Knowledge Graphs – a key component in Bilateral AI

Axel Polleres

WIRTSCHAFTSUNIVERSITÄT WIEN VIENNA UNIVERSITY OF ECONOMICS AND BUSINESS

OPENDATA data.wu.ac.at

Institute for Data , Process and Knowledge Management

EFMD EQUIS ACCREDITED AACSB ACCREDITED ASSOCIATION AMBA ACCREDITED

Great to be back!

WIRTSCHAFTS UNIVERSITÄT WIEN VIENNA UNIVERSITY OF ECONOMICS AND BUSINESS

Logic Programming, "Symbolic" Artificial Intelligence

Semantic Web, Knowledge Representation & Reasoning

Web Standard Query Language (SPARQL)

Linked Data

Data Integration, Open Data

Knowledge Graphs

"Bilateral" AI

2003

2005

2007

2011

2013

~2020

LLMs


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”…

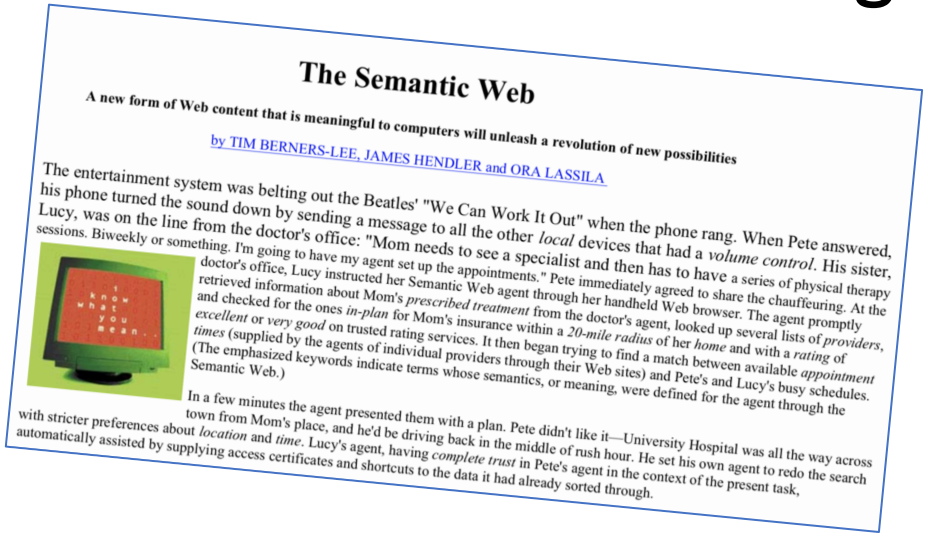

The Semantic Web

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*

Semantic Web: Standard formats, Reasoning & Logics

- (2000s - ca. 2009)

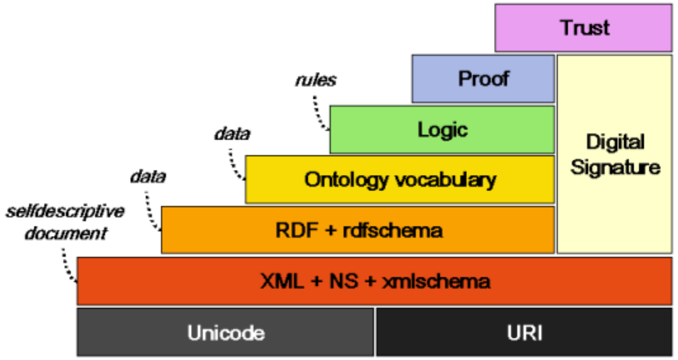

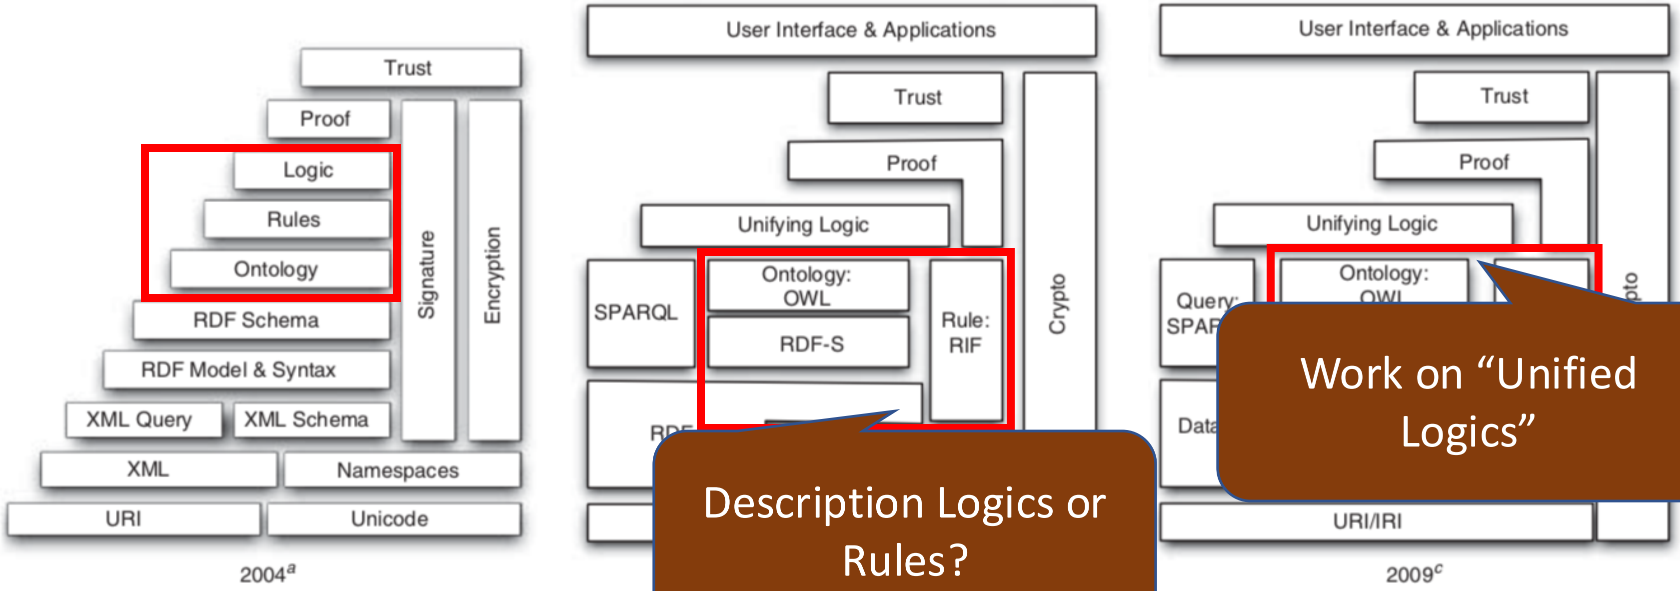

Description Logics or Rules?

Work on "Unified Logics"

(Contextualized Reasoning?)

[a]http://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

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)

…etti, Michael Luck, John-Jules Meyer, and Axel Polleres. Norms, … *The Knowledge Engineering Review*, 28(1):107--116, March 2013.

Focus on Data: Linked Data

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

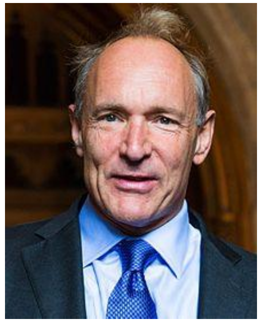

Linked Data Principles

- **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)

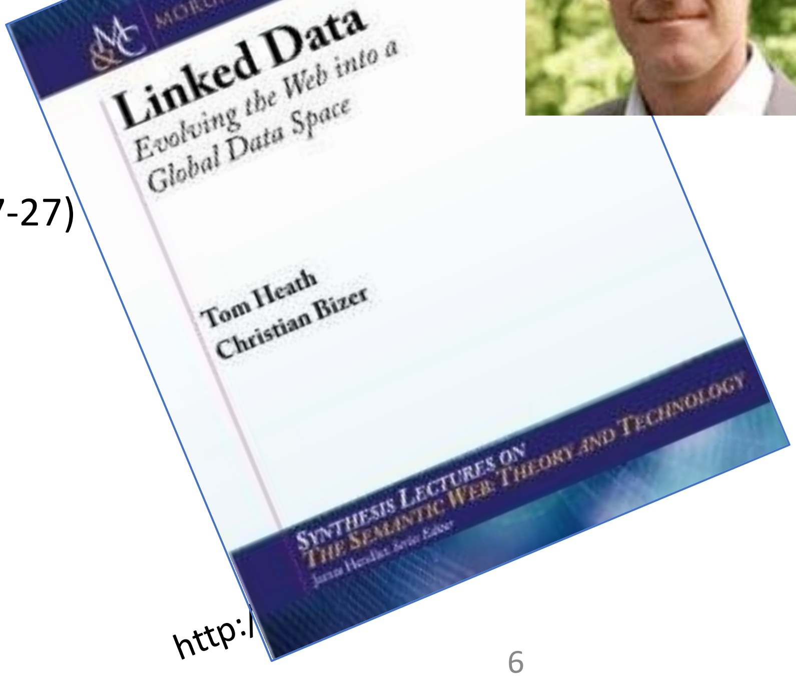

"A Little Semantics Goes a Long Way" (Jim Hendler)

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

http:/

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”)

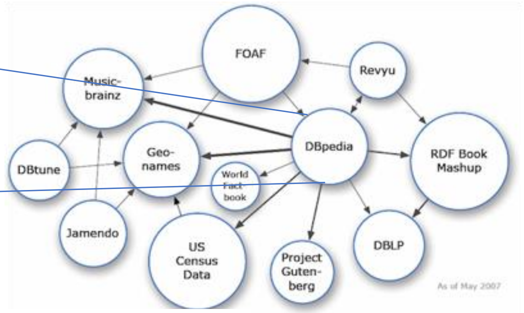

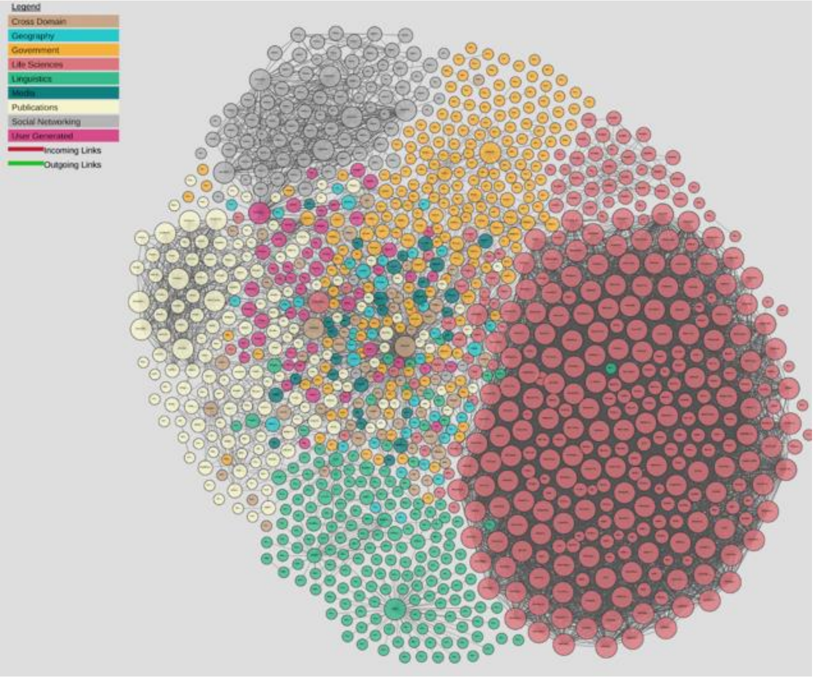

http://lod-cloud.net/

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

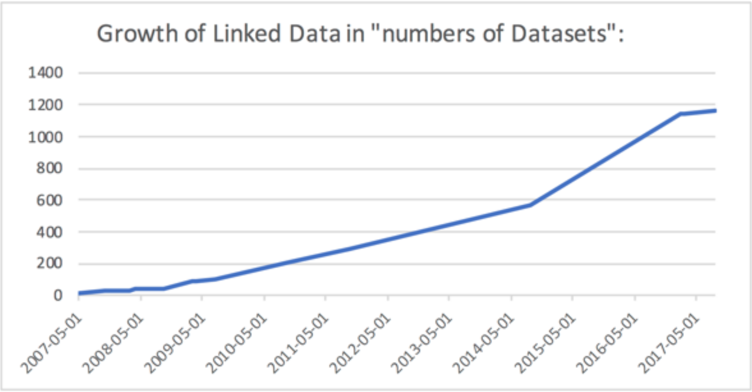

Axel Polleres, Maulik R. Kamdar, Javier D. Fernández, Tania Tudorache, and Mark A. Musen. [A more decentralized vision for linked data](). In *Decentralizing the Semantic Web (Workshop of ISWC2018).*

From Linked Open Data to Knowledge Graphs:

2013: Google adopts Semantic Web ideas under a new name

- Jamie Taylor, Google, Inc., Keynote [ISWC2017]

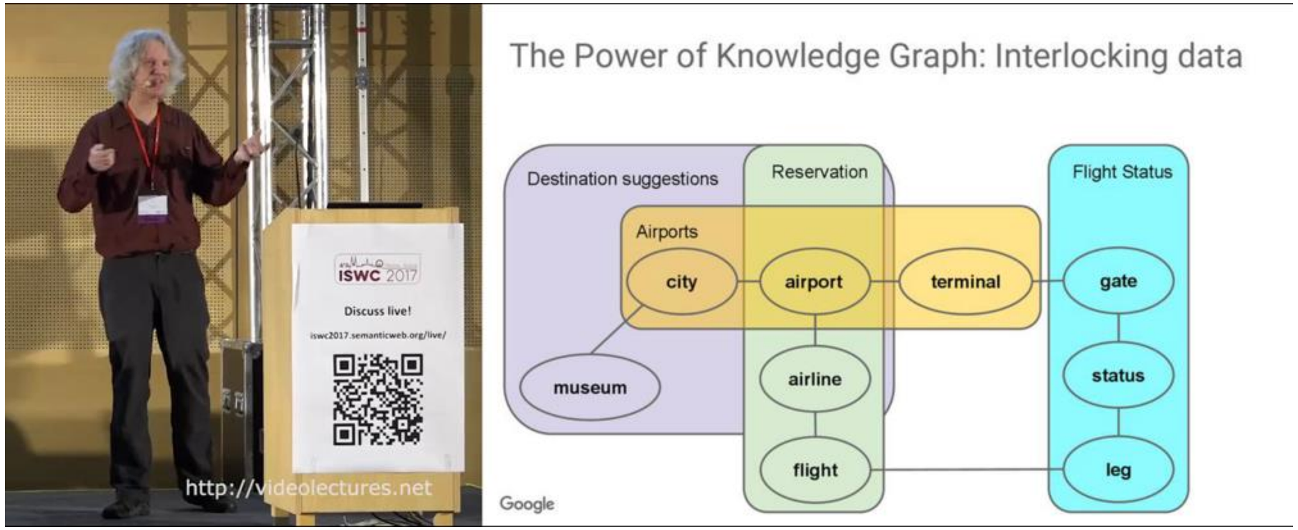

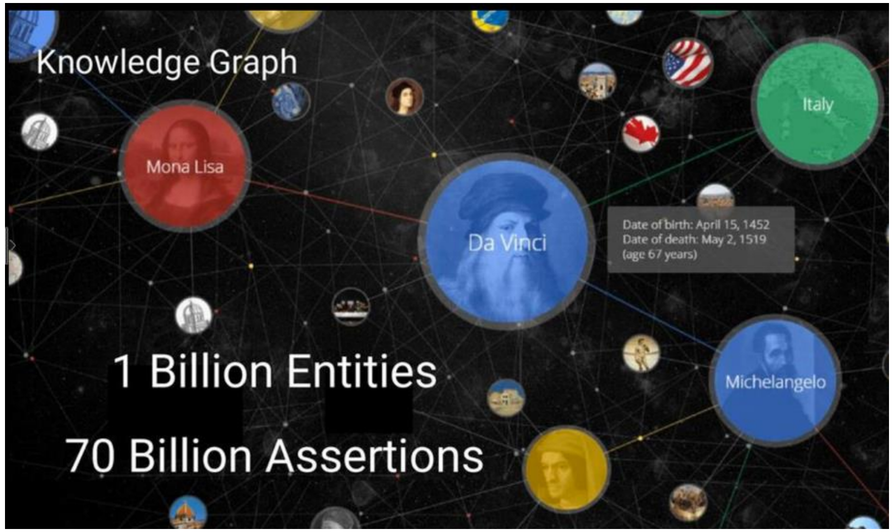

- Actors, Directors, Movies
- Art Works & Museums
- Cities & Countries
- Islands, Lakes, Lighthouses
- Music Albums & Music Groups
- Planets & Spacecraft
- Roller Coasters & Skyscrapers
- Sports Teams

[...]

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

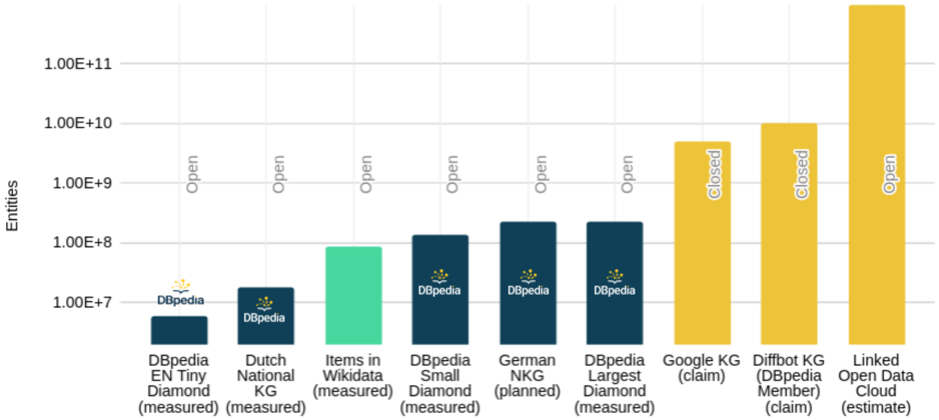

https://www.dbpedia.org 2021

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

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

http://dbpedia.org/resource/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

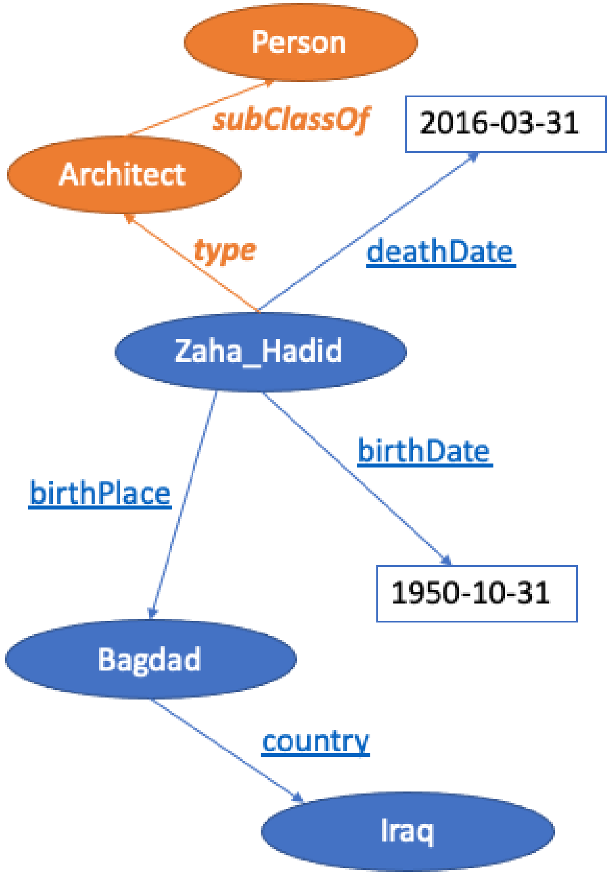

Collaborative, Open Knowledge Graphs:

WIKIDATA

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

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

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

Compare selected revisions

- (cur | prev) 08:56, 8 December 2024 द्वाद्रुन गाजा (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) . . *(Se see autofix на / on Property talk:P12924)*

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

- 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
- Enterprise KGs: knowledge necessary to power applications
- *Ontological expressivity* not central – BUT: ***Expresssing context*** is!

For instance:

- *Provenance*
- *Temporal context*

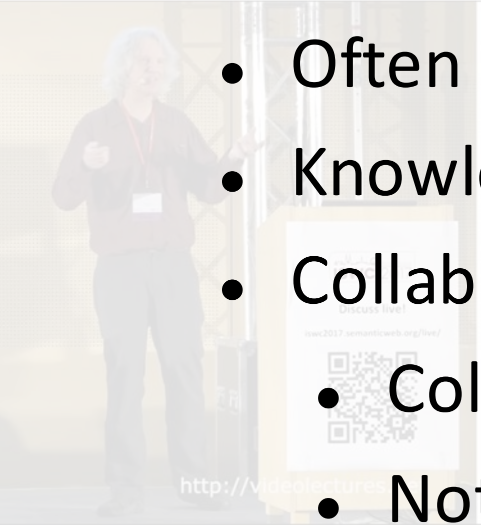

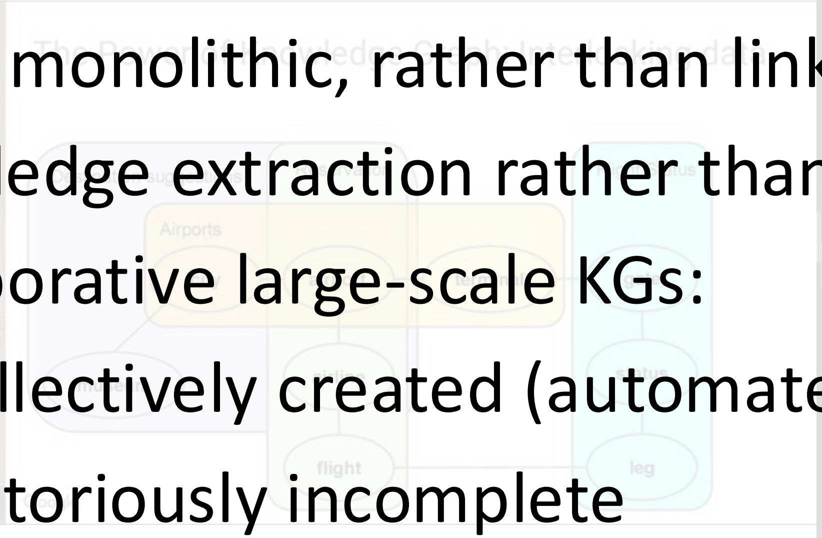

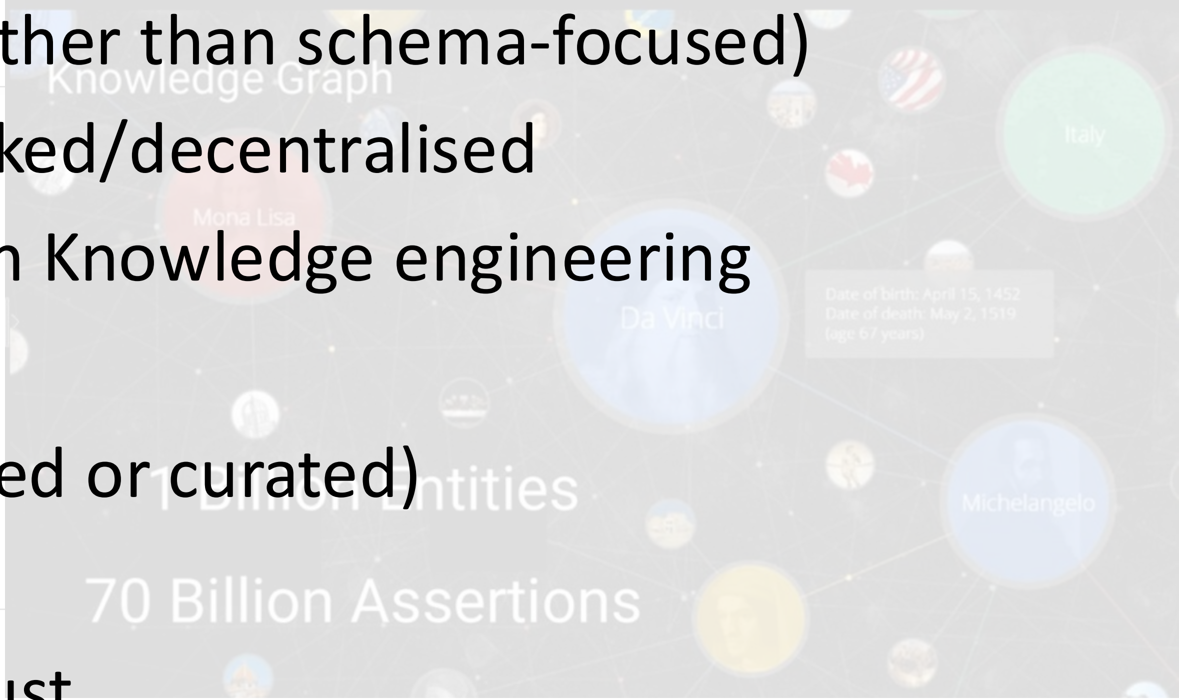

- Actors, Directors, Movies
- Art Works & Museums
- Cities & Countries
- Islands, Lakes, Lighthouses
- Music Albums & Music Groups
- Planets & Spacecraft
- Roller Coasters & Skyscrapers
- Sports Teams

[...]

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

- DBPedia (since 2007) vs. Wikidata (since 2012)

DBpedia

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

DBPedia		Wikidata
✓	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 DBpedia

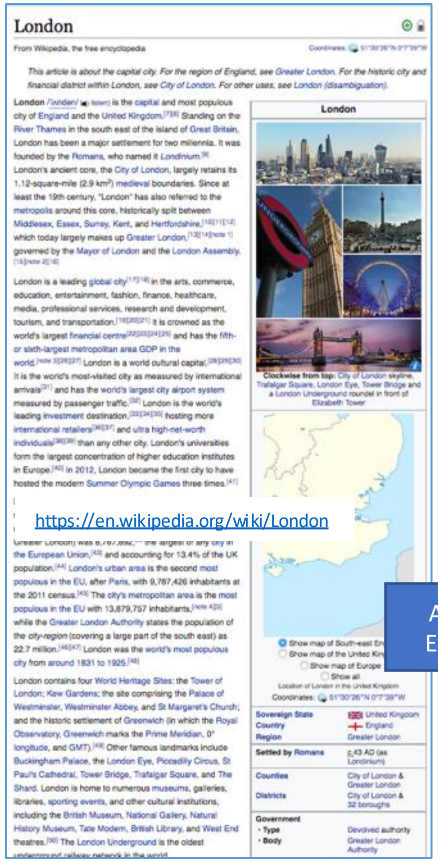

https://en.wikipedia.org/wiki/London

- 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“:

Automatic Exctractors

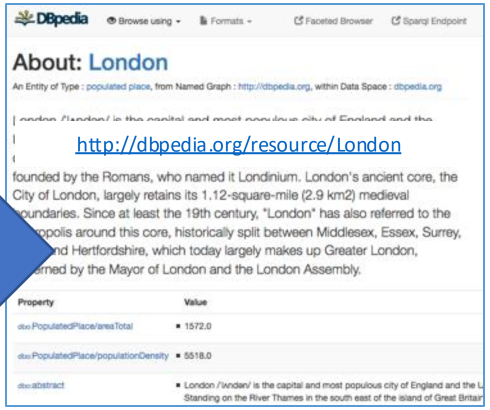

http://dbpedia.org/resource/London

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.

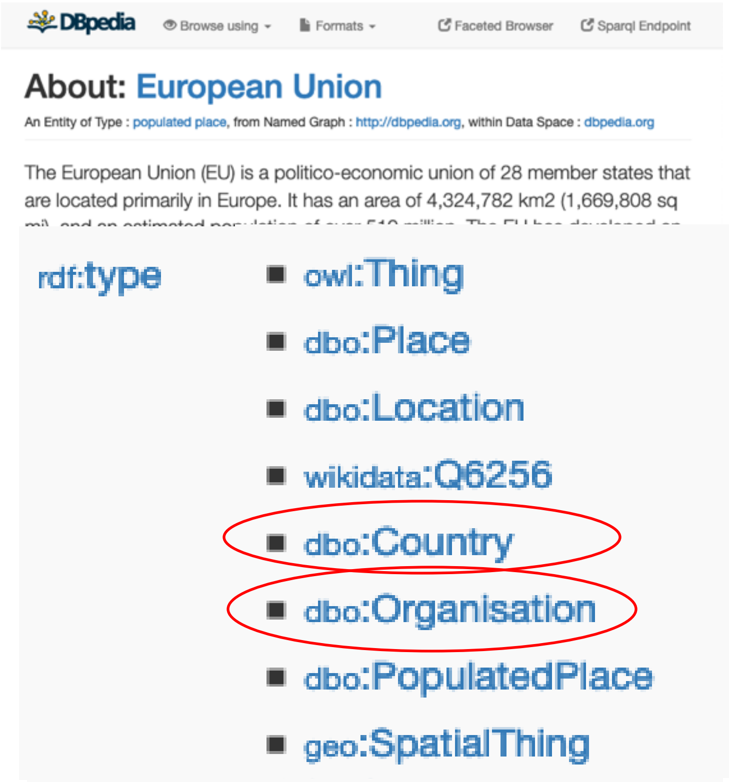

Dbpedia Ontology:

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

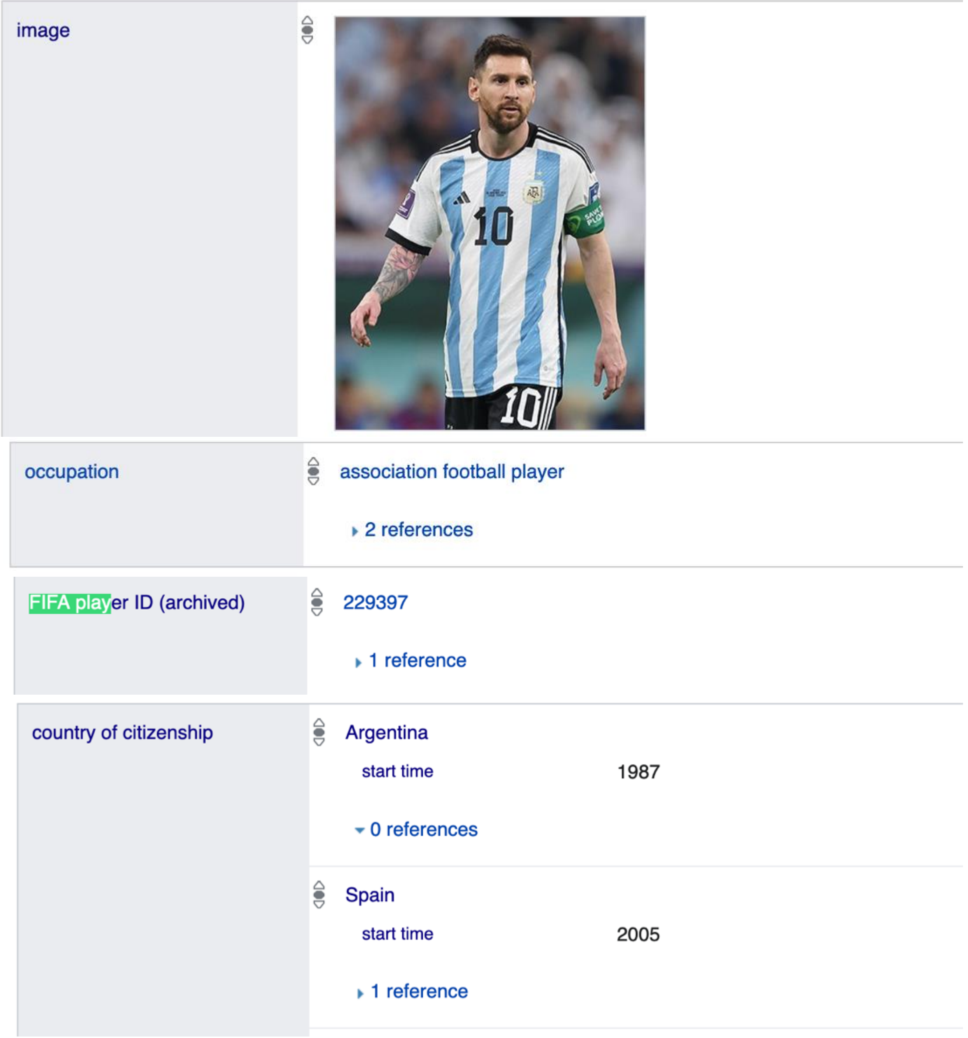

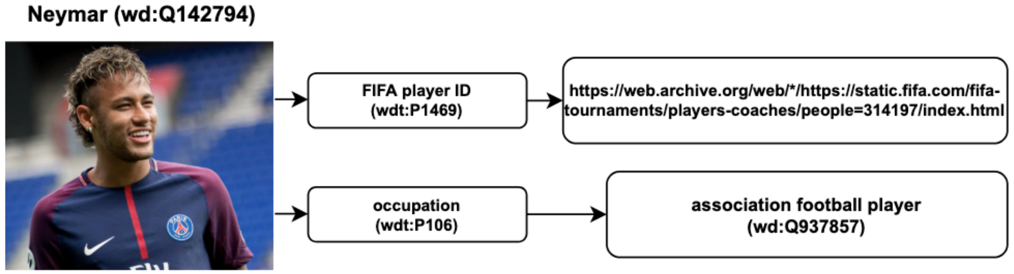

(a) Data graph complying to the constraint

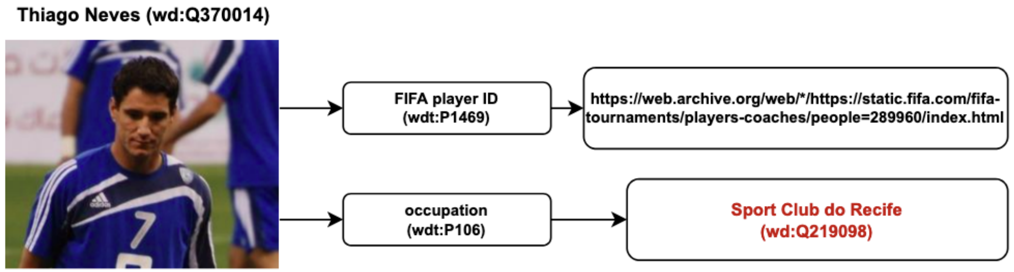

(b) Data graph not complying to the constraint

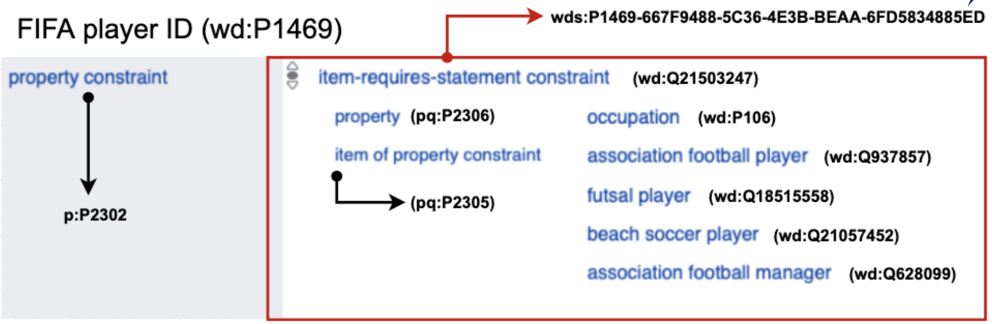

… user defined ***Property Constraints*** (rather than OWL)

The same question as before in Wikidata:

Note: Wikidata does not even use standard OWL

Note: Wikidata uses numeric IDs

- “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) }
```

city (Q515)
large and permanent human settlement

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

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).

country (P17)
sovereign state of this item

United Kingdom (Q145)
country in Europe

- What’s this?

The same question as before in Wikidata:

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

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)
}
```

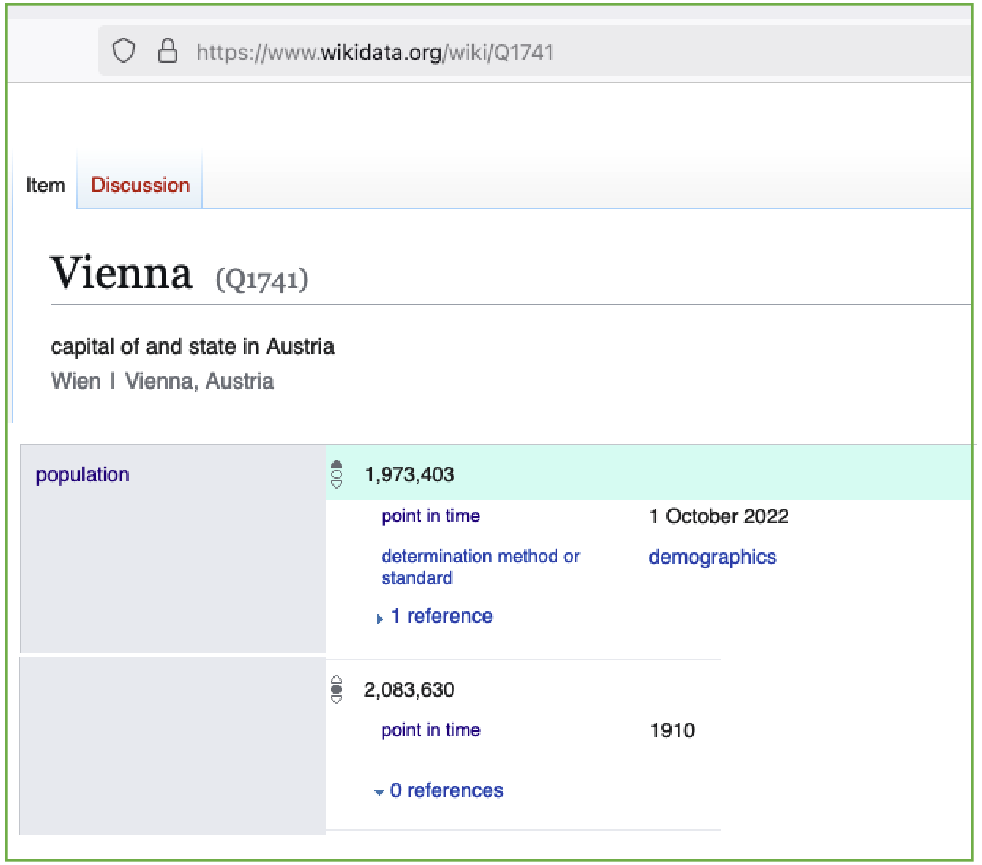

So, WHEN did Vienna have 2M inhabitants?

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:

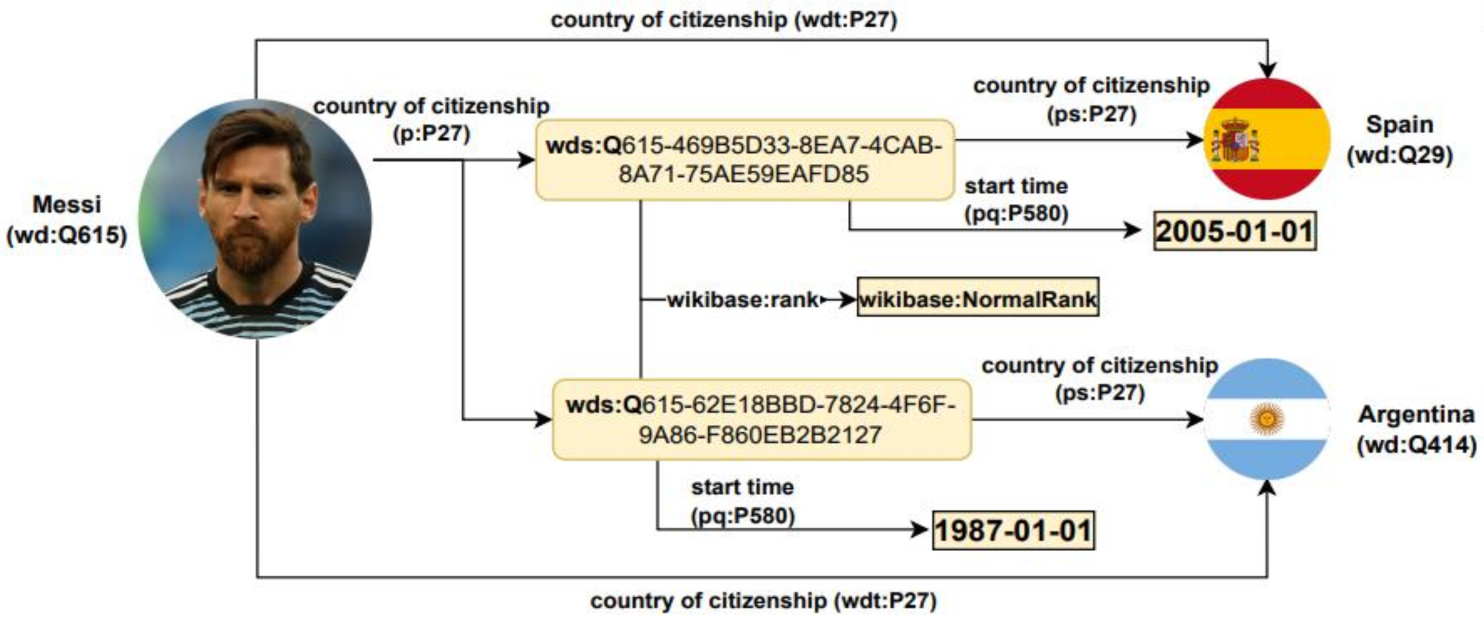

See our recent ISWC2024 tutorial: https://ww101.ai.wu.ac.at/

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!

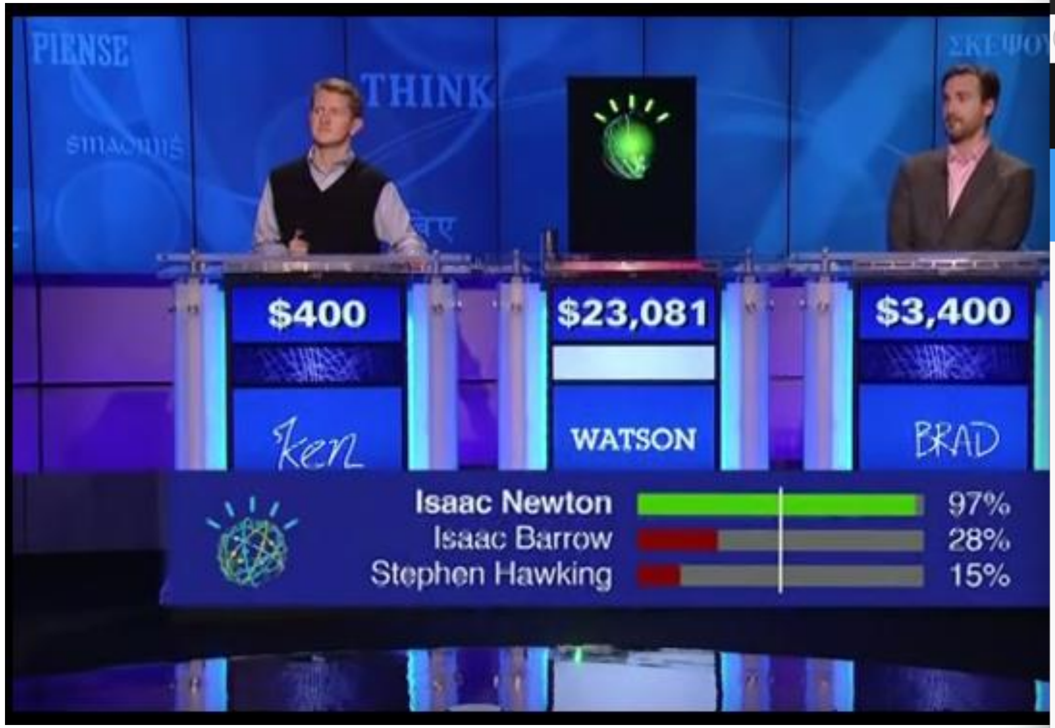

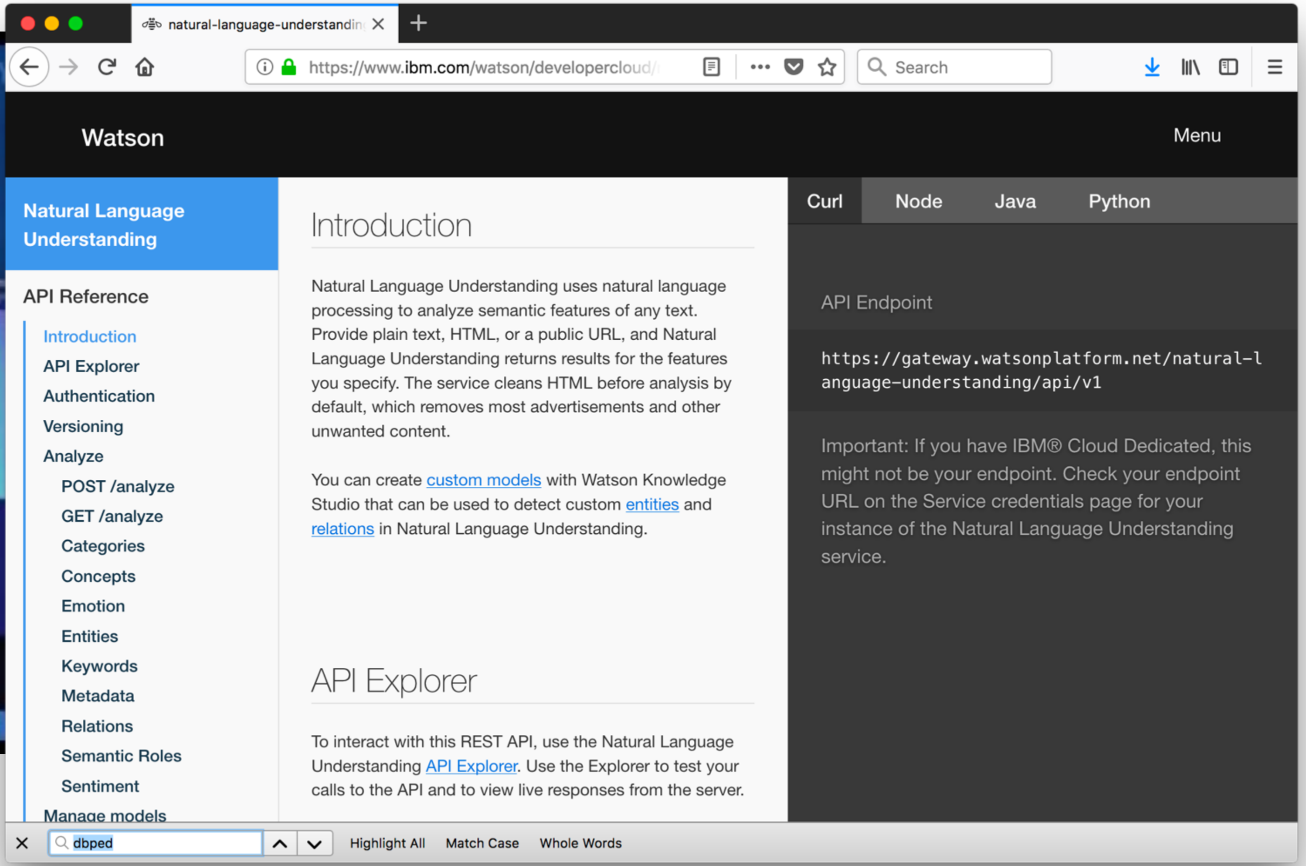

https://youtu.be/P0Obm0DBvwI?t=951

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

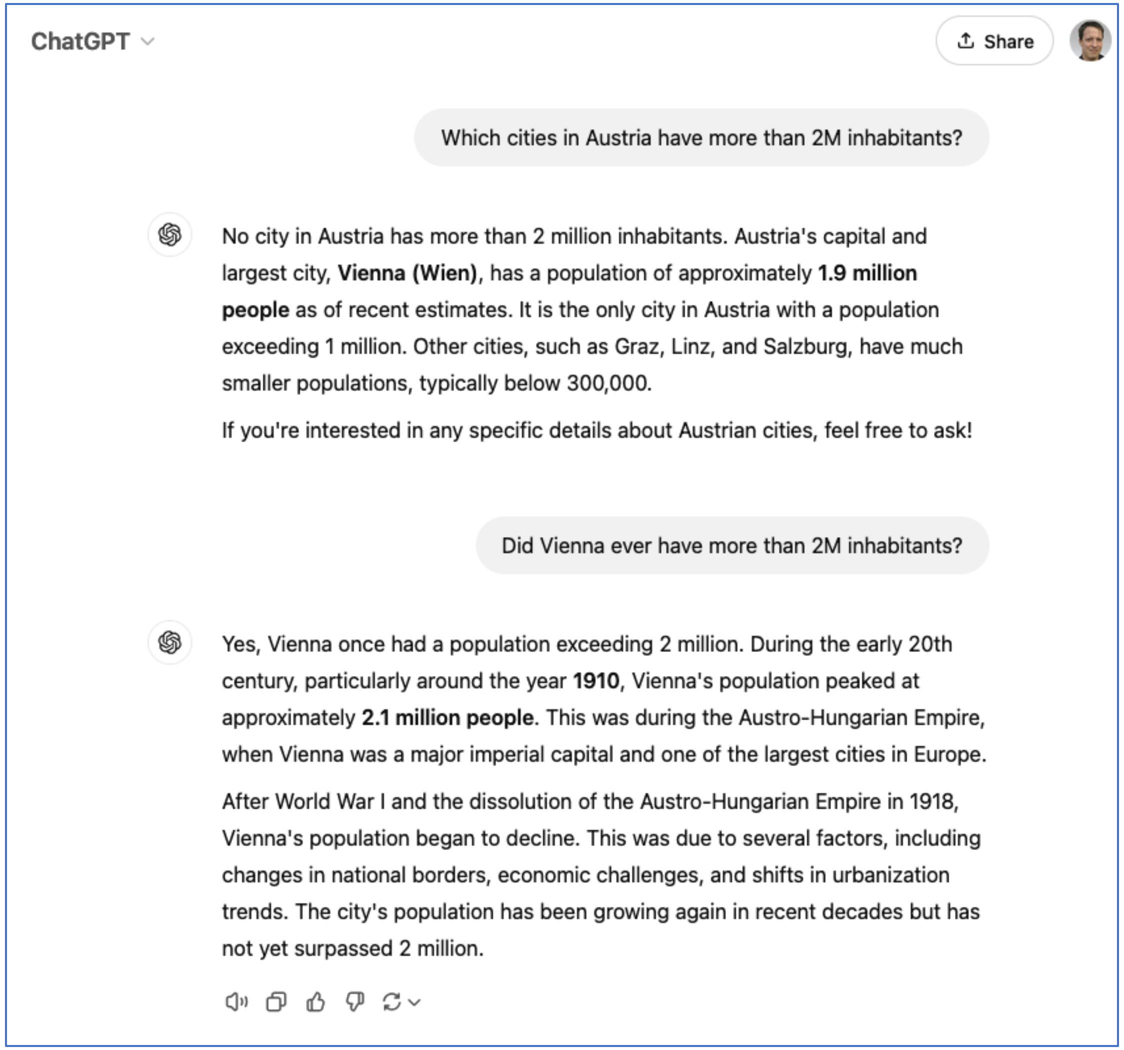

- KGC23 Keynote: “The Future of Knowledge Graphs in a World of LLMs — Denny Vrandečić, Wikimedia”

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

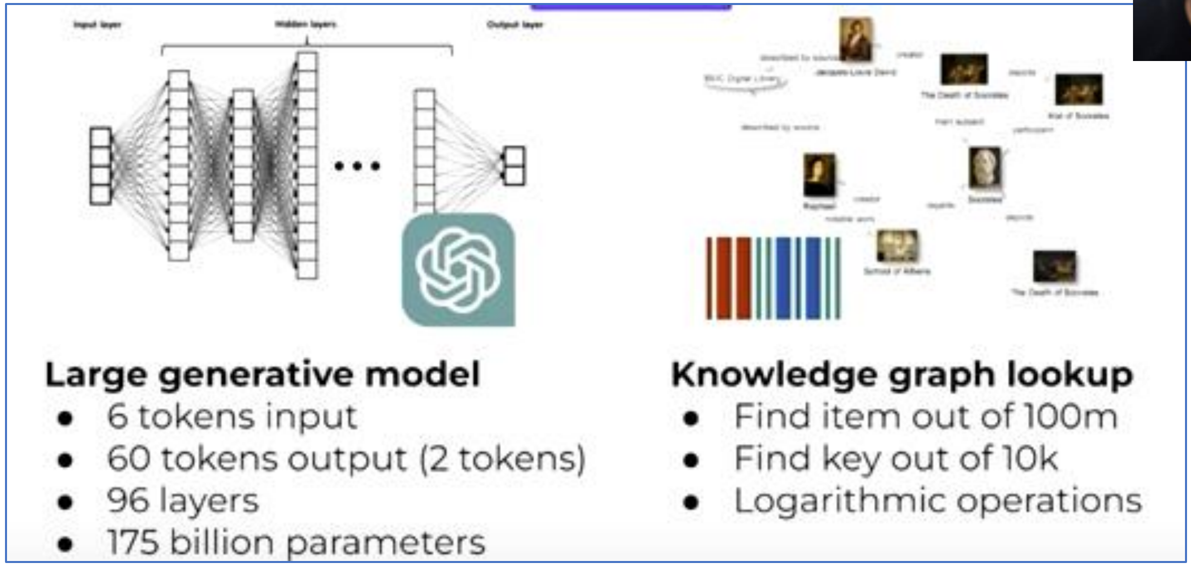

Admittedly, Denny didn't talk about this…

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

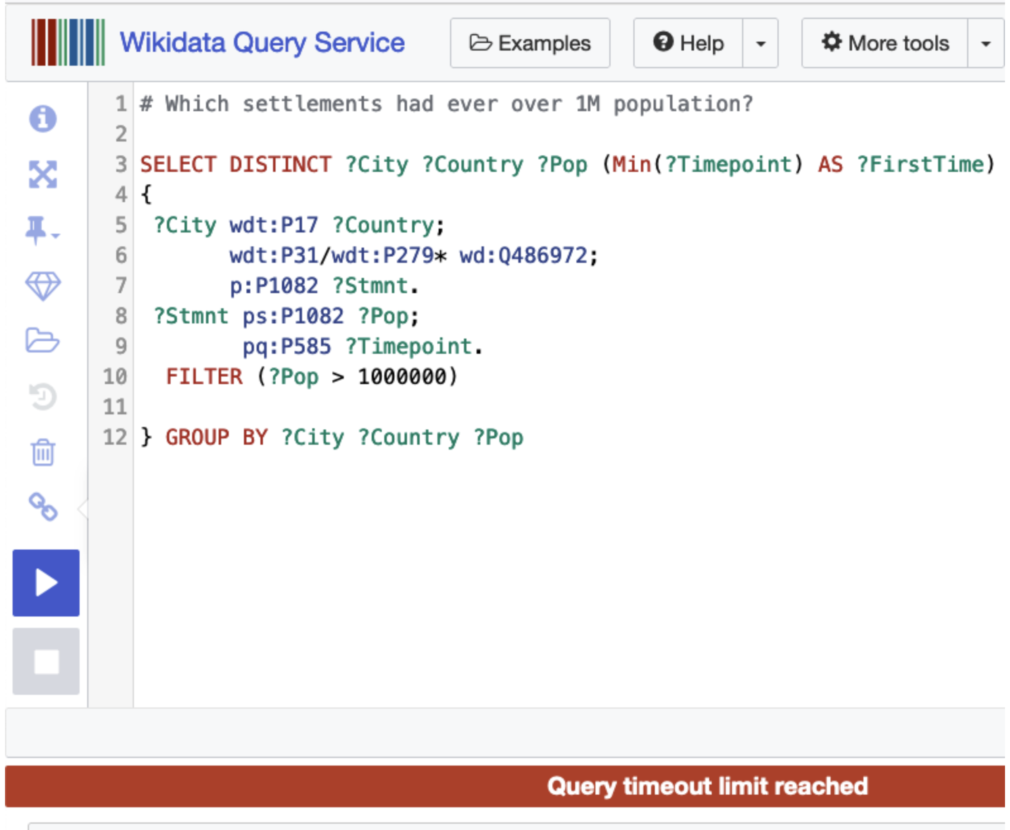

Challenge: scaling queries to large-scale, schemaless KGs

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 KGs deal 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.*

Idea: use ***unsupervised message passing*** *to propagate confidence scores obtained by parsing an input question and matching terms in the knowledge graph to a set of possible answers.*

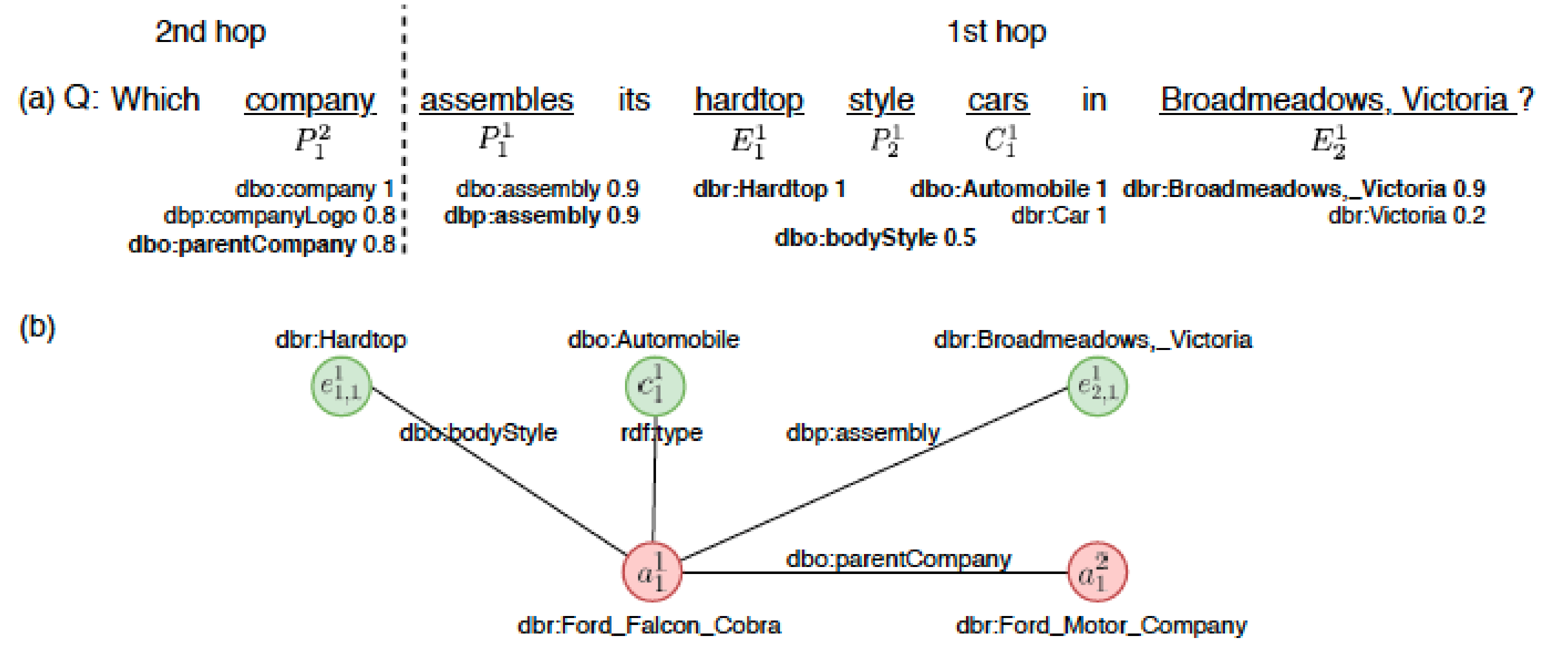

Figure 1: (a) A sample question Q highlighting different components of the question interpretation model: references and matched URIs with the corresponding confidence scores, along with (b) the illustration of a sample KG subgraph relevant to this question. The URIs in bold are the correct matches corresponding to the KG subgraph.

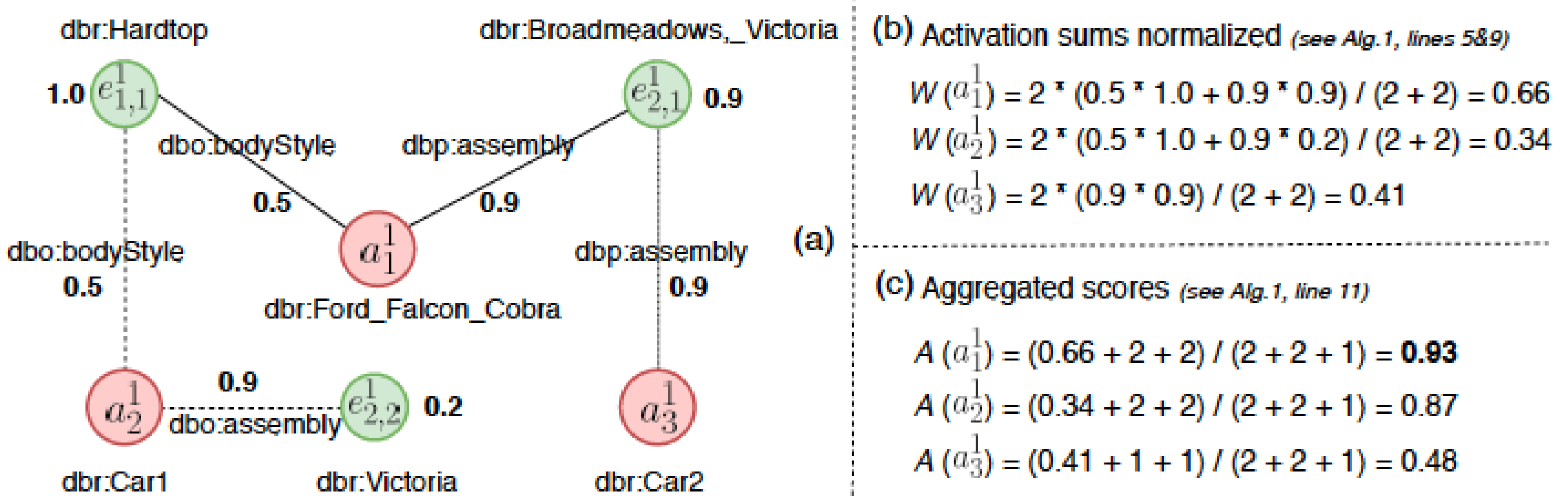

Figure 2: (a) A sample subgraph with three entities as candidate answers, (b) their scores after predicate and entity propagation, and (c) the final aggregated score.

Some of our own research in this area:

- How good or bad can KG swith 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.*
- 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

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

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

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

Category	Subcategory	Example	SE	KG	LM	SE + KG + LM
Facts	Popular	*Who directed the movie Spotlight?*	+ good coverage − noisy content	+ high precision − poor coverage	+ good coverage − noisy content	+ good coverage + high precision
	Long-Tail	*Which galaxy is closest to the Sunflower Galaxy?*	+ good coverage − needle in haystack	+ high precision − sparse coverage	− limited storage − hallucinations	+ good coverage + high precision
	Multi-hop	*Which Turing Award winners were born in Latin America?*	− no reasoning − single-shot search	+ formal reasoning + structured queries	+ latent reasoning − hallucinations	+ formal reasoning + structured queries
	Analytical	*How many U.S. Congress Members are younger than 50?*	− no datatypes − no aggregation	+ rich datatypes + aggregation	− no datatypes − no aggregation	+ rich datatypes + aggregation
Explanations	Commonsense	*How do snakes move?*	+ good coverage + text output	− poor coverage − structured output	+ good coverage + text output	+ good coverage + text output
	Causal	*What caused the dancing plague of 1518?*	+ good coverage + text output	+ long tail − structured output	+ good coverage + text output	+ good coverage + text output
	Exploratory	*Who was Williamina Fleming?*	+ text output + ranked results	+ graph algorithms + browsing + navigation	+ interactive + synthesis	+ hybrid output + hybrid interactivity
Planning	Instructive	*How do I tie a Windsor Knot?*	+ multimedia + diverse results	− poor coverage − non-didactic output	+ interactive − no multimedia	+ interactive + diverse results
	Recommendation	*Should I pack warm clothes for Iceland in June?*	+ diverse results + ranked results	− poor coverage − no recommendations	+ interactive + synthesis	+ interactive + synthesis
	Spatio-temporal	*What kid-friendly Italian restaurants are near Disneyland?*	+ events & maps − no integration	+ integration + s.-t. operators	− lacks freshness − no s.-t. operators	+ s.-t. operators + integration
Advice	Lifestyle	*How can I improve my work/life balance?*	+ diverse results + ranked results	− poor coverage − lacks nuance	+ interactive + synthesis	+ diverse results + interactive
	Cultural	*Should I tip bartenders in Canada?*	+ diverse results + ranked results	− poor coverage − lacks nuance	+ synthesis − cultural bias	+ diverse results + synthesis
	Philosophical	*Is the death penalty ever acceptable?*	+ diverse results + ranked results	− poor coverage − lacks nuance	+ interactive + synthesis	+ diverse results + synthesis

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

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)*”

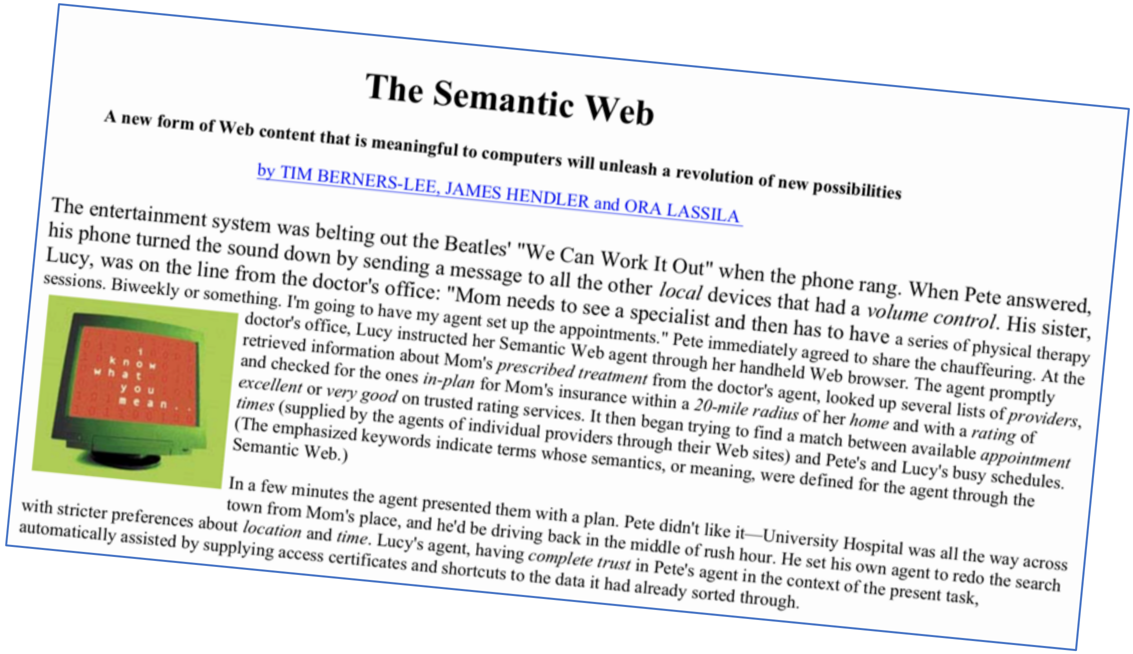

The Semantic Web

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 Exellence” BILAI (funded by FWF):**
  - **Vison 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].

FWF Österreichischer Wissenschaftsfonds

BILAI

BilAI Consortium (~30M EUR, 5Y)

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

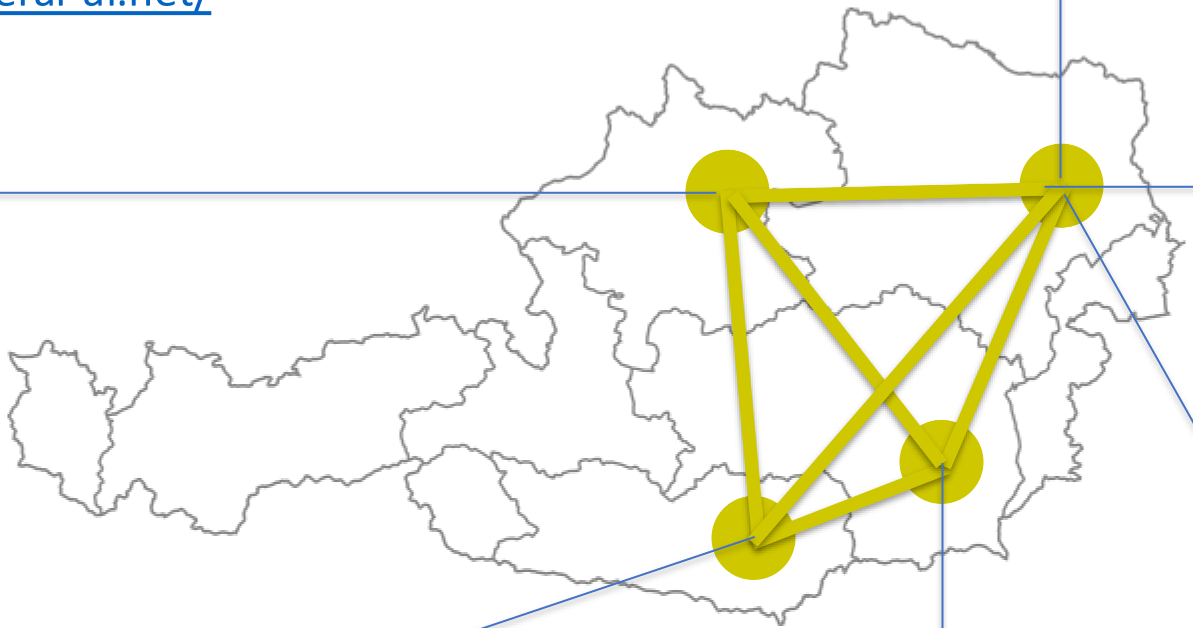

JKU

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 Intelligrnce

Institute of Science and Technology Austria

Christoph Lambert
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

UNIVERSITÄT KLAGENFURT

Gerhard Friedrich
Symbolic AI
Model-based reasoning

- Institute for Artificial Iltelligence and Cybersecurity

TU Graz

Robert Legenstein
Machine Learning
Computational Neuroscience

- Institute of Theoretical Computer Science

WU VIENNA

Axel Polleres
Knowledge Graphs

- Institute for Data Process and Knowledge Management

BILAI Slides: thanks to Günther Klambauer (JKU)

BILAI

General Architecture

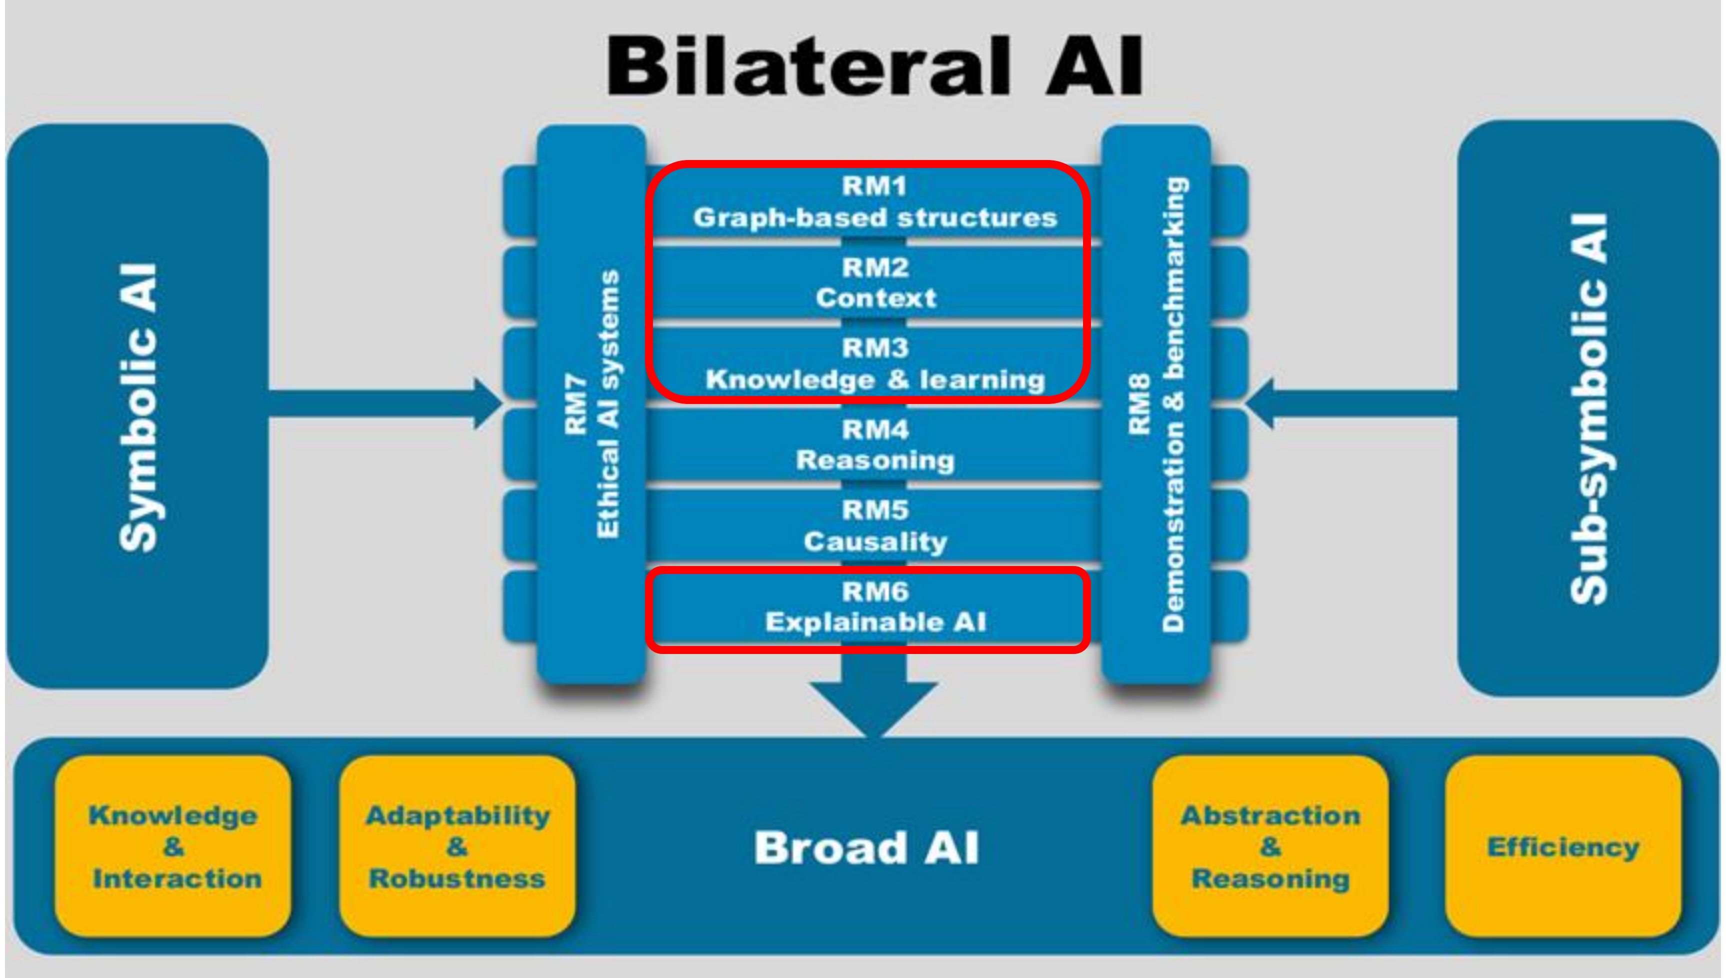

BILAI Slides: thanks to Günther Klambauer (JKU)

Vision: Building a „Broad“ AI

BILAI

Human cognitive abilities

General AI
Human level abilities

Knowledge & Interaction

Adaptability & Robustness

Abstraction & Reasoning

Efficiency

Broad AI
Broad cognitive abilities

Image generation
DALL·E 2

Language processing
ChatGPT

Game playing
AlphaGO

Structure prediction
AlphaFold

Narrow AI
Task specific skills

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.

BILAI Slides: thanks to Günther Klambauer (JKU)

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

    → <u>tight integration</u> of LLMs with symbolic ***[models (KGs) &]*** solvers [...] could be employed to leverage incremental reasoning capabilities (**RM1-4**)

BILAI

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)**

BILAI Slides: thanks to Günther Klambauer (JKU)

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.

BILAI Slides: thanks to Günther Klambauer (JKU)

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
Time and other Contextual information

- Our Starting Points:
  - In order to learn over time, we nee to build Corpora (Crawling) of Evolving KGs

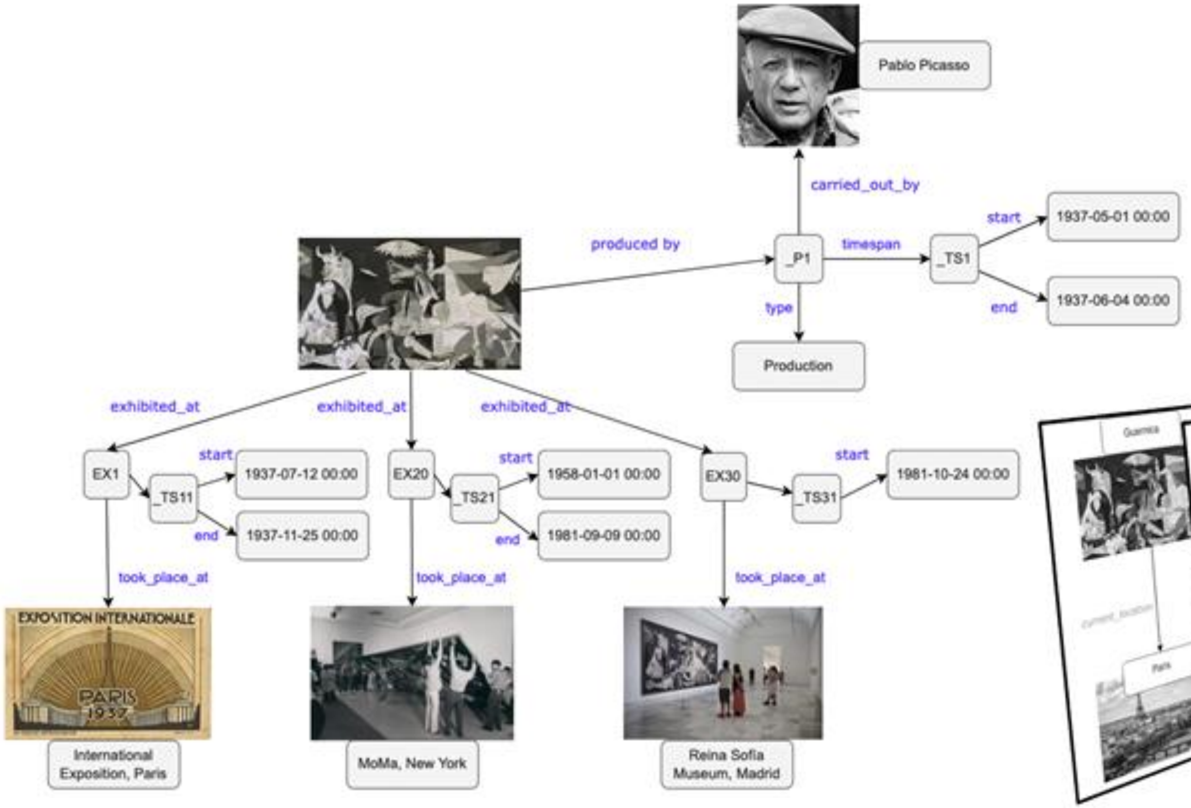

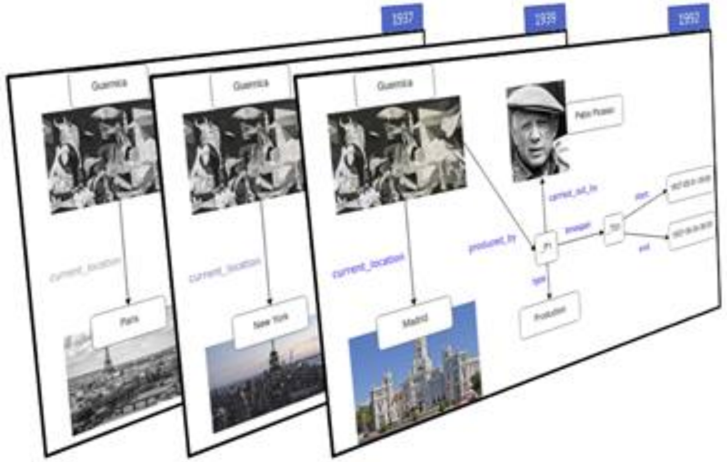

How Does Knowledge Evolve in Open Knowledge Graphs?

Axel Polleres
Vienna University of Economics and Business, Austria
Complexity Science Hub Vienna, Austria
Romana Pernisch[1]
Vrije Universiteit Amsterdam, the Netherlands
Discovery Lab, Elsevier, the Netherlands
Angela Bonifati
Lyon 1 University, CNRS LIRIS, France
IUF, France
Daniele Dell'Aglio
Aalborg University, Denmark
Daniil Dobriy
Vienna University of Economics and Business, Austria
Stefania Dumbrava
ENSIIE, France
SAMOVAR, IP Paris, France
Lorena Etcheverry

SMW Cloud: A Corpus of Domain-Specific Knowledge Graphs from Semantic MediaWikis

Daniil Dobriy[1(✉)], Martin Beno[1], and Axel Polleres[1,2]

[1] Vienna University of Economics and Business, Vienna, Austria
{daniil.dobriy,martin.beno,axel.polleres}@wu.ac.at
[2] Complexity Science Hub, Vienna, Austria

Abstract. Semantic wikis have become an increasingly popular means of collaboratively managing Knowledge Graphs. They are powered by platforms such as Semantic MediaWiki and Wikibase, both of which enable MediaWiki to store and publish structured data. While there are many semantic wikis currently in use, there has been little effort to collect and analyse their structured data, nor to make it available for the research community. This paper seeks to address this gap by systematically collecting structured data from an extensive corpus of Semantic-MediaWiki-powered portals and providing an in-depth analysis of the ontological diversity (and re-use) amongst these wikis using a variety of ontological metrics. Our paper aims to demonstrate that semantic wikis are a valuable and extensive part of Linked Open Data (LOD), and in fact may be considered an own active "sub-cloud" within the LOD ecosystem, which can provide useful insights into the evolution of small and medium-sized domain-specific Knowledge Graphs.

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!*

Formalizing and Validating Wikidata’s Property Constraints using SHACL and SPARQL

Nicolas Ferranti[a,*], Jairo Francisco De Souza[b], Shqiponja Ahmetaj[c] and Axel Polleres[a,d]

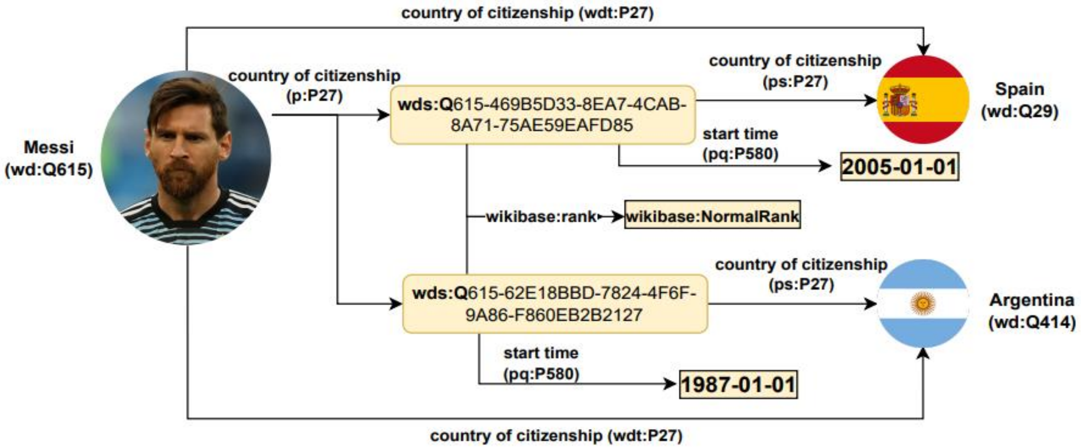

Research Questions & Starting points in my group
Querying large-scale KGs

- Starting Points:
  - Speeding up decentralized SPARQL Querying by Graph partition-shipping 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

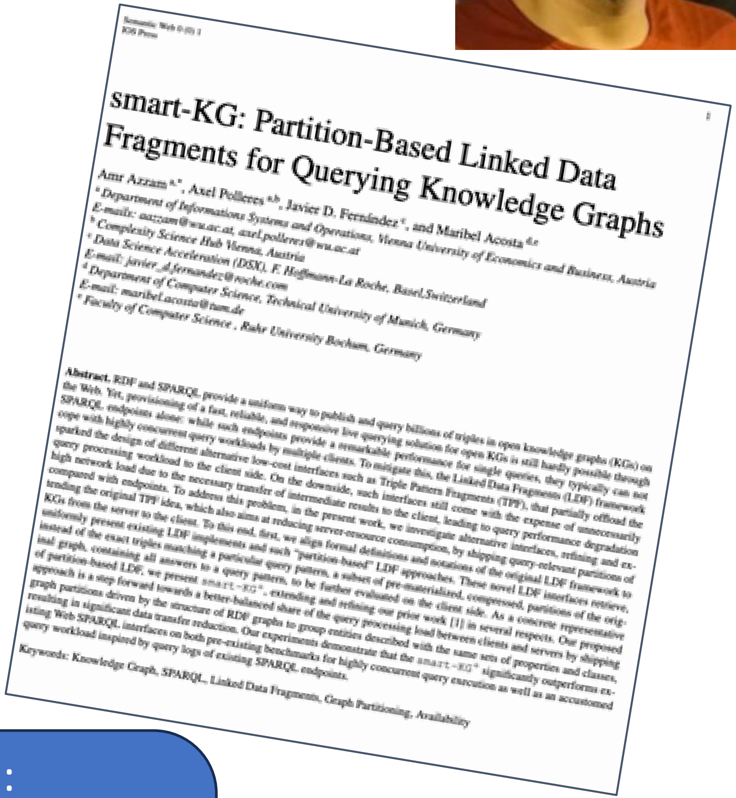
smart-KG: Partition-Based Linked Data Fragments for Querying Knowledge Graphs

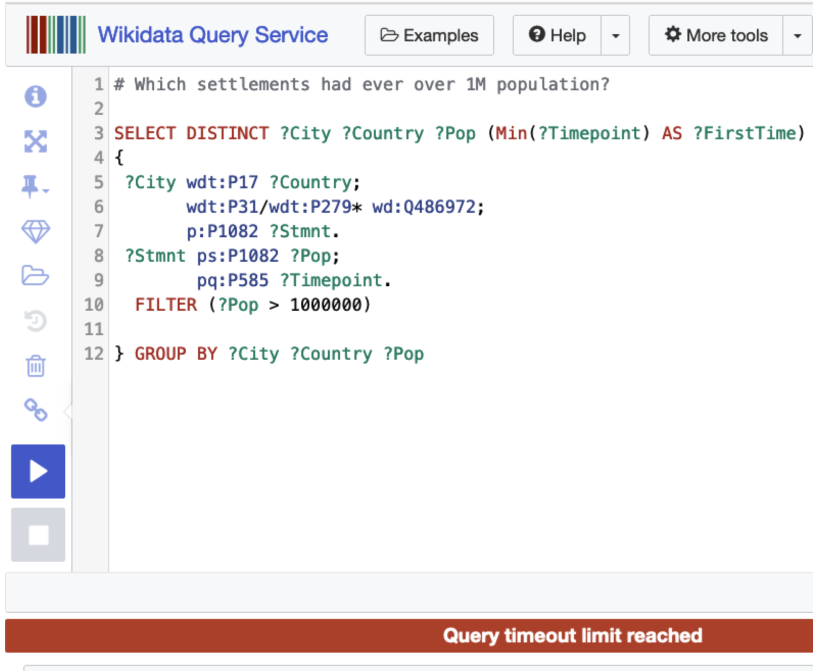

Future work in BILAI:
How can we similarly split work in a decentralized manner for other KG/AI tasks?
e.g. can we similarly modularize Knowledge Graph embeddings?

Starting points for collaboration:

Slide: Emanuel Sallinger

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?

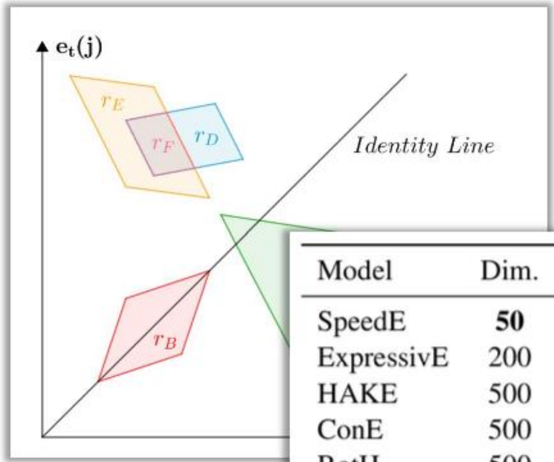

Model	Dim.	MRR	Conv. Time	#Parameters
SpeedE	**50**	**.500**	**6min**	**2M**
ExpressivE	200	**.500**	31min	8M
HAKE	500	.497	50min	41M
ConE	500	.496	1.5h	20M
RotH	500	.496	2h	21M

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

What’s next?

- Austrian National “Cluster of Exellence” BILAI:
  - Vison 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/

BILAI

AI in our Department – at a glance

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

Semantic Systems Group
(Marta Sabou)

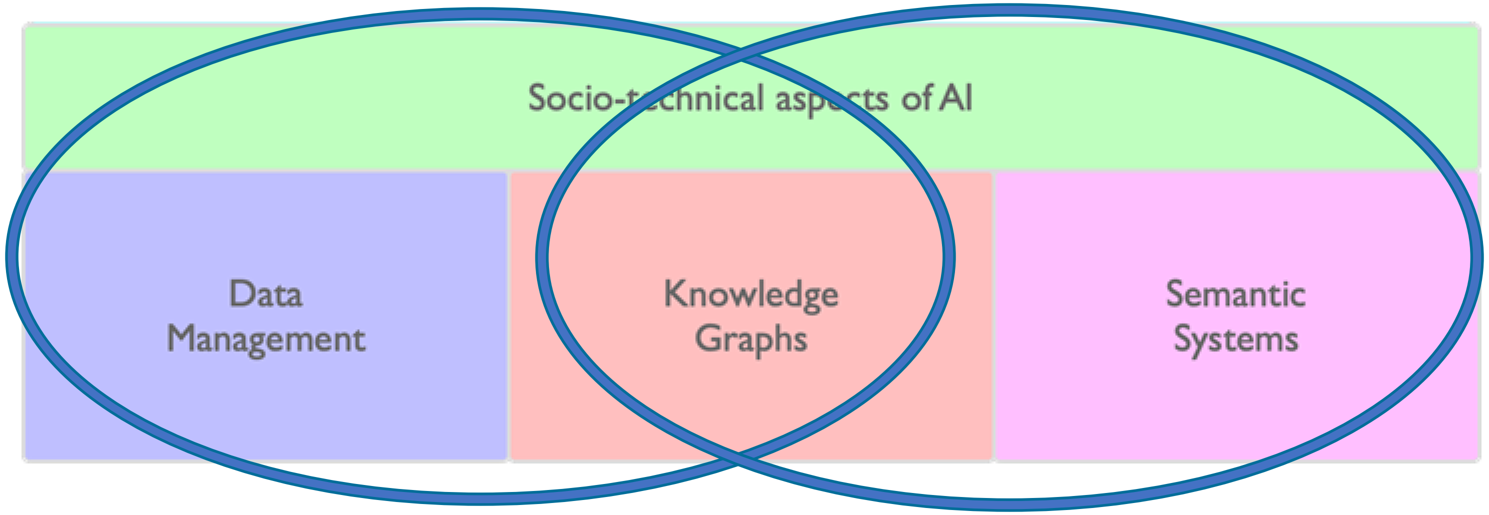

AI areas of interest:

- 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)

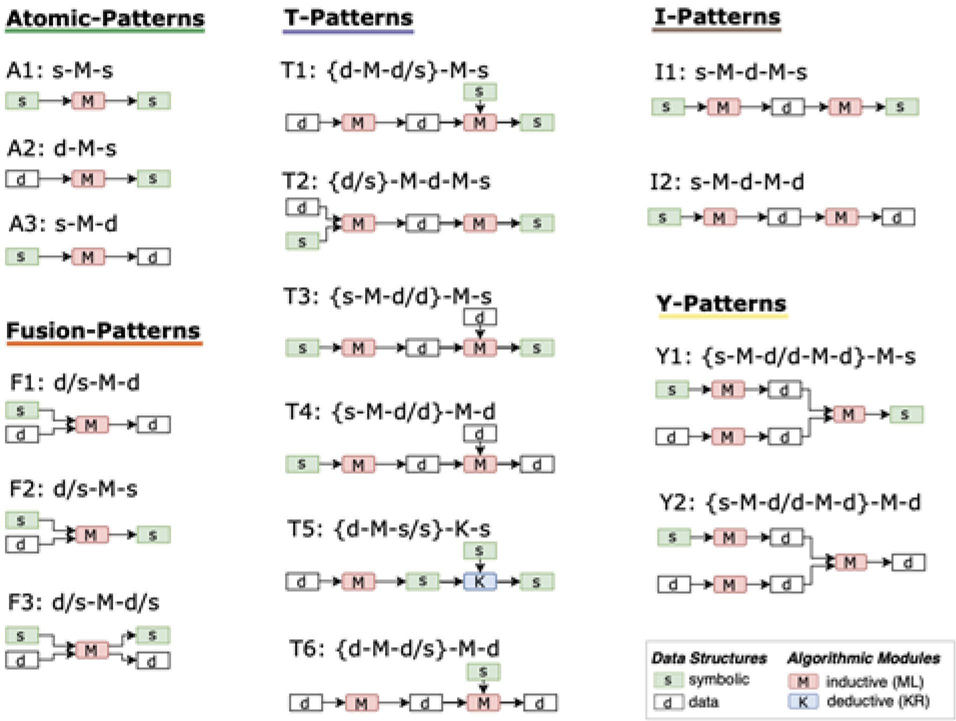

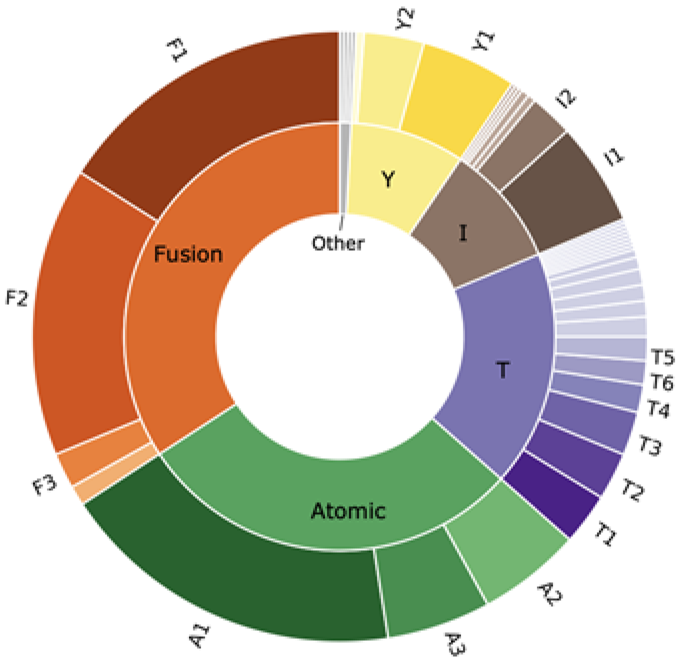

A. Breit, L. Waltersdorfer, F.J. Ekaputra, M. Sabou, A. Ekelhart, A. Iana, H. Paulheim, J. Portisch, A. Revenko, A. ten Teije, and F. van Harmelen. 2023. Combining Machine Learning and Semantic Web: A Systematic Mapping Study. ACM Computing Survey. March 2023.

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

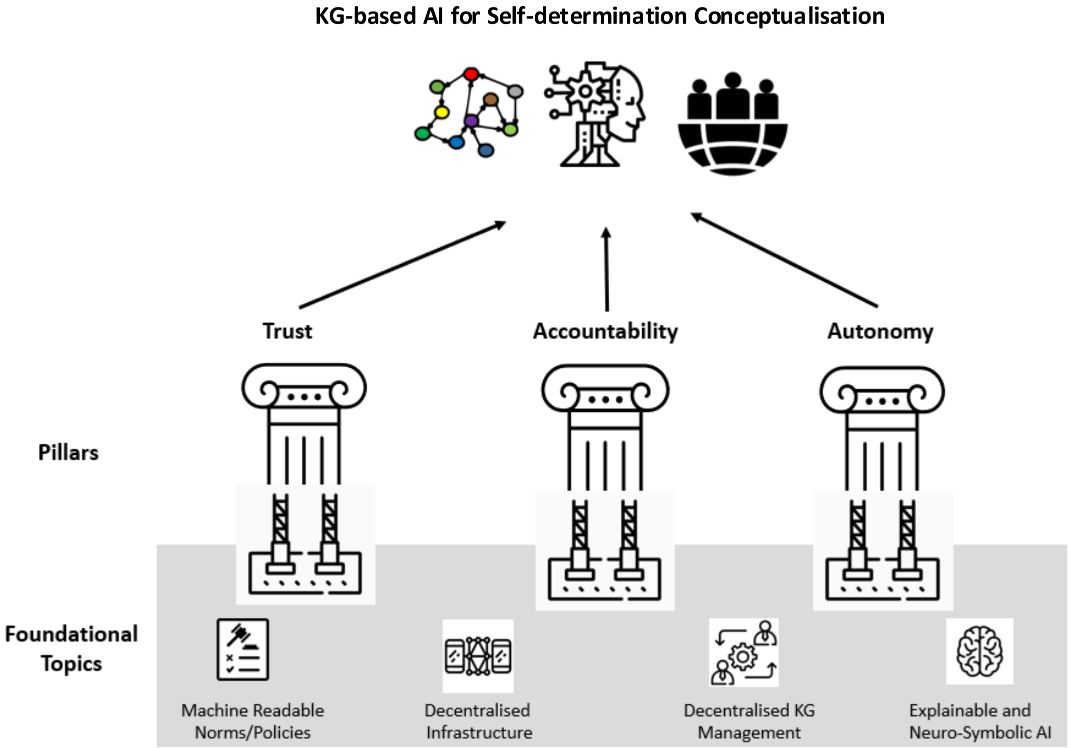

Ibáñez, L., Domingue, J., **Kirrane, S.**, Seneviratne, O., Third, A., Vidal, M., 2023. Trust, Accountability, and Autonomy in Knowledge Graph-based AI for Self-determination. Transactions on Graph Data and Knowledge (TGDK) *(revised and resubmitted)*

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

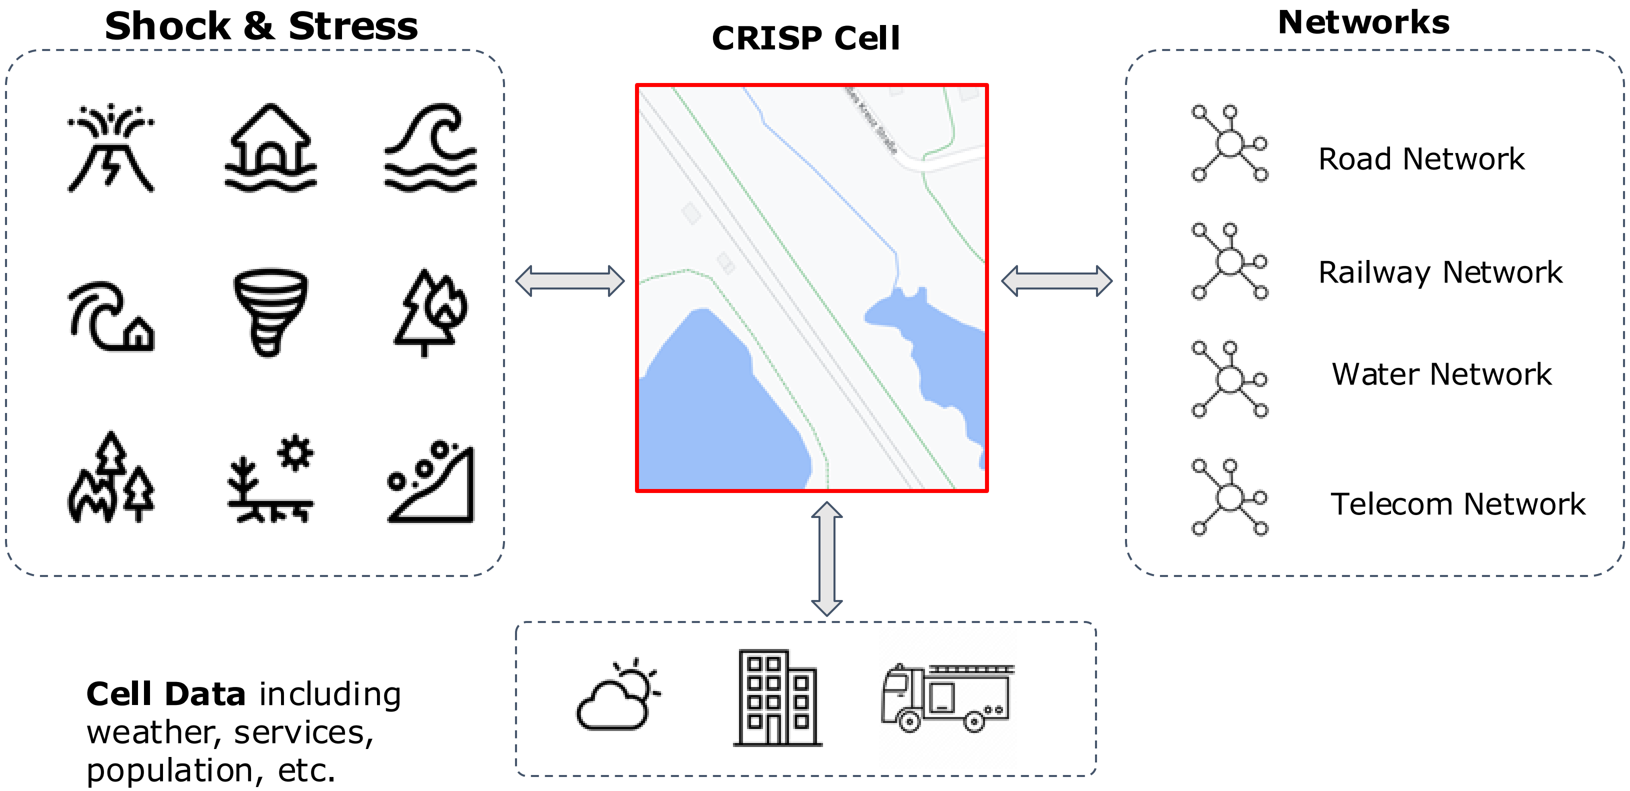

Thank you!

- Summary:
  - (Semantic) Web & Knowledge Graphs play in 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

WU WIRTSCHAFTS UNIVERSITÄT WIEN VIENNA UNIVERSITY OF ECONOMICS AND BUSINESS

P.S.: we're hiring! ☺

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

BILAI

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

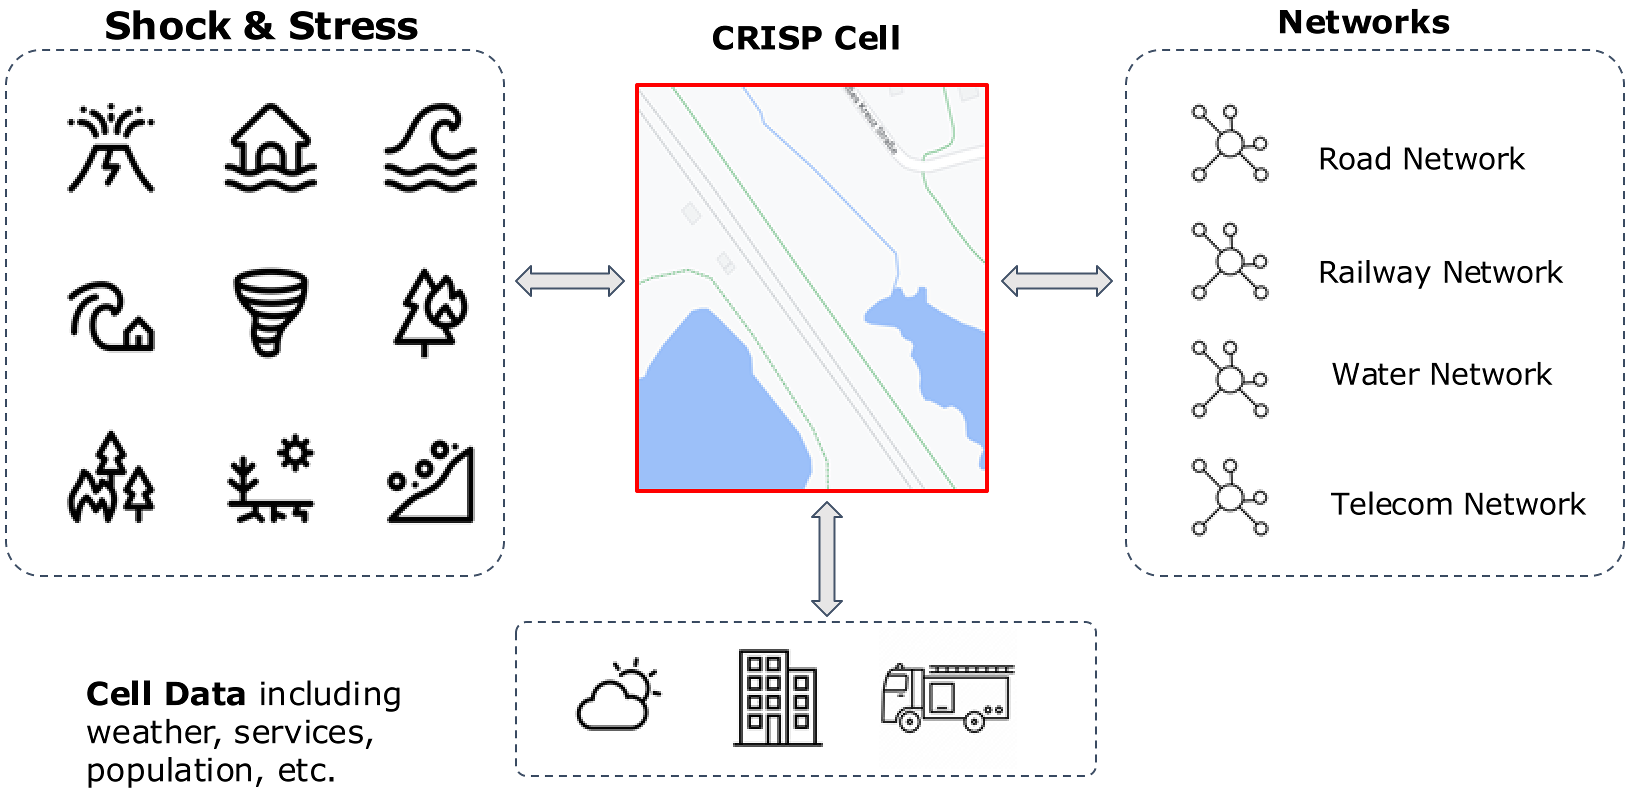

CRISP Semantic Model

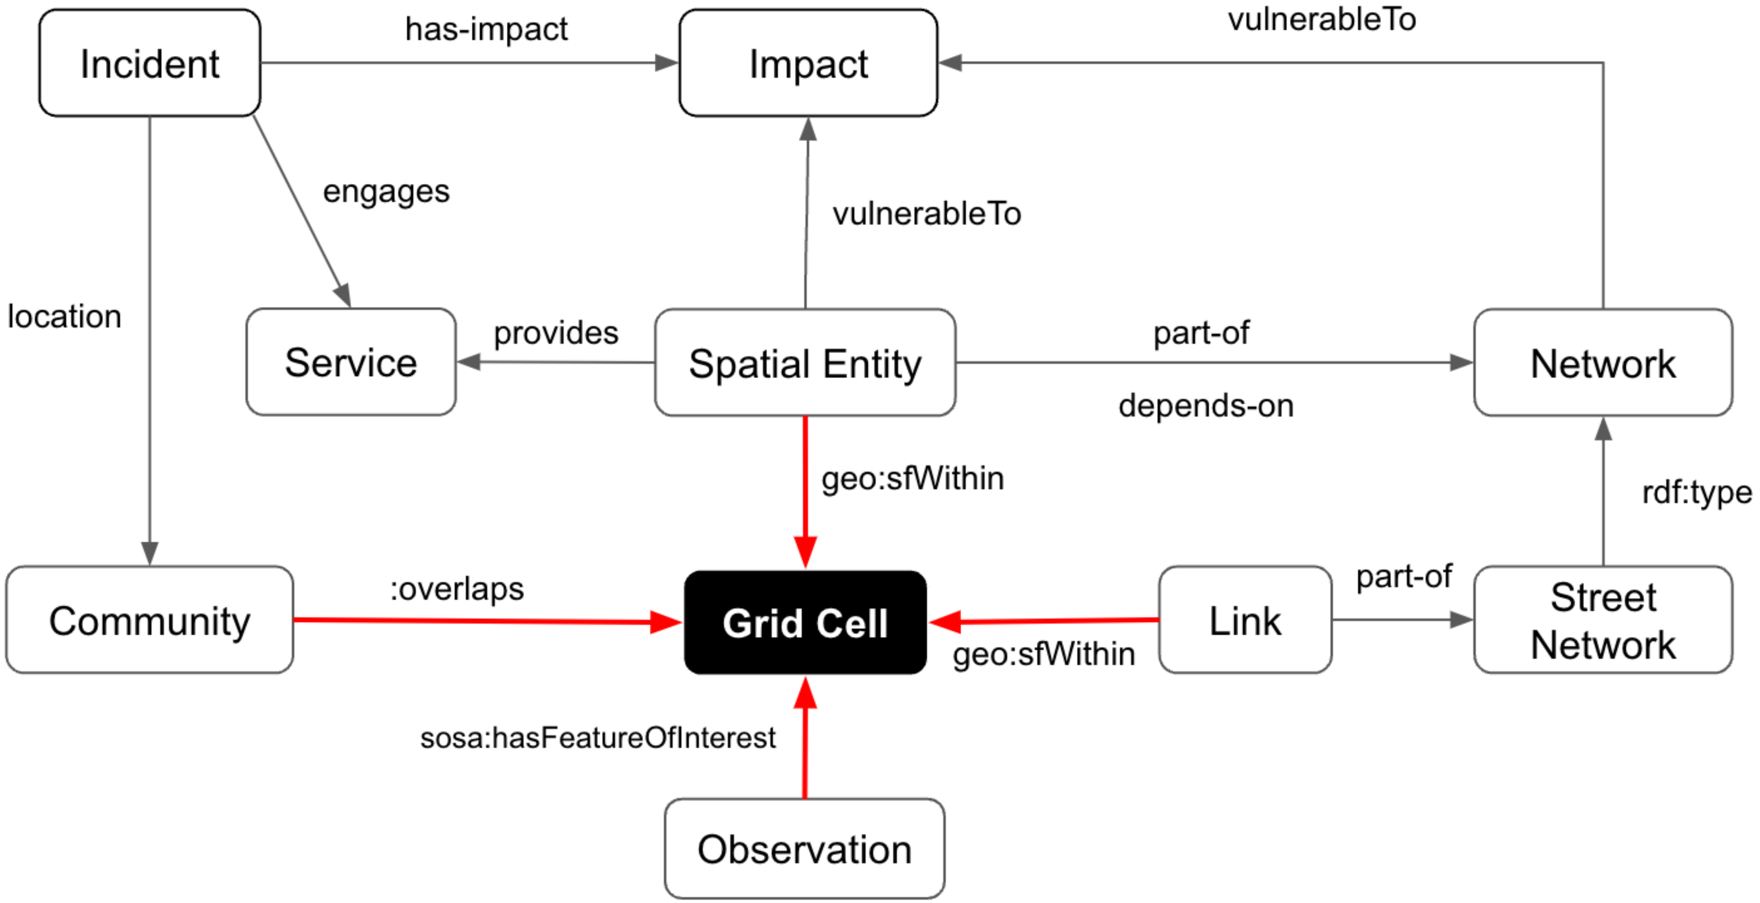

Crisis Management Use Case

Incident

Provide incident's basic information.

Incident type

Gewitter

Community code

40609

Potential Impacts

Incident Impacts

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

Required Services

Emergency Services

Commonly engaged emergency responders for selected impact(s)

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

L1.285.471 [48.43765 , 14.56269]

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.

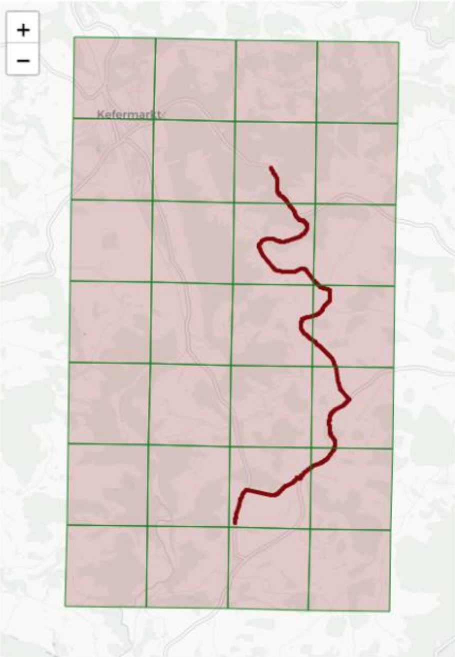

Future Work: Real-time Crisis KG Construction + AI Services

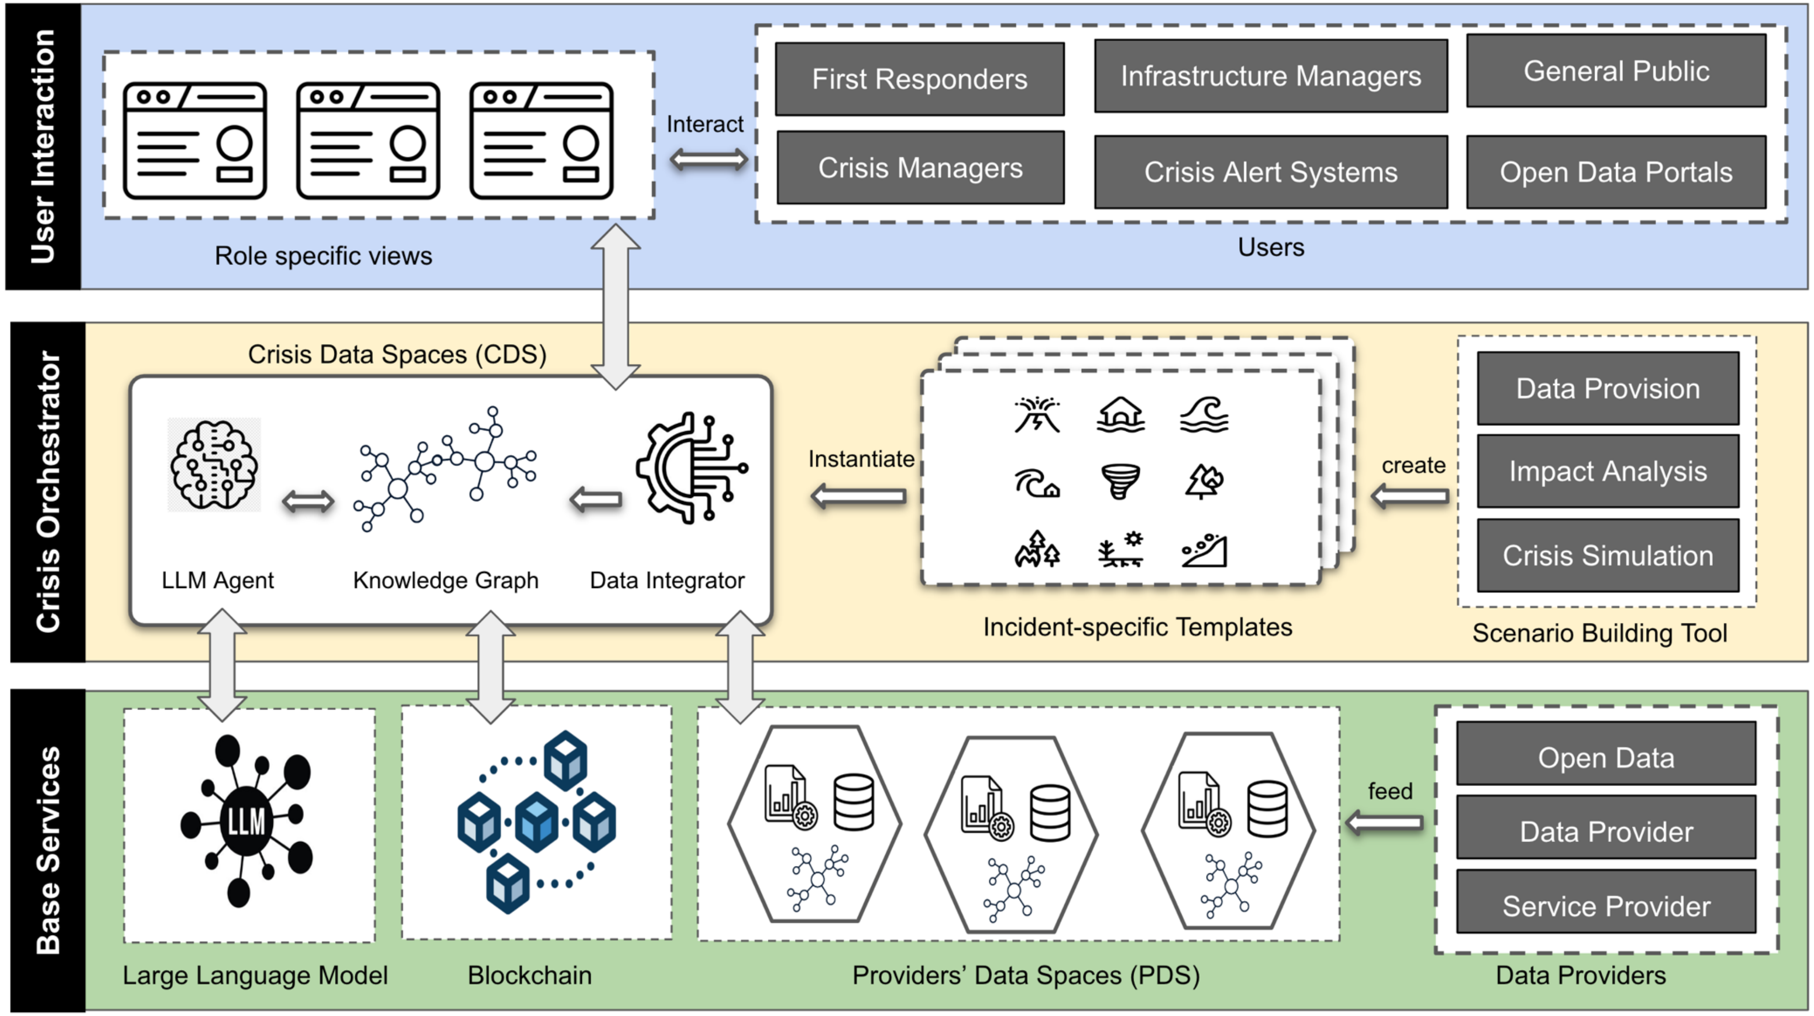

CRISP Portal & SPARQL Endpoint

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

SMWCloud:

Apart from Wikidata, there are many other Semantic Wikis

- Powered by MediaWiki software.
- We know of **60527** currently active wikis. (It's a lot ☺)

MediaWiki

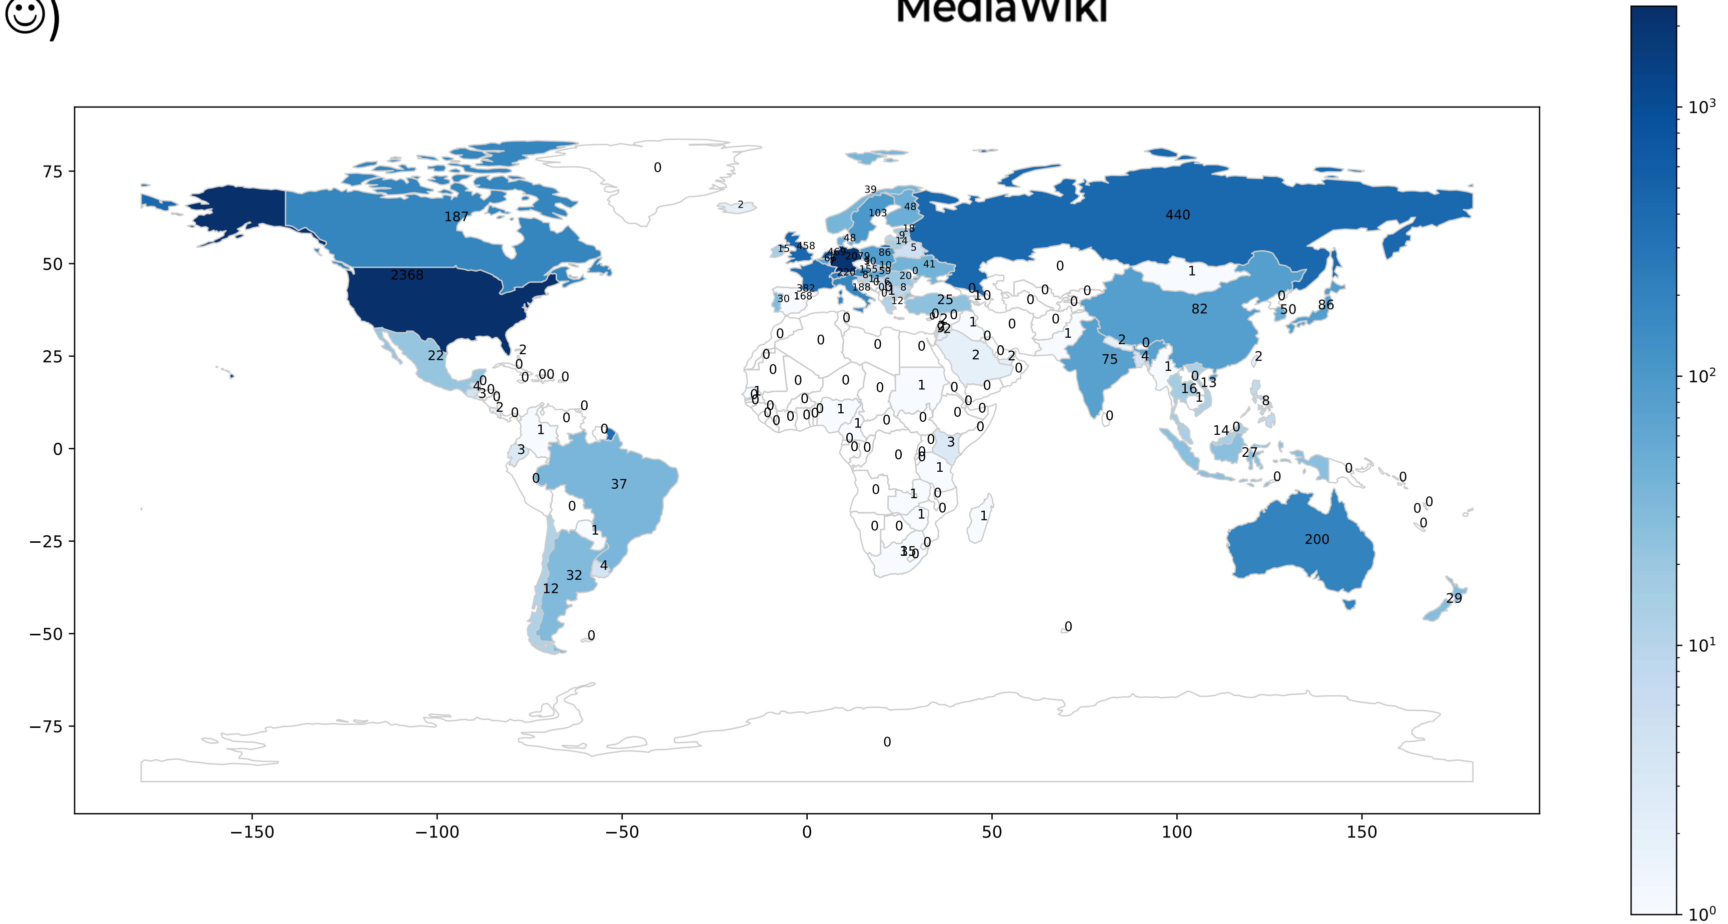

How many **Semantic** MediaWikis?
SMW Cloud (1458 wikis)

Semantic MediaWiki (SMW) is a free, open-source extension to MediaWiki – the wiki software that powers Wikipedia – that lets you store and query data within the wiki's pages.

Semantic MediaWiki is also a full-fledged framework, in conjunction with many spinoff extensions, that can turn a wiki into a powerful and flexible knowledge management system. All data created within Semantic MediaWiki can easily be exported or published via the Semantic Web, allowing other systems to use this data seamlessly.

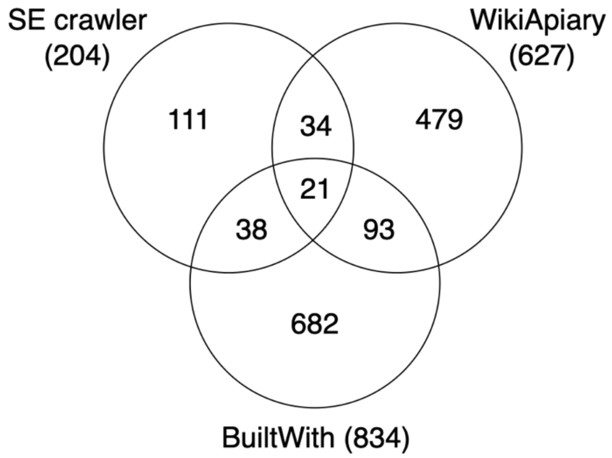

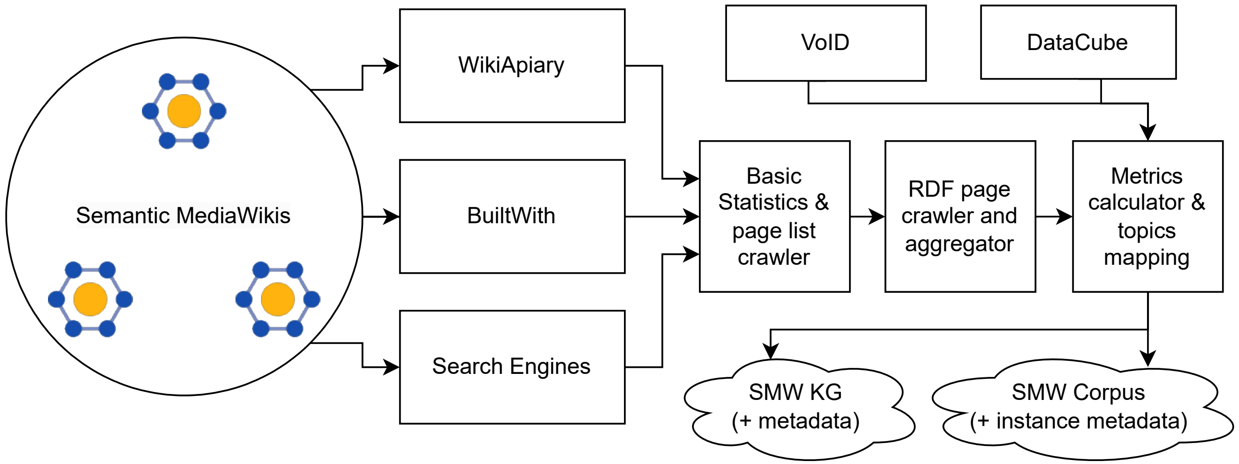

Dataset	#Triples	#Subjects	#Predicates	#Objects	#Literals
LODStats [10]	192,230,648	Not reported	49,916	Not reported	90,261,655
SMW Cloud	236,505,705	24,010,566	52,670	66,052,823	160,108,216
Wikidata 2021 [23]	17,662,800,665	1,625,057,179	38,867	Not reported	Not reported
LOD-a-lot [15]	28,362,198,927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

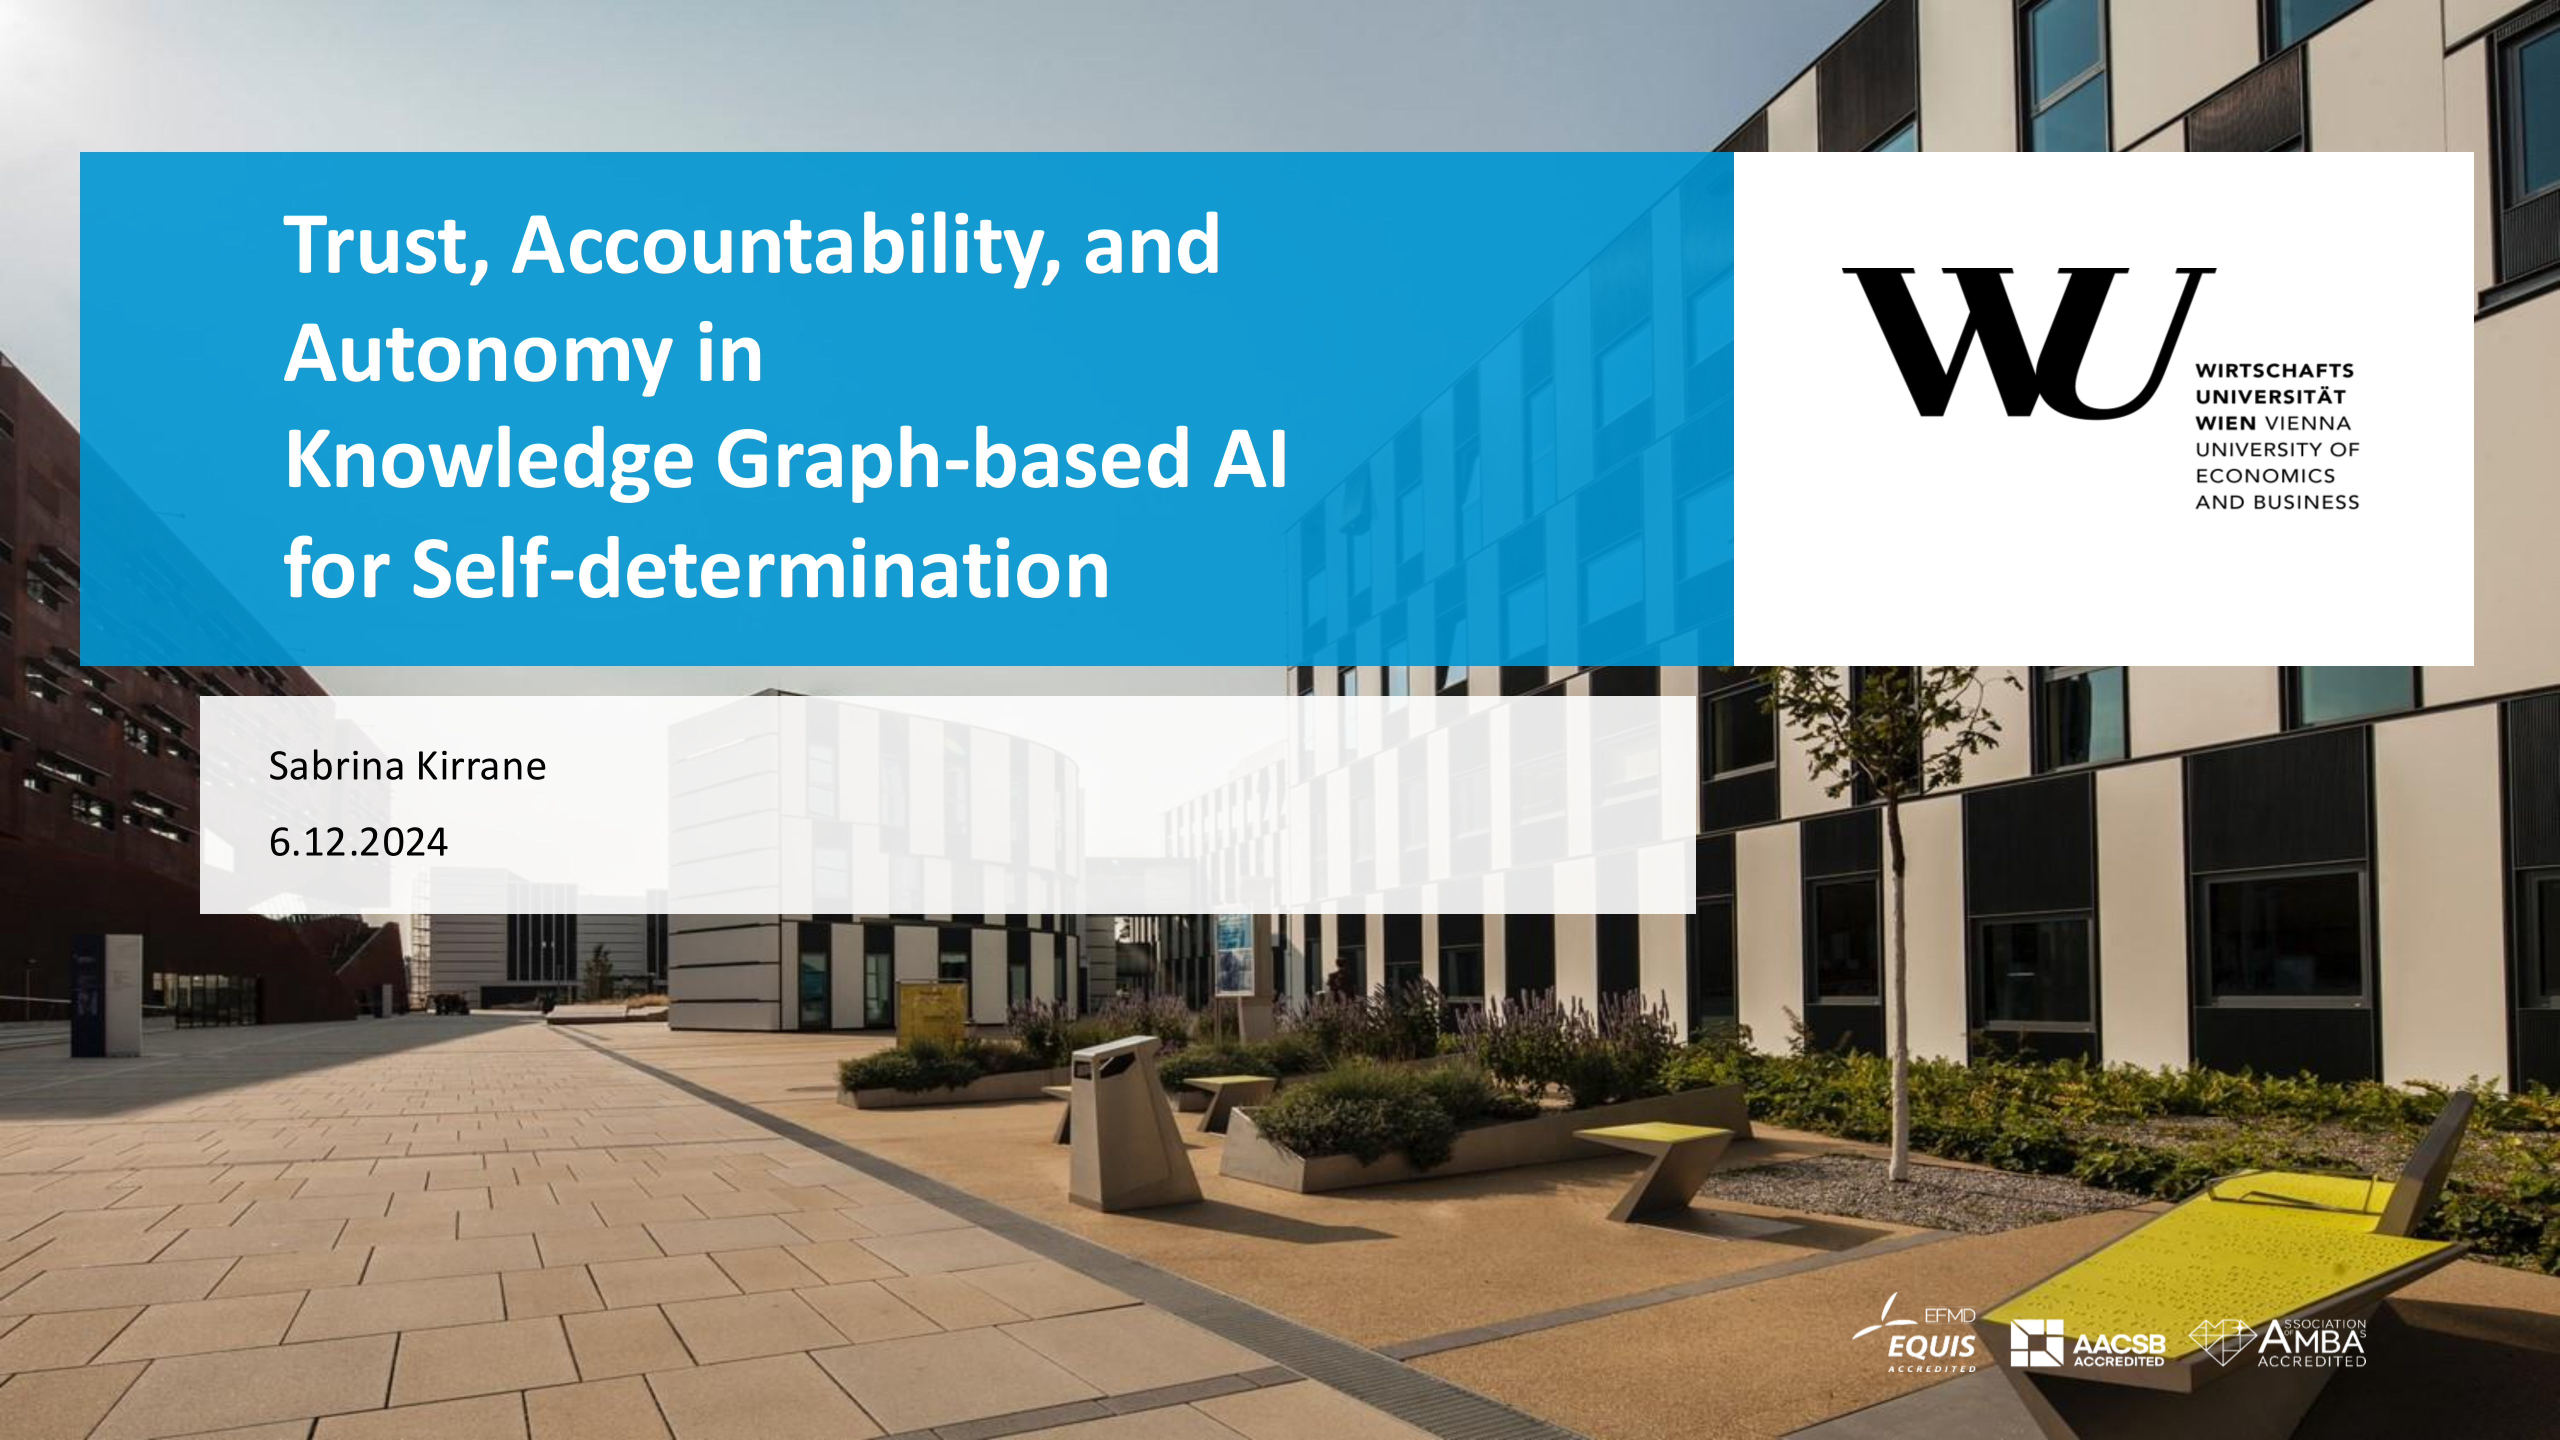

Trust, Accountability, and Autonomy in Knowledge Graph-based AI for Self-determination

WIRTSCHAFTS UNIVERSITÄT WIEN VIENNA UNIVERSITY OF ECONOMICS AND BUSINESS

Sabrina Kirrane

6.12.2024

KG-based AI for Self-Determination

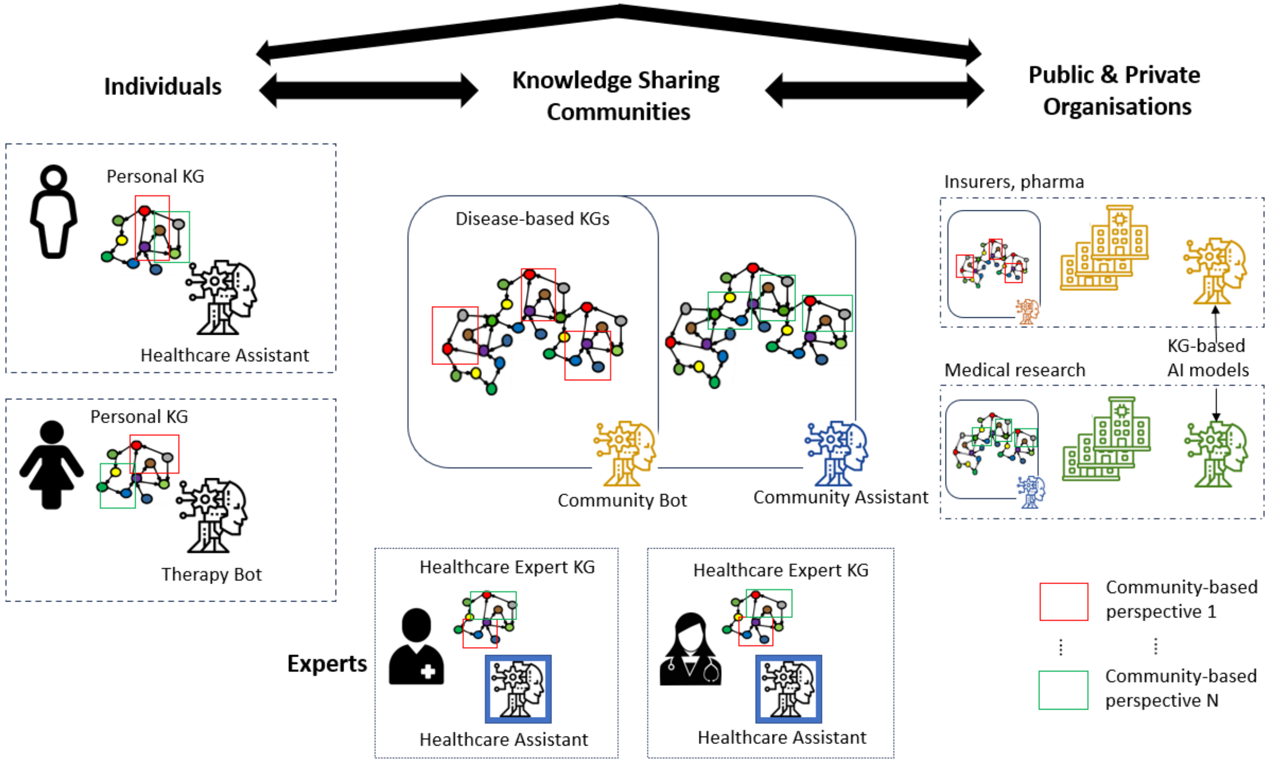

- Individuals use Artificial Intelligence (AI) assistants to make sense of data collected in their Personal Knowledge Graphs (PKGs).
- They may share perspectives of their PKGs with other individuals and healthcare experts in knowledge-sharing communities that aggregate and curate data to power AI services 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.

Ibáñez, L., Domingue, J., **Kirrane, S.**, Seneviratne, O., Third, A., Vidal, M., 2023. Trust, Accountability, and Autonomy in Knowledge Graph-based AI for Self-determination. Transactions on Graph Data and Knowledge (TGDK) *(revised and resubmitted)*

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

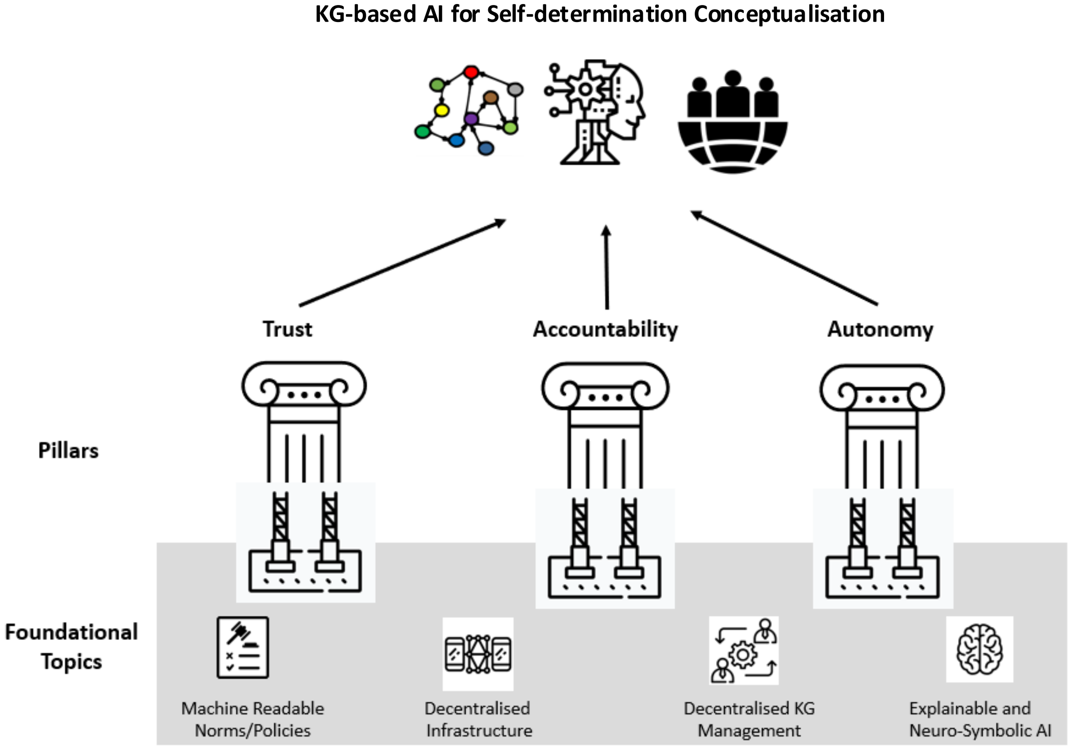

Ibáñez, L., Domingue, J., **Kirrane, S.**, Seneviratne, O., Third, A., Vidal, M., 2023. Trust, Accountability, and Autonomy in Knowledge Graph-based AI for Self-determination. Transactions on Graph Data and Knowledge (TGDK) *(revised and resubmitted)*

Machine-readable norms and policies

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.

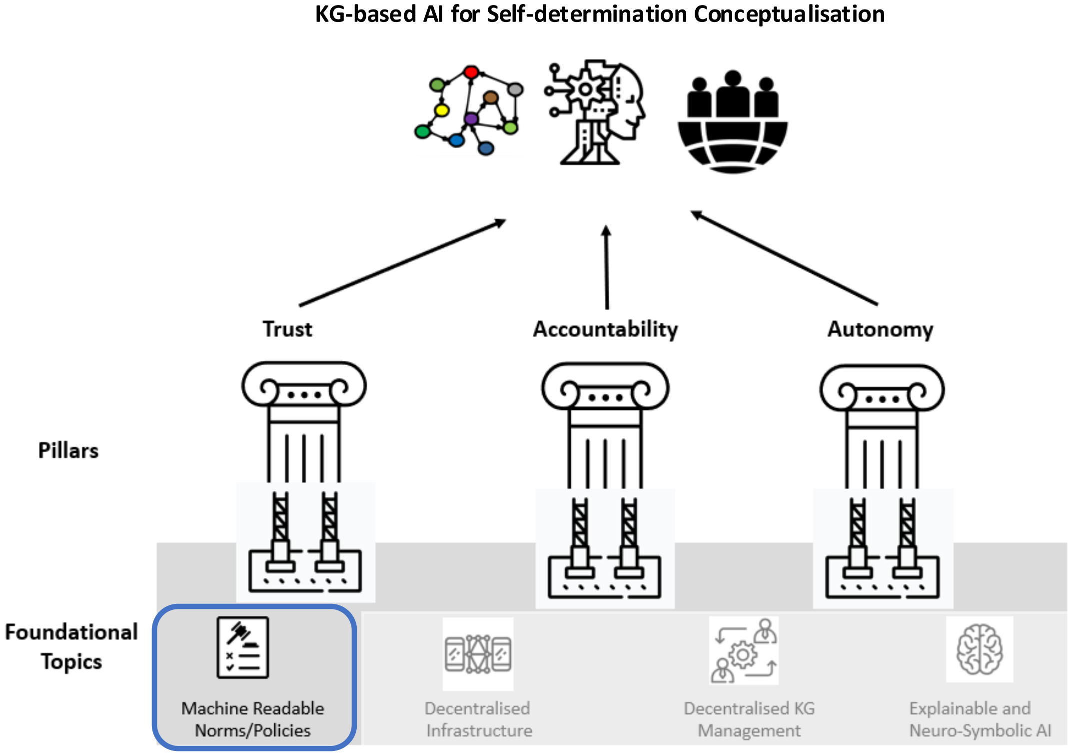

Decentralised infrastructure

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.

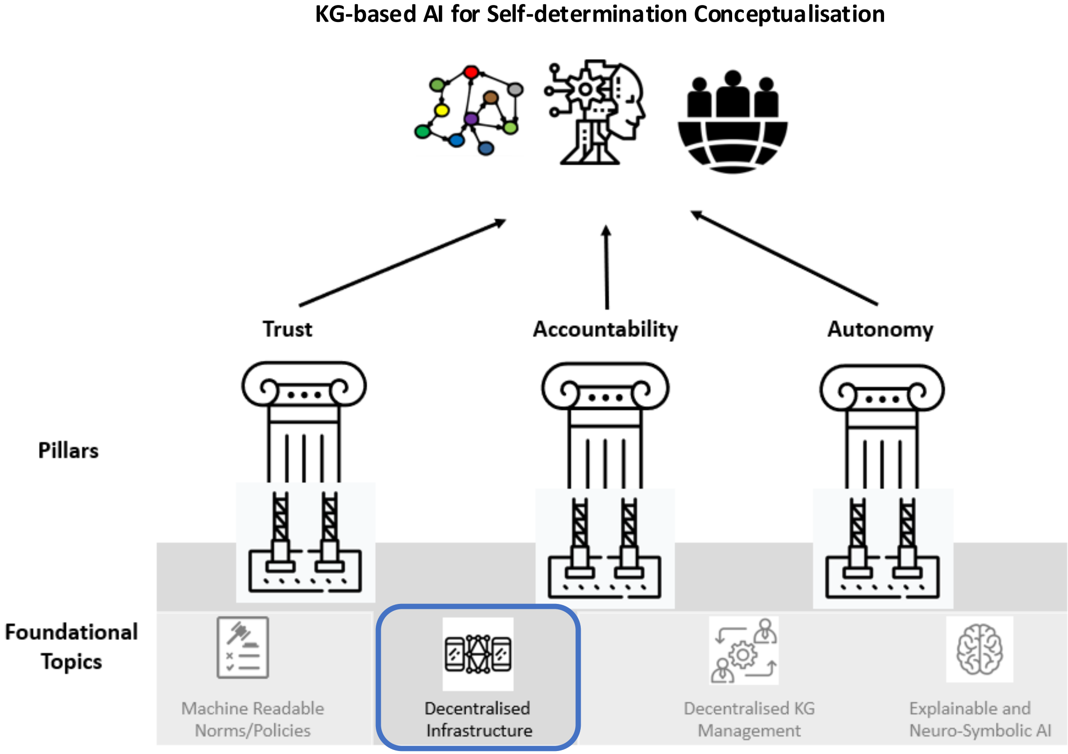

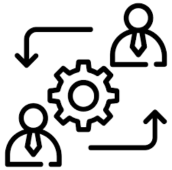

Decentralised KG management

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.

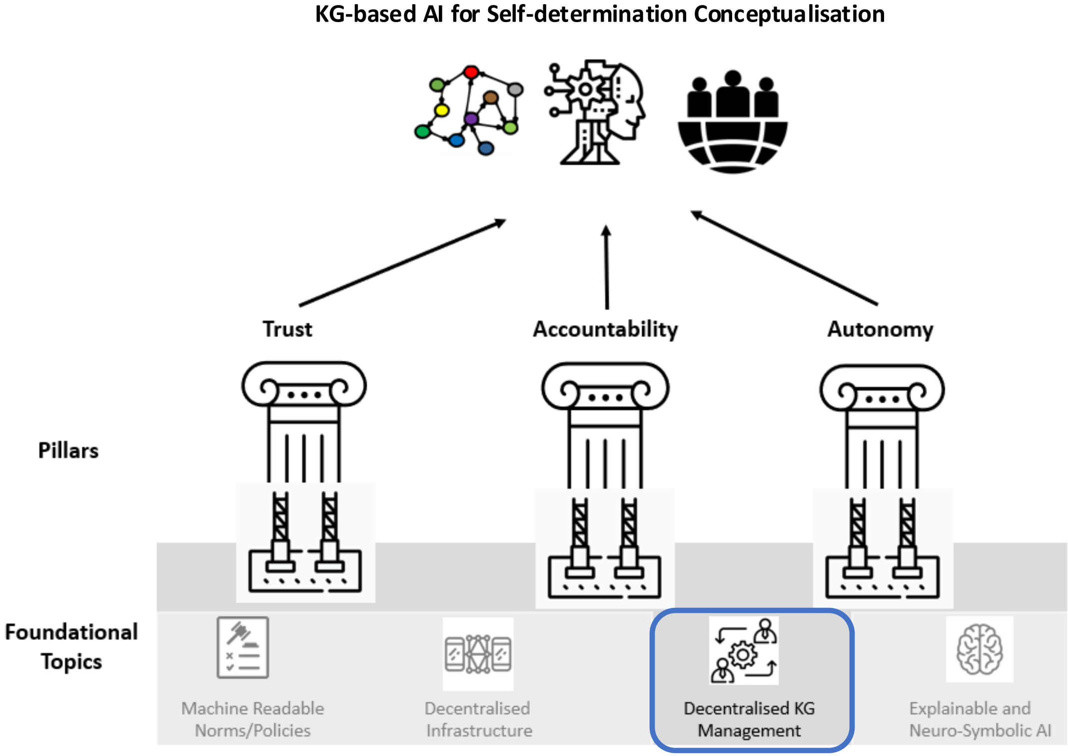

Towards
Explainable and Neuro-Symbolic AI

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.

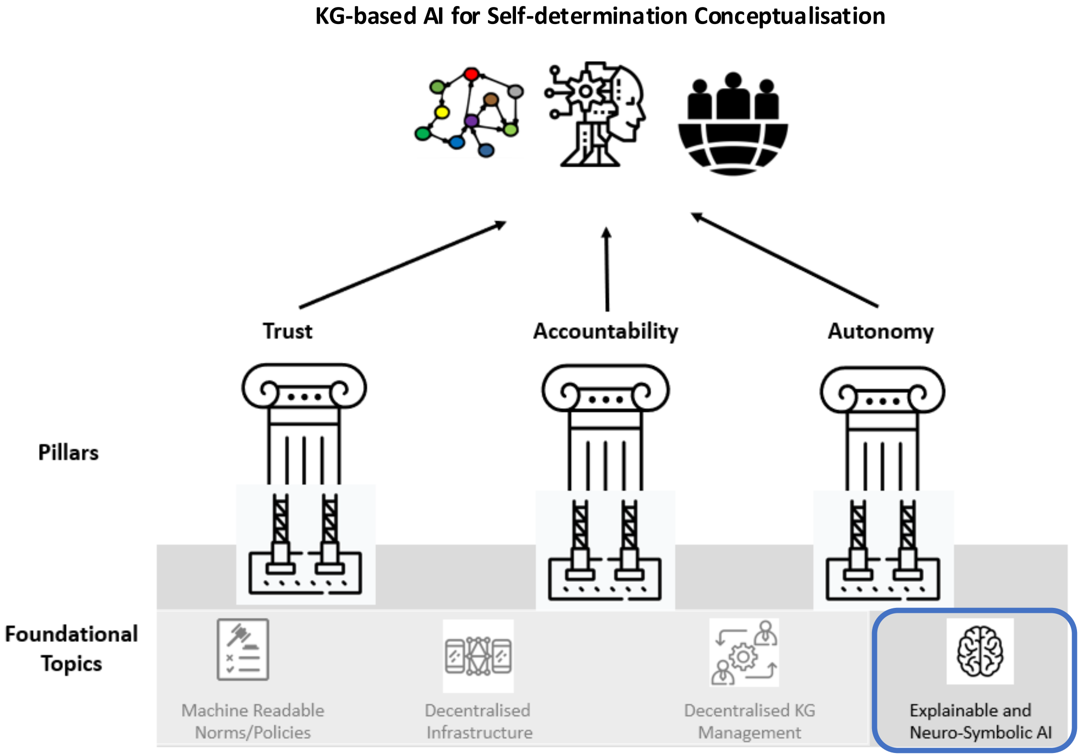

KG-based AI for Self-Determination

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