# Knowledge Graph-based AI for Self-determination

## AI4Industry Summer School 2025

**Trust, Interoperability, Autonomy and Resilience in Industry 4.0**

WIRTSCHAFTS UNIVERSITÄT WIEN VIENNA UNIVERSITY OF ECONOMICS AND BUSINESS

Sabrina Kirrane

22.07.2025

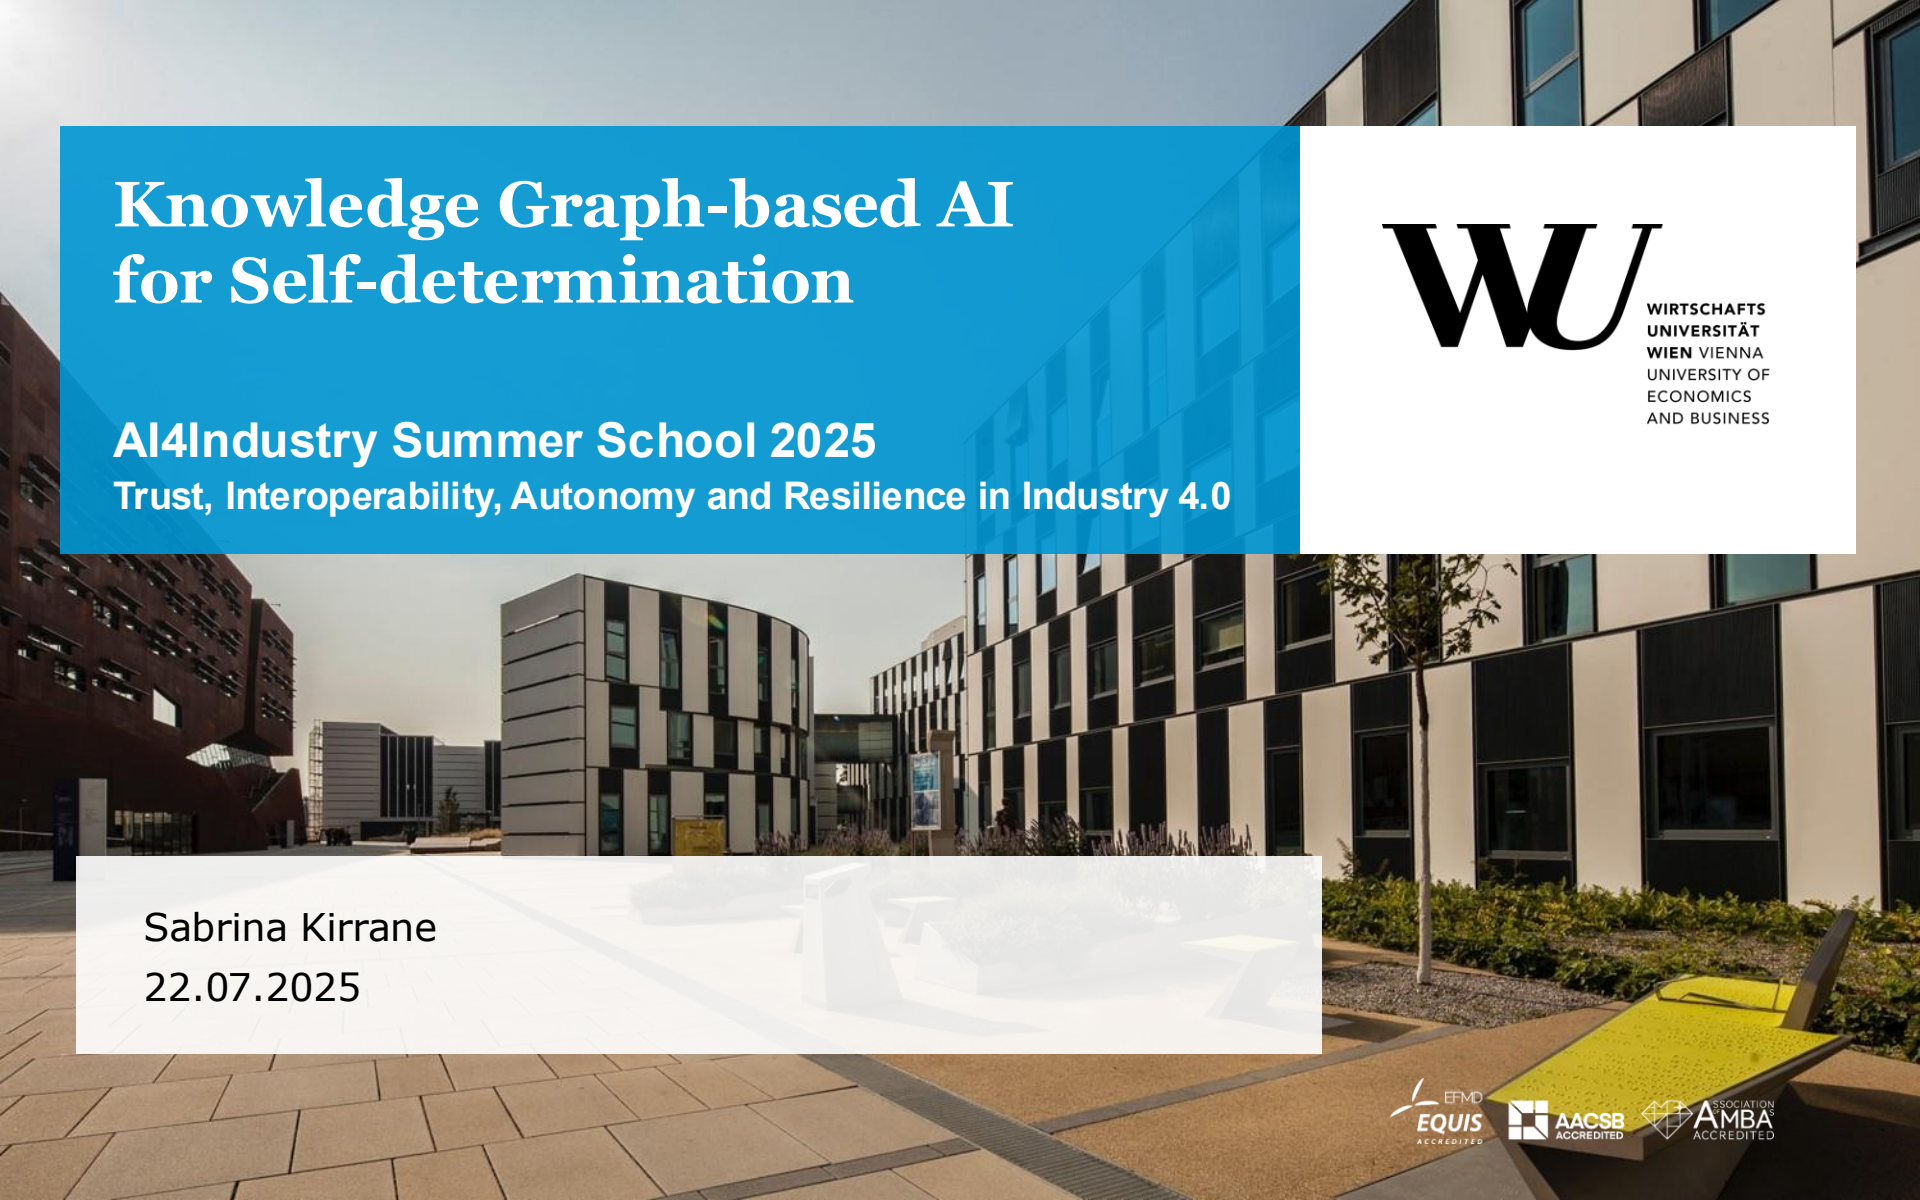

# Knowledge Graphs

A knowledge graph is a graph of data intended to accumulate and convey knowledge of the real world, whose **nodes** represent **entities** of interest and whose **edges** represent **relations** between these entities.

https://kgbook.org/

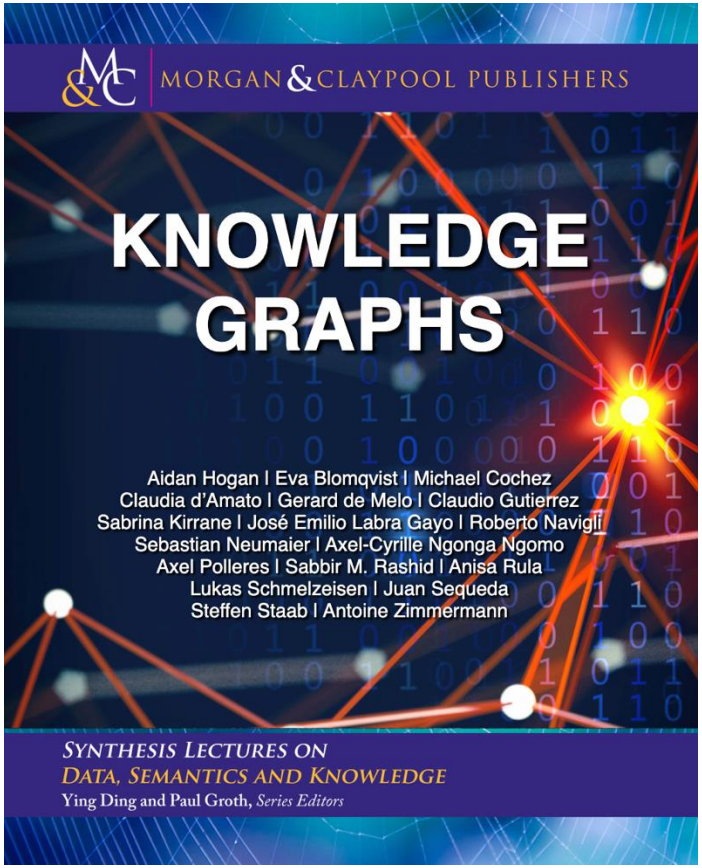




A Hogan, E. Blomqvist, M. Cochez, C. d'Amato, G. de Melo, C. Gutierrez, J. E. Labra Gayo, S. Kirrane, S. Neumaier, A. Polleres, R. Navigli, A. Ngonga Ngomo, S. M. Rashid, A. Rula, L. Schmelzeisen, J. Sequeda, S. Staab, A. Zimmermann. Knowledge Grpahs.

# Directed Labelled Graphs

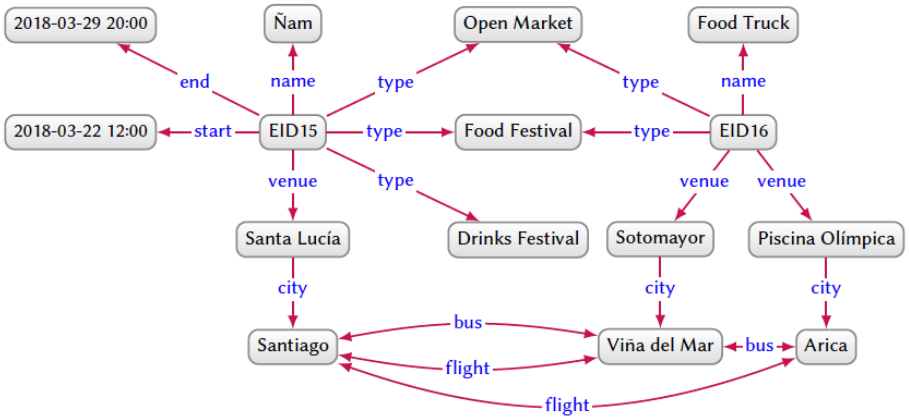


- A directed edge-labelled graph is defined as a set of nodes and a set of directed labelled edges between those nodes
- In the case of knowledge graphs, nodes are used to represent entities and edges are used to represent (binary) relations between those entities.
- Modelling data as a graph in this way offers more flexibility for integrating new sources of data, compared to the standard relational model, where a schema must be defined upfront and followed at each step.



A Hogan, E. Blomqvist, M. Cochez, C. d'Amato, G. de Melo, C. Gutierrez, J. E. Labra Gayo, S. Kirrane, S. Neumaier, A. Polleres, R. Navigli, A. Ngonga Ngomo, S. M. Rashid, A. Rula, L. Schmelzeisen, J. Sequeda, S. Staab, A. Zimmermann. Knowledge Grpahs.

# Graph Datasets

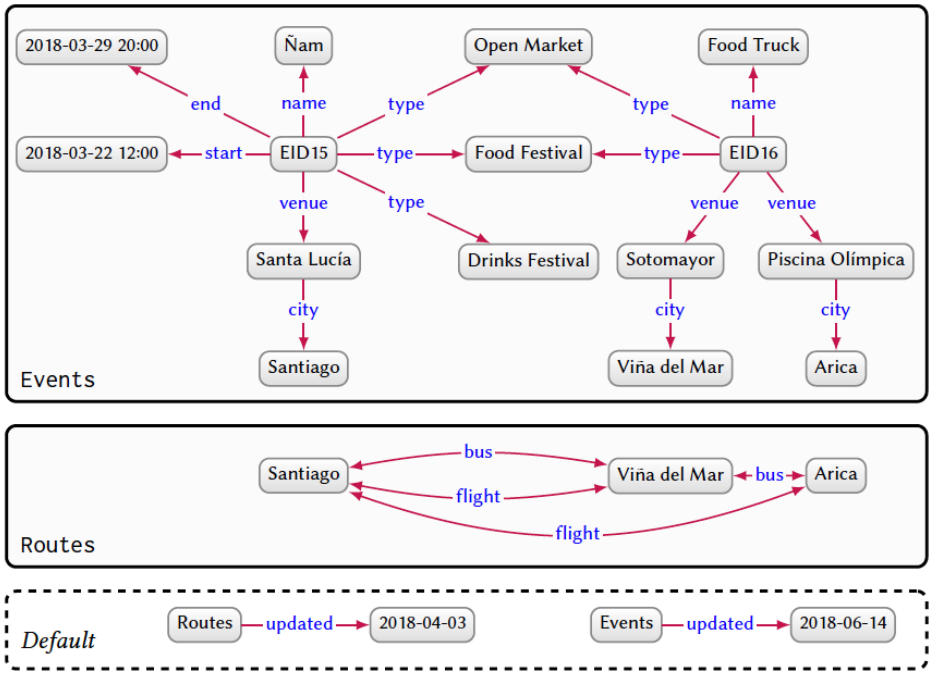


- It is often desirable to manage several graphs rather than one monolithic graph
- A graph dataset consists of a set of named graphs and a default graph. Each named graph is a graph ID and graph pair
- The default graph is a graph without an ID and is referenced "by default" if a graph ID is not specified.



A Hogan, E. Blomqvist, M. Cochez, C. d'Amato, G. de Melo, C. Gutierrez, J. E. Labra Gayo, S. Kirrane, S. Neumaier, A. Polleres, R. Navigli, A. Ngonga Ngomo, S. M. Rashid, A. Rula, L. Schmelzeisen, J. Sequeda, S. Staab, A. Zimmermann. Knowledge Grpahs.

# Graph Querying

- At the core of every structured query language for graphs are (basic) graph patterns, which follow the same model as the data graph being queried, additionally allowing variables as terms
- A graph pattern is then evaluated against the data graph by generating mappings from the variables of the graph pattern to constants in the data graph
- Homomorphism-based semantics allows multiple variables to be mapped to the same term such that all mappings shown would be considered results

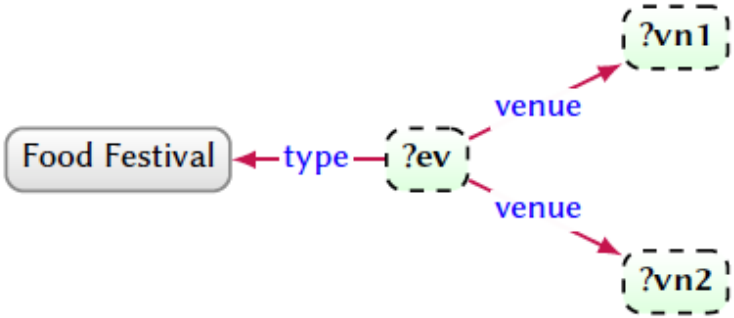


| ?ev | ?vn1 | ?vn2 |
|---|---|---|
| EID16 | Piscina Olímpica | Sotomayor |
| EID16 | Sotomayor | Piscina Olímpica |
| EID16 | Piscina Olímpica | Piscina Olímpica |
| EID16 | Sotomayor | Sotomayor |
| EID15 | Santa Lucía | Santa Lucía |



A Hogan, E. Blomqvist, M. Cochez, C. d'Amato, G. de Melo, C. Gutierrez, J. E. Labra Gayo, S. Kirrane, S. Neumaier, A. Polleres, R. Navigli, A. Ngonga Ngomo, S. M. Rashid, A. Rula, L. Schmelzeisen, J. Sequeda, S. Staab, A. Zimmermann. Knowledge Grpahs.

# Graph Querying

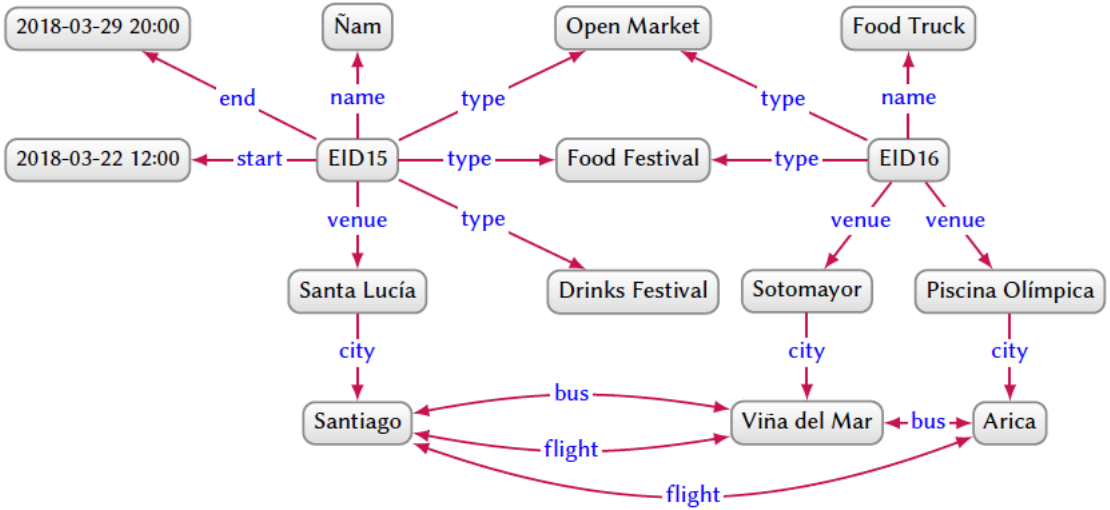


Fig. 1. Directed edge-labelled graph describing events and their venues.

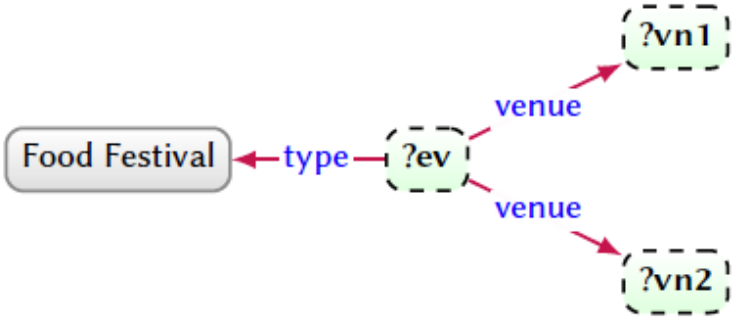


| ?ev | ?vn1 | ?vn2 |
|---|---|---|
| EID16 | Piscina Olímpica | Sotomayor |
| EID16 | Sotomayor | Piscina Olímpica |
| EID16 | Piscina Olímpica | Piscina Olímpica |
| EID16 | Sotomayor | Sotomayor |
| EID15 | Santa Lucía | Santa Lucía |



A Hogan, E. Blomqvist, M. Cochez, C. d'Amato, G. de Melo, C. Gutierrez, J. E. Labra Gayo, S. Kirrane, S. Neumaier, A. Polleres, R. Navigli, A. Ngonga Ngomo, S. M. Rashid, A. Rula, L. Schmelzeisen, J. Sequeda, S. Staab, A. Zimmermann. Knowledge Grpahs.

# Schema Definition & Reasoning

| Feature | Definition | Condition | Example |
|---|---|---|---|
| SUBCLASS | $c$ — subc. of → $d$ | $x$ — type → $c$ implies $x$ — type → $d$ | City — subc. of → Place |
| SUBPROPERTY | $p$ — subp. of → $q$ | $x$ — $p$ → $y$ implies $x$ — $q$ → $y$ | venue — subp. of → location |
| DOMAIN | $p$ — domain → $c$ | $x$ — $p$ → $y$ implies $x$ — type → $c$ | venue — domain → Event |
| RANGE | $p$ — range → $c$ | $x$ — $p$ → $y$ implies $y$ — type → $c$ | venue — range → Venue |

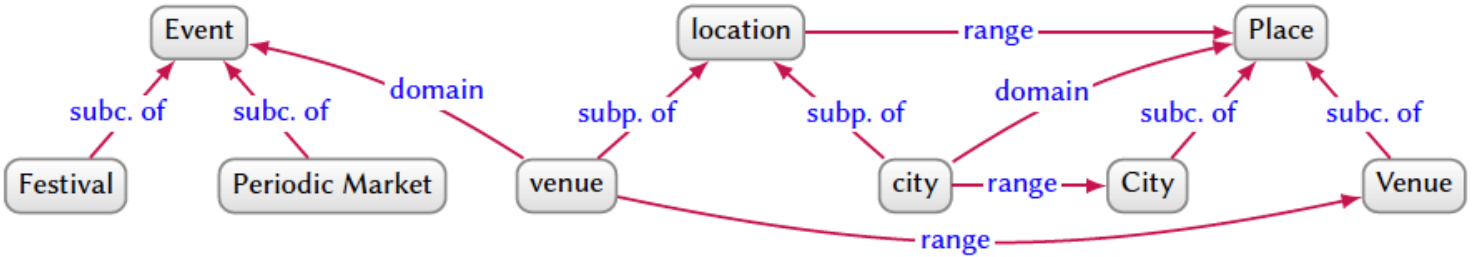




A Hogan, E. Blomqvist, M. Cochez, C. d'Amato, G. de Melo, C. Gutierrez, J. E. Labra Gayo, S. Kirrane, S. Neumaier, A. Polleres, R. Navigli, A. Ngonga Ngomo, S. M. Rashid, A. Rula, L. Schmelzeisen, J. Sequeda, S. Staab, A. Zimmermann. Knowledge Grpahs.

# Web Standards

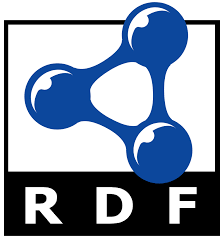


Data model for encoding triples

**Disclaimer:** This is not a complete list!


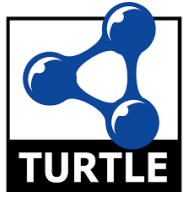




Textual syntax for RDF


Query language


Data modelling vocabularies

# Vocabularies

FOAF Vocabulary Specification 0.99

Namespace Document 14 January 2014 - *Paddington Edition*

**This version:**
http://xmlns.com/foaf/spec/20140114.html (rdf)
**Latest version:**
http://xmlns.com/foaf/spec/ (rdf)
**Previous version:**
http://xmlns.com/foaf/spec/20100809.html (rdf)
**Authors:**
Dan Brickley, Libby Miller
**Contributors:**
Members of the FOAF mailing list (foaf-dev@lists.foaf-project.org) and the wider RDF and Semantic Web developer community. See acknowledgements.


Regarding underlying technology, FOAF uses W3C's RDF technology, an open Web standard that can be freely used by anyone.

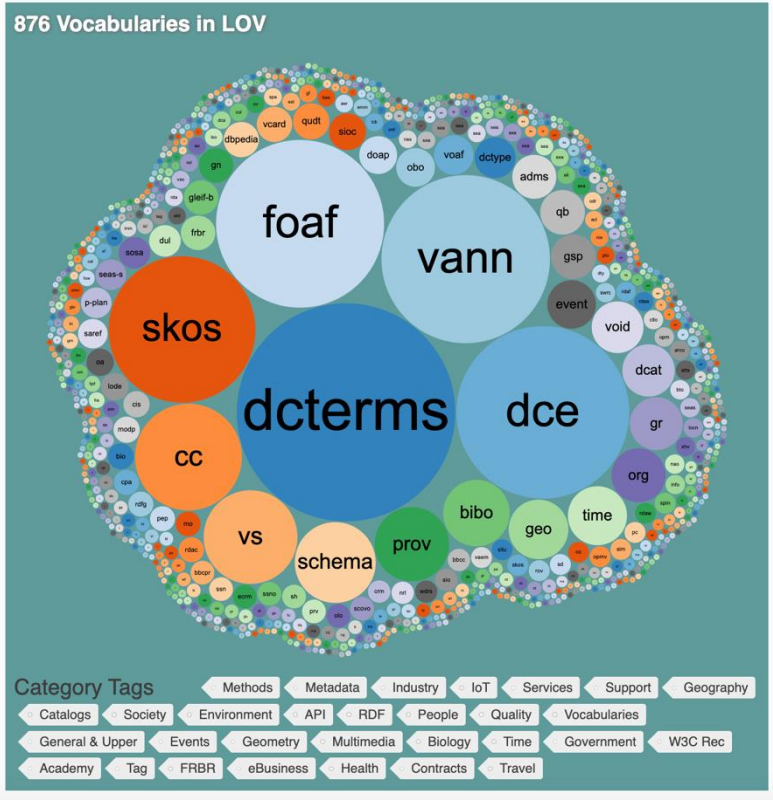


W3C

SKOS Simple Knowledge Organization System Reference

W3C Recommendation 18 August 2009

**This version:**
http://www.w3.org/TR/2009/REC-skos-reference-20090818/
**Latest version:**
http://www.w3.org/TR/skos-reference
**Previous versions:**
http://www.w3.org/TR/2009/PR-skos-reference-20090615/
**Editors:**
Alistair Miles, STFC Rutherford Appleton Laboratory / University of Oxford
Sean Bechhofer, University of Manchester

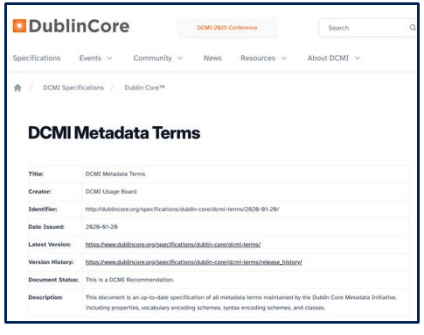


W3C

PROV-O: The PROV Ontology

W3C Recommendation 30 April 2013

**This version:**
http://www.w3.org/TR/2013/REC-prov-o-20130430/
**Latest published version:**
http://www.w3.org/TR/prov-o/
**Implementation report:**
http://www.w3.org/TR/2013/NOTE-prov-implementations-20130430/
**Previous version:**
http://www.w3.org/TR/2013/PR-prov-o-20130312/
**Editors:**
Timothy Lebo, Rensselaer Polytechnic Institute, USA
Satya Sahoo, Case Western Reserve University, USA
Deborah McGuinness, Rensselaer Polytechnic Institute, USA
**Contributors:**
(in alphabetical order)
Khalid Belhajjame, University of Manchester, UK
James Cheney, University of Edinburgh, UK
David Corsar, University of Aberdeen, UK
Daniel Garijo, Ontology Engineering Group, Universidad Politécnica de Madrid, Spain
Stian Soiland-Reyes, University of Manchester, UK
Stephan Zednik, Rensselaer Polytechnic Institute, USA
Jun Zhao, University of Oxford, UK

https://lov.linkeddata.es/dataset/lov/

# The Semantic Web

WU WIRTSCHAFTS UNIVERSITÄT WIEN VIENNA UNIVERSITY OF ECONOMICS AND BUSINESS

Sabrina Kirrane

Sabrina Kirrane
Researcher
Vienna, Austria
sabrinakirrane@gmail.com

**Academic Profile**

ORCID
SCOPUS
Google Scholar
DBLP
Turtle
LDF Server

```
ttp://sabrinakirrane.com/> <http://xmlns.com/foaf/0.1/primaryTopic> <http://sabrinakirrane.com/#me> ,
ttp://sabrinakirrane.com/#me> ;
<http://xmlns.com/foaf/0.1/maker> <http://sabrinakirrane.com/#me> ;
<http://ogp.me/ns#image> "https://www.sabrinakirrane.com/images/sabrinakirrane.jpg"@en ;
<http://ogp.me/ns#title> "Sabrina Kirrane"@en ;
<http://ogp.me/ns#description> "Computer scientist, researcher, programmer\n"@en ;
<http://ogp.me/ns#url> "https://www.sabrinakirrane.com/"@en ;
<http://ogp.me/ns#locale> "en_UK"@en ;
<http://ogp.me/ns#site_name> "Sabrina Kirrane"@en ;
<http://ogp.me/ns#type> "website"@en .

ttp://sabrinakirrane.com/#me> a <http://
<http://xmlns.com/foaf/0.1/name> "Sabrir
<http://www.w3.org/2000/01/rdf-schema#la
<http://xmlns.com/foaf/0.1/givenName> "S
<http://xmlns.com/foaf/0.1/familyName> "
<http://www.w3.org/2006/vcard/ns#title>
<http://xmlns.com/foaf/0.1/logo> <http:/
<http://xmlns.com/foaf/0.1/img> <http://
<http://xmlns.com/foaf/0.1/homepage> <ht
<http://xmlns.com/foaf/0.1/workplaceHome
<http://xmlns.com/foaf/0.1/mbox> <mailto
<http://xmlns.com/foaf/0.1/topic_interes

ttps://scholar.google.com/citations?hl=e
ttp://sabrinakirrane.com/#me> <http://xm
ttps://scholar.google.com/citations?hl=e
  <http://xmlns.com/foaf/0.1/accountName

ttp://orcid.org/0000-0002-6955-7718> a <
ttp://sabrinakirrane.com/#me> <http://xm
ttp://orcid.org/0000-0002-6955-7718> <ht
  <http://xmlns.com/foaf/0.1/accountName

ttps://www.scopus.com/authid/detail.uri?
ttp://sabrinakirrane.com/#me> <http://xm
ttps://www.scopus.com/authid/detail.uri?
  <http://xmlns.com/foaf/0.1/accountName
```

## Sabrina's Linked Data Fragments server

#LD Linked Data Fragments

### My Turtle file

**Search My Turtle file by triple/quad pattern**

**subject:**
**predicate:**
**object:**
**graph:**

Search

**Matches in My Turtle file for *{ ?s ?p ?o ?g. }***

Showing items 1 to 100 of ±5,981 with 100 items per page. **next**

```
sabrinakirrane.com   primaryTopic   me   @default.
sabrinakirrane.com   maker   me   @default.
sabrinakirrane.com   image   "https://www.sabrinakirrane.com/images/sabrinakirrane.jpg"   @de…
sabrinakirrane.com   title   "Sabrina Kirrane"   @default.
sabrinakirrane.com   description   "Computer scientist, researcher, programmer "   @default.
sabrinakirrane.com   url   "https://www.sabrinakirrane.com/"   @default.
sabrinakirrane.com   locale   "en_UK"   @default.
sabrinakirrane.com   site_name   "Sabrina Kirrane"   @default.
sabrinakirrane.com   type   "website"   @default.
```

http://xmlns.com/foaf/0.1/primaryTopic

http://sabrinakirrane.com/

# The Semantic Web

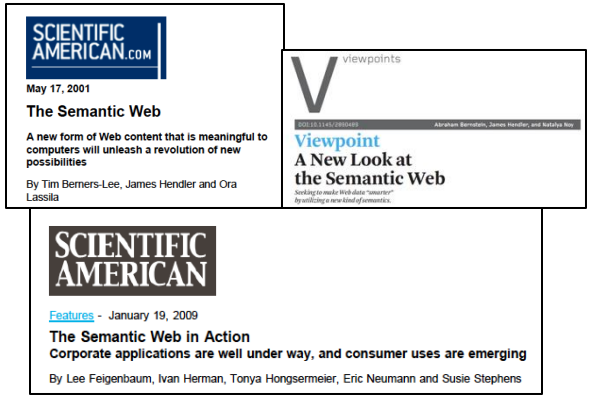


Knowledge structures and modeling

social semantic web

knowledge representation languages and standards

logic and reasoning

linked data

search, retrieval, ranking, question answering

matching and data integration

visualization, user interfaces and annotation

semantic web services

query languages and mechanisms

knowledge extraction, discovery and acquisition

streaming & sensor data

Rexplore

Saffron
Knowledge Extraction Framework

poolparty

distribution, decentralization, federation

semantic web databases

intelligent software agents

multilingual intelligent agents

data quality

change management and propagation

privacy, trust, security, provenance

scalability, efficiency, robustness



***A decade of Semantic Web research through the lenses of a mixed methods approach**, Sabrina Kirrane, Marta Sabou, Javier D. Fernández, Francesco Osborne, Cécile Robin, Paul Buitelaar, Enrico Motta, and Axel Polleres, The Semantic Web Journal, 2020*

# The Original Semantic Web Vision

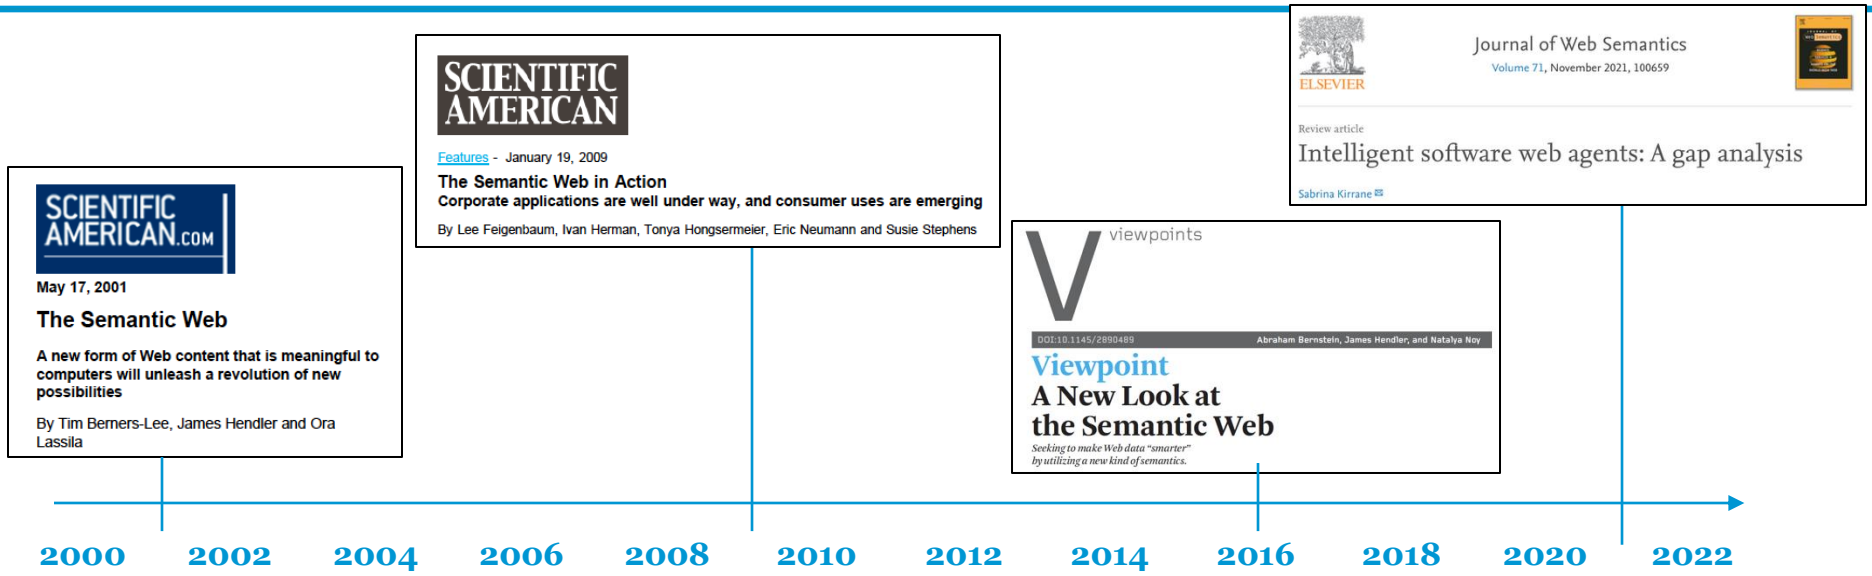


- Develop a **reference architecture** that could serve to bridge the gap between theory and practice
- Better understand the specific requirements relating to the cross cutting **behavioural functions** (i.e., benevolence, rationality, and mobility), **code of conduct functions** (i.e., identification, security, privacy, trust, and ethics), and **basic functions** (i.e., autonomy, and social ability)
- Develop/extend existing **benchmarks** for assessing the performance and scalability

# The Semantic Web



## Summary

The Semantic Web is about AI: KR + agents

People thought it was about something else

The vision has now matured into modern knowledge graphs (also: I declare victory)

We still do not have agents, but can build them, and LLMs will help

People (still) trying to reinvent RDF (poorly)

## Semantic Web and AI: Can we finally realize the vision?

Dr. Ora Lassila
Principal Technologist, Amazon Neptune
Co-chair, W3C RDF-star WG

Keynote Address
International Semantic Web Conference, November 2024

aws

https://videolectures.net/videos/iswc2024_lassila_web_and_ai

# Knowledge Graph-based AI for Self-determination


## Special Issue

Special Issue on Trends in Graph Data and Knowledge

## Editors


Aidan Hogan
› DCC, Universidad de Chile, IMFD, Chile

Ian Horrocks
› University of Oxford, U.K.

Andreas Hotho
› Department of Informatics, University of Würzburg, Germany

Lalana Kagal
› Massachusetts Institute of Technology, Cambridge, MA, USA


## Trust, Accountability, and Autonomy in Knowledge Graph-Based AI for Self-Determination


**Luis-Daniel Ibáñez**
Department of Electronics and Computer Science, University of Southampton, UK

**John Domingue**
Knowledge Media Institute, The Open University, Milton Keynes, UK

**Sabrina Kirrane**
Institute for Information Systems & New Media, Vienna University of Economics and Business, Austria

**Oshani Seneviratne**
Department of Computer Science, Rensselaer Polytechnic Institute, USA

**Aisling Third**
Knowledge Media Institute, The Open University, Milton Keynes, UK

**Maria-Esther Vidal**
Leibniz University of Hannover, Germany
TIB-Leibniz Information Centre of Science and Technology, Hannover, Germany
L3S Research Centre, Hannover, Germany



**Abstract**

Knowledge Graphs (KGs) have emerged as fundamental platforms for powering intelligent decision-making and a wide range of Artificial Intelligence (AI) services across major corporations such as Google, Walmart, and AirBnb. KGs complement Machine Learning (ML) algorithms by providing data context and semantics, thereby enabling further inference and question-answering capabilities. The integration of KGs with neuronal learning (e.g., Large Language Models (LLMs)) is currently a topic of active research, commonly named neuro-symbolic AI. Despite the numerous benefits that can be accomplished with KG-based AI, its growing ubiquity within online services may result in the loss of self-determination for citizens as a fundamental societal issue. The more we rely on these technologies, which are often centralised, the less citizens will be able to determine their own destinies. To counter this threat, AI regulation, such as the European Union (EU) AI Act, is being proposed in certain regions. The regulation sets what technologists need to do, leading to questions concerning How the output of AI systems can be trusted? What is needed to ensure that the data fuelling and the inner workings of these artefacts are transparent? How can AI be made accountable for its decision-making? This paper conceptualises the foundational topics and research pillars to support KG-based AI for self-determination. Drawing upon this conceptual framework, challenges and opportunities for citizen self-determination are illustrated and analysed in a real-world scenario. As a result, we propose a research agenda aimed at accomplishing the recommended objectives.




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)

# Knowledge Graphs (KGs)

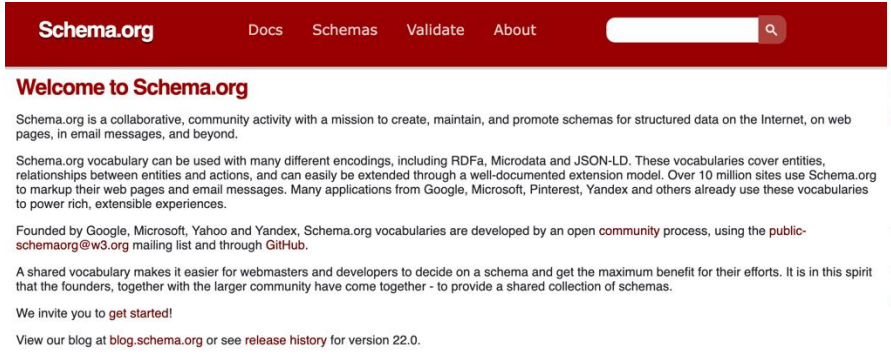



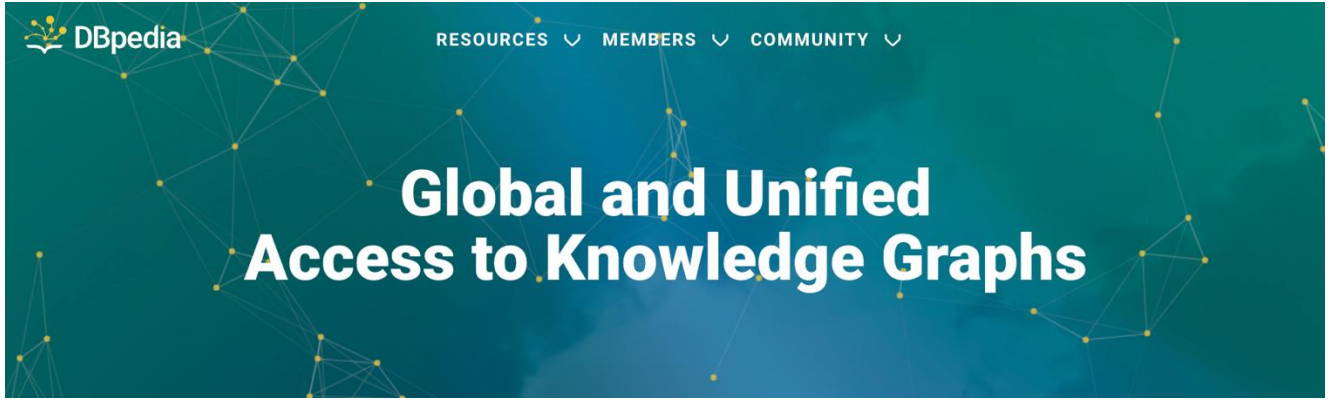

# Artificial Intelligence (AI)

## ChatGPT vs. Bing Chat vs. Google Bard: Which is the best AI chatbot?


Maria Diaz/ZDNET

https://www.zdnet.com/article/chatgpt-vs-bing-chat-vs-google-bard-which-is-the-best-ai-chatbot/

*"Artificial intelligence (AI) has transformed how we work and play in recent months, giving almost anyone the ability to write code, create art, and even make investments."*

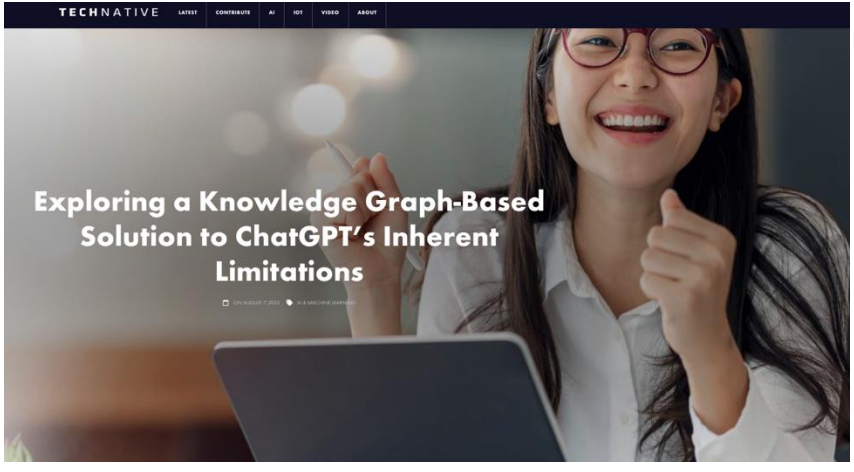


https://technative.io/exploring-a-knowledge-graph-based-solution-to-chatgpts-inherent-limitations/

*It can complete writing tasks and some coding challenges, but in both cases human expertise is still required since its output is not always precise. Its expertise is general, and it lacks deep knowledge in your domain."*

# Self-determination

*How can we ensure that individuals are aware of who knows what about them and can influence data processing that concerns them (i.e.self-determination)?*

Document 32016R0679

**Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance)**

Document 32024R1689

**Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) (Text with EEA relevance)**

PE/24/2024/REV/1

A key characteristic of AI systems is their capability to infer.

The techniques that enable inference while building an AI system include machine learning approaches that learn from data how to achieve certain objectives, and logic- and knowledge-based approaches that infer from encoded knowledge or symbolic representation of the task to be solved.

# Motivating Scenario
## European Health Data Spaces




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)

# KG-based AI for Self-Determination
## The Three Pillars

- 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**.

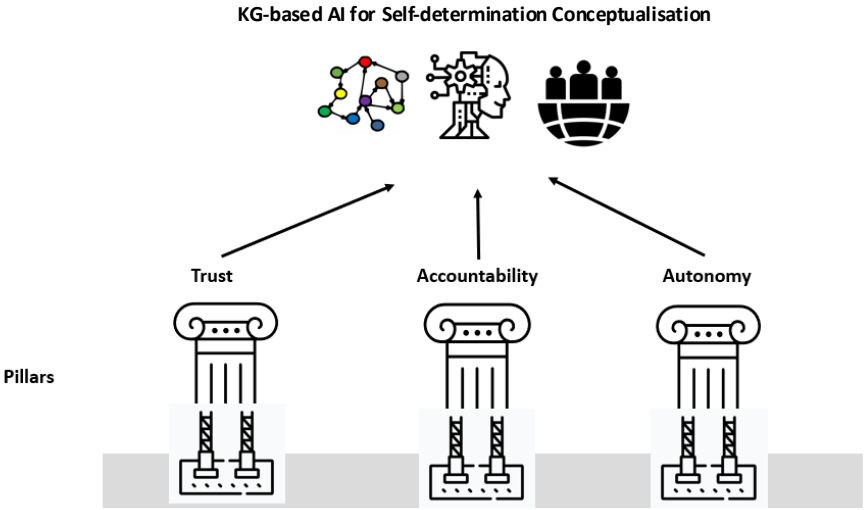




Ibáñez, L., Domingue, J., Kirrane, S., Seneviratne, O., Third, A., Vidal, M., 2023. Trust, Accountability, and Autonomy in Knowledge Graph-based AI for Self-determination. Transactions on Graph Data and Knowledge (TGDK)

# KG-based AI for Self-Determination
## The Three Pillars

### Trust

- Machine-readable policies must **faithfully represent** human policies and norms
- **Policy Enforcement** and **compliance** checking:
  - (semi-)automated techniques
  - auditing and tracing
  - trusted execution environments
  - certification mechanisms



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)

# KG-based AI for Self-Determination
## The Three Pillars

### Trust

- Machine-readable policies must **faithfully represent** human policies and norms
- **Policy Enforcement** and **compliance** checking:
  - (semi-)automated techniques
  - auditing and tracing
  - trusted execution environments
  - certification mechanisms

### Accountability

- Integrating, querying, and aggregating knowledge from **disparate sources**
- Detecting if any party **violated** policies and norms
- Facilitating learning **transparency**
- Providing **explanations** for recommendations and decisions



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)

# KG-based AI for Self-Determination
## The Three Pillars

### Trust

- Machine-readable policies must **faithfully represent** human policies and norms
- **Policy Enforcement** and **compliance** checking:
  - (semi-)automated techniques
  - auditing and tracing
  - trusted execution environments
  - certification mechanisms

### Accountability

- Integrating, querying, and aggregating knowledge from **disparate sources**
- Detecting if any party **violated** policies and norms
- Facilitating learning **transparency**
- Providing **explanations** for recommendations and decisions

### Autonomy

- **Controlling** who has access to our personal data
- **Negotiating** terms of use
- Fostering collaboration via **aggregation and strong privacy** guarantees (e.g., anonymisation)
- **Continuous monitoring** via auditing, tracing, and certification
- **Self-sovereign identities**



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)

# KG-based AI for Self-Determination
## The Four Foundations

- 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

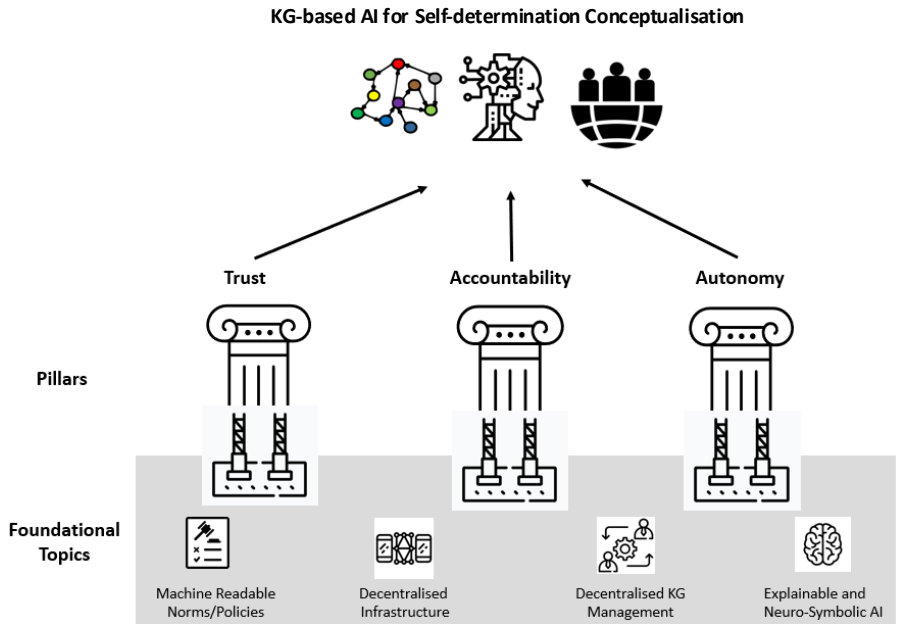




Ibáñez, L., Domingue, J., Kirrane, S., Seneviratne, O., Third, A., Vidal, M., 2023. Trust, Accountability, and Autonomy in Knowledge Graph-based AI for Self-determination. Transactions on Graph Data and Knowledge (TGDK)

# KG-based AI for Self-Determination
## The Four Foundations

**Disclaimer:** Today's talk is a self reflection framed according to the foundational topics identified in our vision paper!

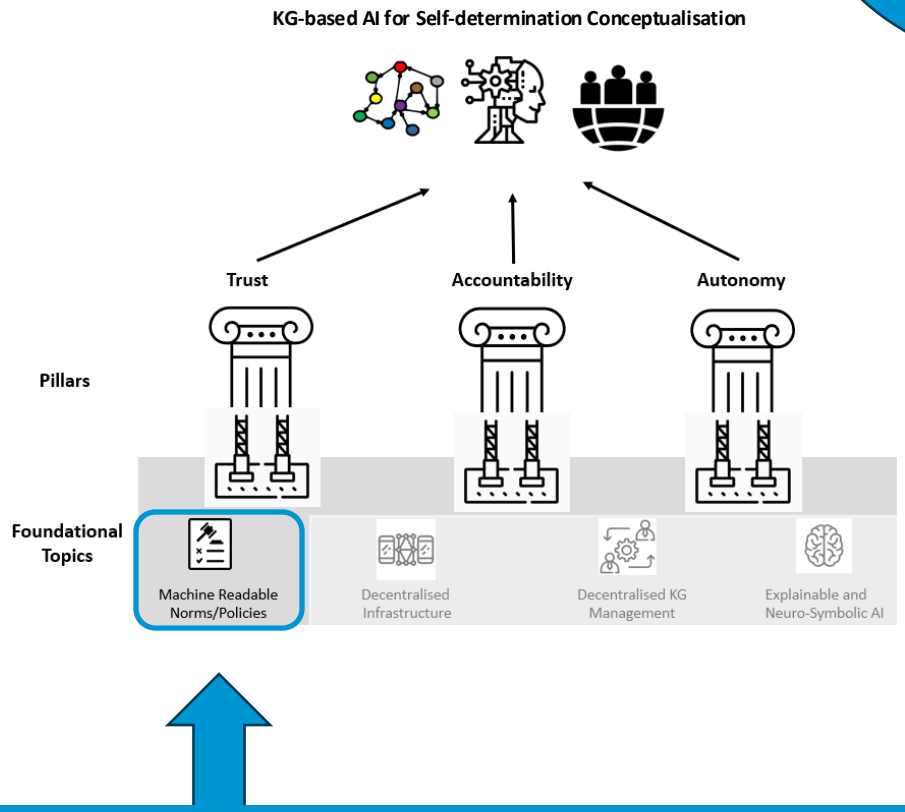




Ibáñez, L., Domingue, J., Kirrane, S., Seneviratne, O., Third, A., Vidal, M., 2023. Trust, Accountability, and Autonomy in Knowledge Graph-based AI for Self-determination. Transactions on Graph Data and Knowledge (TGDK)

# 1 Machine-readable norms and policies

*Consent, Transparency & Compliance Checking*

# 1 Machine-readable norms and policies
## Consent, Transparency & Compliance Checking

WU WIRTSCHAFTS UNIVERSITÄT WIEN VIENNA UNIVERSITY OF ECONOMICS AND BUSINESS

KI - Künstliche Intelligenz (2020) 34:303–315
https://doi.org/10.1007/s13218-020-00677-4

TECHNICAL CONTRIBUTION

Check for updates

### Machine Understandable Policies and GDPR Compliance Checking

Piero A. Bonatti[1] · Sabrina Kirrane[2] · Iliana M. Petrova[1] · Luigi Sauro[1]

Received: 4 November 2019 / Accepted: 18 June 2020 / Published online: 8 July 2020


**Abstract**
The European General Data Protection Regulation (GDPR) calls for technical and organizational measures to support its implementation. Towards this end, the SPECIAL H2020 project aims to provide a set of tools that can be used by data controllers and processors to automatically check if personal data processing and sharing complies with the obligations set forth in the GDPR. The primary contributions of the project include: (i) a policy language that can be used to express consent, business policies, and regulatory obligations; and (ii) two different approaches to automated compliance checking that can be used to demonstrate that data processing performed by data controllers/processors complies with consent provided by data subjects, and business processes comply with regulatory obligations set forth in the GDPR.

**Keywords** GDPR · Policies · Compliance checking

Personal and Ubiquitous Computing (2020) 24:465–486
https://doi.org/10.1007/s00779-019-01330-0

ORIGINAL ARTICLE

Check for updates

### User consent modeling for ensuring transparency and compliance in smart cities

Javier D. Fernández[1] · Marta Sabou[2] · Sabrina Kirrane[1] · Elmar Kiesling[2] · Fajar J. Ekaputra[3] · Amr Azzam[1] · Rigo Wenning[4]

Received: 4 December 2018 / Accepted: 19 September 2019 / Published online: 14 January 2020


**Abstract**
Smart city infrastructures such as transportation and energy networks are evolving into so-called cyber physical social systems (CPSSs), which collect and leverage citizens' data in order to adapt services to citizens' needs. The privacy implications of such systems are, however, significant and need to be addressed. Current systems either try to escape the privacy challenge via anonymization or use very rigid, hard-coded workflows that have been agreed with a data protection authority. In the case of the latter, there is a severe impact on data quality and richness, whereas in the former, only these hard-coded flows are permitted resulting in diminished functionality and potential. We address these limitations via *user modeling* in terms of investigating how to model and semantically represent user consent, preferences, and data usage policies that will guide the processing of said data in the data lake. Data protection is a horizontal field and consequently very wide. Therefore, we focus on a concrete setting where we extend the domain-agnostic SPECIAL policy language for a smart mobility use case supplied by Vienna's largest utility provider. To that end, (1) we create an extension of SPECIAL in terms of a core CPSS vocabulary that lowers the semantic gap between the domain agnostic terms of SPECIAL and the vocabulary of the use case; (2) we propose a workflow that supports defining domain-specific vocabularies for complex CPSSs; and (3) show that these two contributions allow successfully achieving the goals of our setting.

**Keywords** Cyber physical (social) systems · Smart mobility · User consent modeling · Privacy · GDPR · Linked data



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.

26

# 1 Machine-readable norms and policies
## Consent as a Legal Basis for Data Processing

**Data Subject**

Personal Data

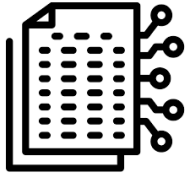

Bob

Bob makes his data available only for **medical purposes**

**Data Controller**

Consent UI


Medical KG

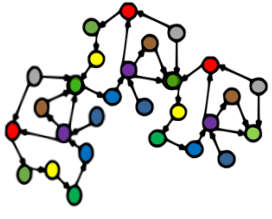

Medical Organisations

Usage Policies

The Data Controller processes data in accordance with the **data subjects consent**

# 1 Machine-readable norms and policies
## OWL based Usage Policy Language

**SPECIAL's Usage Policy Language Grammar**

**UsagePolicy** :='ObjectUnionOf' '(' **BasicUsagePolicy BasicUsagePolicy** { **BasicUsagePolicy** } ')'
| **BasicUsagePolicy**

**BasicUsagePolicy** :='ObjectIntersectionOf' '(' **Data Purpose Processing Recipients Storage** ')'

**Data** :='ObjectSomeValueFrom' '(' 'spl:hasData' **DataExpression** ')'

**Purpose** :='ObjectSomeValueFrom' '(' 'spl:hasPurpose' **PurposeExpression** ')'

**Processing** :='ObjectSomeValueFrom' '(' 'spl:hasProcessing' **ProcessingExpression** ')'

**Recipients** :='ObjectSomeValueFrom' '(' 'spl:hasRecipient' **RecipientExpression** ')'

**Storage** :='ObjectSomeValueFrom' '(' 'spl:hasStorage' **StorageExpression** ')'

**DataExpression** :='spl:AnyData' | **DataVocabExpression**

**PurposeExpression** :='spl:AnyPurpose' | **PurposeVocabExpression**

**ProcessingExpression** :='spl:AnyProcessing' | **ProcessingVocabExpression**

**RecipientsExpression** :='spl:AnyRecipient' | 'spl:Null' | **RecipientVocabExpression**

**StorageExpression** :='spl:AnyStorage' | 'spl:Null' |
'ObjectIntersectionOf' '(' **Location Duration** ')'

**Location** :='ObjectSomeValueFrom' '(' 'spl:hasLocation' **LocationExpression** ')'

**Duration** :='ObjectSomeValueFrom' '(' 'spl:hasDuration' **DurationExpression** ')'
| 'DataSomeValueFrom' '(' 'spl:durationInDays' **IntervalExpression** ')'

- The SPECIAL usage policy language can be used to express:
  - data subject **consent**
  - data controllers **usage requests**
  - fragments of the **GDPR**
  - processing requirements as **business policies**
- We extensively **re-uses standards** based privacy-related vocabularies
- Policies are expressed using the **Web Ontology Language (OWL)**, thus we are able to **leverage existing OWL reasoners** out of the box



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.

# 1 Machine-readable norms and policies
## Log Vocabulary

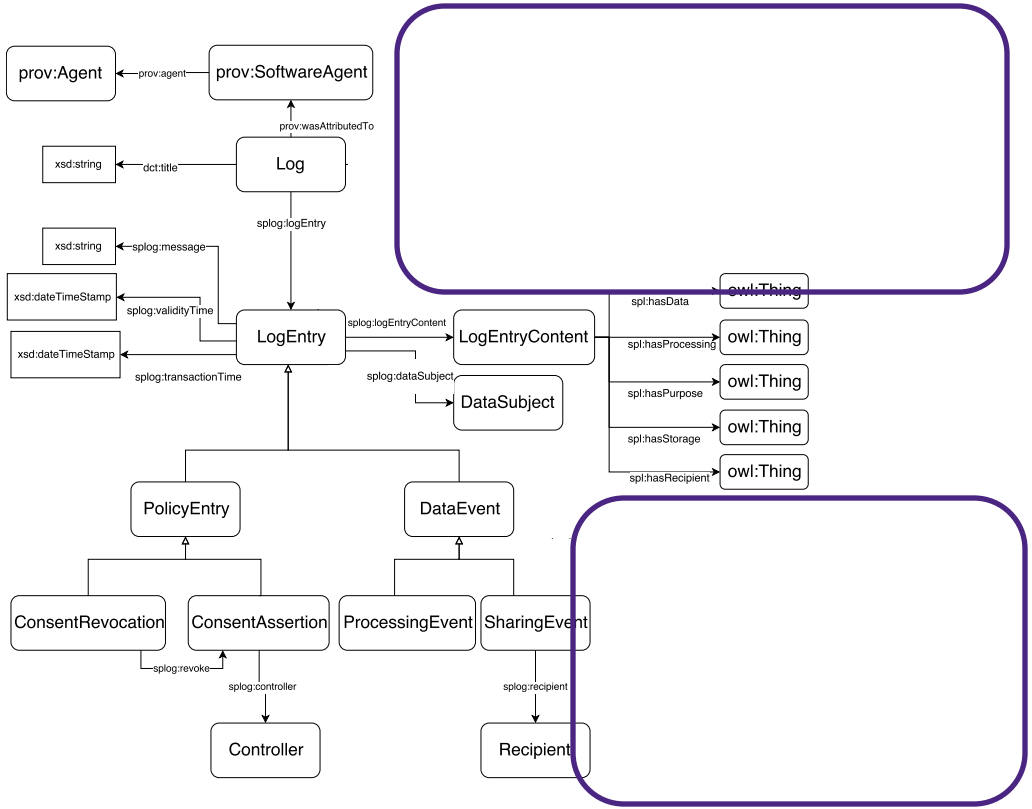


- The SPECIAL **log vocabulary** reuses well-known vocabularies such as **PROV** for representing provenance metadata
- Log entries are used to represent:
  - Data processing events
  - Policy events
- Optional components are provided for:
  - Immutability
  - Business process management (BPM)



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.

# 1 Machine-readable norms and policies
## Compliance Checking Architecture

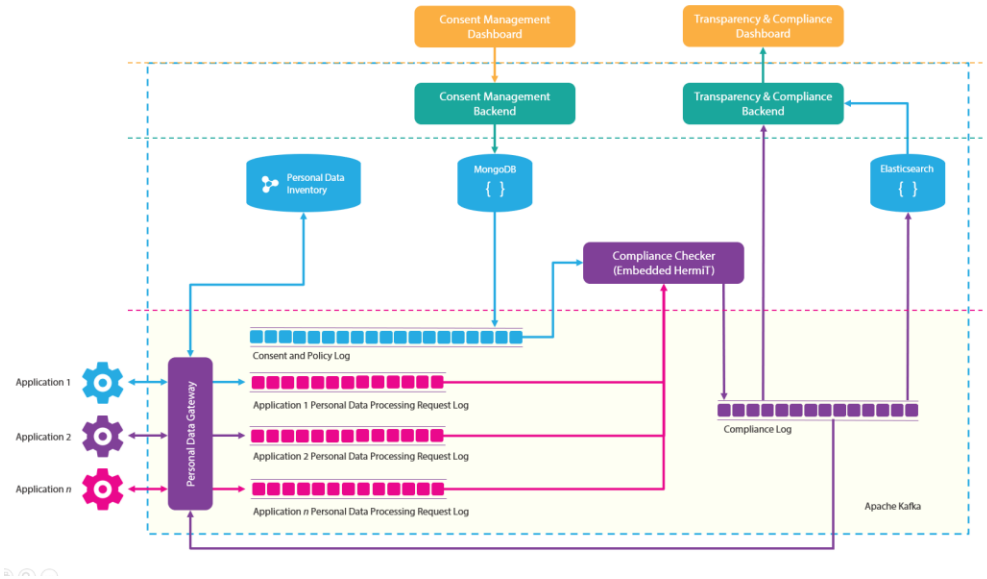


- Data processing and sharing event logs are stored in the **Kafka** distributed streaming platform
- We assume that consent updates are infrequent and as such usage policies and the respective vocabularies are represented in a **virtuoso triple store**
- The compliance checker, which includes an embedded **HermiT reasoner** uses the consent saved in Virtuoso together with the application logs provided by Kafka to check that data processing and sharing complies with the relevant usage control policies
- As logs can be serialized using JSON-LD, it is possible to benefit from the faceting browsing capabilities of **Elasticsearch** and the out of the box visualization capabilities provided by **Kibana**



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.

# 1 Machine-readable norms and policies
## Assessing Effectiveness

WU WIRTSCHAFTS UNIVERSITÄT WIEN VIENNA UNIVERSITY OF ECONOMICS AND BUSINESS

### Telco and Financial Services Pilots


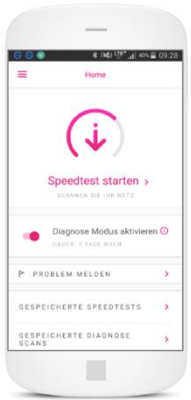



### Cyber-Physical Social Systems Project

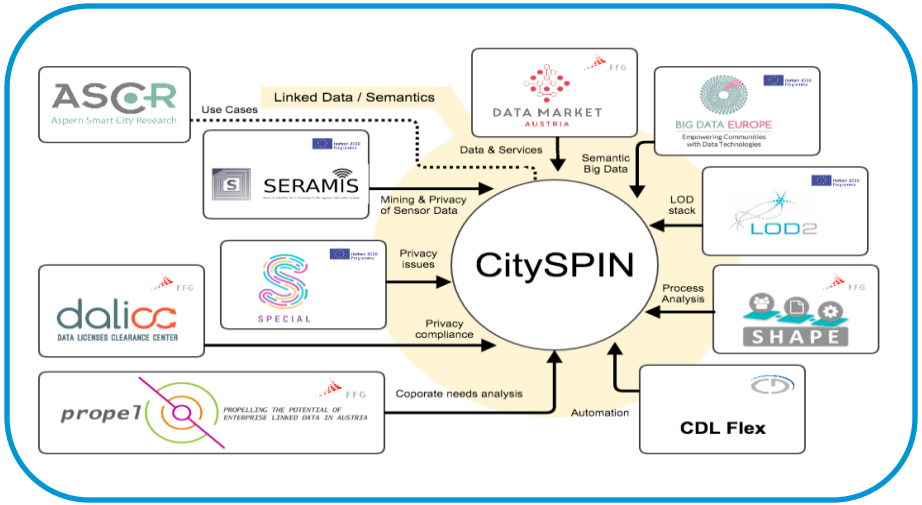




Bonatti, P.A., Kirrane, S., Petrova, I.M. and Sauro, L., 2020. Machine understandable policies and GDPR compliance checking. KI-Künstliche Intelligenz.

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.

# 1 Machine-readable norms and policies

*Usage Control in Decentalised Settings*

# 1 Machine-readable norms and policies
## Usage Control in Decentalised Settings

Contents lists available at ScienceDirect

Computer Science Review

journal homepage: www.elsevier.com/locate/cosrev

Review article

### A comprehensive review of usage control frameworks

Ines Akaichi *, Sabrina Kirrane

*Institute for Complex Networks, Vienna University of Economics and Business, Welthandelspl. 1, 1020, Vienna, Austria*

ARTICLE INFO

*Keywords:*
Usage control
Policy languages
Enforcement frameworks
Decentralized systems

ABSTRACT

The sharing of data and digital assets in a decentralized settling is entangled with various legislative challenges, including, but not limited to, the need to adhere to legal requirements with respect to privacy and copyright. In order to provide more control to data and digital asset owners, usage control could be used to make sure that consumers handle data according to privacy, licenses, regulatory requirements, among others. However, considering that many of the existing usage control frameworks were designed to cater for different use cases (e.g., networking, operating systems, and industry 4.0), there is a need to better understand the existing proposals and how they compare to one another. In this paper, we provide a holistic overview of existing usage control frameworks and their support for a broad set of requirements. We systematically collect requirements that are routinely used to guide the development of usage control solutions, which are classified according to three broad dimensions: *specification*, *enforcement*, and *system*. We use these requirements to conduct a qualitative comparison of the most prominent usage control frameworks found in the literature. Finally, we identify existing gaps, challenges, and opportunities in the field of usage control in general, and in decentralized environments in particular.

Poster

### Implementing Usage Control Policies Using Reification with RDF-Star and SPARQL-Star

Ines Akaichi[1,*], Giorgos Flouris[2], Irini Fundulaki[2] and Sabrina Kirrane[1]

[1]*Institute for Complex Networks, WU, Vienna, Austrian*
[2]*Institute of Computer Science, FORTH, Heraklion, Greece*

**Abstract**

In decentralized environments usage control (UC) is crucial for governing asset and resource usage. In an earlier work, we proposed GUCON, a generic graph pattern based policy framework for UC enforcement, which offers a formal semantics for specifying and reasoning over UC policies. Although GUCON caters for the expression of various rules using graph patterns, the incorporation of temporal constraints is cumbersome. In this paper, we propose an instantiation of the GUCON framework that demonstrates how RDF-star can be used for rule representation and SPARQL-star for compliance checking. Additionally, we extend the original policy language to cater for temporal constraints.

**Keywords**

RDF-star, SPARQL-star, Usage Control, Policies



Ines Akaichi and Sabrina Kirrane. A Comprehensive Review of Usage Control Frameworks. 2025. Computer Science Review Journal.

Ines Akaichi, Giorgos Flouris, Irini Fundulaki and Sabrina Kirrane. Implementing Usage Control Policies Using Reification with RDF-Star and SPARQL-Star. 2024. Proceedings of the Posters and Demos Track of the 23rd International Semantic Web Conference (ISWC),

# 1 Machine-readable norms and policies
## Usage Control in Decentalised Settings

**Data Producer**

Bob

Usage Policies

Knowledge Graph

Bob makes his data available only for **medical purposes**

**Data Consumer**

Knowledge Graph

Usage Policies

Alice

The **Data Consumer** processes data in accordance with the **Data Producers** policies

# 1 Machine-readable norms and policies
## SPARQL* based Usage Policy Language

WU WIRTSCHAFTS UNIVERSITÄT WIEN VIENNA UNIVERSITY OF ECONOMICS AND BUSINESS

### A GUCON Rule $cond \rightsquigarrow d\ a$

- *cond* is a *graph pattern*
- *d a* is a *deontic pattern*
  - $d \in \{A, P, O, D\}$, such as A (permission) P (prohibition), O (obligation), and D (dispensation)
  - $a = \{(s, p, o) \mid s \in E, p \in A, o \in R\}$, such as an (s, p, o) corresponds to a SPARQL triple pattern, E is the set of all **entities**; A is the set of all **actions**; and R is the set of all **resources**
- $\rightsquigarrow$ denotes if *cond* is satisifed , then *da* may (A)/must not (P)/must (O)/need not (D) be satisfied.

### A Serialization of GUCON using SPARQL

- *cond* is a ***SPARQL*** *graph pattern*
- *a* is a ***SPARQL*** *triple pattern*

```
{?share rdf:type splog:SharingEvent ;
    rdf:subject ?x ;
    rdf:predicate eg:share ;
    rdf:object egmc:l.premierleague.com-2020-coverage ;
    splog:occurs  ?t .
?x  rdf:type eg:Person ;
    eg:position eg:Journalist .
FILTER (?t > "2021-02-25T21:00:00Z"^^xsd:dateTime)}
~ {A {?x  :share  egmc:l.premierleague.com-2020-coverage }}
```

### A Serialization of GUCON using SPARQL-star

- *cond* is a ***SPARQL-star*** *graph pattern*
- *constraints are expressed using* ***quoted action patterns***
- *a* is a ***SPARQL*** *triple pattern*

```
{<<?x eg:share egmc:l.premierleague.com-2020-coverage>>
   splog:occurs  ?t .
?x  rdf:type eg:Person;
    eg:position eg:Journalist .
FILTER (?t > "2021-02-25T21:00:00Z"^^xsd:dateTime )}
~ {A {?x eg:share egmc:l.premierleague.com-2020-coverage }}
```



Ines Akaichi, Giorgos Flouris, Irini Fundulaki and Sabrina Kirrane. Implementing Usage Control Policies Using Reification with RDF-Star and SPARQL-Star. 2024. Proceedings of the Posters and Demos Track of the 23rd International Semantic Web Conference (ISWC),

# 1 Machine-readable norms and policies
## Usage Control Requirements

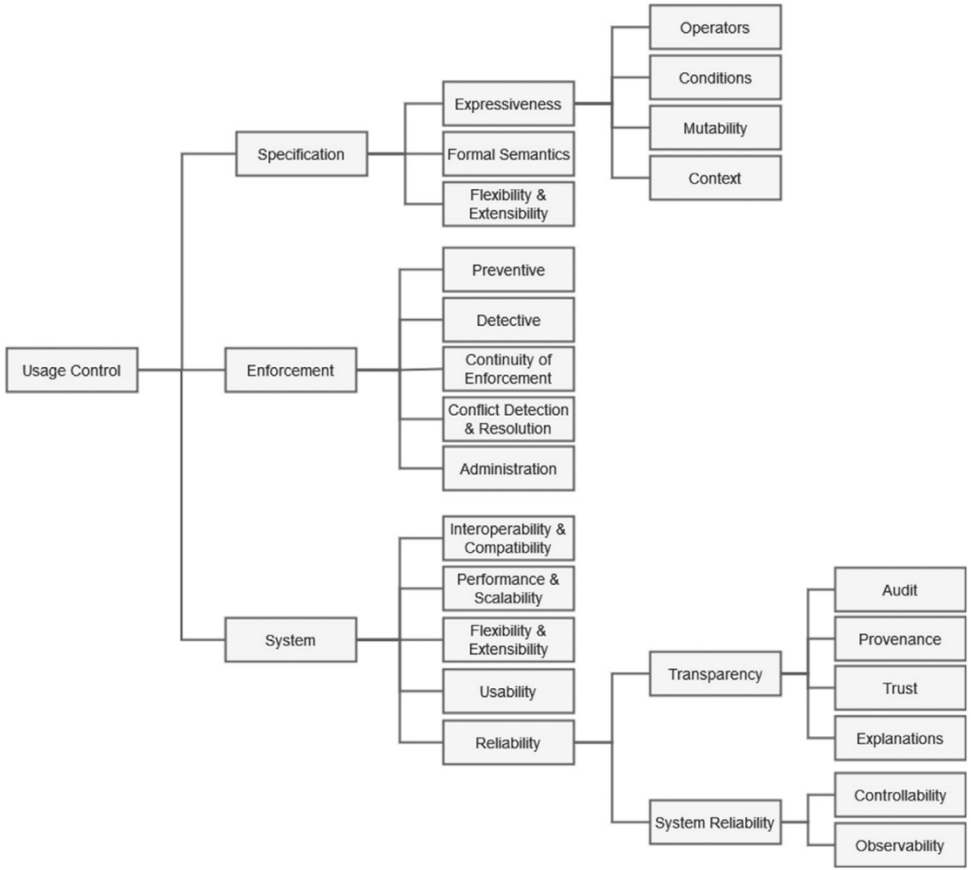


- **Trust**
  - Concept of trust in the context of distributed usage control remains relatively unexplored
  - Heterogeneous entities and systems
- **Control**
  - Different copies and derivations
  - Distribute and synchronize policies across decentralized systems
- **Governance**
  - No centralized authority and oversight
  - Need to handle diverse regulatory requirements



Ines Akaichi and Sabrina Kirrane. A Comprehensive Review of Usage Control Frameworks. 2025. Computer Science Review Journal.

# 2 Decentralised infrastructure

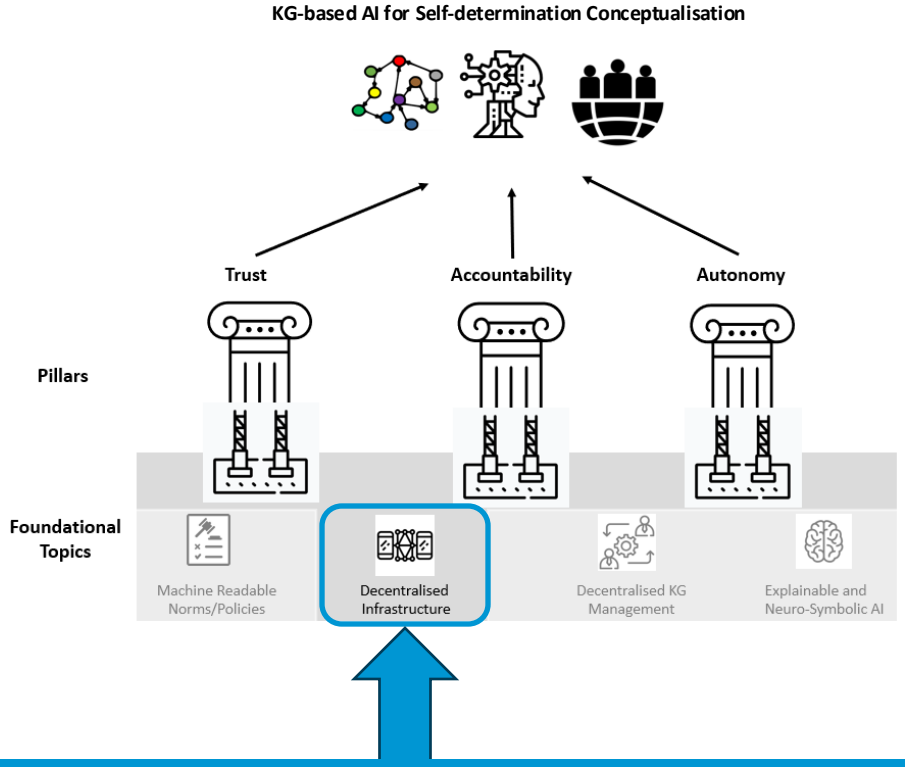



37

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)

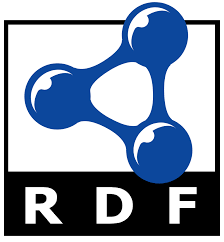


Data model for encoding triples

**Disclaimer:** This is not a complete list!


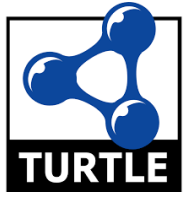




Textual syntax for RDF


Query language


Data modelling vocabularies

# Solid is an ecosystem of data and apps that work seamlessly together.

**data pods**

profile, photos, comments, likes, ...

**applications**

photo album, meeting invites, document collaboration, ...

**standards**

HTTP, URL, RDF, LDP, LDN, ...



Slides from https://github.com/RubenVerborgh/Solid-DeSemWeb-2018

# 2 Decentralised infrastructure
## Solid: your data, your choice

**You can grant apps and people access to very specific parts of your data.**

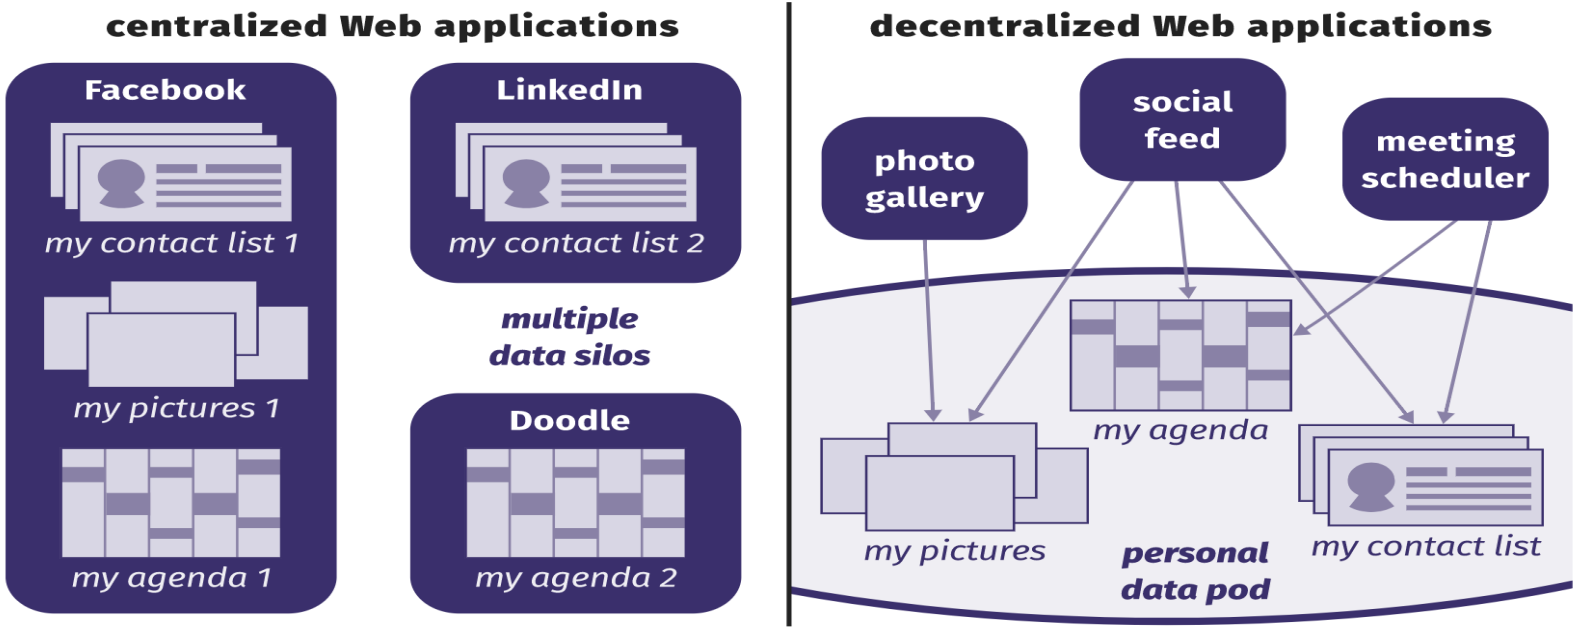




Slides from https://github.com/RubenVerborgh/Solid-DeSemWeb-2018

# Building apps over decentralized data requires different app techniques.

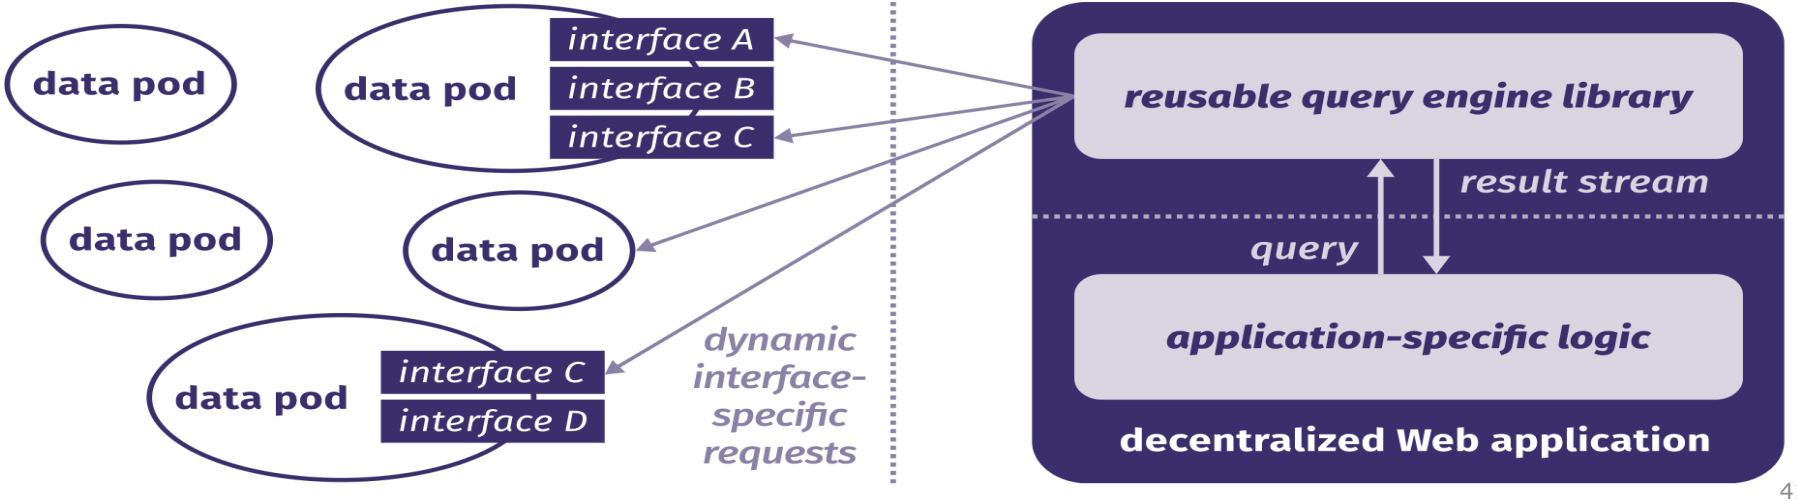




Slides from https://github.com/RubenVerborgh/Solid-DeSemWeb-2018

# 2 Decentralised infrastructure
## Solid: your data, your choice

WU WIRTSCHAFTS UNIVERSITÄT WIEN VIENNA UNIVERSITY OF ECONOMICS AND BUSINESS

**You can choose where you store every single piece of data you produce.**

**Author's name and latest profile picture**
*stored in author's personal data pod*

**Work-related opinion about an article**
*stored in data pod of author's company*

**Discussed article title and photo**
*stored in news website's data pod*

**Likes on this post**
*each one in different individuals' data pods*

**Comments on this post**
*each one in different individuals' data pods*



Slides from https://github.com/RubenVerborgh/Solid-DeSemWeb-2018

# *Ensuring Compliance with Policies in Decentalised Environments*

# 2 Decentralised infrastructure
## Solid Resource Governance

frontiers | Frontiers in Blockchain

TYPE Original Research
PUBLISHED 09 May 2023
DOI 10.3389/fbloc.2023.1141909

Check for updates

OPEN ACCESS

EDITED BY
Marco Comuzzi,
Ulsan National Institute of Science and Technology, Republic of Korea

REVIEWED BY
Petros Kavassalis,
University of the Aegean, Greece
Paulo Rupino Cunha,
University of Coimbra, Portugal

*CORRESPONDENCE
Valerio Goretti,
valerio.goretti@uniroma1.it

RECEIVED 11 January 2023
ACCEPTED 04 April 2023
PUBLISHED 09 May 2023

CITATION
Basile D, Di Ciccio C, Goretti V and Kirrane S (2023), Blockchain based resource governance for decentralized web environments.
*Front. Blockchain* 6:1141909.
doi: 10.3389/fbloc.2023.1141909



# Blockchain based resource governance for decentralized web environments

Davide Basile[1], Claudio Di Ciccio[1], Valerio Goretti[1]* and Sabrina Kirrane[2]

[1]Department of Computer Science, Sapienza University of Rome, Rome, Italy, [2]Institute for Information Systems and New Media, Vienna University of Economics and Business, Vienna, Austria

Decentralization initiatives such as Solid, Digi.me, and ActivityPub aim to give data owners more control over their data and to level the playing field by enabling small companies and individuals to gain access to data, thus stimulating innovation. However, these initiatives typically use access control mechanisms that cannot verify compliance with usage conditions after access has been granted to others. In this paper, we extend the state of the art by proposing a resource governance conceptual framework, entitled ReGov, that facilitates usage control in decentralized web environments. We subsequently demonstrate how our framework can be instantiated by combining blockchain and trusted execution environments. Through blockchain technologies, we record policies expressing the usage conditions associated with resources and monitor their compliance. Our instantiation employs trusted execution environments to enforce said policies, inside data consumers' devices. We evaluate the framework instantiation through a detailed analysis of requirments derived from a data market motivating scenario, as well as an assessment of the security, privacy, and affordability aspects of our proposal.

KEYWORDS

**decentralization, usage control, governance, blockchain, trusted execution environment (TEE)**



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

# 2 Decentralised infrastructure
## A Decentralised Data Market

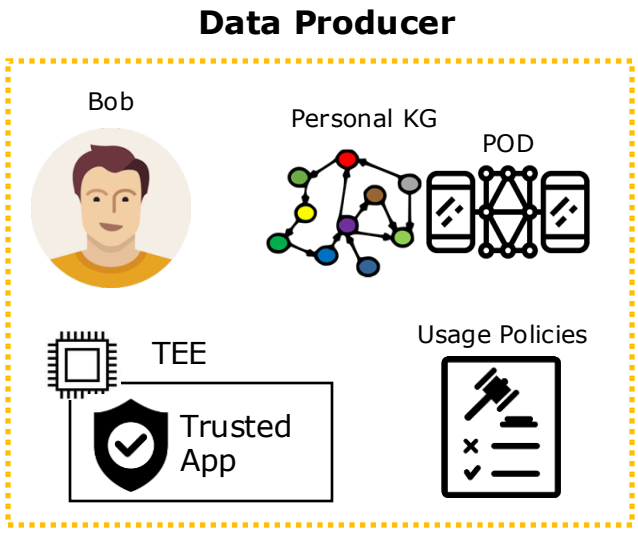


Bob sets up a **personal online datastore (POD)**

Makes his resources available only for **medical purposes**

Gets a **remuneration** according to the **number of accesses**

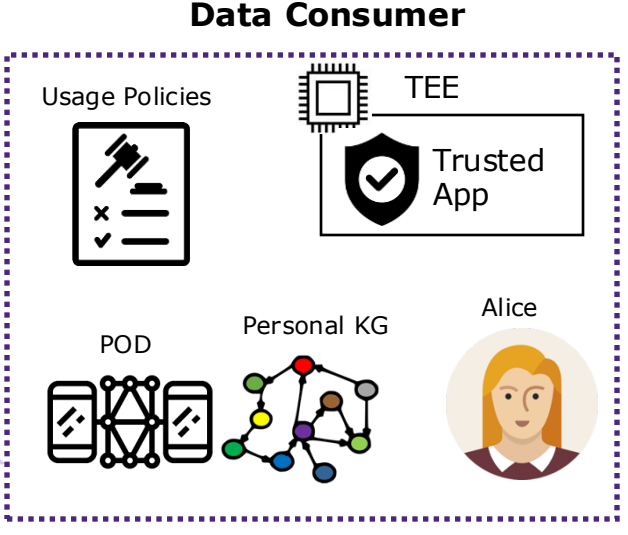


Alice asks the market for **medical data**

**Contacts** the personal online datastore (POD)

Uses the retrieved resources in **her trusted execution environment (TEE)**

# 2 Decentralised infrastructure
## Solid Resource Governance

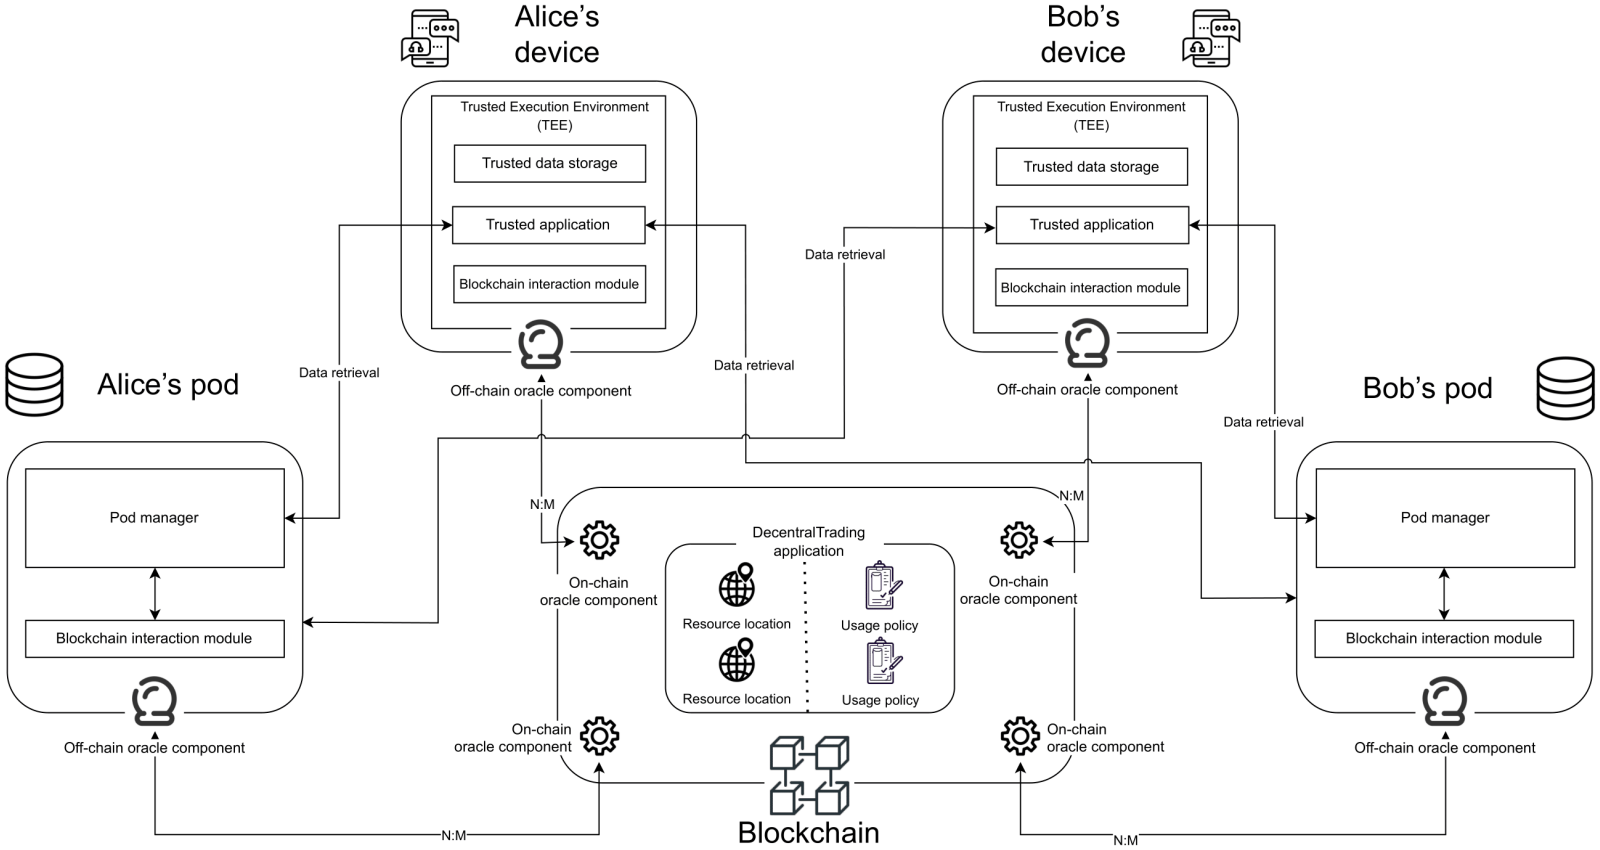




Basile, D., Di Ciccio, C., Goretti, V. and **Kirrane, S.,** 2023. Blockchain based Resource Governance for Decentralized Web Environments. Frontiers in Blockchain.

# 3 Decentralised KG management

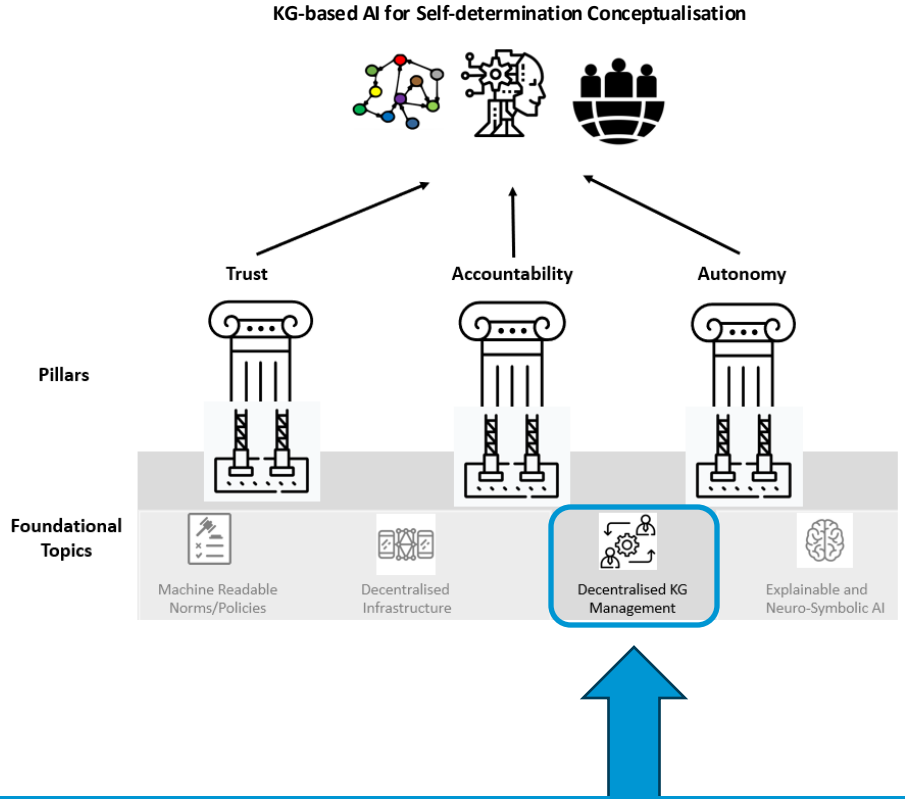




Ibáñez, L., Domingue, J., **Kirrane, S.**, Seneviratne, O., Third, A., Vidal, M., 2023. Trust, Accountability, and Autonomy in Knowledge Graph-based AI for Self-determination. Transactions on Graph Data and Knowledge (TGDK)

# *Federated Data Querying with Policies*

# 3 Decentralised KG management
## Federated Querying with Policies

Report from Dagstuhl Seminar 24061

# Are Knowledge Graphs Ready for the Real World? Challenges and Perspective

### 5.1 Access and Usage Control for Federations of Knowledge Graphs

*Piero Andrea Bonatti (University of Naples, IT)*
*Pierre-Antoine Champin (INRIA – Sophia Antipolis, FR)*
*Anastasia Dimou (KU Leuven, BE)*
*Peter Haase (Metaphacts GmbH – Walldorf, DE)*
*Olaf Hartig (Linköping University, SE)*
*Katja Hose (TU Wien, AT)*
*Sabrina Kirrane (Wirtschaftsuniversität Wien, AT)*
*Edelmira Pasarella (UPC Barcelona Tech, ES)*





Chaves-Fraga, D., Corcho, O., Dimou, A., Vidal, M.E., Iglesias-Molina, A. and Van Assche, D., 2024. Are Knowledge Graphs Ready for the Real World? Challenges and Perspective.

# Access and Usage Control for Federated Querying with Policies

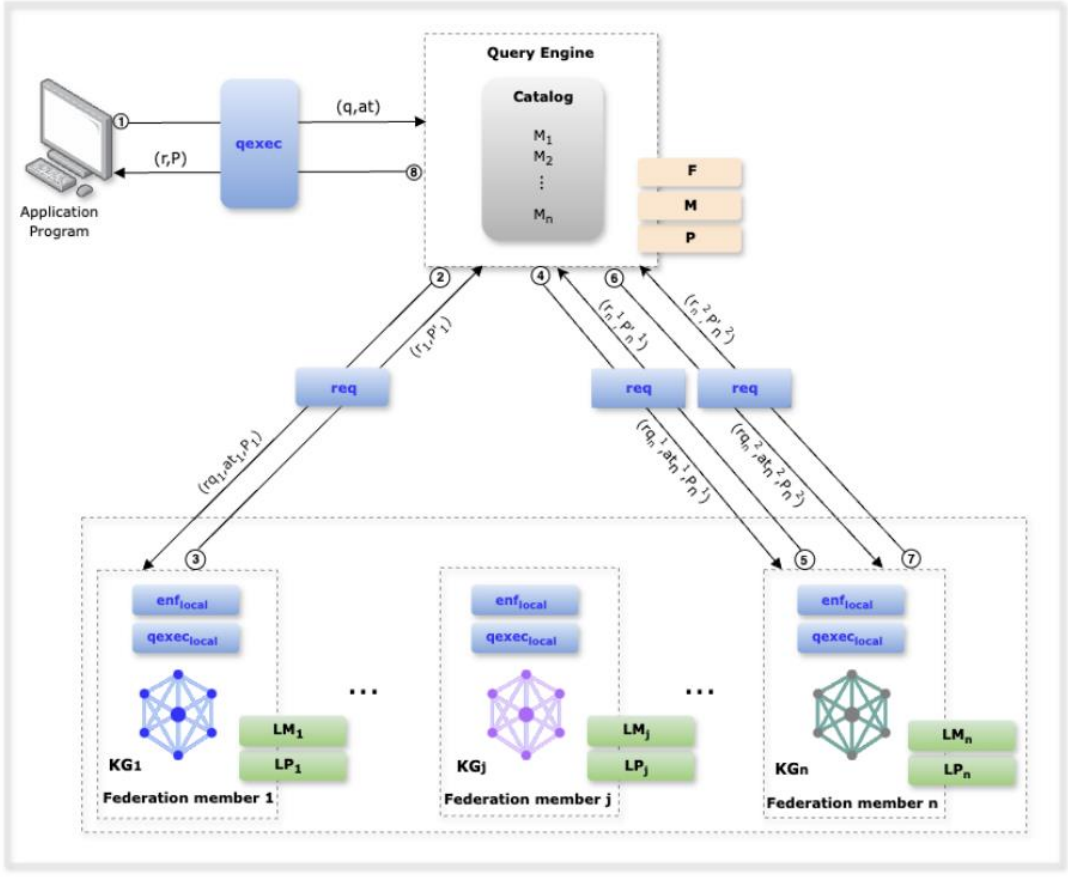


- Federation members are under the control of a **single organization**
- Different organizations (or different units within the same organization) share and integrate some of their respective data within a **collaboration or consortium** that has been created to achieve a common goal (e.g., in a Data Space)
- Federation members that have not been created explicitly for participating in that particular federation and that **may not even be aware of their participation**

# *Autonomous Agents with Policies and Norms*

# 3 Decentralised KG management
## Governance of Autonomous Agents

### Governance of Autonomous Agents on the Web: Challenges and Opportunities

TIMOTHEUS KAMPIK, Umeå University, Sweden
ADNANE MANSOUR and OLIVIER BOISSIER, Mines Saint-Étienne, France
SABRINA KIRRANE, Wirtschaftsuniversität Wien, Austria
JULIAN PADGET, University of Bath, UK
TERRY R. PAYNE, University of Liverpool, UK
MUNINDAR P. SINGH, North Carolina State University, USA
VALENTINA TAMMA, University of Liverpool, UK
ANTOINE ZIMMERMANN, Mines Saint-Étienne, France

The study of autonomous agents has a long history in the Multiagent System and the Semantic Web communities, with applications ranging from automating business processes to personal assistants. More recently, the Web of Things (WoT), which is an extension of the Internet of Things (IoT) with metadata expressed in Web standards, and its community provide further motivation for pushing the autonomous agents research agenda forward. Although representing and reasoning about norms, policies, and preferences is crucial to ensuring that autonomous agents act in a manner that satisfies stakeholder requirements, normative concepts, policies, and preferences have yet to be considered as first-class abstractions in Web-based multiagent systems. Towards this end, this article motivates the need for alignment and joint research across the Multiagent Systems, Semantic Web, and WoT communities, introduces a conceptual framework for governance of autonomous agents on the Web, and identifies several research challenges and opportunities.

CCS Concepts: • **Computing methodologies** → **Multiagent Systems**; **Intelligent agents**; • **Information systems** → **web services;**

Additional Key Words and Phrases: Autonomous agents, norms, policies, preferences, governance

# 3 Decentralised KG management
## Vaccination Appointment Scheduling

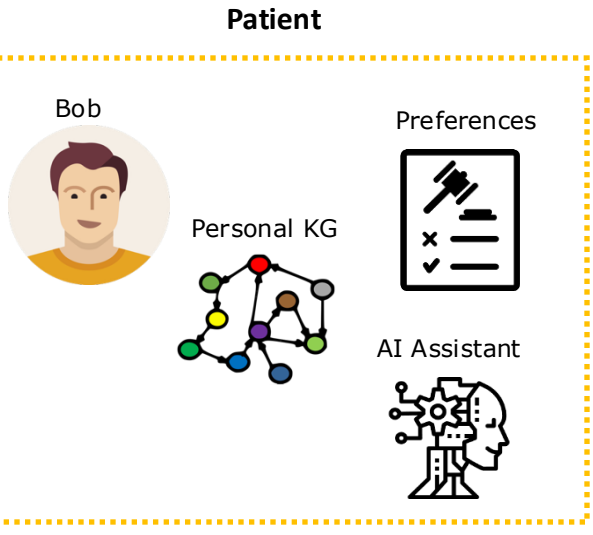


Bob's assistant agent is in charge of **managing personal data** on his behalf.

Bob's assistant agent uses his **appointment calendar** and his **medical history** to schedule a vaccination appointment.

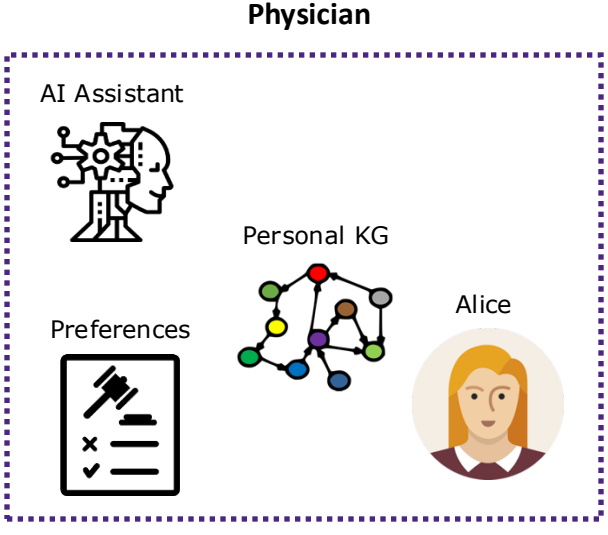


Alice's physician agent is in charge of **managing administrative tasks**.

Alice's physician agent provides **recommendations on vaccinations** and makes **appointments based on preference classes**.

# 3 Decentralised KG management
## Governance Conceptual Framework

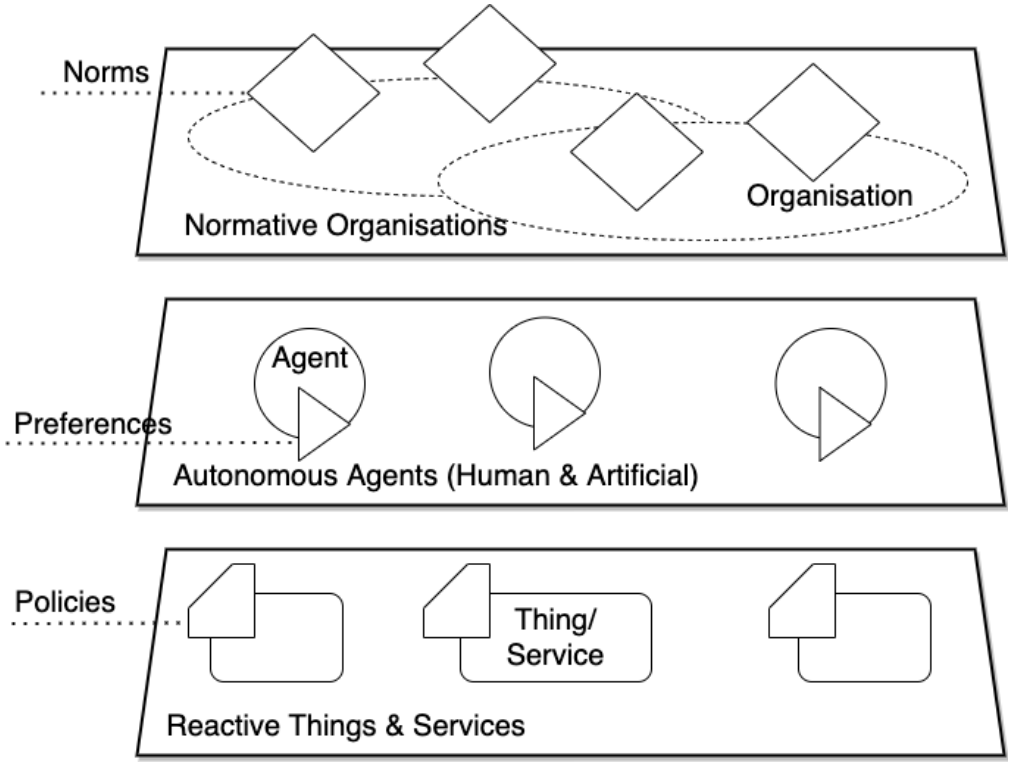


**Norms**

- Organisations are first-class abstractions that group agents and their governance (i.e., norms)
- Logical grouping of agents with a particular purpose, and the provision of legal, regulatory and social norms that may possibly span multiple organisations

**Preferences**

- Entities that autonomously perceive and act upon their environment (i.e., things and services) and interact with the other entities
- Agents have preferences that inform and constrain their actions with respect to things, web services and other agents.

**Policies**

- Non-autonomous entities in the environment
- Policies state who can access things/services and constraints on their usage



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.

54

# 3 Decentralised KG management
## Assessing Effectiveness

Vaccination Motivating Scenario

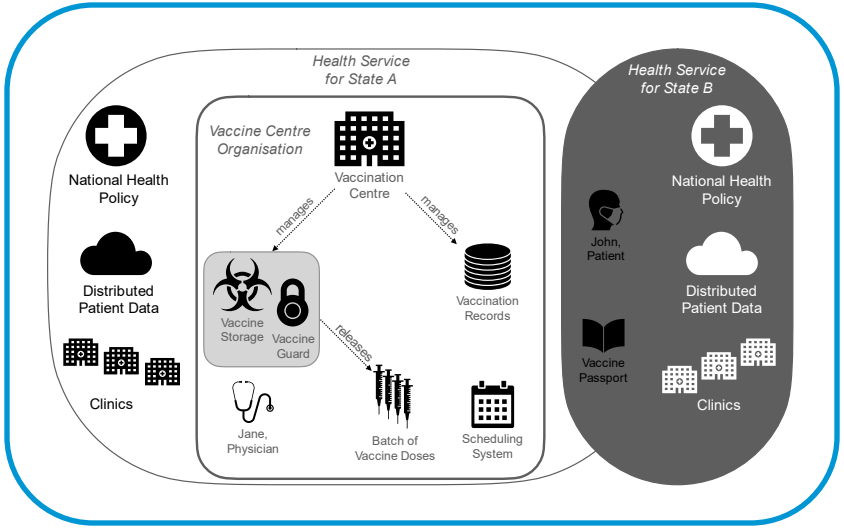


Instantiation of the Conceptual Framework

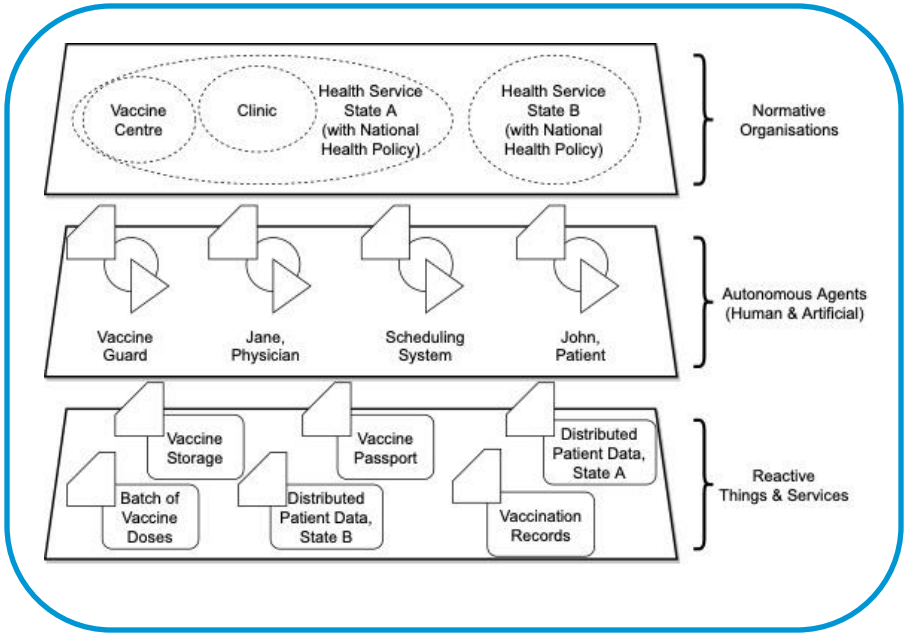




Kampik, T., Mansour, A., Boissier, O., Kirrane, S., Padget, J., Payne, T.R., Singh, M.P., Tamma, V. and Zimmermann, A., 2022. Governance of Autonomous Agents on the Web: Challenges and Opportunities. ACM Transactions on Internet Technology.

55

# 4 Explainable and Neuro-Symbolic AI

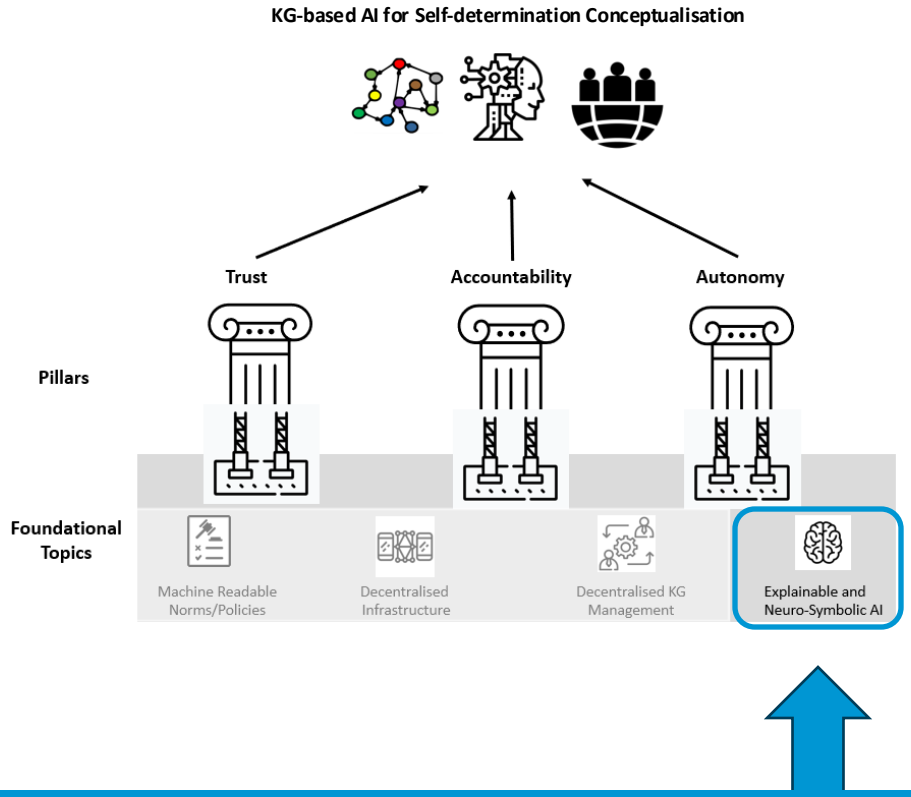




Ibáñez, L., Domingue, J., Kirrane, S., Seneviratne, O., Third, A., Vidal, M., 2023. Trust, Accountability, and Autonomy in Knowledge Graph-based AI for Self-determination. Transactions on Graph Data and Knowledge (TGDK)

# *Encoding Legal Knowledge*

Artificial Intelligence and Law (2021) 29:485–539
https://doi.org/10.1007/s10506-021-09282-8

ORIGINAL RESEARCH

### The linked legal data landscape: linking legal data across different countries

Erwin Filtz[1,2] · Sabrina Kirrane[1] · Axel Polleres[1,3]

Accepted: 6 January 2021 / Published online: 25 February 2021


**Abstract**
The European Union is working towards harmonizing legislation across Europe, in order to improve cross-border interchange of legal information. This goal is supported for instance via standards such as the European Law Identifier (ELI) and the European Case Law Identifier (ECLI), which provide technical specifications for Web identifiers and suggestions for vocabularies to be used to describe metadata pertaining to legal documents in a machine readable format. Notably, these ECLI and ELI metadata standards adhere to the RDF data format which forms the basis of Linked Data, and therefore have the potential to form a basis for a pan-European legal Knowledge Graph. Unfortunately, to date said specifications have only been partially adopted by EU member states. In this paper we describe a methodology to transform the existing legal information system used in Austria to such a legal knowledge graph covering different steps from modeling national specific aspects, to population, and finally the integration of legal data from other countries through linked data. We demonstrate the usefulness of this approach by exemplifying practical use cases from legal information search, which are not possible in an automated fashion so far.

**Keywords** Linked data · Legal knowledge graph · Legal ontology · Law identifier

### Events Matter: Extraction of Events from Court Decisions

Erwin FILTZ [a,b], María NAVAS-LORO [c], Cristiana SANTOS [d], Axel POLLERES [a] and Sabrina KIRRANE [a]
[a] *Vienna University of Economics and Business*
[b] *Siemens AG Österreich*
[c] *Universidad Politécnica de Madrid – Ontology Engineering Group, Madrid, Spain*
[d] *Utrecht University*

**Abstract.** The analysis of court decisions and associated events is part of the daily life of many legal practitioners. Unfortunately, since court decision texts can often be long and complex, bringing all events relating to a case in order, to understand their connections and durations is a time-consuming task. Automated court decision timeline generation could provide a visual overview of what happened throughout a case by representing the main legal events, together with relevant temporal information. Tools and technologies to extract events from court decisions however are still underdeveloped. To this end, in the current paper we compare the effectiveness of three different extraction mechanisms, namely deep learning, conditional random fields, and rule-based method, to facilitate automated extraction of events and their components (i.e., the event type, who was involved, and when it happened). In addition, we provide a corpus of manually annotated decisions of the European Court of Human Rights, which shall serve as a gold standard not only for our own evaluation, but also for the research community for comparison and further experiments.

**Keywords.** event extraction, named entity recognition, court decisions

# 4 Explainable and Neuro-Symbolic AI
## Linked Legal Data

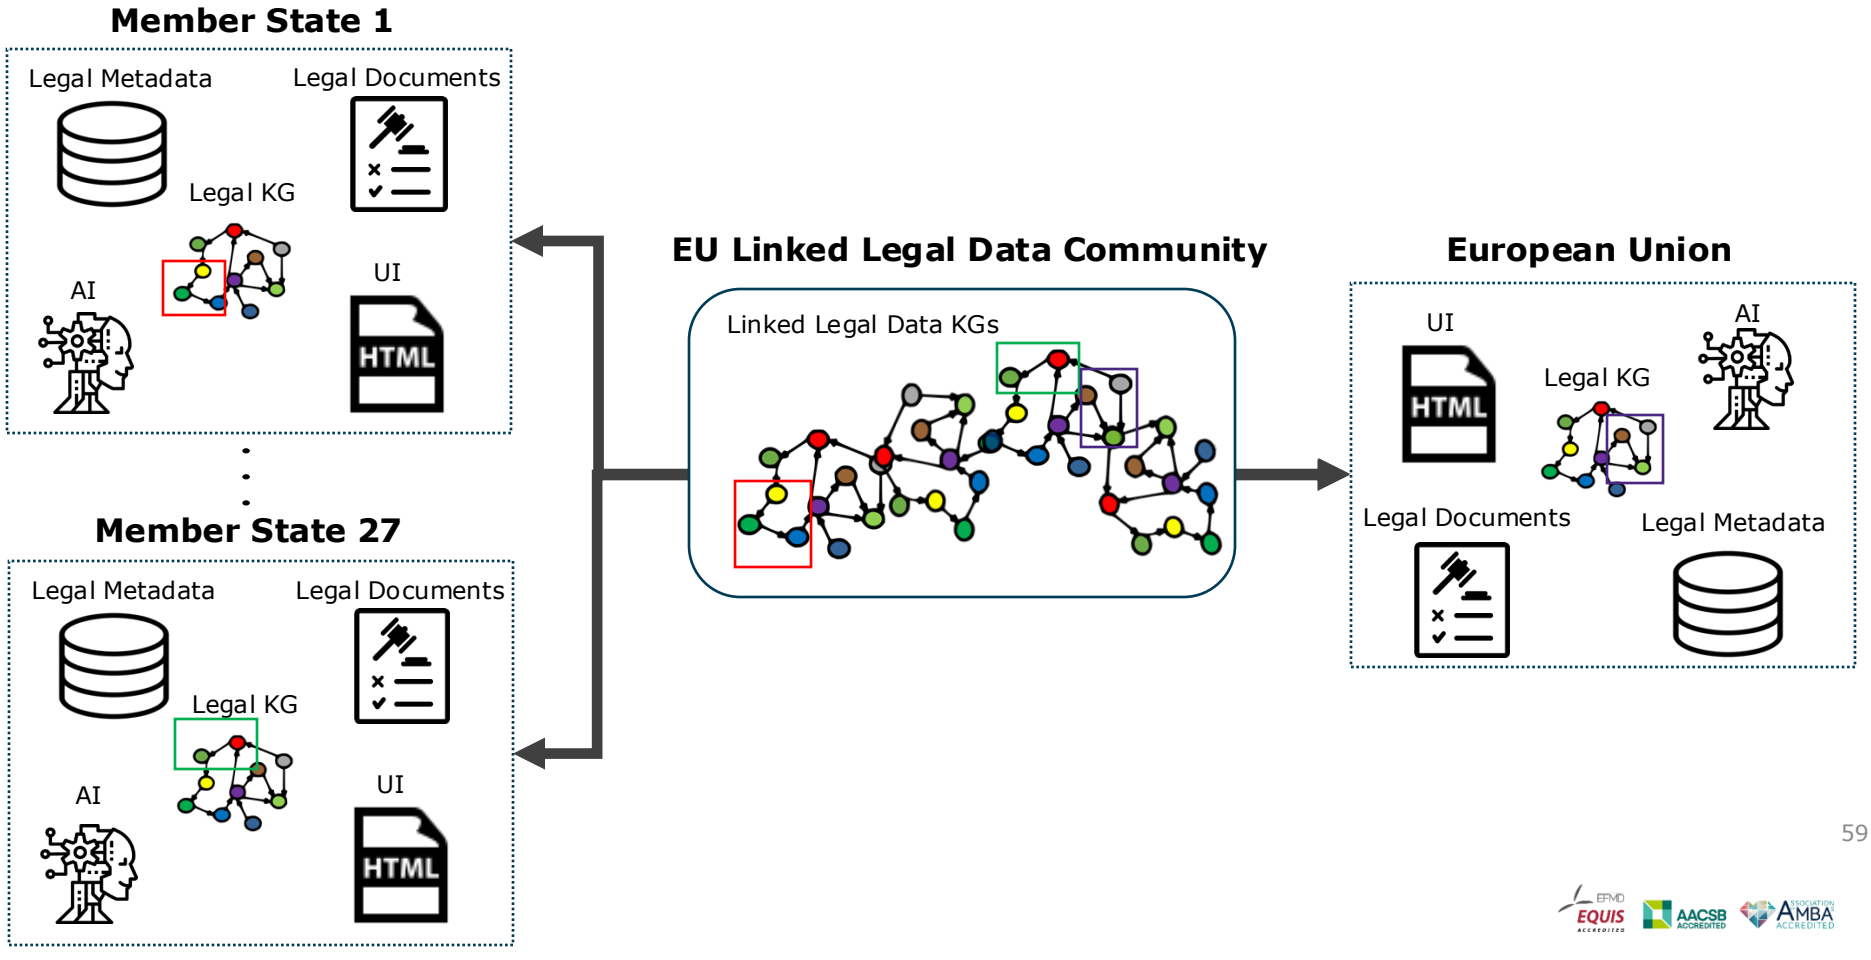

# 4 Explainable and Neuro-Symbolic AI
## EU Linked Legal Data Community

= court

= legal rule (e.g., RS0053483)

= reference to a specific article/paragraph (e.g., §1323 ABGB = civil code)

= other cases that have been decided by a court, in this snippet all of them are from the Supreme Court (e.g., 6Ob161/10k)

= literature (eg ZLB 2013/38 should be the journal "Österreichische Zeitschrift für Liegenschaftsbewertung")

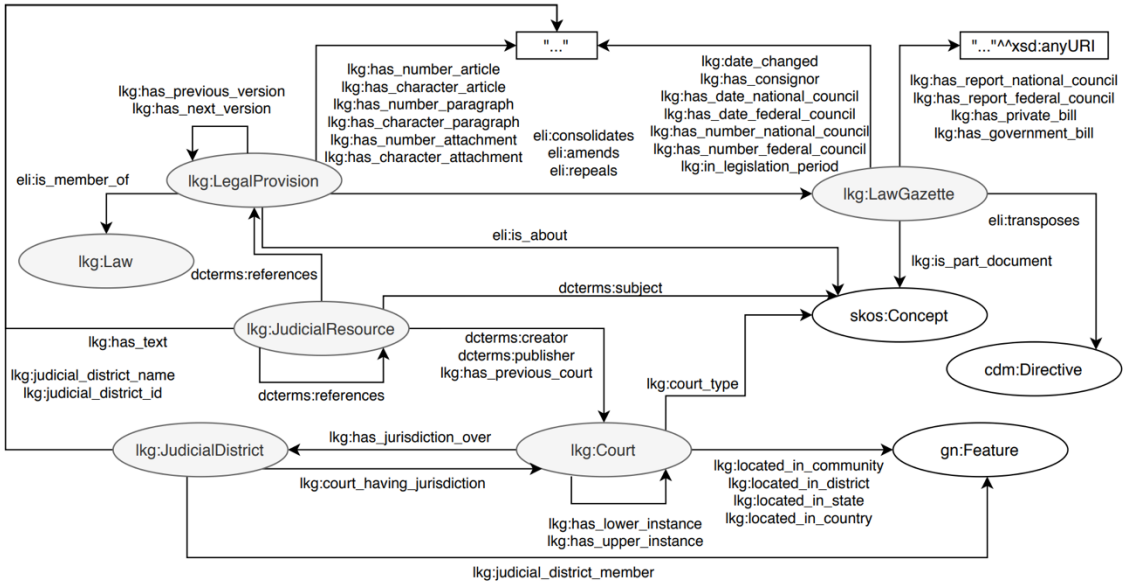


**Austrian Supreme Court Decision**

1. Gemäß § 18 Abs 1 erster Satz BStG 1971 gebührt dem Enteigneten für alle durch die Enteignung verursachten vermögensrechtlichen Nachteile Schadloshaltung (§ 1323 ABGB). Die dem Enteigneten gebührende Entschädigung muss alle durch die Enteignung verursachten vermögensrechtlichen Nachteile erfassen, wobei der Verkehrswert der entzogenen Liegenschaft den wichtigsten Faktor für dessen Bemessung darstellt (6 Ob 161/10k mwN). Auch eine nachträgliche Entwicklung der tatsächlichen Verhältnisse, die im Zeitpunkt der Enteignung als wahrscheinlich vorausgesehen werden konnte, kann die Höhe des zu ersetzenden Verkehrswerts (hier: nach § 18 BStG) beeinflussen (RIS-Justiz RS0053483). Für die Bewertung eines Grundstücks sind nach ständiger Rechtsprechung des Obersten Gerichtshofs daher neben der bestehenden Widmung auch realistisch beurteilte künftige Verwendungsmöglichkeiten samt ihrer Auswirkung auf den Marktwert entscheidend, soferne die reale Möglichkeit einer solchen Verwendung bereits im Zeitpunkt der Enteignung gegeben war und nicht bloß für eine noch unbestimmte Zukunft erhofft worden ist (RIS Justiz RS0058043 [T3]; RS0057977; RS0057981; zuletzt 3 Ob 46/11b; 1 Ob 201/13k je mwN [krit zur Einstufung als Bauerwartungsland bei bloßer Möglichkeit der Umwidmung: Kerschner , Bauerwartungsland, ZLB 2013/38]).



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

# 4 Explainable and Neuro-Symbolic AI
## Assessing Effectiveness

**50 Austrian Supreme Court decisions**

*„[...] zugunsten des obsiegenden Klägers (RIS-Justiz **RS0079624** [T14]). Ein berechtigtes Interesse des obsiegenden Beklagten an der Urteilsveröffentlichung ist dann gegeben, wenn der Rechtsstreit eine gewisse Publizität erlangt hat (RIS-Justiz **RS0079511**), etwa wenn [...]"*

*„[...] vgl. **Mayrhofer/Tangl in Fenyves/Kerschner/Vonkilch, Klang3 § 6 Abs 1 Z 2 KSchG** Rz 1 [...]"*

F-Scores

| | Approach | Case Reference | Contributor | Court | Legal Provision | Law Gazette | Legal Rule | Literature |
|---|---|---|---|---|---|---|---|---|
| Rules | Rules (lenient) | **0.9824** | 0.8426 | **0.9801** | 0.9090 | 0.9460 | **1** | 0.8674 |
| ML / DL | CRF | 0.9787 | 0.9328 | 0.9616 | 0.9459 | **0.9473** | 0.9997 | **0.8866** |
| | BERT | 0.9712 | **0.9583** | 0.9616 | 0.9489 | 0.9396 | 0.9986 | 0.8448 |
| | DistilBERT | 0.9772 | 0.9551 | 0.9586 | **0.9521** | 0.9437 | 0.9989 | 0.8626 |
| Δ | | 0.0112 | 0.1157 | 0.0215 | 0.0431 | 0.0077 | 0.0014 | 0.0418 |



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

# 4 Explainable and Neuro-Symbolic AI
## Extraction of Events from English Legal Text

THE FACTS
I. THE CIRCUMSTANCES OF THE CASE
5. The applicant was born in 1965 and lives in Szajol.
6. On 6 October 1990 he married Ms N.R. The couple had one daughter, R.Ó, born on 23 March 1994.
7. On 24 March 2006 the couple divorced and agreed on custody of the child and other parental rights. Their agreement, which was approved by the Szolnok District Court, placed the daughter with her mother and granted the applicant contact every other weekend, during the school holidays, Easter Sunday, Whit Monday, the first holiday following Christmas Eve and three weeks during the summer holidays. The applicant was also to pay child allowance, amounting to 20% of his salary, but at least 20,000 Hungarian forints (HUF approximately 66 euros (EUR)) per month.
8. Ms N.R. moved out of the former family home and took up residence in Budapest.
9. The applicant could exercise contact on 25 and 26 June and 24 July 2006. However, his daughter was reluctant to stay with him during the summer holidays and preferred to spend her holidays with her grandparents.
10. The applicant's further attempts to exercise contact failed, seemingly because in the mother's view it was up to the child to decide whether she wanted to see her father. The applicant applied to the Szolnok County Guardianship Authority to have the contact agreement enforced. In a decision of 11 September 2006 it ordered Ms N.R. to comply with the agreement, warning her that failure to do so could result in her being fined and ordered to reimburse the applicant's travel costs. Given the strained relationship between the parents, a notary initiated child protection proceedings ( védelembe vétel ) on 27 December 2006. Subsequently, on a number of occasions Ms N.R. was fined for not complying with the agreement. It appears from the case file that the applicant could exercise contact between 10 July 2007 and 10 November 2008.
11. In 2007 Ms N.R. brought an action seeking to change the applicant's contact rights and an increase in the amount of child allowance. In a counterclaim the applicant requested that the child be placed in his custody.
12. The applicant also filed a criminal complaint against Ms N.R. alleging child endangerment. The parties' attempt to stabilise the relationship between the applicant and his daughter during the suspension of the custody proceedings were unsuccessful and the Szolnok District Court ordered that any meetings between them take place at a child protection centre.
13. In a judgment of 10 June 2008 the Szolnok District Court reduced the applicant's contact with his daughter to every first and third Saturday of the month from 9 a.m. to 6 p.m. It held that the previously agreed form of contact was unlikely to be implemented and would only lead to further proceedings before the guardianship authority, which would be to the detriment of the child. The court based its judgment on an expert opinion finding that the child should have had an adaptation period to re-establish her relationship with her father. The court dismissed the applicant's claim concerning custody, stating that the child's wishes had to be taken into account, given that she was now fourteen years of age.
14. On 19 November 2008 the Jász-Nagykun-Szolnok County Regional Court upheld the first-instance judgment in essence but amended the applicant's contact rights to every other Saturday between 9 a.m. and 4 p.m. until 31 May 2009, and all weekend-long visits every other weekend as of 1 June 2009.
15. In 2009 the applicant failed to turn up at numerous scheduled meetings for months, for which he was fined HUF 10,000 (approximately EUR 35).

- European Court of Human Rights (ECHR)
- Quick overview over case
  - Dates (“when“)
  - Subjects (“who“)
  - Core (“what“)



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.

# 4 Explainable and Neuro-Symbolic AI
## Assessing Effectiveness

**30 court decisions of European Court of Human Rights (ECHR)**

*„The applicant **refused to see his son** in June 1999."*

*„On 15 January 2000 the District Court **upheld the Judgment**. "*

F-Scores

| | Approach | What | When | Who |
|---|---|---|---|---|
| Rules | Rules (lenient) | 0.2195 | 0.6977 | 0.6857 |
| ML/DL | CRF | **0.8050** | 0.8658 | 0.7834 |
| | BERT | 0.6583 | **0.9022** | **0.9044** |
| | DistilBERT | 0.6237 | 0.8823 | 0.8998 |
| Δ | | 0.5855 | 0.2045 | 0.2187 |



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.

# *Tackling Information Disorders*

## Seven types of mis- and disinformation

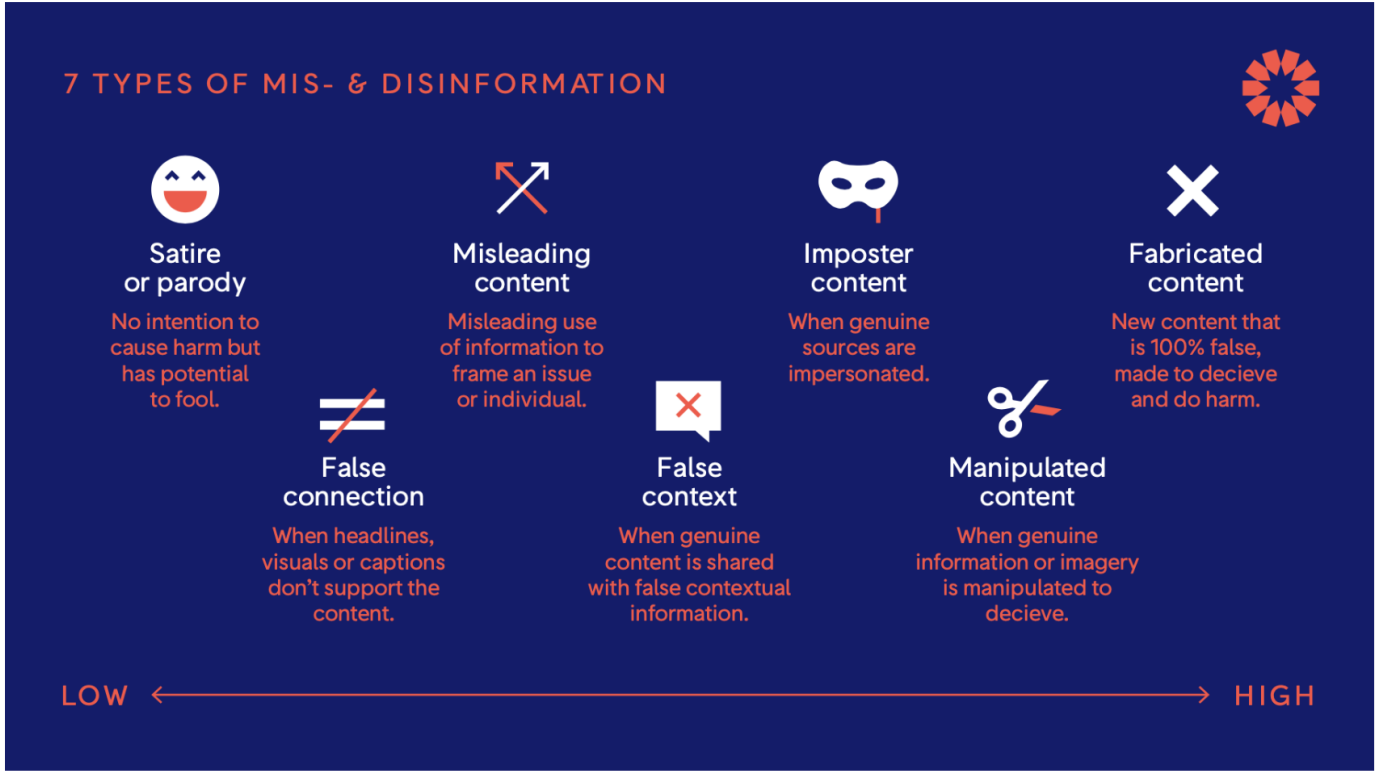


**Claire Wardle**
https://firstdraftnews.org/long-form-article/understanding-information-disorder/

# 4 Explainable and Neuro-Symbolic AI
## Information Disorder Detection

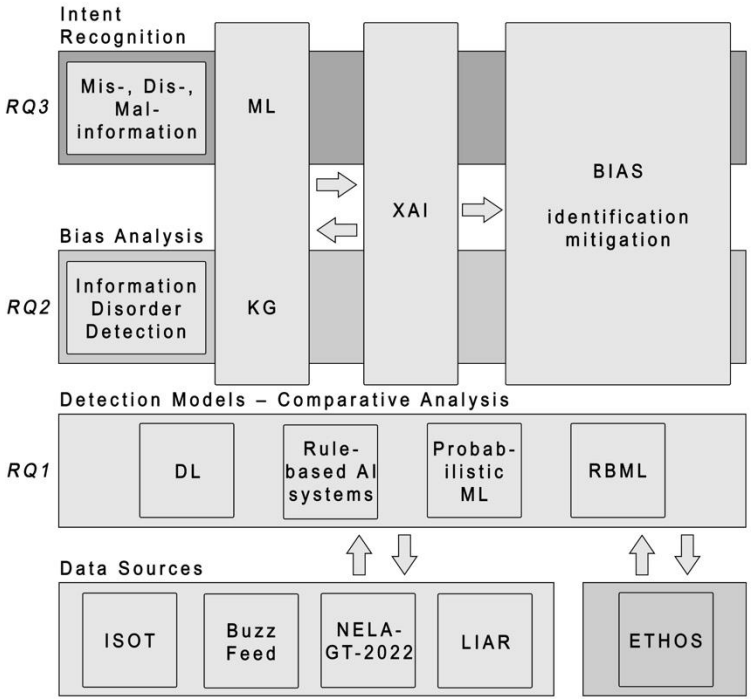


**RQ1** Which machine and deep learning algorithms are most effective when it comes to information disorder detection?

**RQ2** To what extent can machine and deep learning explainability and bias detection be facilitated via knowledge graph-based enrichment?

**RQ3** How can hybrid AI-based approaches be used to better distinguish between the three information disorder types (misinformation, disinformation, and malformation)? 4

# Bilateral AI
## Combining Symbolic and Sub-symbolic AI

BilAI

### BilAI Consortium

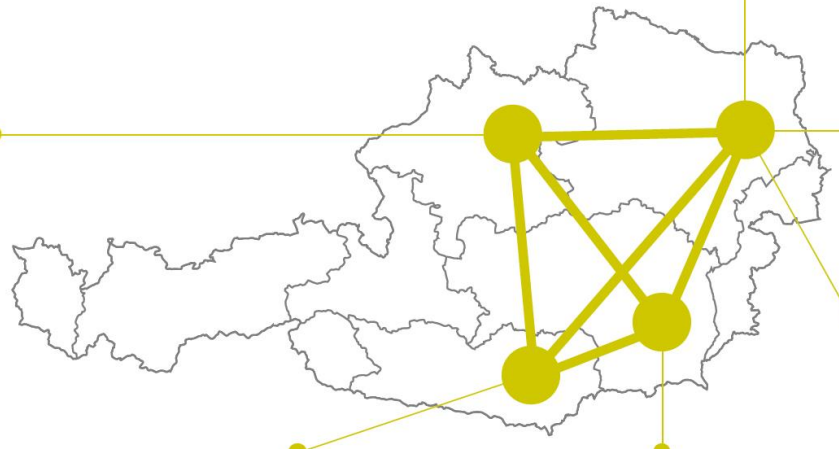

**Martina Seidl**
Symbolic AI
SAT Solving
Formal methods

**Sepp Hochreiter**
Machine Learning
LSTM
Vanishing gradient

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

**Gerhard Friedrich**
Symbolic AI
Model-based reasoning

- Institute for Artificial Intelligence and Cybersecurity

**Robert Legenstein**
Machine Learning
Computational Neuroscience

- Institute of Machine Learning and Neural Computation

**Christoph Lampert**
Machine Learning
Trustworthy Learning

- Machine Learning and Computer Vision group
- ELLIS Unit ISTA

**Agata Ciabattoni**
Logic Reasoning

**Thomas Eiter**
Symbolic AI
Knowledge representation

- Institute for Logic and Computation

**Axel Polleres**
Knowledge Graphs

- Institute for Data Process and Knowledge Management

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

# Bilateral AI
## Combining Symbolic and Sub-symbolic AI

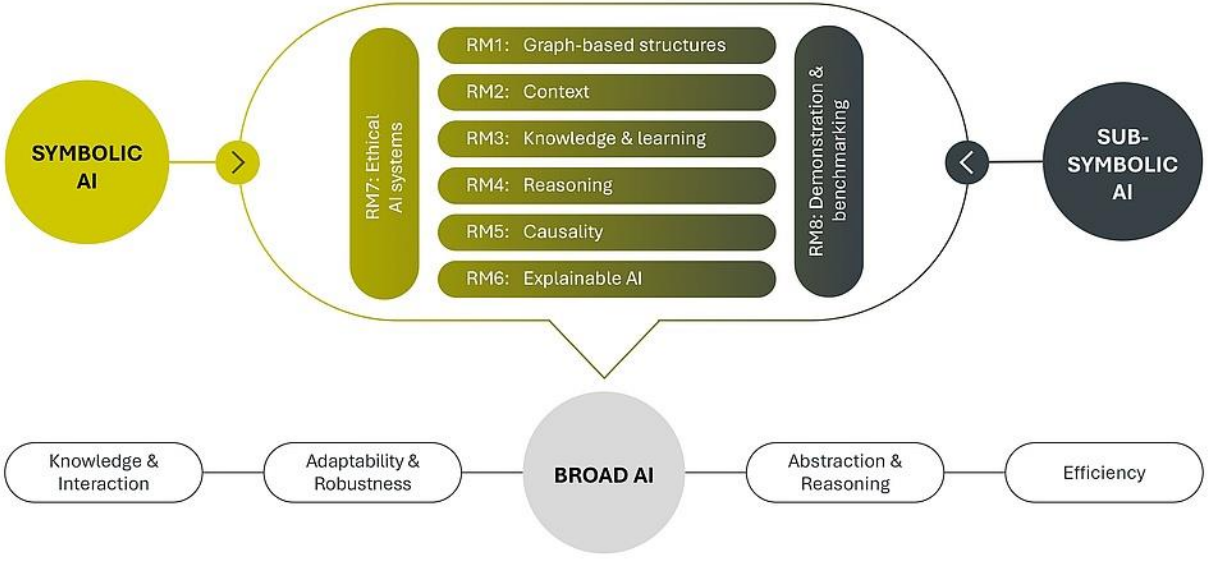


- Providing legal and ethical foundations for AI technologies and systems, via machine readable policies and norms
- Policy aware federated querying and learning
- Developing normative autonomous agents in general and autonomous web agents in particular

# Thank you / contact details

*This research was funded in whole or in part by the Austrian Science Fund (FWF) 10.55776/COE12.*

VIENNA UNIVERSITY OF ECONOMICS AND BUSINESS

**Department of Information Systems & Operations**
Institute for Information Systems & New Media
Welthandelsplatz 1, 1020 Vienna, Austria

**Dr. Sabrina Kirrane**

T +43-1-313 36-4494
F +43-1-313 36-90 4494
sabrina.kirrane@wu.ac.at
www.wu.ac.at
www.sabrinakirrane.com

The PENNI project is funded by the FWF Austrian Science Fund and the Internet Foundation Austria under the FWF Elise Richter and netidee SCIENCE programmes as project number V 759-N.

The project SPECIAL (Scalable Policy-awarE linked data arChitecture for prIvacy, trAnsparency and compLiance) has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731601 as part of the ICT-18-2016 topic Big data PPP: privacy-preserving big data technologies.