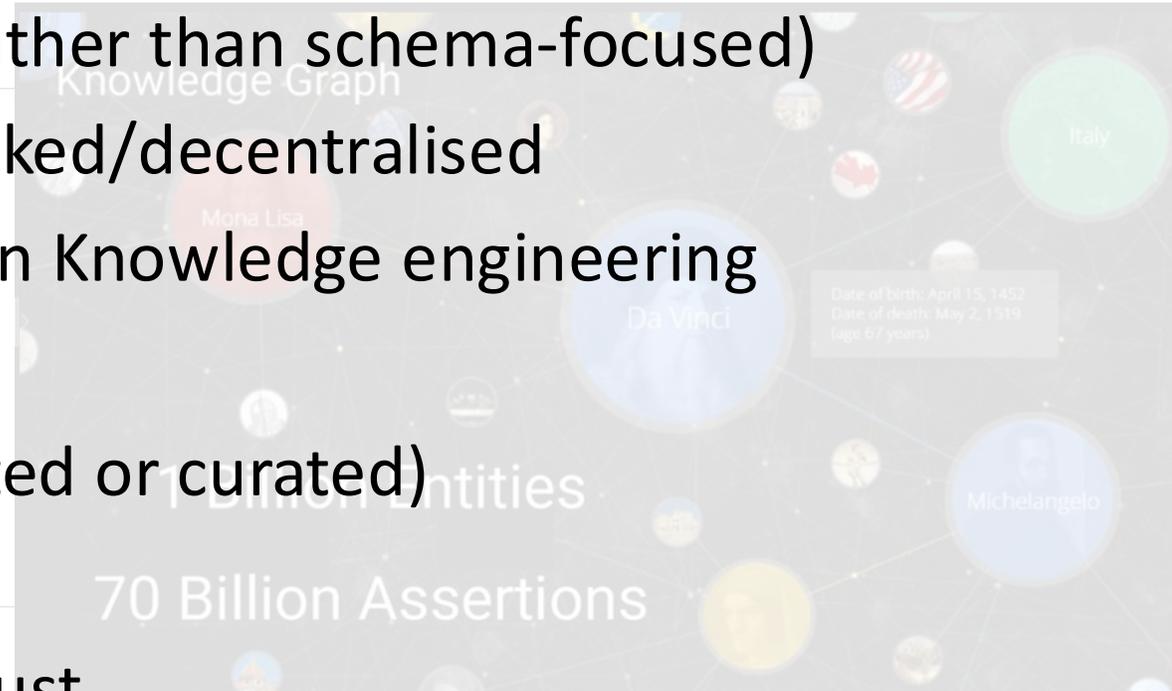
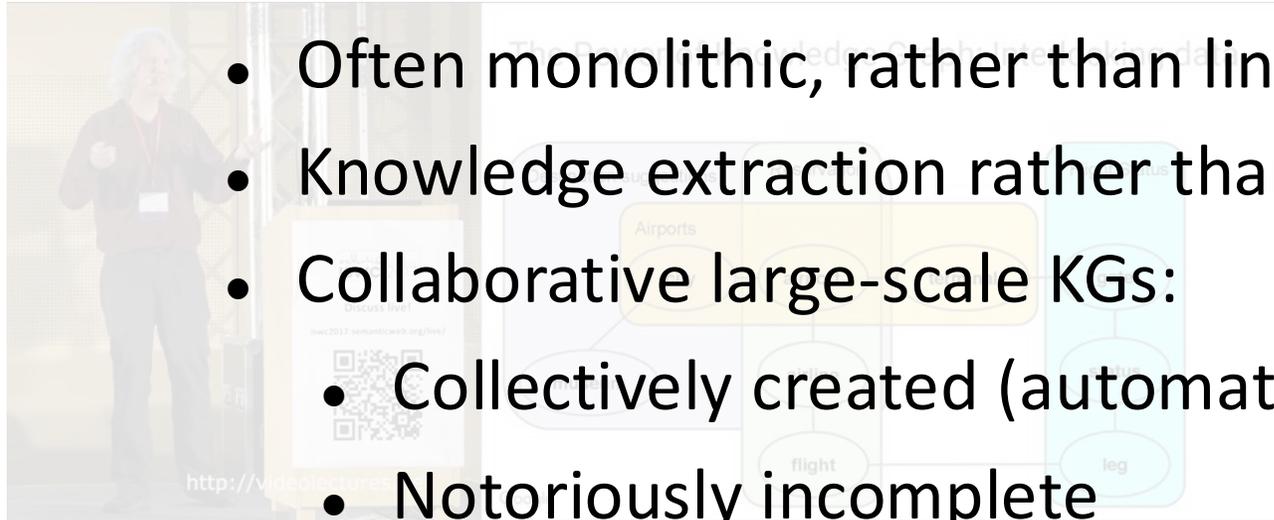
[...]

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

Wikidata



- ✓
- ✓
- ✗
- ✗
- ✓

Wikidata	
Screenshot [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,^{[1][2][13][14]} which today largely makes up Greater London,^{[12][14][note 1]} governed by the Mayor of London and the London Assembly.^{[14][note 2][15]}

London is a leading global city,^{[16][17]} in the arts, commerce, education, entertainment, fashion, finance, healthcare, media, professional services, research and development, tourism, and transportation.^{[18][19][21]} It is crowned as the world's largest financial centre^{[22][23][24][25]} and has the fifth- or sixth-largest metropolitan area GDP in the world.^{[16][19][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] 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,^{[16][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, 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>

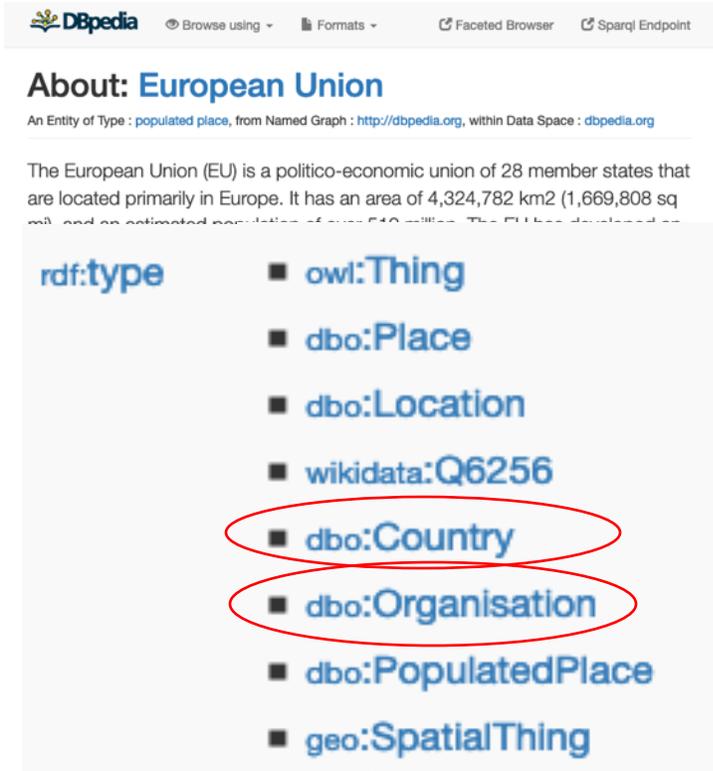
```
PREFIX : <http://dbpedia.org/resource/>
PREFIX dbo: <http://dbpedia.org/ontology/>
PREFIX yago: <http://dbpedia.org/class/yago/>

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

    FILTER ( ?pop > 1000000 )
}
```

Dbpedia is not logically consistent! ☹️ [1]

• E.g. 



DBpedia

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

association football player

► 2 references

FIFA player ID (archived)

229397

► 1 reference

country of citizenship

Argentina

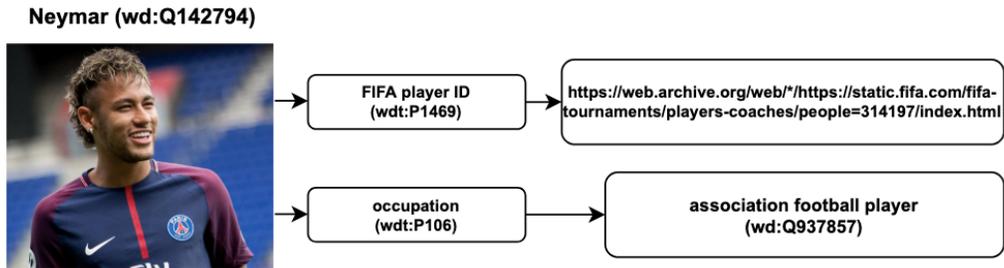
start time 1987

► 0 references

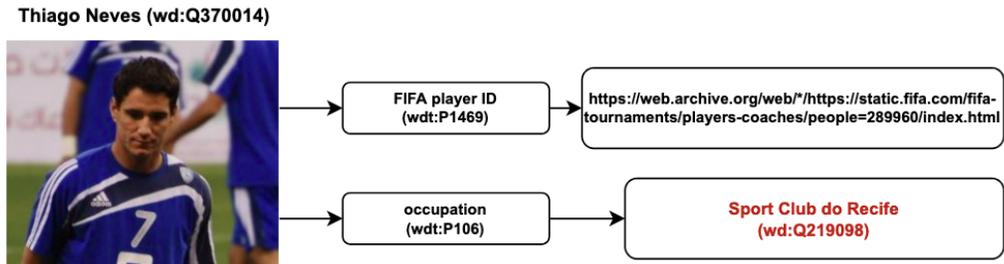
Spain

start time 2005

► 1 reference

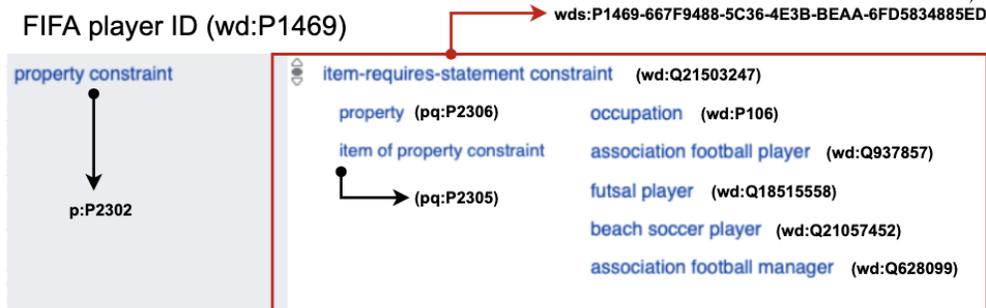


(a) Data graph complying to the constraint



(b) Data graph not complying to the constraint

... user defined **Property Constraints** (rather than OWL)



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:

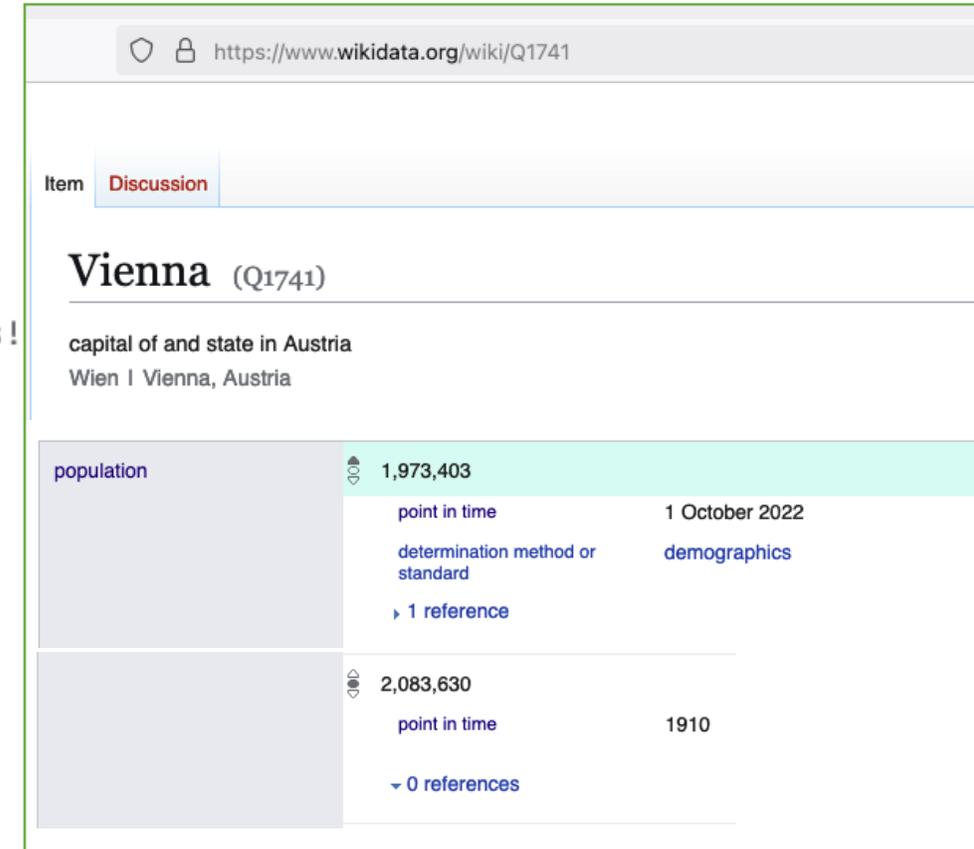
<https://w.wiki/BqRX>

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
	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)
}
```

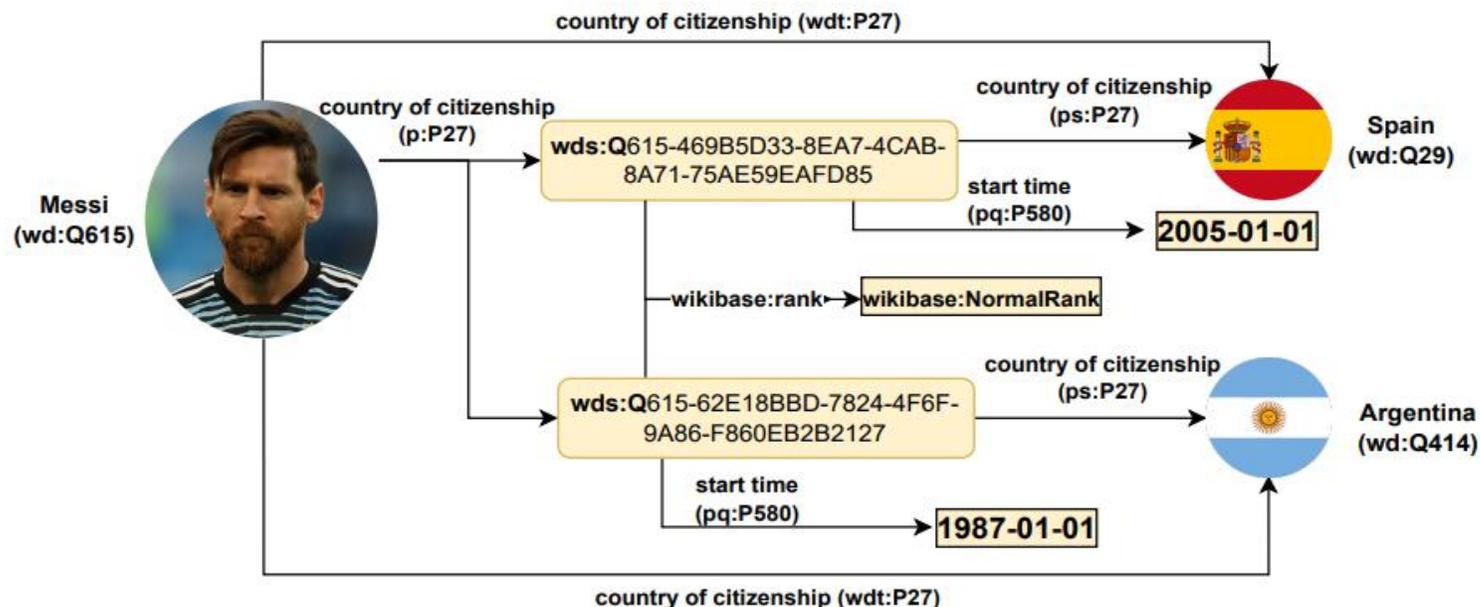
So, WHEN did Vienna have 2M inhabitants? Works!!!!

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/>

Wikidata's proprietary RDF reification model

- Wikidata's internal Data Model rather is a **Labelled Property Graph** than fitting into "flat" RDF:



<https://www.wikidata.org/wiki/Q615>

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

<https://www.wikidata.org/wiki/Q30>

capital	Washington, D.C.	17 November 1800 <i>Gregorian</i>
	start time	17 November 1800 <i>Gregorian</i>
	end time	no value
	▼ 0 references	
	Philadelphia	6 December 1790 <i>Gregorian</i>
	start time	6 December 1790 <i>Gregorian</i>
	end time	14 May 1800 <i>Gregorian</i>
	▼ 0 references	

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

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

i.e.,

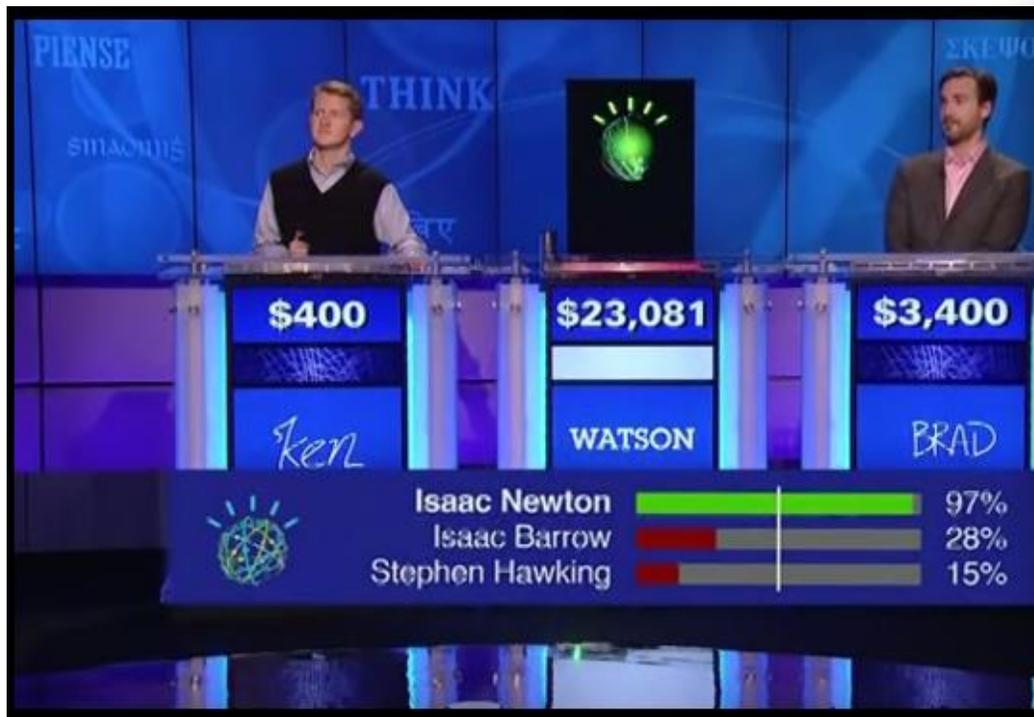
Large-scale, partially incomplete,
inconsistent, labelled property graphs

(rather than curated ontologies)

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 the "API Reference" section. The "Introduction" section explains that Natural Language Understanding uses natural language processing to analyze semantic features of any text. The "API Explorer" section provides instructions on how to interact with the REST API. The page also includes a "Menu" and a "Search" bar.

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

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

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.

🔊 📄 👍 🗨️ ↻

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

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

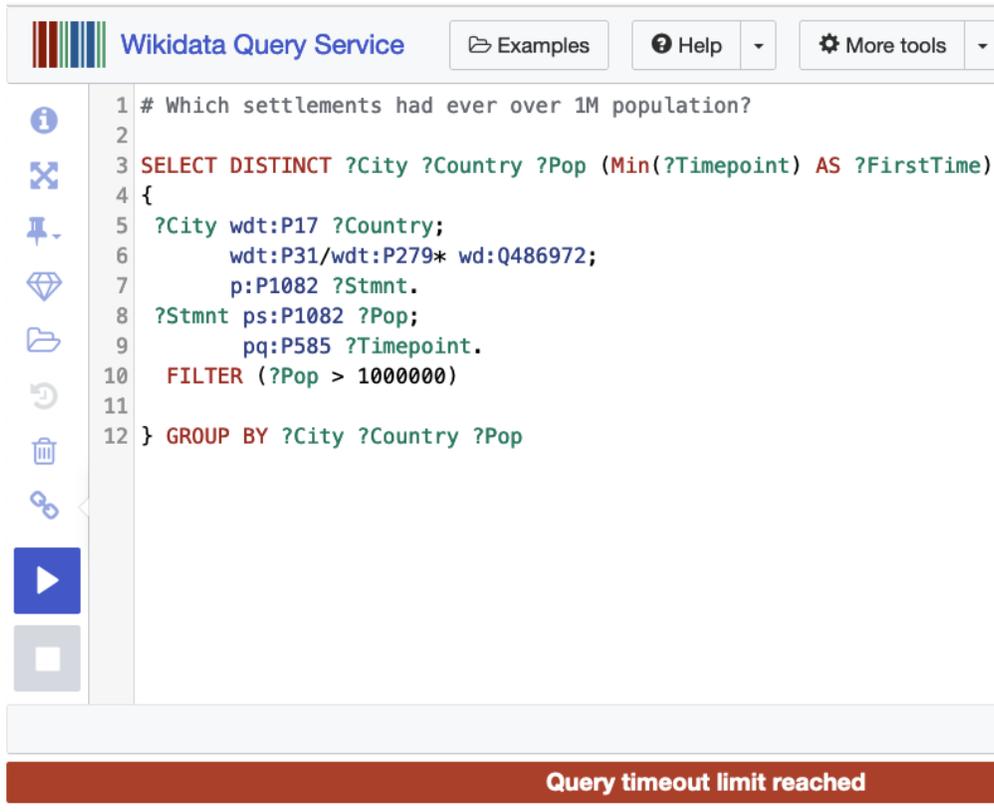


The diagram illustrates the difference between a Large generative model and a Knowledge graph lookup. On the left, a neural network diagram shows an input layer with 6 tokens, followed by multiple hidden layers, and an output layer with 2 tokens. Below this is the OpenAI logo. On the right, a knowledge graph visualization shows a network of nodes and edges, with a search bar and a list of results. Below this is a bar chart with vertical bars of varying heights.

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	

Admittedly, Denny didn't talk about this...

... at least Wikidata also struggles on some questions: <https://w.wiki/CLw9>



The screenshot shows the Wikidata Query Service interface. At the top, there are navigation buttons for 'Examples', 'Help', and 'More tools'. Below this is a text area containing a SPARQL query:

```
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
```

At the bottom of the interface, a red error bar displays the message: "Query timeout limit reached".

Challenge:
scaling queries
to large-scale,
schemaless KGs

For the records: comparison with GPT ;-)

<https://chatgpt.com/share/675585c7-04cc-8006-a20e-c70d75619e13>

Some of our own research in this area:



- How good or bad can KG with Question can answering?

- **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 KG with Question answering?

- *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 and what do they struggle with?

- *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.*

- Lessons learnt:

- One of LLMs main problem: **recency**
- Can we use LLMs to generate SPARQL queries?
 - Main problem: “training” (**identifiers** in the database) is hard...

Other main trends in our community:

- **(Graph)RAG** – Retrieval augmented generation leveraging Knowledge Graphs
(a significant share at this year's ISWC)
- **Knowledge Graph Embeddings** – similar to word embeddings use vector space embeddings to predict missing information in KGs
- **Neurosymbolic Systems** that involve KGs

→ Trend is to combine:

Search Engines (SE)

Querying KGs (KG)

LLMs (LM)

So... *What's good for what? And What's next?*

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



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



Forthcoming work by :

Aidan Hogan, Xin Luna Dong,

Denny Vrandečić, Gerhard Weikum

<https://aidanhogan.com/talks/2024-09-04-wuwien-invited-talk.pdf>

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



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



Category	Subcategory	Example	SE	KG	LM	SE + KG + LM
Facts	POPULAR	<i>Who directed the movie Spotlight?</i>	+ good coverage - noisy content	+ high precision - poor coverage	+ good coverage - noisy content	+ good coverage + high precision
	LONG-TAIL	<i>Which galaxy is closest to the Sunflower Galaxy?</i>	+ good coverage - needle in haystack	+ high precision - sparse coverage	- limited storage - hallucinations	+ good coverage + high precision
	MULTI-HOP	<i>Which Turing Award winners were born in Latin America?</i>	- no reasoning - single-shot search	+ formal reasoning + structured queries	+ latent reasoning - hallucinations	+ formal reasoning + structured queries
	ANALYTICAL	<i>How many U.S. Congress Members are younger than 50?</i>	- no datatypes - no aggregation	+ rich datatypes + aggregation	- no datatypes - no aggregation	+ rich datatypes + aggregation
Explanations	COMMONSENSE	<i>How do snakes move?</i>	+ good coverage + text output	- poor coverage - structured output	+ good coverage + text output	+ good coverage + text output
	CAUSAL	<i>What caused the dancing plague of 1518?</i>	+ good coverage + text output	+ long tail - structured output	+ good coverage + text output	+ good coverage + text output
	EXPLORATORY	<i>Who was Williamina Fleming?</i>	+ text output + ranked results	+ graph algorithms + browsing + navigation	+ interactive + synthesis	+ hybrid output + hybrid interactivity
Planning	INSTRUCTIVE	<i>How do I tie a Windsor Knot?</i>	+ multimedia + diverse results	- poor coverage - non-didactic output	+ interactive - no multimedia	+ interactive + diverse results
	RECOMMENDATION	<i>Should I pack warm clothes for Iceland in June?</i>	+ diverse results + ranked results	- poor coverage - no recommendations	+ interactive + synthesis	+ interactive + synthesis
	SPATIO-TEMPORAL	<i>What kid-friendly Italian restaurants are near Disneyland?</i>	+ events & maps - no integration	+ integration + s.-t. operators	- lacks freshness - no s.-t. operator	+ s.-t. operators + integration
Advice	LIFESTYLE	<i>How can I improve my work/life balance?</i>	+ diverse results + ranked results	- poor coverage - lacks nuance	+ interactive + synthesis	+ diverse results + interactive
	CULTURAL	<i>Should I tip bartenders in Canada?</i>	+ diverse results + ranked results	- poor coverage - lacks nuance	+ synthesis - cultural bias	+ diverse results + synthesis
	PHILOSOPHICAL	<i>Is the death penalty ever acceptable?</i>	+ diverse results + ranked results	- poor coverage - lacks nuance	+ interactive + synthesis	+ diverse results + synthesis

What's missing?



Ora Lassila (Keynote ISWC2024)

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

*“Agents! The Semantic Web vision is predicated on the idea that we can converse with our **agents** and give them tasks to perform. Using LLMs, sufficiently flexible and open-ended conversational user interfaces are finally possible. Through curated and audited **knowledge graphs**, we get trusted sources of information for the agents to consume (and avoid LLM hallucinations)”*

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 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

What's next (from our side)?

- **Austrian National “Cluster of Excellence” BILAI (funded by FWF):**
 - **Vision of Broad AI**
 - **Role of (Knowledge) Graph-Based AI in BILAI**
- Ongoing Research in our Institute/Department

This research is funded in whole or in part by the Austrian Science Fund (FWF) [10.55776/COE12].

BiAI Consortium (~30M EUR, 5Y)

<https://www.bilateral-ai.net/>

JKU

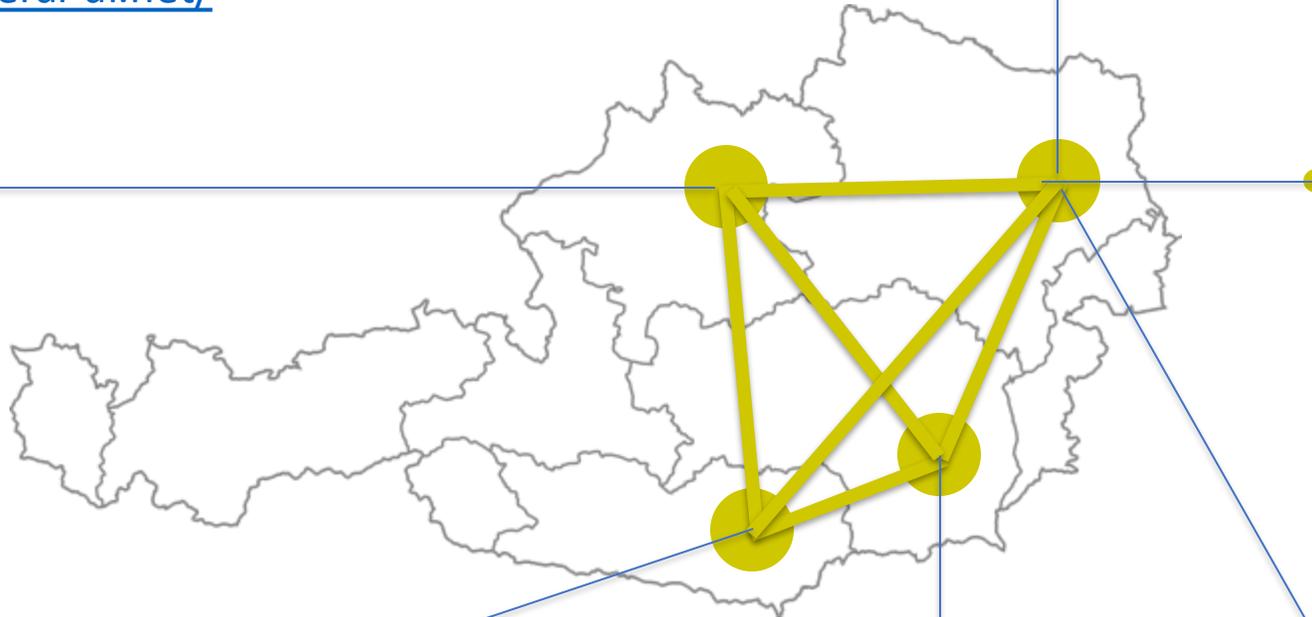


Martina Seidl
Symbolic AI
SAT Solving
Formal methods



Sepp Hochreiter
Machine Learning
LSTM
Vanishing gradient

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



**UNIVERSITÄT
KLAGENFURT**



Gerhard Friedrich
Symbolic AI
Model-based reasoning

- Institute for Artificial Intelligence and Cybersecurity

**TU
Graz**



Robert Legenstein
Machine Learning
Computational Neuroscience

- Institute of Theoretical Computer Science

ISTA Institute of
Science and
Technology
Austria



Christoph Lambert
Machine Learning
Trustworthy Learning

- Machine Learning and Computer Vision group
- ELLIS Unit ISTA

**TU
WIEN** TECHNISCHE
UNIVERSITÄT
WIEN



Agata Ciabattoni
Logic Reasoning



Thomas Eiter
Symbolic AI
Knowledge representation

- Institute for Logic and Computation

**WU
VIENNA**

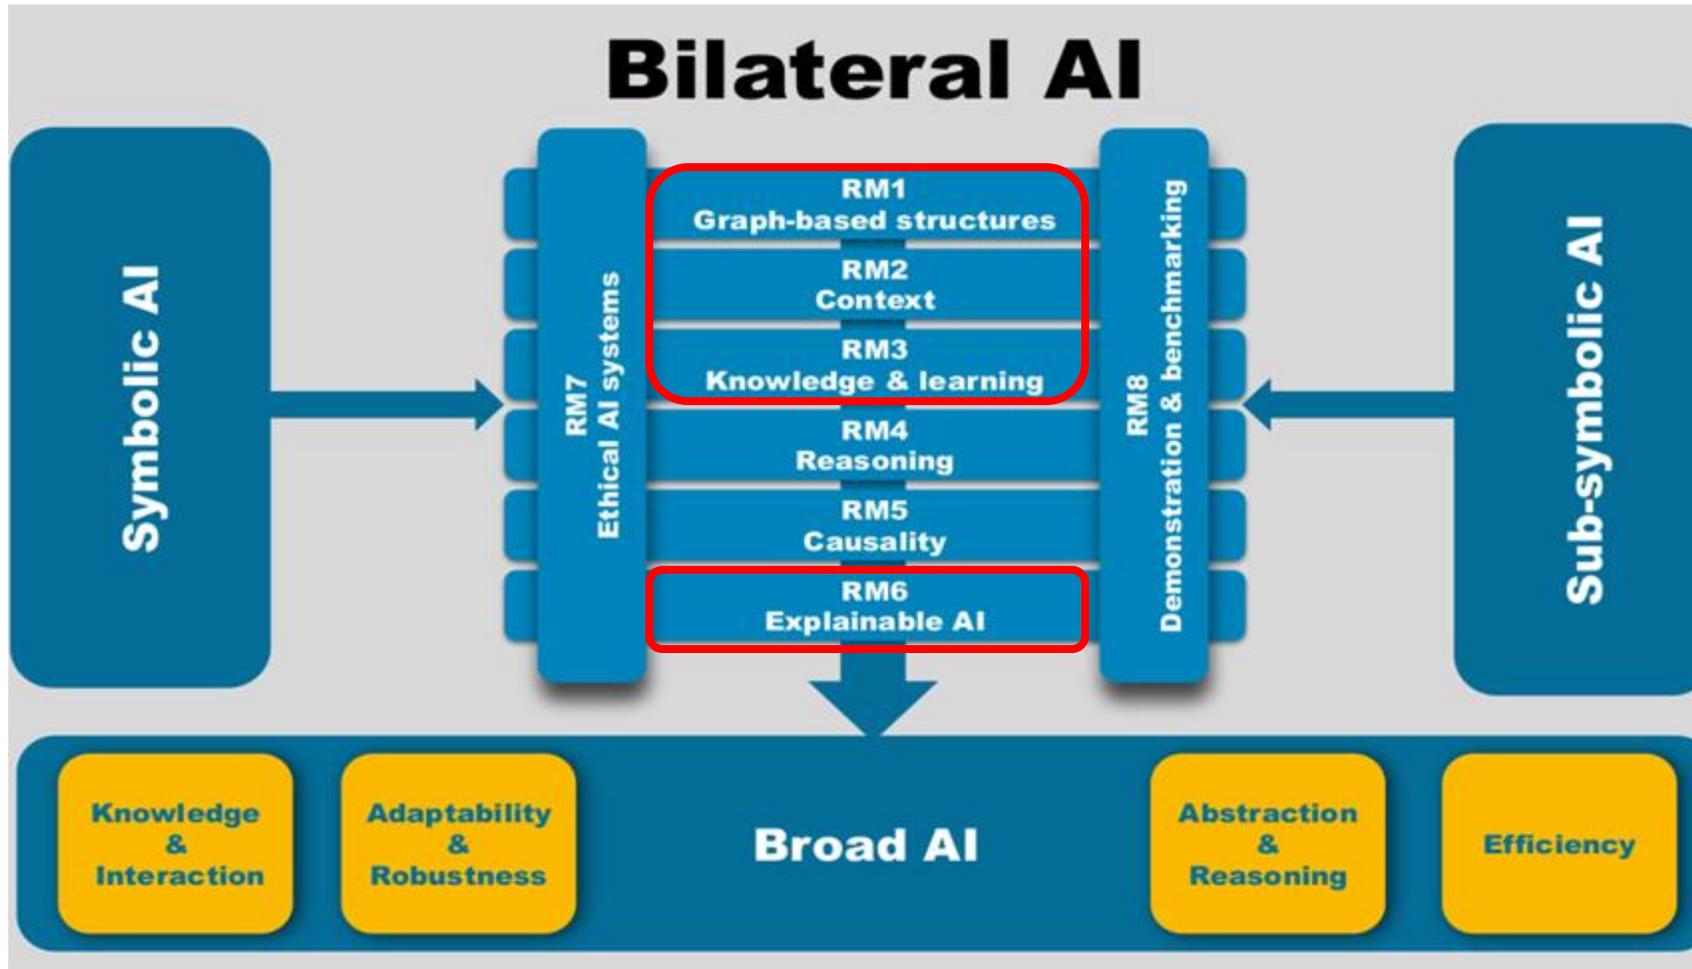


Axel Polleres
Knowledge Graphs

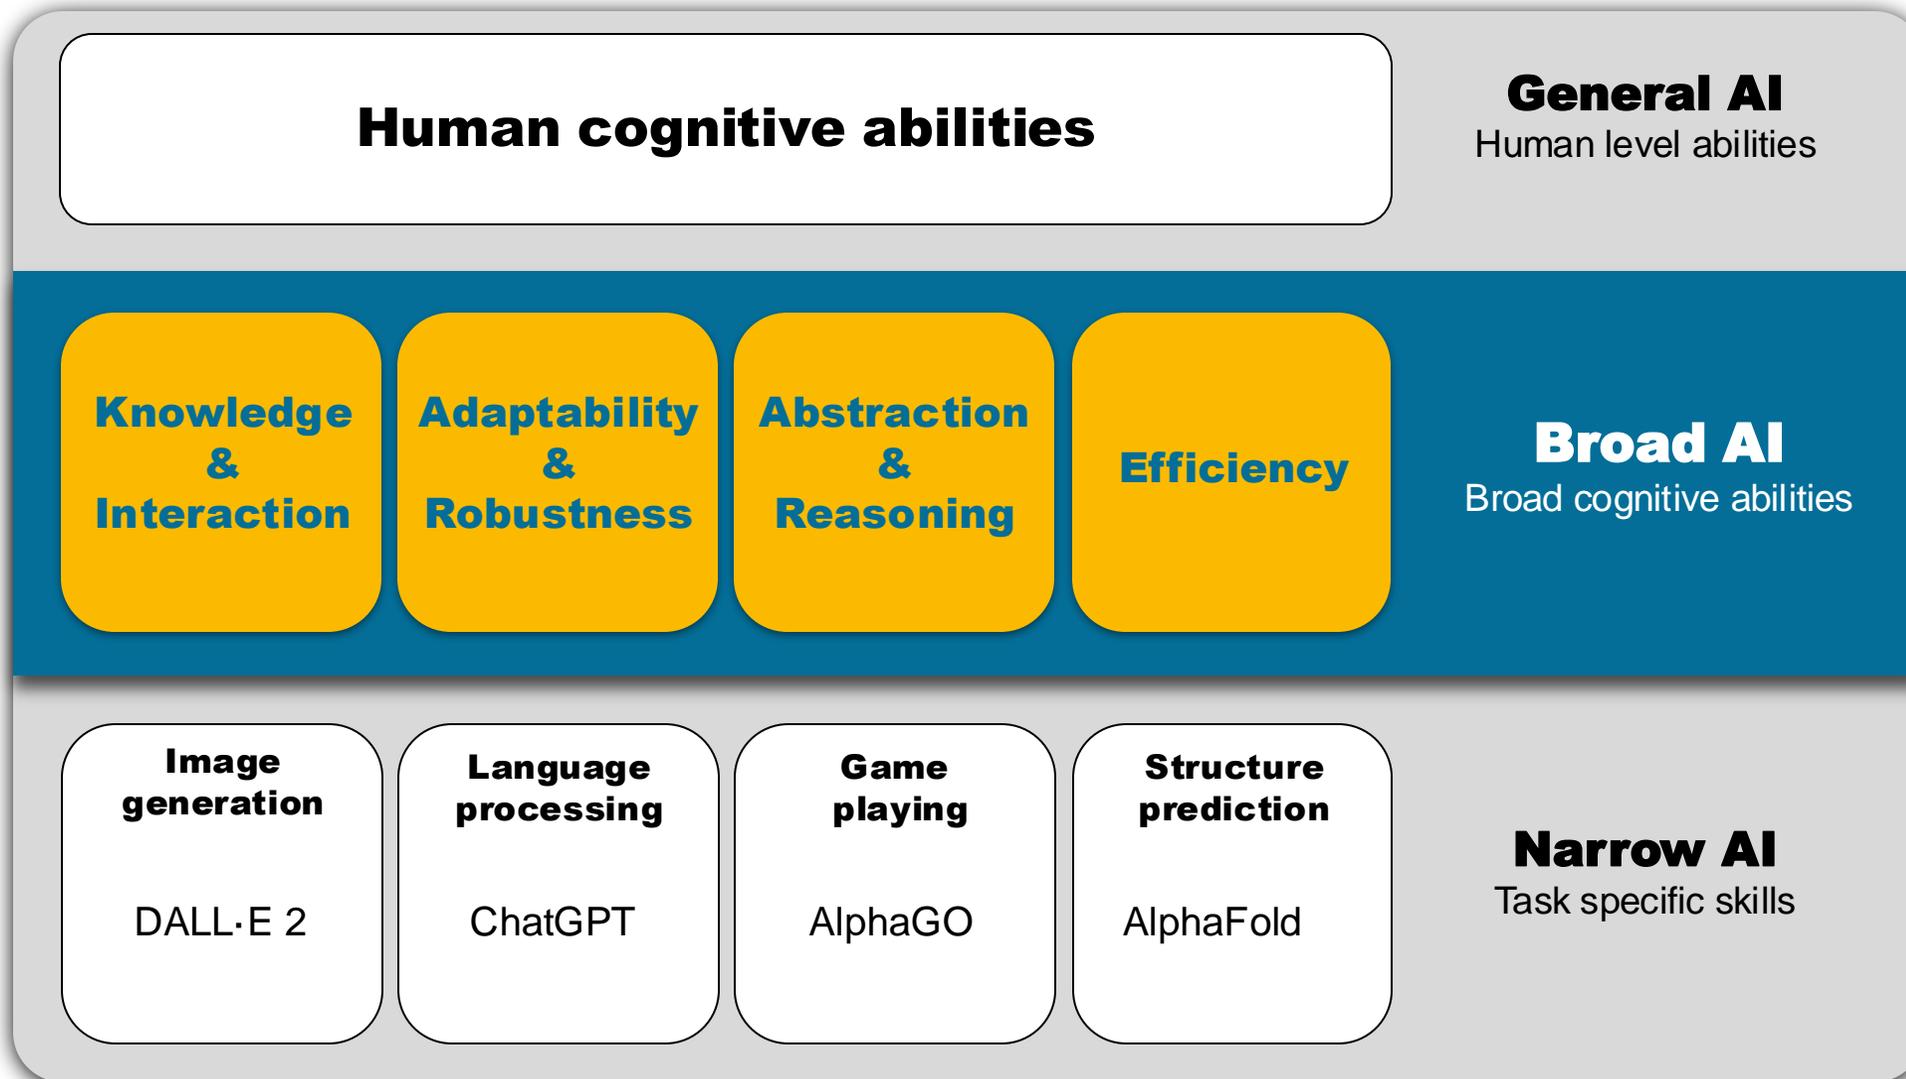
- Institute for Data Process and Knowledge Management



General Architecture



Vision: Building a „Broad“ AI



Chollet, F. (2019). On the measure of intelligence. *arXiv preprint arXiv:1911.01547*.

Hochreiter, S. (2022). Toward a broad AI. *Communications of the ACM*, 65(4), 56-57.



Large language models and the essential properties of broad AI

Auto-Regressive LLMs can't plan

(and can't really reason). — Yann LeCun (September 2023)

- **Challenge1 (Knowledge): LLMs hallucinate.**
 - difficult to remove or delete particular knowledge or suppress particular examples from the training set
 - knowledge that is gathered or collected after training, is difficult to integrate into LLMs (recency)
 - questions that cannot be factually answered as the corresponding information is missing in the training data, LLMs hallucinate
- tight integration of LLMs with symbolic **[models (KGs) &]** solvers [...] could be employed to leverage incremental reasoning capabilities (RM1-4)



Large language models and the essential properties of broad AI

Auto-Regressive LLMs can't plan

(and can't really reason). — Yann LeCun (September 2023)

- **Challenge2 (adaptability and robustness):** current LLMs lack adaptability and robustness.
 - low adversarial robustness
 - simple adversarial attacks can lead to critical threats, such as the extraction of training data
 - underscores the pressing need for advancements in AI that can enhance the adaptability and robustness of LLMs

→ ensure reliability and safety of LLMs in diverse contexts with approaches ***[leveraging context and again tight integration of Symbolic and Sub-symbolic inference]*** (e.g., by verification) **(RM2+3)**



Large language models and the essential properties of broad AI

Auto-Regressive LLMs can't plan

(and can't really reason). — Yann LeCun (September 2023)

- **Challenge3 (abstraction & reasoning):**

- LLMs are weak at reasoning and causality.
- if successful in causal inference, there is usually sufficiently close training data.
- Usual case: LLMs fail and, thus, they could be considered as weak “causal parrots”.
- LLMs are far from reasoning reliably about causality

→ Research Module on **Causality (RM5)** suggests how to approach such issues.



Research Questions & Starting points:

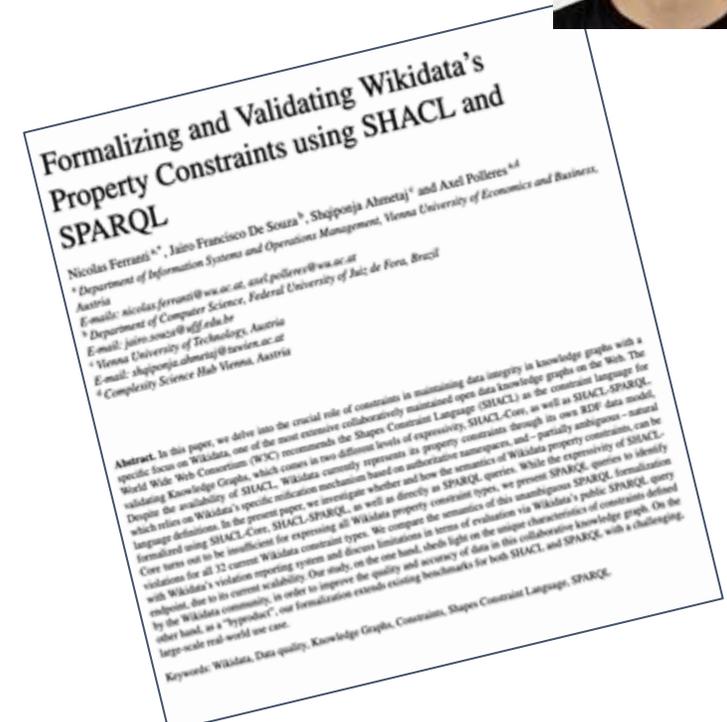
- **Time and other Contextual information:** Temporal Evolution of Graphs (and their **quality**) vs. Evolution of Embeddings – Constraints & Repairs (TGDK)
- **Knowledge at scale - Modularization and Decentralization of Knowledge** federated graph querying techniques and partitioning techniques vs. federated models/learning
- **Integrating vector representation vs graph representation** ... what's good for what?
 - a. graph pattern matching and isomorphism* → obviously great for symbolic processing, modularization, etc.
 - b. vector representation, embeddings* → obviously great for modeling similarity, semantic closeness, link prediction, but also dissimilarity/inconsistency/outliers
 - c. Different graph representations: RDF vs. Property Graphs*
 - d. How could we integrate these representations and their processing?*

Research Questions & Starting points in my group

Automatically Repairing KGs



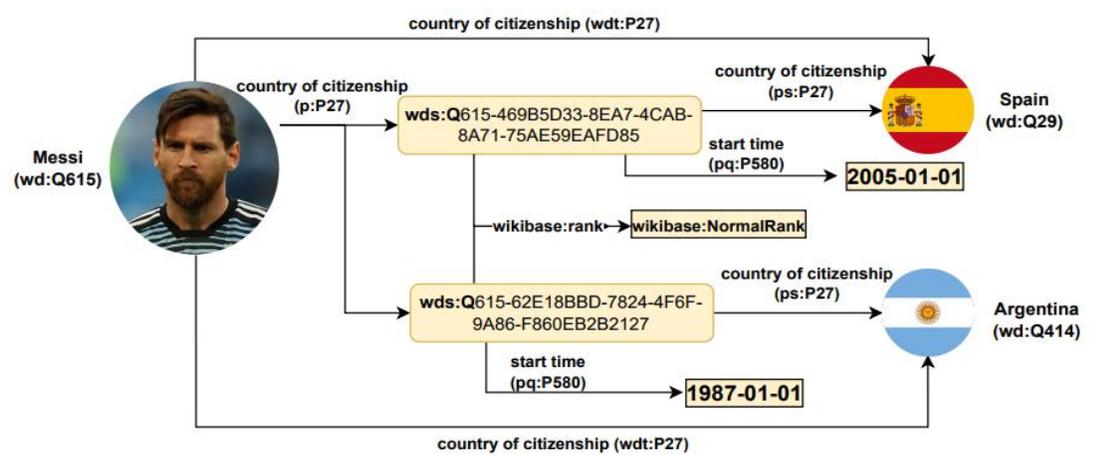
- Starting Points:
 - Formalizing the proprietary Integrity Constraint “Language” of Wikidata & Observing violations over time**
 - Wikidata does not rely on OWL or SHACL, but uses a community-defined way to define constraints:*
 - We formulated all these constraints in SPARQL, to extract all violations*
 - We now investigate which constraints have been repaired how to learn patterns!*



Property Discussion

country of citizenship (P27)

property constraint	value-type constraint	former or current state
	class	country in a fiction work
		nation
		dependent territory
subject type constraint		human
	class	character that may or may not be fictional
		fictional character



Research Questions & Starting points in my group

Querying large-scale KGs

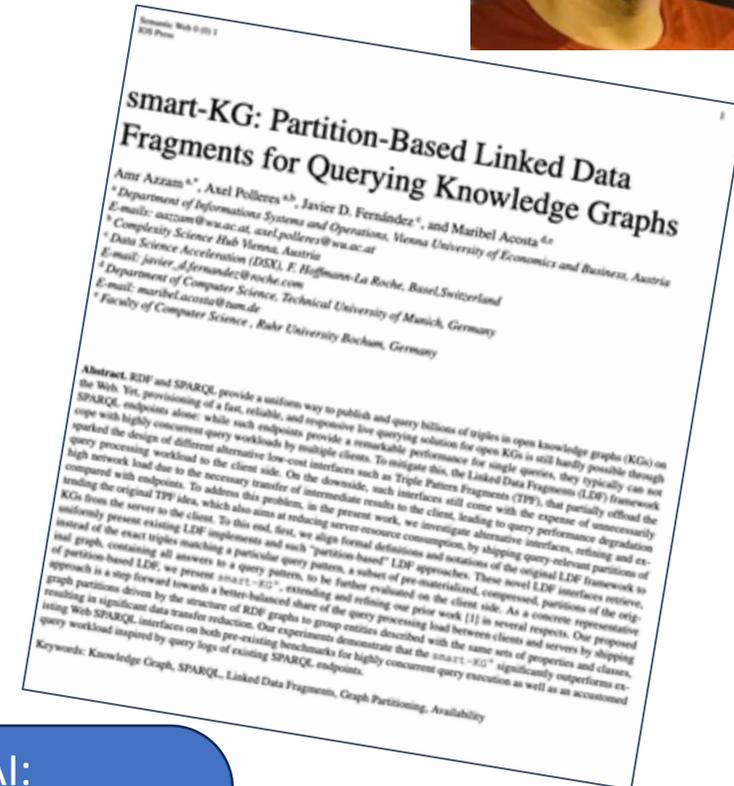


- Starting Points:
 - Speeding up decentralized SPARQL Querying by Graph partition-shiping to avoid query time-outs
 - In Amr's thesis we demonstrated that by clever **graph partitioning** and **splitting processing between clients and SPARQL endpoints** the performance, the problems of central query endpoints can be significantly alleviated

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



Future work in BILAI:
How can we similarly split work in a decentralized manner for other KG/AI tasks?
e.g. can we similarly modularize Knowledge Graph embeddings?



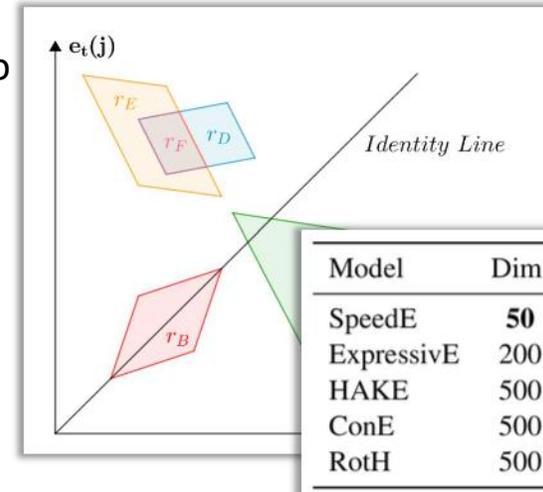
Slide: Emanuel Sallinger



Starting points for collaboration:

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

- Similar to word embeddings and LLMs, Knowledge Graph Embeddings allow to
 - predict missing edges in incomplete KGs
 - predict inconsistencies
 - ... I.e. predict possible repairs?



Model	Dim.	MRR	Conv. Time	#Parameters
SpeedE	50	.500	6min	2M
ExpressivE	200	.500	31min	8M
HAKKE	500	.497	50min	41M
ConE	500	.496	1.5h	20M
RotH	500	.496	2h	21M

Open Problems:

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

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: $r_1(X, Y) \wedge r_2(X, Y) \Rightarrow \perp$	✓	✓	✓	✓	✓	✓

What's next?

- Austrian National “Cluster of Excellence” BILAI:
 - Vision of Broad AI
 - Role of (Knowledge) Graph-Based AI in BILAI
- **Other Ongoing Research in our Institute/Department**

P.S.: we're hiring! 😊

<https://www.bilateral-ai.net/jobs/>



AI in our Department – at a glance

Data Management Group
(Axel Polleres, Elmar Kiesling,
Amin Anjomshoaa)

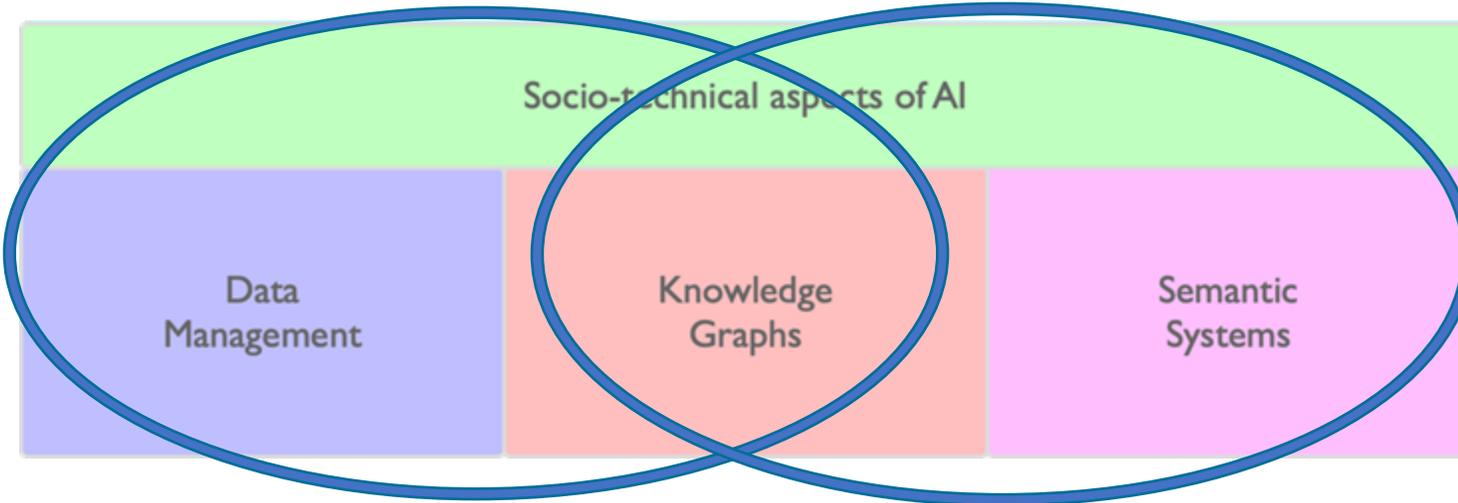


Semantic Systems Group
(Marta Sabou)



AI areas of interest:

-  Neurosymbolic AI Systems
-  Digital Humanism and AI
-  AI for Engineering
-  AI for Data Ecosystems/
Data Ecosystems for AI



Inst. for Complex Networks
(Sabrina Kirrane)

- => AI based policy representation and reasoning (e.g., regulatory obligations)
- => AI Transparency and trust

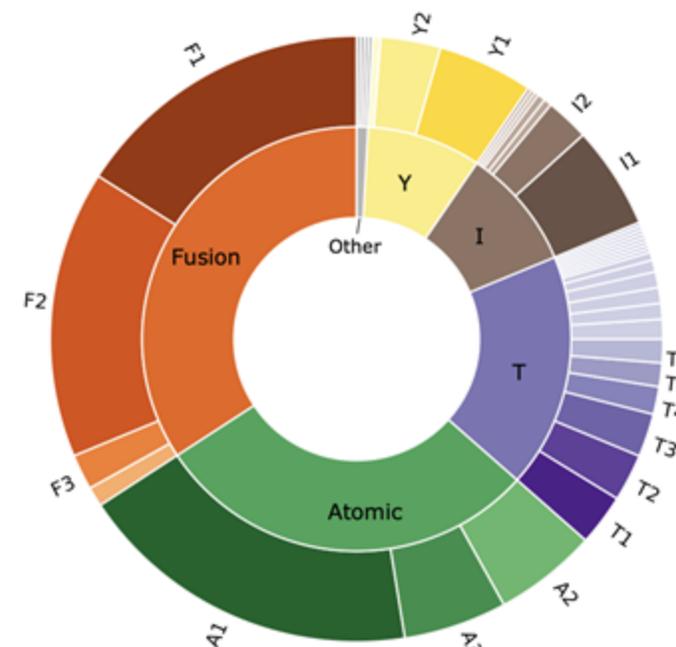
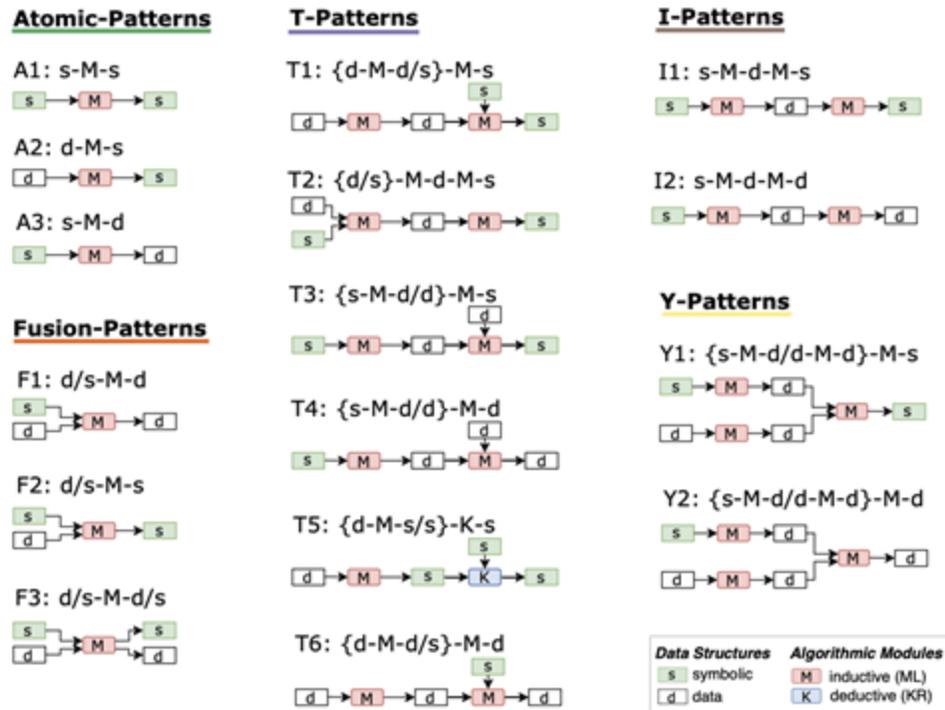


Neurosymbolic AI Systems

Prof. Marta Sabou



= Semantic Web and Machine Learning systems
(a type of neural-symbolic systems)

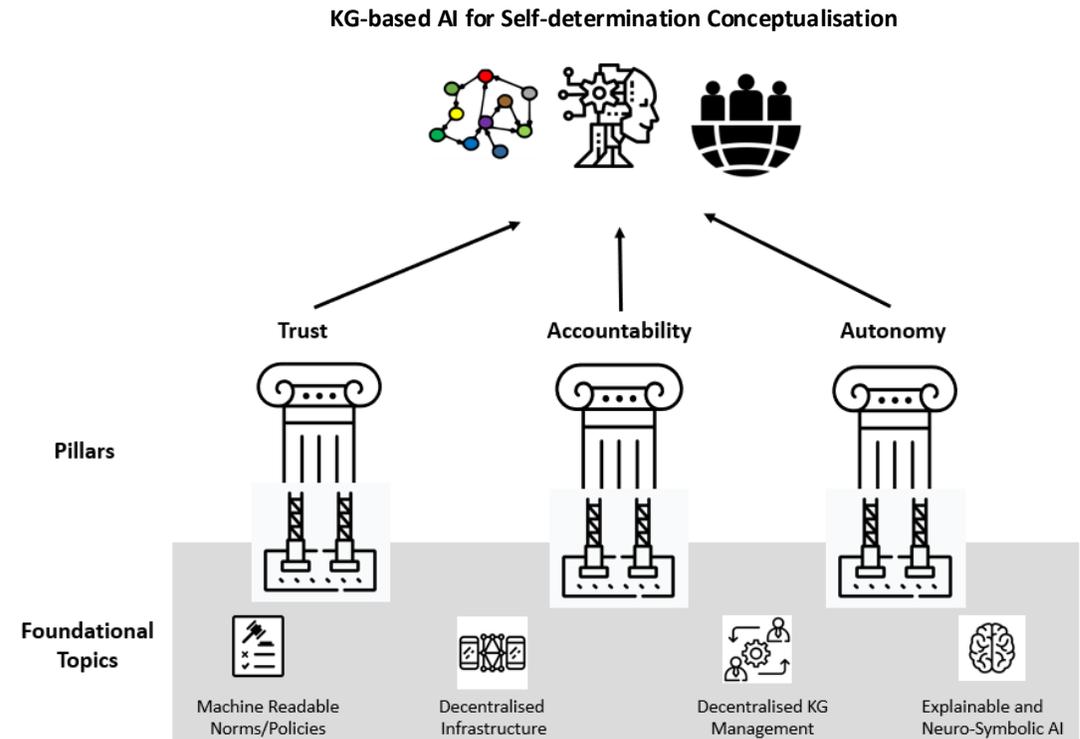


A. Breit, L. Waltersdorfer, F.J. Ekaputra, M. Sabou, A. Ekelhart, A. Iana, H. Paulheim, J. Portisch, A. Revenko, A. ten Teije, and F. van Harmelen. 2023. Combining Machine Learning and Semantic Web: A Systematic Mapping Study. ACM Computing Survey. March 2023.

KG-based AI for Self-Determination



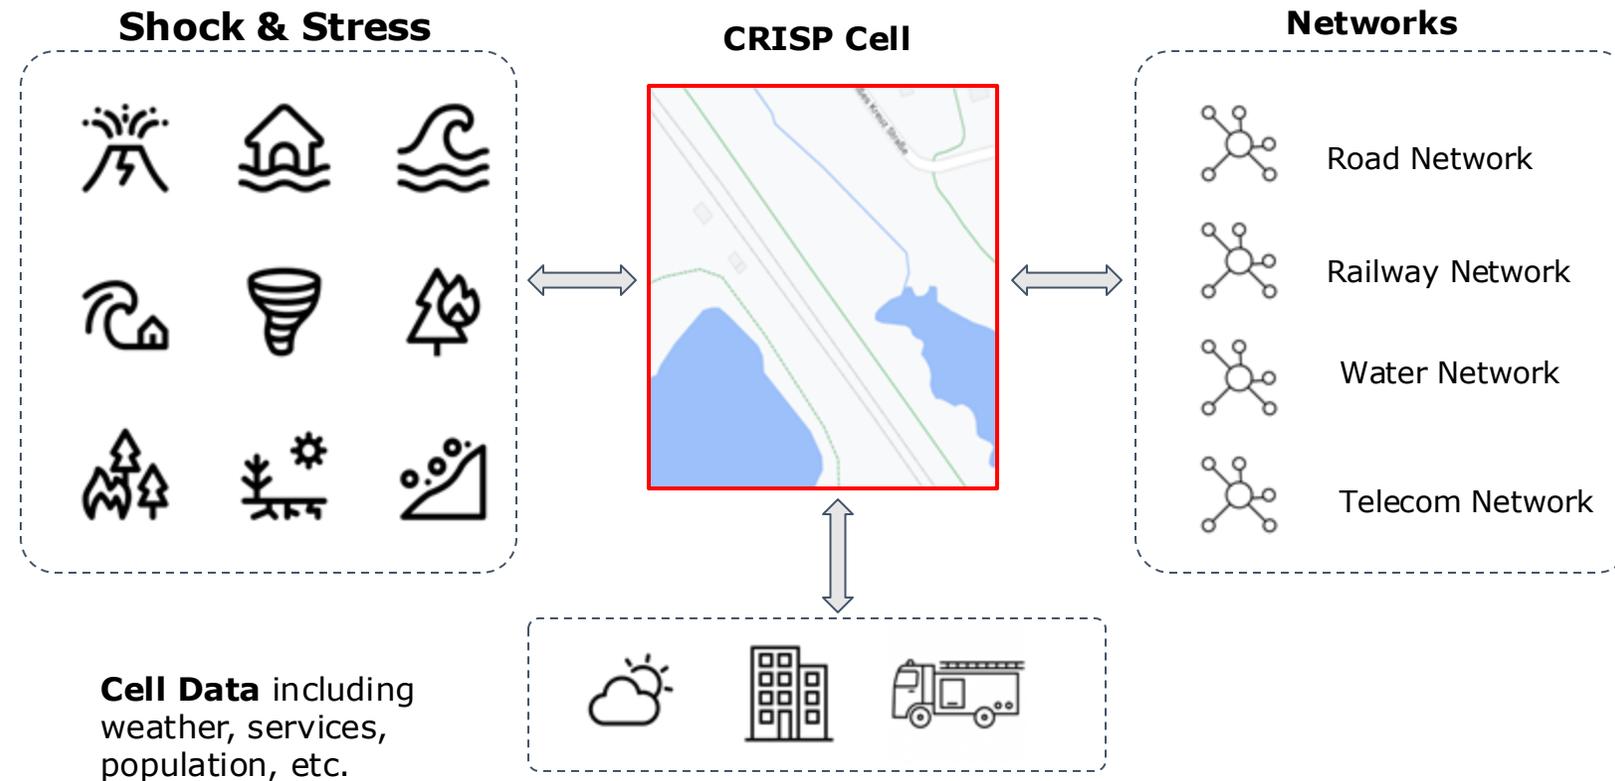
- The three pillar research topics - trust, accountability, and autonomy - represent the **desired goals for how AI can benefit society and facilitate self-determination**
- The pillars combine **fundamental principles of the proposed EU AI Act and self-determination theory**.
- The pillars are supported via four foundational research topics that represent the **tools and techniques needed to support the three research pillars**:
 - machine-readable norms and policies
 - decentralised infrastructure
 - decentralised KG management
 - explainable and neuro-symbolic AI



Building application specific Knowledge Graphs: CRISP Knowledge Graph



- Aims to establish the backbone of information integration for gathering Austrian infrastructure systems pertinent for crisis management.
- Is built on the foundation of three core elements: **event of hazards**, **geographical regions**, and **infrastructure networks**.
- Some statistics
 - **6,375,118** Triples collected from different open data resources.
 - **3,887** First Responders Organizations involved in crisis management.
 - **249,781** Observations of properties associated with specific features of interest



Thank you!

- Summary:
 - (Semantic) Web & Knowledge Graphs play an important role in latest trends in AI
 - GraphRAG, NeuroSymbolic Systems powered by KGs, etc.
 - GenAI could help to create, improve and curate KGs (but symbolic constraints will be needed!)
 - Collaborative, Open Knowledge Graphs like Wikidata are a particularly interesting subject of study (observable!)
 - evolution, repairs, etc.
 - embeddings
 - but also: collaboration patterns
 - My guess: agents will play an important role!
 - Getting back to decentralized approaches needed to scale & democratize AI
 - Also for trends I didn't talk about, e.g. Data Spaces

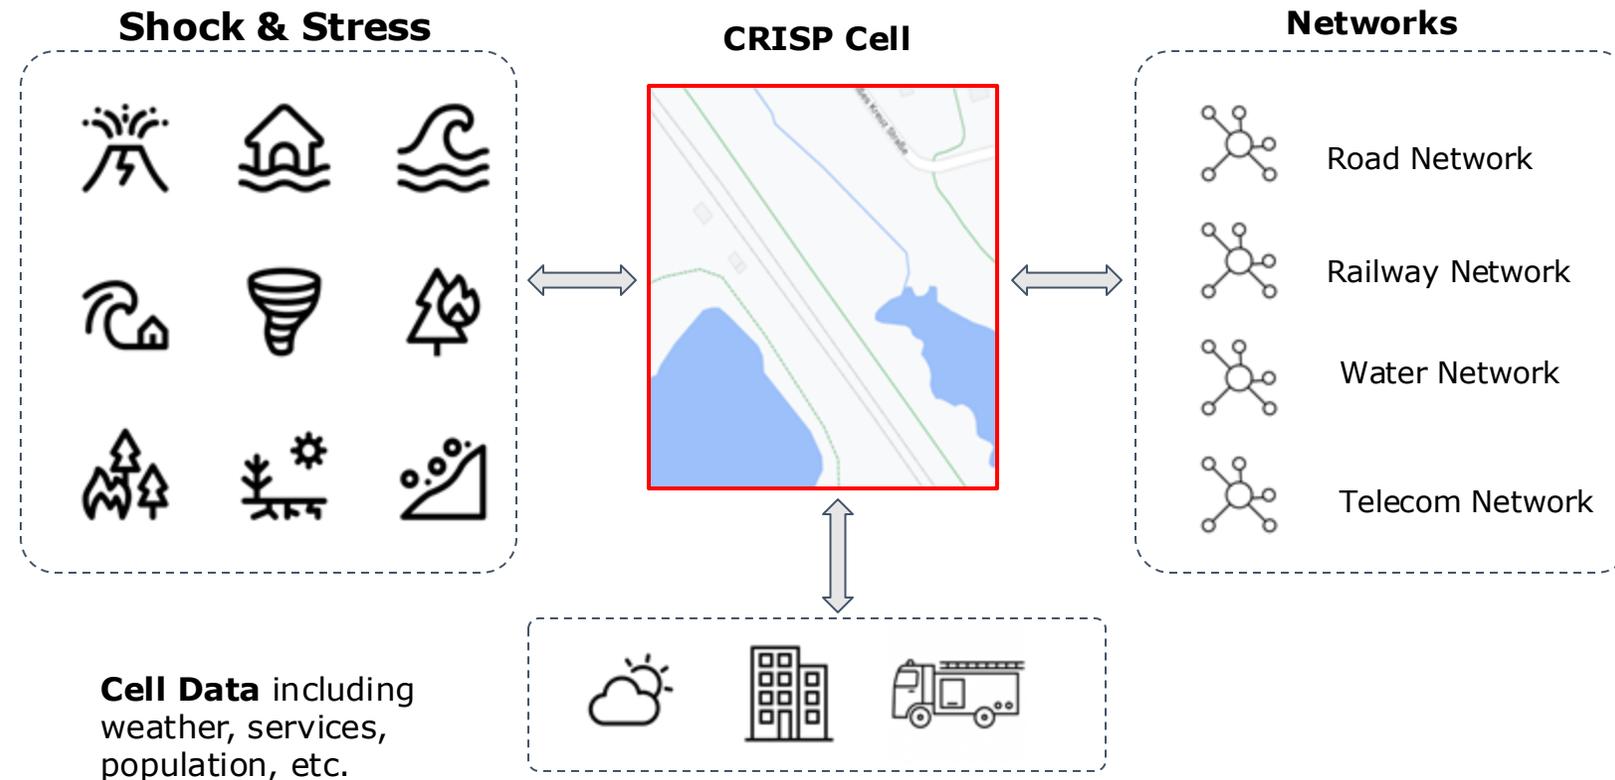
Backup Slides

CRISP KG

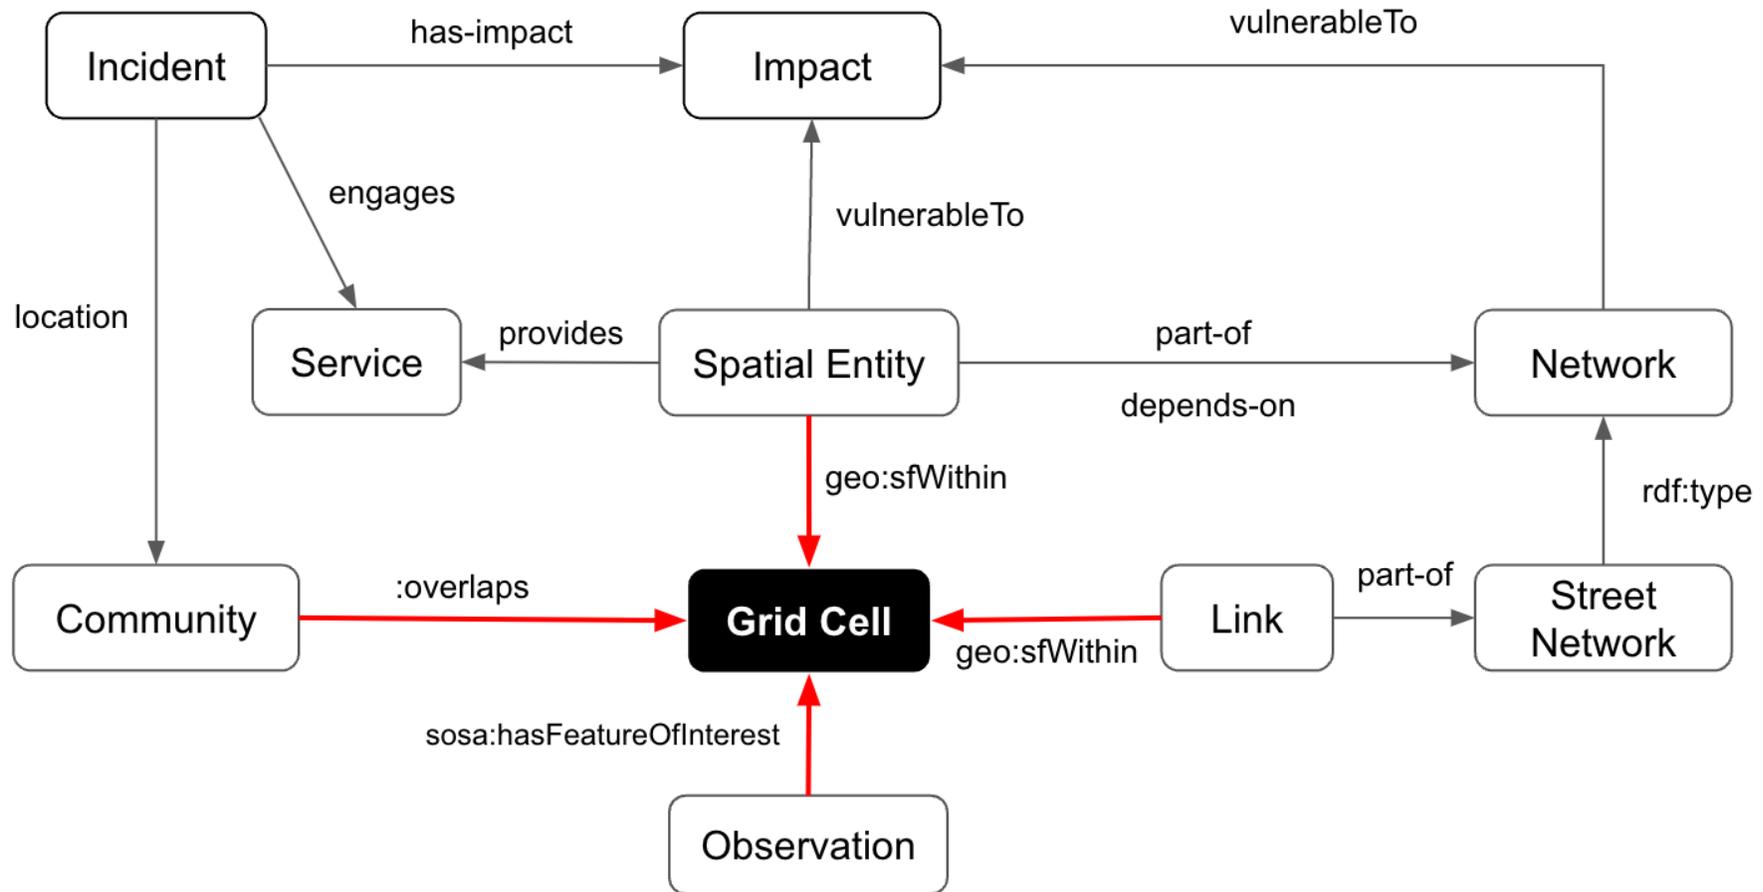
Building application specific Knowledge Graphs: CRISP Knowledge Graph



- Aims to establish the backbone of information integration for gathering Austrian infrastructure systems pertinent for crisis management.
- Is built on the foundation of three core elements: **event of hazards**, **geographical regions**, and **infrastructure networks**.
- Some statistics
 - **6,375,118** Triples collected from different open data resources.
 - **3,887** First Responders Organizations involved in crisis management.
 - **249,781** Observations of properties associated with specific features of interest



CRISP Semantic Model



Crisis Management Use Case



Incident

Provide incident's basic information.

Incident type

Community code

Potential Impacts

Incident Impacts

- Flood (2)
- Basement Flooding (1)
- Hail (1)
- Fire (1)
- Mudslides (1)
- Road/Railway Closed (1)
- River Lake Flooding (1)

Required Services



Emergency Services

Commonly engaged emergency responders for selected impact(s)

- Fire Department Service (5)
- Rescue Service (2)
- Infrastructure Restore Service (1)
- Police Service (1)
- Hospital Service (-)

Providers

Local providers for rendering selected services (10 km radius).

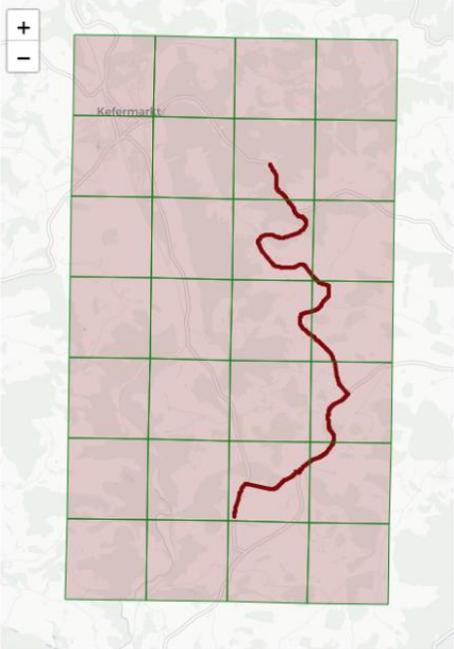
FireDepartmentService 

- [FF Selker-Neustadt](#)
- [FF Dingdorf](#)
- [FF Gutau](#)
- [FF Kefermarkt](#)
- [FF Lasberg](#)

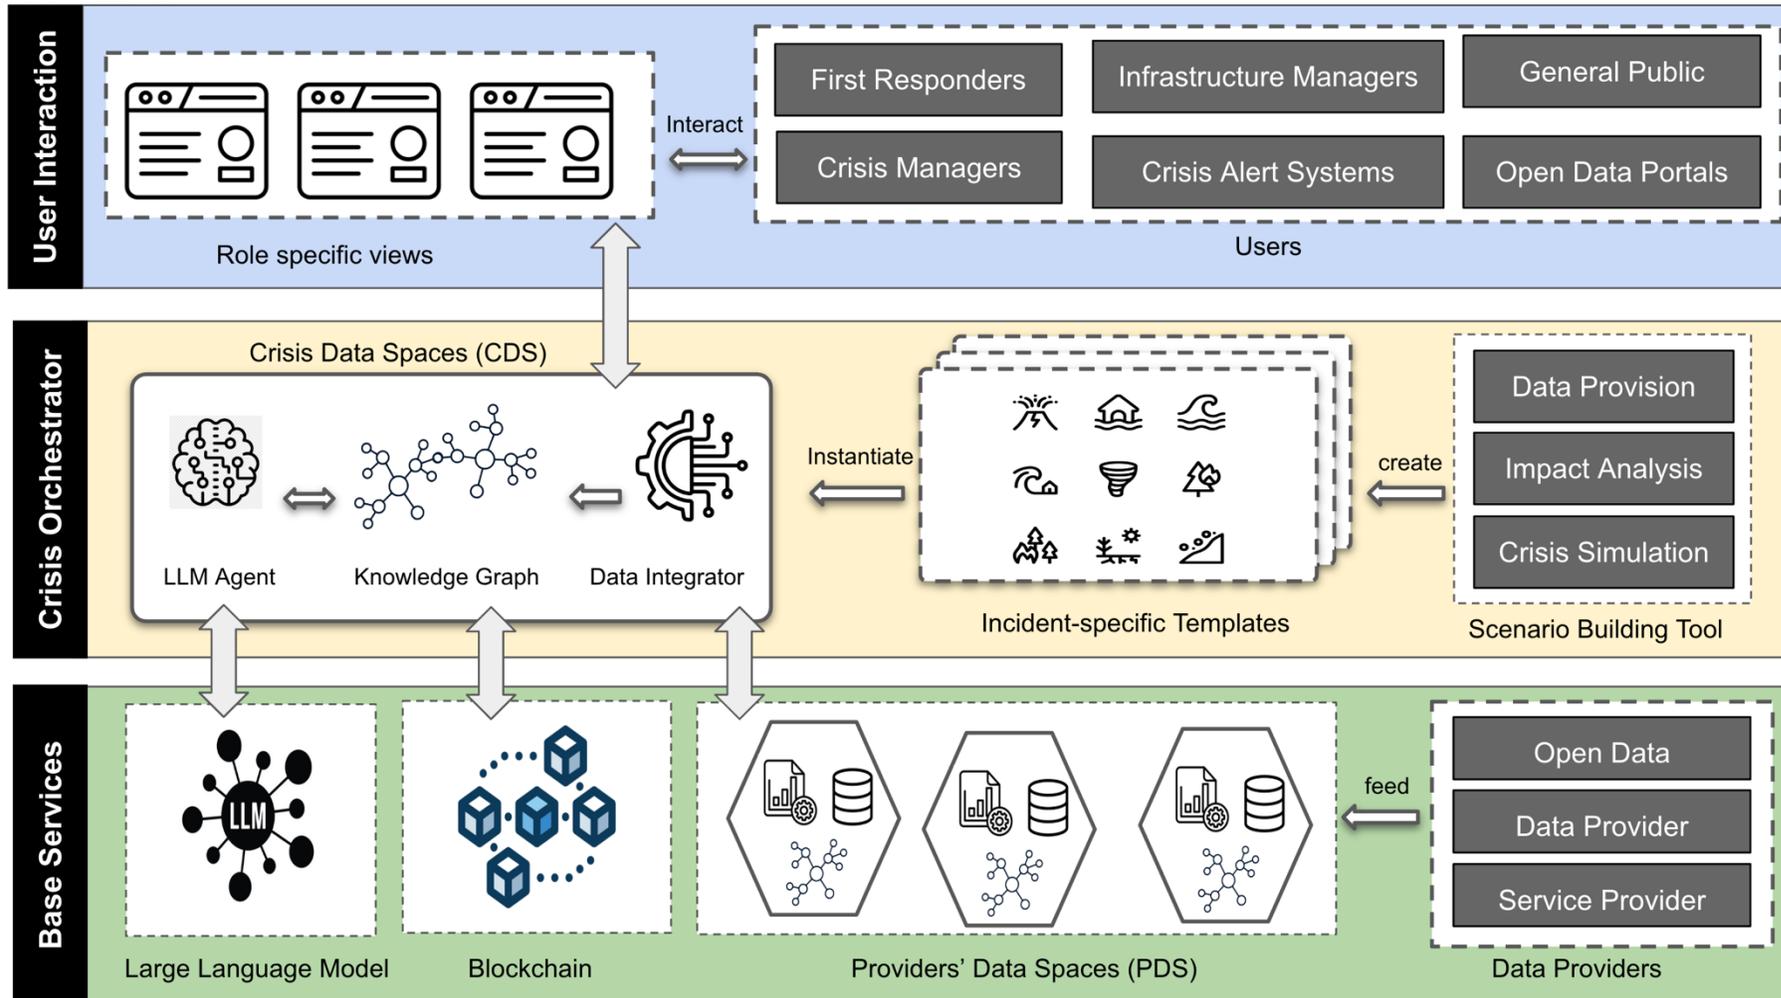


Networks

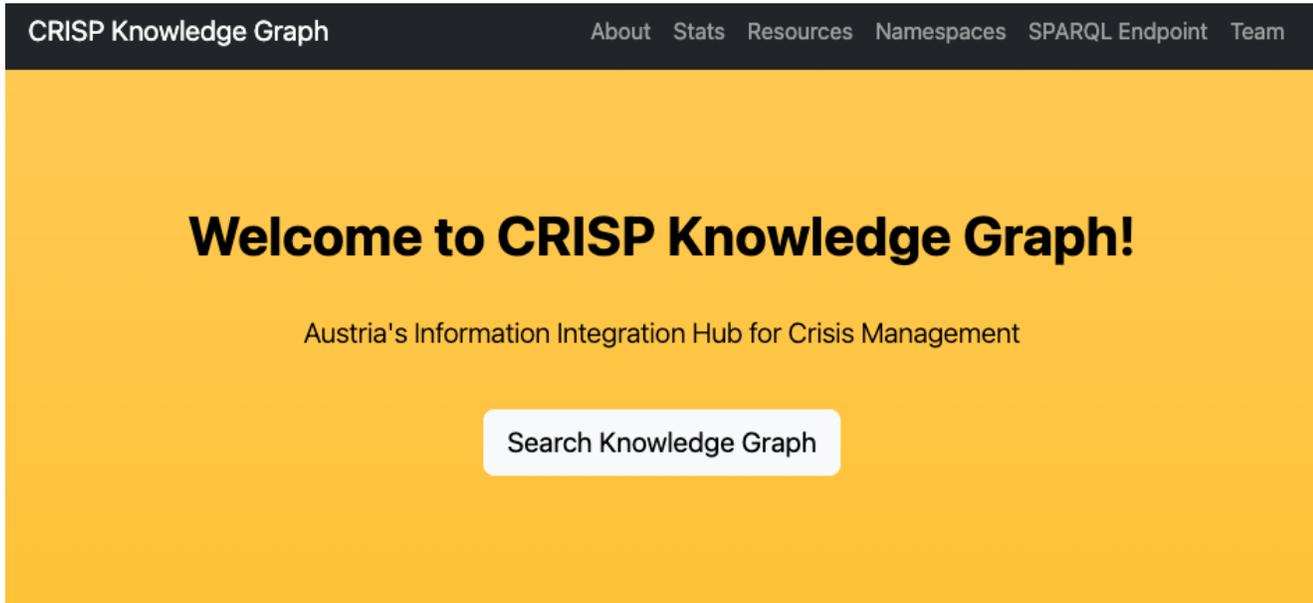
Networks and resources used by selected providers.



Future Work: Real-time Crisis KG Construction + AI Services



CRISP Portal & SPARQL Endpoint



The screenshot shows the homepage of the CRISP Knowledge Graph. At the top, there is a dark navigation bar with the text "CRISP Knowledge Graph" on the left and a menu of links: "About", "Stats", "Resources", "Namespaces", "SPARQL Endpoint", and "Team". The main content area has a bright yellow background. In the center, it says "Welcome to CRISP Knowledge Graph!" in large, bold, black letters. Below this, in smaller text, it reads "Austria's Information Integration Hub for Crisis Management". At the bottom of the yellow area, there is a white rounded rectangular button with the text "Search Knowledge Graph".



<http://crisp.ai.wu.ac.at/>

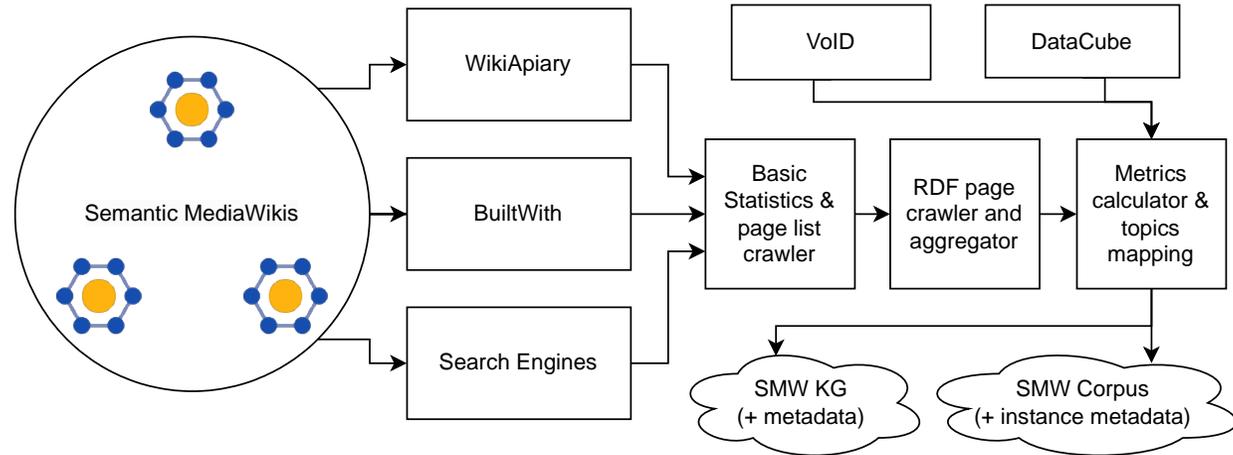
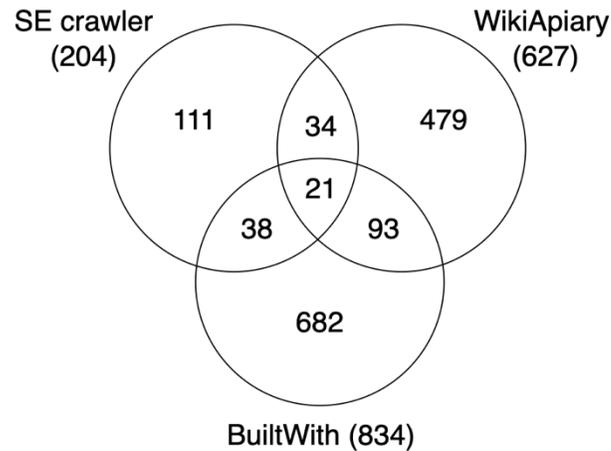
About CRISP Project

The CRISP Knowledge Graph aims to establish the backbone of information integration for gathering Austrian infrastructure systems pertinent for crisis management. It offers a comprehensive and collective view of urban infrastructure, service networks, and diverse environmental indicators. CRISP KG is built on the foundation of three

SMWCloud:

How many Semantic MediaWikis?

SMW Cloud (1458 wikis)



Dataset	#Triples	#Subjects	#Predicates	#Objects	#Literals
LODStats [10]	192,230,648	Not reported	49,916	Not reported	90,261,655
SMW Cloud	236,505,705	24,010,566	52,670	66,052,823	160,108,216
Wikidata 2021 [23]	17,662,800,665	1,625,057,179	38,867	Not reported	Not reported
LOD-a-lot [15]	28,362,198,927	3,214,347,198	1,168,932	3,178,409,386	1,302,285,394

Crawled RDF data available at semantic-data.cluster.ai.wu.ac.at/smwcloud/

- Currently ongoing work/next steps:
- also crawl historic data (Semantic MediaWiki edit history)
 - also crawl Wikiba.se instances!

Trust, Accountability, and Autonomy in Knowledge Graph-based AI for Self-determination

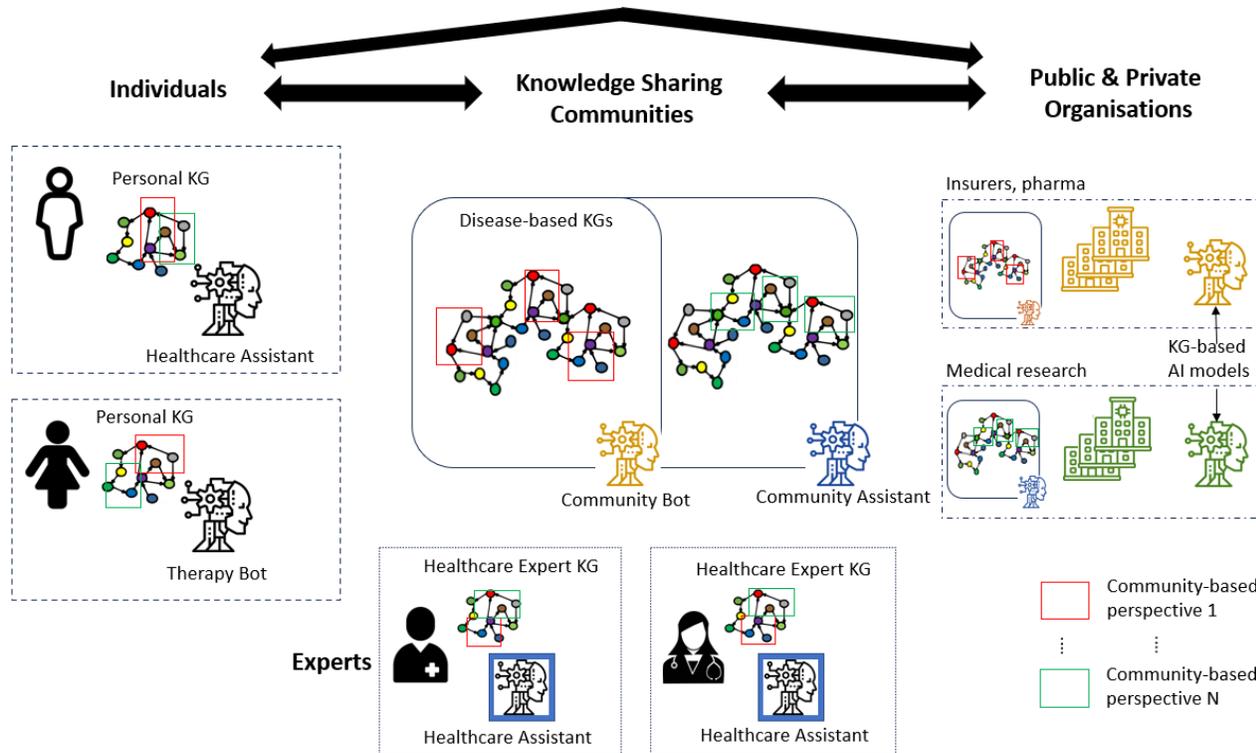


Sabrina Kirrane

6.12.2024



KG-based AI for Self-Determination

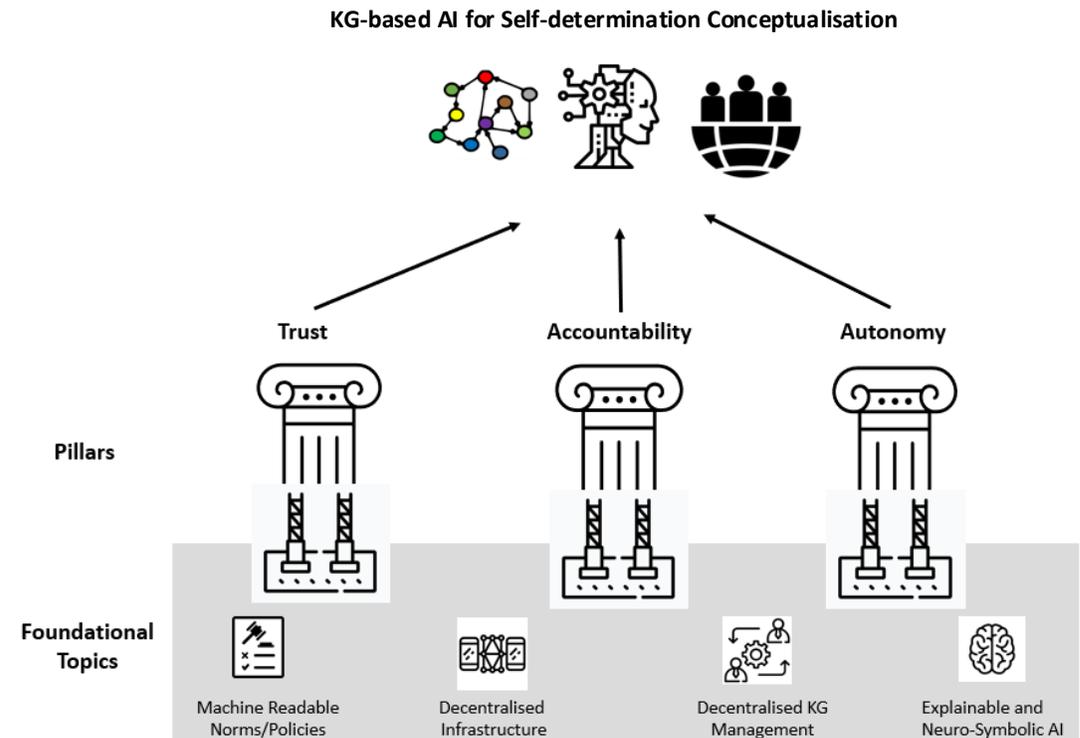


- Individuals use **Artificial Intelligence (AI) assistants** to make sense of data collected in their **Personal Knowledge Graphs (PKGs)**.
- They may **share perspectives of their PKGs** with other individuals and healthcare experts in knowledge-sharing communities that aggregate and curate data to **power AI services** for the benefit of all members.
- Public and private organisations can **negotiate access to data** from communities and individuals to train **KG-based AI models**, which in turn are used to build services for them.

KG-based AI for Self-Determination



- The three pillar research topics - trust, accountability, and autonomy - represent the **desired goals for how AI can benefit society and facilitate self-determination**
- The pillars combine **fundamental principles of the proposed EU AI Act and self-determination theory**.
- The pillars are supported via four foundational research topics that represent the **tools and techniques needed to support the three research pillars**:
 - machine-readable norms and policies
 - decentralised infrastructure
 - decentralised KG management
 - explainable and neuro-symbolic AI





Machine-readable norms and policies



KG-based AI for Self-determination Conceptualisation



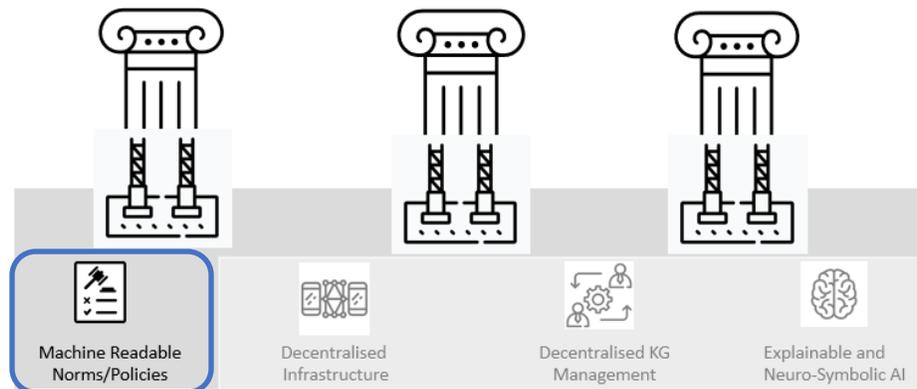
Trust

Accountability

Autonomy

Pillars

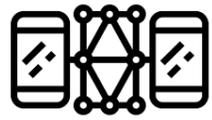
Foundational Topics



Bonatti, P.A., **Kirrane, S.**, Petrova, I.M. and Sauro, L., 2020. Machine understandable policies and GDPR compliance checking. KI-Künstliche Intelligenz.

Fernández, J.D., Sabou, M., **Kirrane, S.**, Kiesling, E., Ekaputra, F.J., Azzam, A. and Wenning, R., 2020. User consent modeling for ensuring transparency and compliance in smart cities. Personal and Ubiquitous Computing.

Kirrane, S., Fernández, J.D., Dullaert, W., Milosevic, U., Polleres, A., Bonatti, P.A., Wenning, R., Drozd, O. and Raschke, P., 2018. A scalable consent, transparency and compliance architecture. In The Semantic Web: ESWC 2018 Satellite Events.



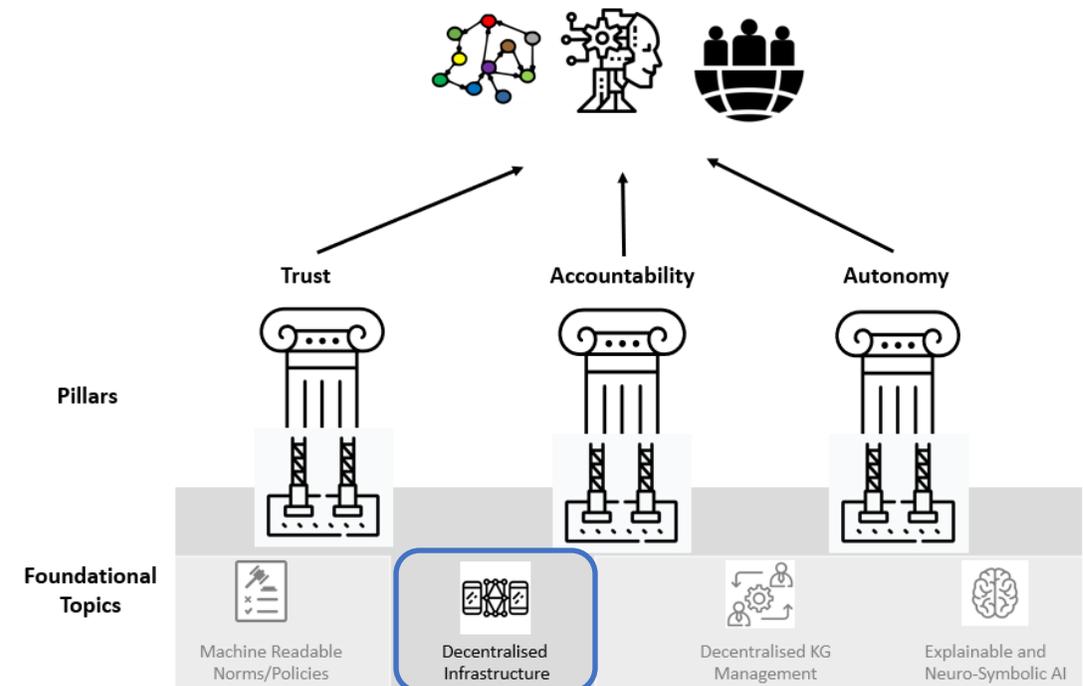
Decentralised infrastructure



Basile, D., Di Ciccio, C., Goretti, V. and **Kirrane**, S., 2023. Blockchain based Resource Governance for Decentralized Web Environments. *Frontiers in Blockchain*.

Basile, D., Di Ciccio, C., Goretti, V. and **Kirrane**, S., 2023. A Blockchain-driven Architecture for Usage Control in Solid. *Proceedings of the 1st Workshop on Fintech and Decentralized Finance (FiDeFix) @ the 43rd IEEE International Conference on Distributed Computing Systems*.

KG-based AI for Self-determination Conceptualisation





Decentralised KG management

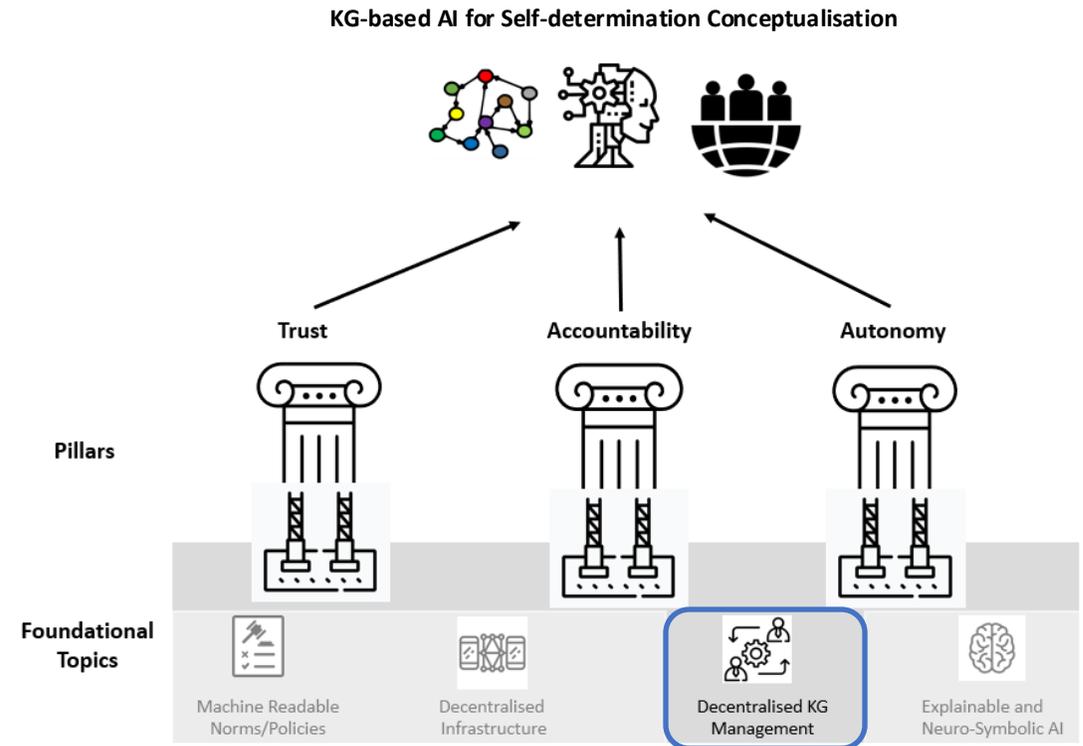


Kirrane, S., 2021. Intelligent software web agents: A gap analysis. *Web Semantics*.

Kampik, T., Mansour, A., Boissier, O., **Kirrane, S.**, Padget, J., Payne, T.R., Singh, M.P., Tamma, V. and Zimmermann, A., 2022. Governance of Autonomous Agents on the Web: Challenges and Opportunities. *ACM Transactions on Internet Technology*.

Fernández, J.D., **Kirrane, S.**, Polleres, A. and Steyskal, S., 2020. HDTcrypt: Compression and encryption of RDF datasets. *Semantic Web*.

Fernández, J.D., **Kirrane, S.**, Polleres, A. and Steyskal, S., 2017. Self-Enforcing Access Control for Encrypted RDF. *Proceedings of the 14th Extended Semantic Web Conference*.





Towards Explainable and Neuro-Symbolic AI

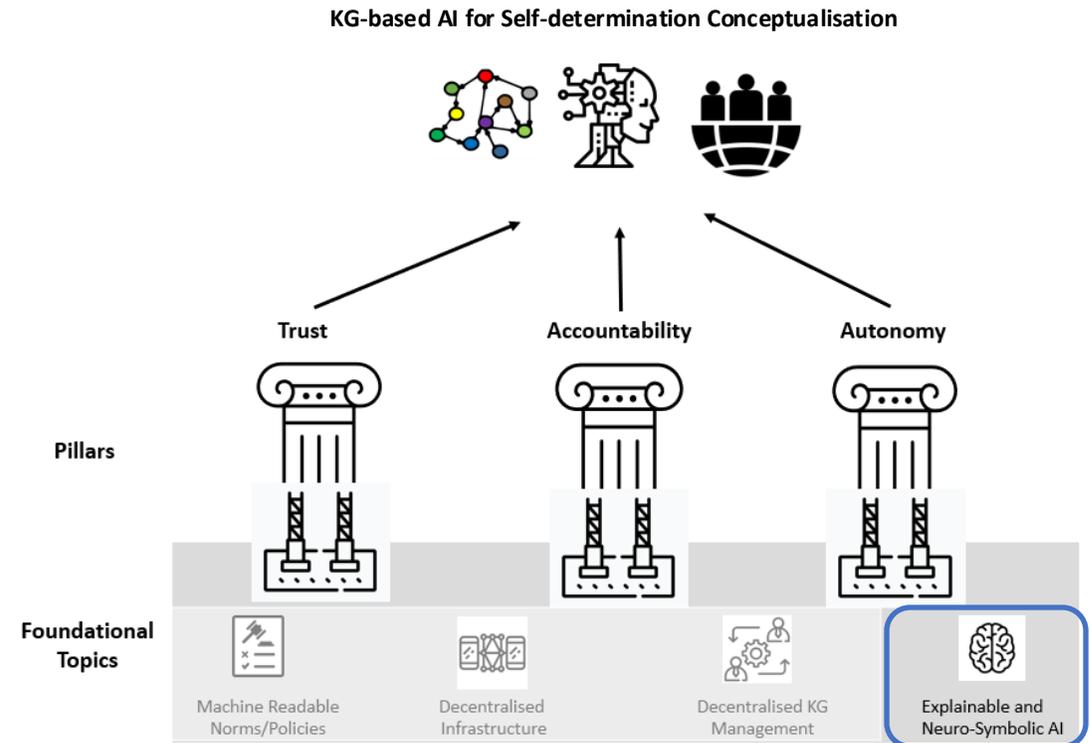


Filtz, E., **Kirrane, S.** and Polleres, A., 2021. The linked legal data landscape: linking legal data across different countries. Artificial Intelligence and Law.

Navas-Loro, M., Filtz, E., Rodríguez-Doncel, V., Polleres, A. and **Kirrane, S.**, 2019. TempCourt: evaluation of temporal taggers on a new corpus of court decisions. The Knowledge Engineering Review.

Filtz, E., Navas-Loro, M., Santos, C., Polleres, A. and **Kirrane, S.**, 2020. Events Matter: Extraction of Events from Court Decisions. Proceedings of the 33rd International Conference on Legal Knowledge and Information Systems.

Dieber, J. and **Kirrane, S.**, 2022. A novel model usability evaluation framework (MUSE) for explainable artificial intelligence. Information Fusion.



KG-based AI for Self-Determination

Challenges & Opportunities



- General-purpose policy languages could be used for **risk-based conformance checking** such as that envisaged in the proposed EU AI Act
- **Policy profiles with well-defined semantics and complexity classes** are needed for (semi)automatic compliance checking and to facilitate negotiation



- **Performance and scalability** are major challenges as applications will need to interact with multiple distributed data sources
- Self Sovereign Identity (SSI) technologies are relatively new and may suffer from **vulnerabilities (e.g., identity theft)**



- The W3C recommendations for decentralized provenance management provides a mechanism for **attributing data to its sources or contributors**.
- For approaches involving the interaction between LLM and KGs, the **transparency of the LLM** itself still depends on the owner



- Studies report limitations of **LLMs in human-like tasks (e.g., explanations, memories, and reasoning over factual statements)**
- Neuro-symbolic systems play a vital role in enhancing trustworthiness by enabling communication between modules and **facilitating tracing**