

RM1: Graph-based structures

Marta Sabou, Axel Polleres
Inst. for Data, Process & Knowledge Management, WU Wien

contributions:
Robert Ganian, Silvia Miksch,
Magdalena Ortiz, Emanuel Sallinger (TU Wien)
Elmar Kiesling (WU Wien)



Funded by:
FWF Austrian Science Fund



bilateral-ai.net





Agenda

- Recall initial RM1 goals...
- ... & what happened since & BILAI “Lighthouses”
- Reported activities
- Discussion about further collaboration potential





RM1: Graph-based structures

Proposal:

RM1 Summary: “Graph-based structures are highly relevant to all the essential properties of a broad AI. [...] navigation **maps**, [...] social or physical interaction **networks**, [...] **object relations**. [...] Because of the **inherent symbolic nature** and their **suitability for learning and sub-symbolic** elements, they naturally constitute a promising starting point as a core component for a bilateral AI [...].”

Essential properties (EP) of Broad AI:

EP1 Knowledge Transfer and Interaction. [...]

A broad AI should be able to integrate existing knowledge from other AI systems or humans. [...]

A broad AI system should be able to interact with other AI systems and humans. [...]

RM1.1: Graph representations: scaling by semantic and visual abstractions

(O1.1) monitoring and understanding the schema development of collaboratively developed KGs,

(O1.2) combining KGs and large corpora of structured raw data,

(O1.3) leveraging and extracting constraint knowledge and automated quality improvement of KGs, and

(O1.4) designing Visual Analytics methods for dynamic graphs to support collaborative exploration and curation of KGs

RM1.2: Understanding and improving learning on graph structures

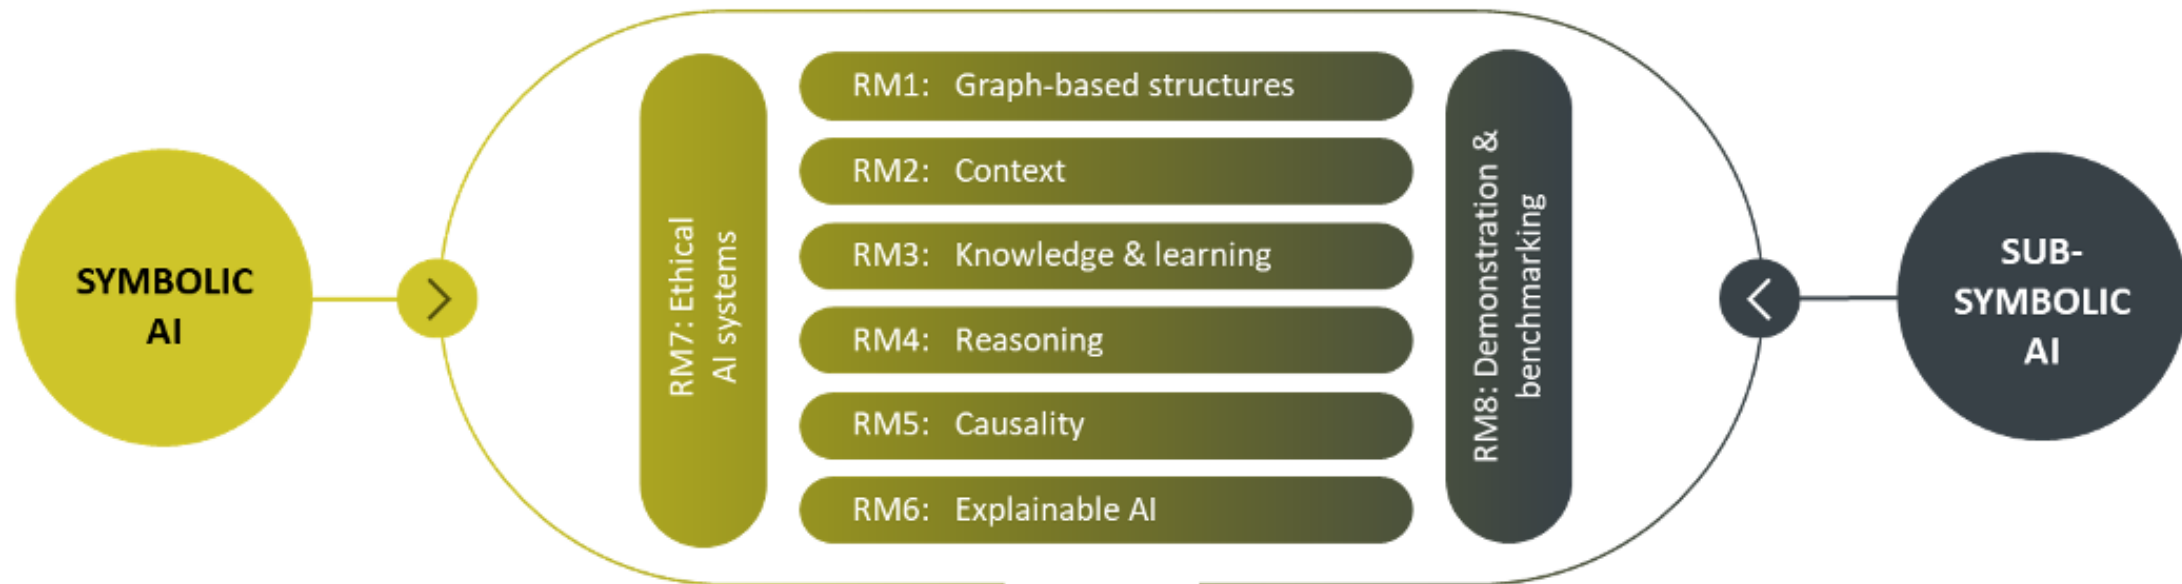
(O2.1) [...] exploiting symbolic methods and advanced tools from the theoretical analysis of deep neural networks and high-dimensional inference problems, and

(O2.2) [...] identifying reasonable subgraph-abstractions and (semantic) modularizations of graphs to improve learning of GNNs under distribution shifts



Bigger Picture: overall BILAI RMs & “Lighthouse” Proposals :

- “BILAI LLM”
- “Agentic BILAI”
- “BILAI Solver”
- (“BILAI Knowledge Graph”?)





What's next?



Neurosymbolic AI Systems


Prof. Marta Sabou




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

Atomic-Patterns


A1: s-M-s



A2: d-M-s




A3: s-M-d




T-Patterns

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




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





I-Patterns

I1: s-M-d-M-s



I2: s-M-d-M-d




Ongoing extension: work on *Systematic understanding of neuro-symbolic knowledge engineering* (Fajar Ekaputra)


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.



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



T6: {d-M-d/s}-M-d



Data Structures	Algorithmic Modules
s symbolic	M inductive (ML)
d data	K deductive (KR)

Marta Sabou (WUW)

- **RM5** Collaboration with Lukas Kook on
 - (1) acquiring causal knowledge through the combination of neural/statistical models; and
 - (2) using symbolic formalism to represent, synthesise and analyse causal knowledge derived by different causal learning methods.
- **RM8** KG-embedding based knowledge infusion methods into use cases from marketing with Niels Wlömert and his colleagues (early stage);
- “**DORSET**” new WWTF project with Emanuel Sallinger & Katja Hose (TUW):
 - impact of knowledge graph characteristics on the performance of embedding methods
 - meaning for current knowledge engineering practices



Magdalena Ortiz (TUW)

(1) Logic-based explanations for GNNs

Sufficient and counterfactual notions of “motif-based formal explanations for graph classifiers expressed as logical formulas. key ingredient: background knowledge to describe the space of relevant graphs to be explored. We consider both plain minimal (sufficient/counterfactual) motifs, and also motifs with respect to background knowledge.

Collaboration with Sagar Malhotra from the Machine Learning group at TU Wien (expert in explainability for GNNs). publication planned by Summer!

Distillation of trained GNNs into logical formulas is a key ingredient, and we would like to have more input/collaborations.

(2) Learning SHACL shapes from examples.

Collaboration with an external partner (symbolic, theory-centric). We already have plans to integrate this with their more practical work on PAC learning for SHACL shapes, but are still trying to understand how to connect it to sub-symbolic learning approaches in meaningful ways.

Currently working on a workshop paper, and hope to submit a conference version before the summer.



Silvia Miksch (TUW) – Ongoing:

(+) **Raphael Buchmüller** started as PhD candidate: 1.12.2025

(<https://www.cvast.tuwien.ac.at/team/raphael-buchmueller>)

(+) work in the scope of graphs/networks, Visual Analytics, and generative/agent AI : Raphael has one WS publication [2] and 3 submissions [1, 3, 4]



[1] IEEE VIS Full Paper (submitted): Raphael Buchmüller, Dennis Collaris, Angelos Chatzimpampas . Evaluating LLM Summaries for Projection Explanation.

[2] VisxAI Workshop at Eurovis Conference 25: Raphael Buchmüller, Dennis Collaris, Linhao Meng, Angelos Chatzimpampas. LangLasso: Interactive Cluster Descriptions through LLM Explanation

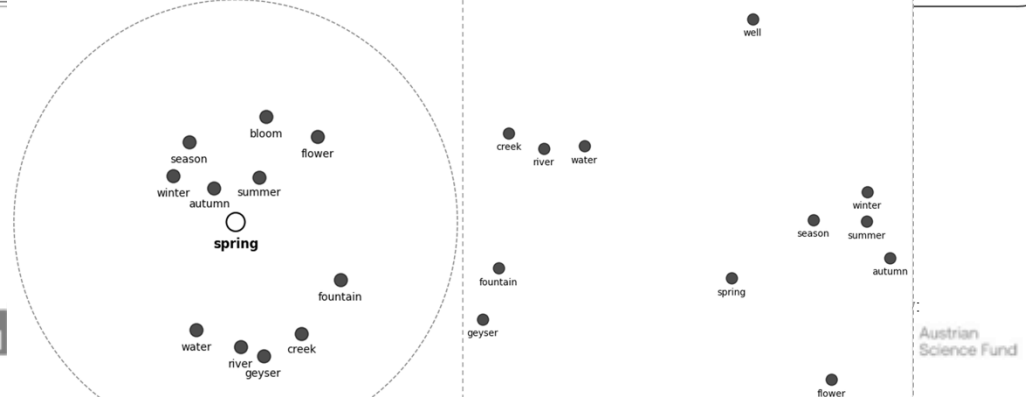
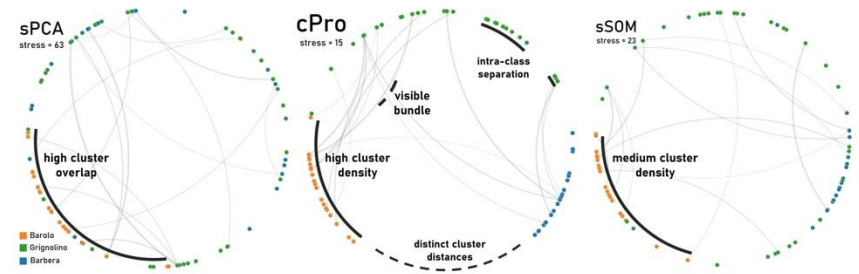
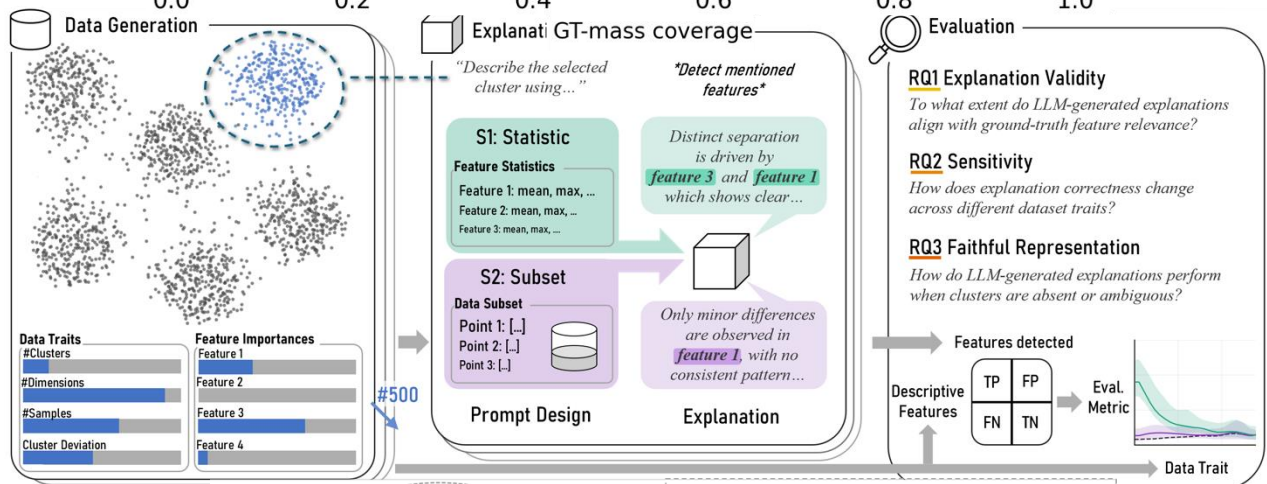
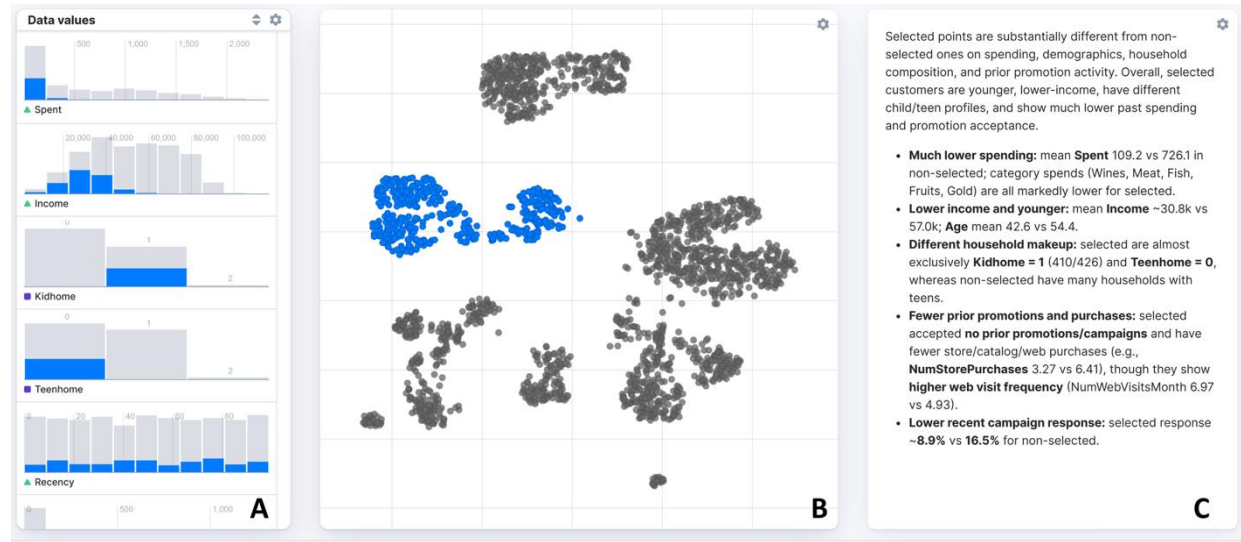

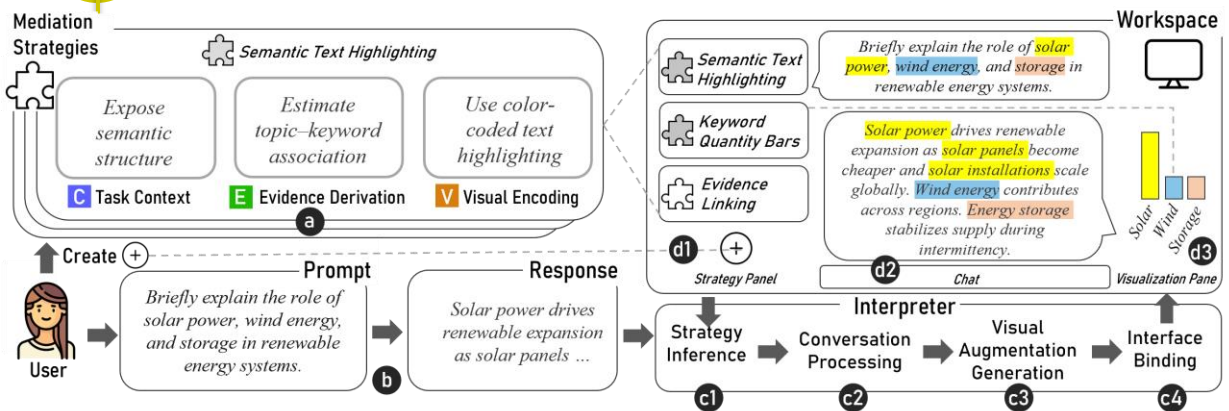
[3] EuroVA Workshop at Eurovis Conference 26 (submitted): Raphael Buchmüller. Evaluating LLM Summaries for Projection Explanation

[4] TVCG Journal - Full Paper (submitted): Raphael Buchmüller, Bastian Jäckl, Michael Behrisch, Daniel A Keim, Frederik L Dennig. Evaluating Circular Projections for High-Dimensional Data Visualization

(+) collaboration with TU Graz (Elisabth Lex Group) and JKU (Markus Schedl) : Visual Analytics and Recommender Systems taking a bilateral perspective”



Silvia Miksch (TUW) – Ongoing:



Silvia Miksch (TUW) - Planned

(B) What I know is planned next

=====

- (+) proceed with the collaboration with TU Graz (Elisabth Lex Group) and JKU (Markus Schedl)
- (+) Visual Analytics for graph representation learning (in collaboration with [Natkamon Tovanich](#) and [Velitchko Filipov](#) (CVASt group))
- (+) Visual Analytics to support exploring and curating of Knowledge Graphs

(C) facilitate discussion on further possible collaborations

=====

- (+) Visual Analytics support for Knowledge Graphs
- (+) Visual Analytics for graph representation learning (also known as graph embedding)
- (+) Neuro-symbolic approaches and Visual Analytics
- (+) Causal representation and learning and Visual Analytics



Robert Galian (TUW, also active in RM5)

New symbolic-based algorithmic techniques for subsymbolic computational problems.

- 1) Exact algorithms + complexity lower bounds for training quantized neural networks
Robert Galian, Frank Sommer, Manuel Sorge: Tractability via Low Dimensionality: The Parameterized Complexity of Training Quantized Neural Networks. **ICLR 2026**
(<https://iclr.cc/virtual/2026/poster/10010967>)
- 2) Computational complexity of training graph neural networks:
Robert Galian, Mathis Rocton, Simon Wietheger: Training One-Dimensional Graph Neural Networks is NP-Hard. **ICLR 2025**
(<https://openreview.net/forum?id=7BESdFZ7YA>)
- 3) Algorithms + lower bounds for non-clashing teaching (restriction to graphs is w.l.o.g.):
Robert Galian, Liana Khazaliya, Fionn Mc Inerney, Mathis Rocton: The Computational Complexity of Positive Non-Clashing Teaching in Graphs. **ICLR 2025**
(<https://openreview.net/forum?id=Jd3Vd7GCyq>)



(1) "Neuro-symbolic Query Answering on Knowledge Graphs"

Paolo Atzeni, Teodoro Baldazzi, Luigi Bellomarini, Eleonora Laurenza, Emanuel Sallinger: Semantic-aware query answering with Large Language Models. *Data Knowl. Eng.* 161: 102494 (2026)

(2) "Neuro-symbolic Explanations in High-Stakes Financial KGs"

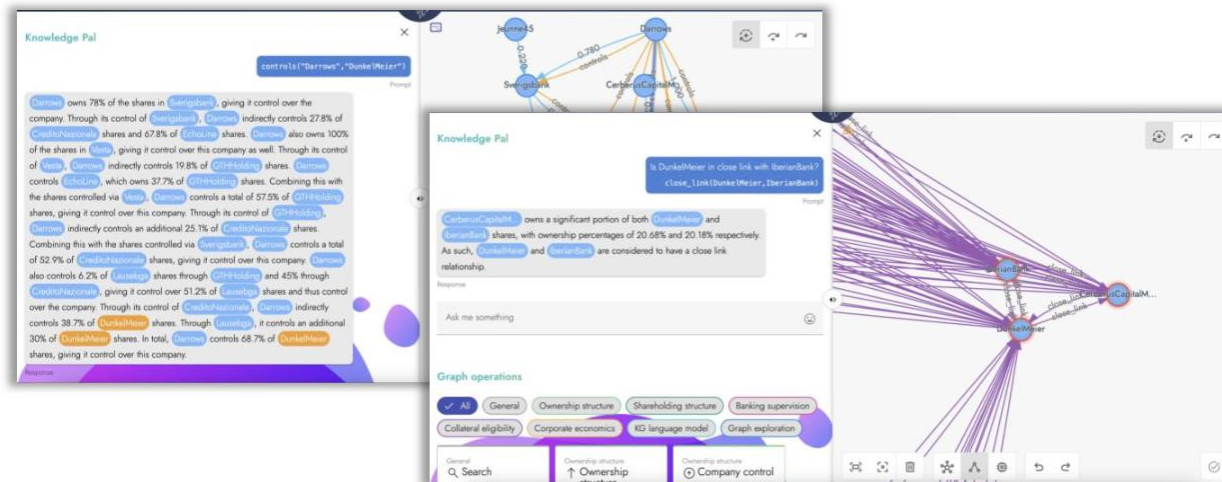
Andrea Colombo, Teodoro Baldazzi, Luigi Bellomarini, Emanuel Sallinger, Stefano Ceri: Template-based Explainable Inference over High-Stakes Financial Knowledge Graphs. *EDBT 2025*: 503-515

(3) "Please, Vadalog, tell me why" - Using LLMs for Explanations

Teodoro Baldazzi, Luigi Bellomarini, Stefano Ceri, Andrea Colombo, Andrea Gentili, Emanuel Sallinger: "Please, Vadalog, tell me why": Interactive Explanation of Datalog-based Reasoning. *EDBT 2024*

(4) "KGEs and GNNs to Capture Logical Rules"

Aleksandar Pavlovic, Emanuel Sallinger, Steven Schockaert: Faithful Differentiable Reasoning with Reshuffled Region-based Embeddings. *KR 2025*



Axel Polleres, Elmar Kiesling (WUW)



1. **Miguel Vázquez**, Kevin Innerebner, Alexander Prock, Günter Klambauer, Elisabeth Lex, Johannes Schimunek, and Axel Polleres. Structure is the signal: Graph encodings and GNNs for constraint repair in collaborative KGs. In *23rd European Semantic Web Conference (ESWC)*, May 2026. to appear. [[.pdf](#)] → **separate presentation!** 😊
2. **Nicolas Ferranti**, Dayane Guiamaraes, Jairo Francisco de Souza, and Axel Polleres. Who can fix this? user recommendation for knowledge graph repair via embedding-based clustering. In *Proceedings Of The 13th Knowledge Capture Conference 2025 (K-CAP 2025)*, December 2025. [[DOI](#)] | [[.pdf](#)]
3. **Ongoing**: *Graph Sampling for evaluating Graph Embeddings in Evolution* (Diego Rincon-Yanez, Elmar Kiesling)

“Agentic BILAI”:

4. Axel Polleres, **Frederik Bauer**, Daniil Dobriy, Tobias Käfer, Lukas Kubelka, Andreas Harth, and Thomas Wehr. On the historic roots of Agentic AI in Semantic Web Services. In *The ACM Web Conference 2026*, 2026. to appear, extended abstract (invited, History of the Web track). [[.pdf](#)]
5. Andreas Ekelhart, Kabul Kurniawan, Fajar Ekaputra and Elmar Kiesling. AgentO: An Ontology for Modeling Agentic AI Systems. In *23rd Extended Semantic Web Conference (ESWC)*, May 2026. to appear.
6. **Daniel Dobriy**, Frederik Bauer, Amr Azzam, Debayan Banerjee, and Axel Polleres. Agentic SPARQL: Evaluating SPARQL-MCP-powered intelligent agents on the federated KGQA benchmark. (submitted, [Tech Report](#) available)
7. Axel Polleres, Diego Rincon-Yanez, Amin Anjomshoaa, Daniel Dobriy and Emanuel Sallinger. From Data Spaces to Agent Spaces (submitted)

“BILAI LLM”:

Discussions with Anamaria Roberta-Hartl & Günter Klambauer to leverage BILAI LLM in works continuing on KG Repairs [1,2], or Agentic SPARQL

other Synergies:

FAIR-AI project (FFG): work on (symbolic) AI risk modeling (Elmar Kiesling)



Axel Polleres (WUW) – SPARQL-MCP BILAI



search.ai.wu.ac.at

SEARCH ENGINE FOR THE WEB OF DATA

Search Catalogue Federated Wiki Documentation About Submit Endpoint

Global Federation Active

Search Engine for the Web of Data

Search Engine powered by Federated Agentic SPARQL over Public Endpoints and Wikis

Agent Thinking & Executing... Change Provider New Task

USER REQUEST
Who is the director of the Bilateral AI project?

SPARQL-MCP STARTUP SPARQL-MCP
Starting a new SPARQL-MCP session...

TECHNICAL STEP SYSTEM
Loading endpoint catalogue...

Project lead: Daniil Dobriy (BILAI PhD)

<https://search.ai.wu.ac.at/>

Automatic Catalog generation
Federated query generation/optimization
Benchmark

Ongoing/Planned:
Collaboration with La Sapienza (R. Navigli) on KG-based
Question answering
Collaboration with Ugent (R. Taelman) on Federated
query Optimization and reducing Server load

**Ambitious goal:
"Game-changer" application!
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 *travelling* to her home and with a *rating* of excellent or very good on trusted rating services. It then began *travelling* to her home and with a *rating* of excellent or very good on trusted rating services.



Axel Polleres (WUW) – “Agent Spaces”

Vision/Position paper under submission for ESWC workshop, together with Emanuel Sallinger (submitted)

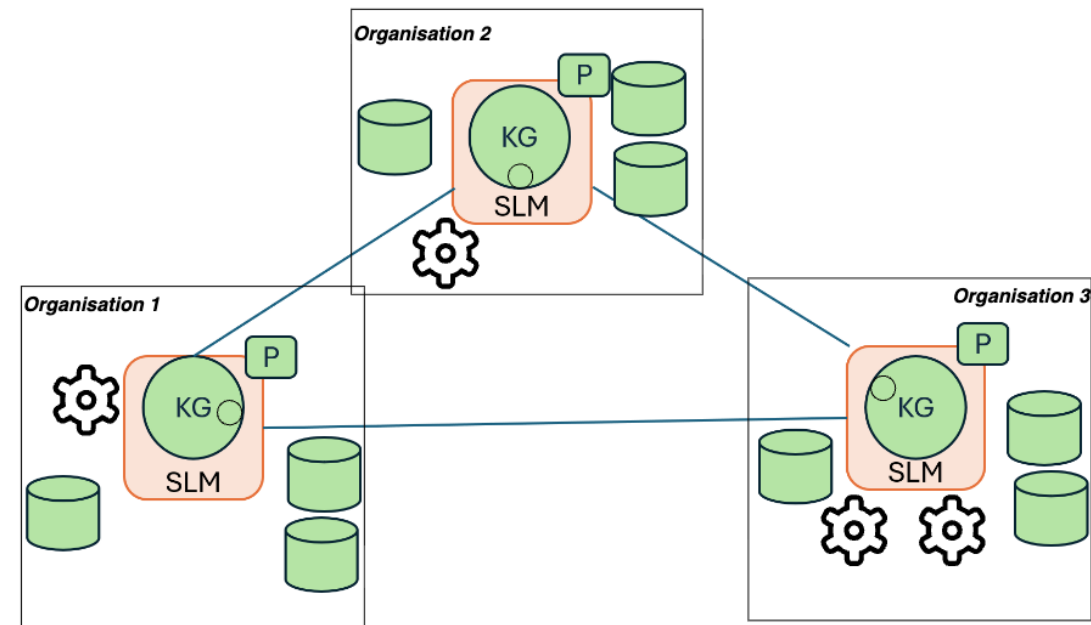
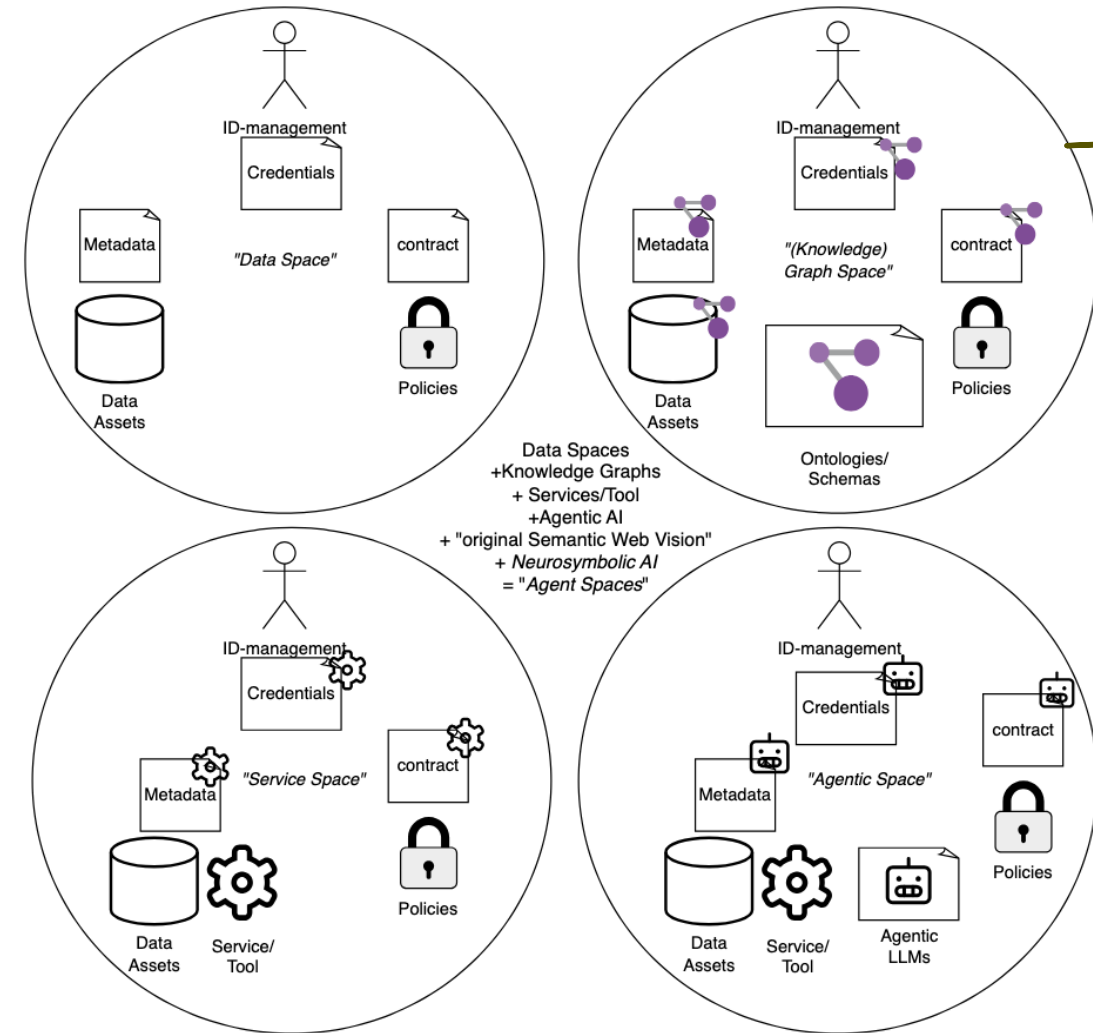


Figure 2: An Agent Space with Linked Knowledge Graphs (KGs), wrapped by a sovereign, agentic language that is governed by symbolic policy guardrails (P).

From Data Spaces to Agent Spaces*

Axel Polleres^{1,2,*,†}, Diego Rincon-Yanez^{1,*,†}, Amin Anjomshoa^{1,†}, Daniel Dobriy^{1,*,†} and Emanuel Sallinger^{3,†}

Further announcements:

- ESSAI summer school lecture
- ISWC2026 all tracks (Research and Resource) open for submission!
workshops (with BILAI involvement): **KG-NeSy, RAGE-KG, Wikidata Workshop**
- ***New BILAI Praedoc position forthcoming (with Svitlana Vakulenko, WWTF VRG, new BILAI KR!)***



The European Summer School on Artificial Intelligence

6 - 10 July 2026, Vienna, Austria



5. Wikidata: A backbone for Hybrid/Bilateral AI




Lecturers: Axel Polleres (WU Wien) , Diego Rincon-Yanez (WU Wien)

Course Type: Introductory

Keywords: Knowledge Representation and Reasoning, Knowledge Graphs, Bilateral AI

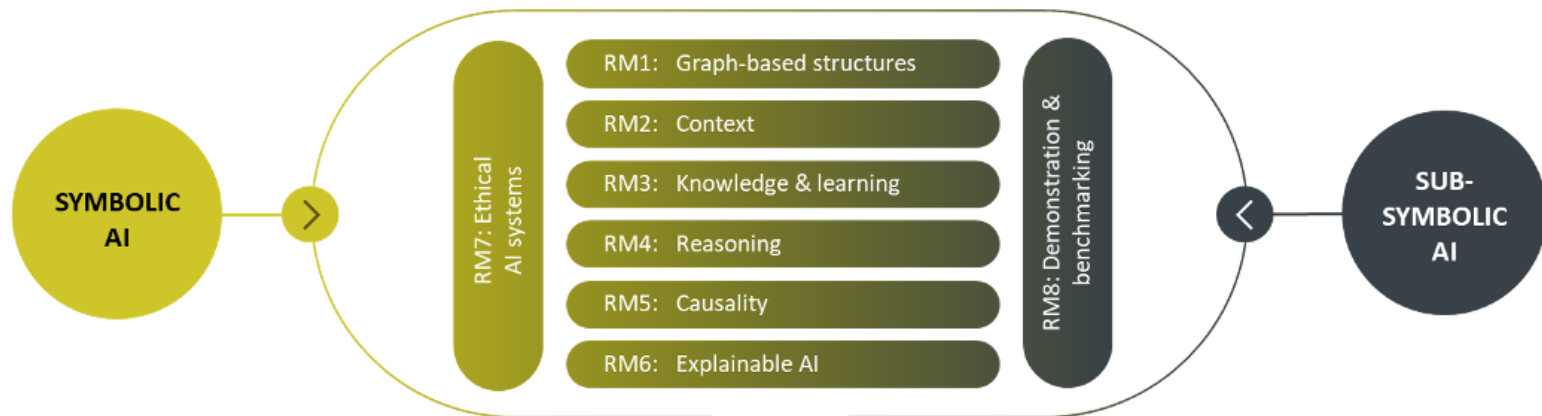
Abstract: The importance of structured, accessible knowledge in Knowledge Graphs like Wikidata is paramount in AI applications, as a basis for reliable, curated facts. This introductory course will introduce you to how to use and leverage large-scale collaboratively edited Knowledge Graphs such as Wikidata, both for driving "bilateral" (aka "hybrid", i.e. combining symbolic and sub-symbolic) AI applications on the one hand, but also in terms of discussing how hybrid/bilateral AI can help to build and consolidate such large collections of structured knowledge. We will focus on a hands-on introduction and examples.



Discussion

- Recall initial RM1 goals...
- ... & what happened since & BILAI “Lighthouses”
- Reported activities
- **Discussion about further collaboration potential**

- *BILAI LLM*
- *“Agentic BILAI”*
- *“Bilateral” Optimizer*
- *(“BILAI Knowledge Graph”?)*



Bigger Picture: overall BILAI RMs & “Lighthouse” Proposals :

- “BILAI LLM”
- “Agentic BILAI”
- “BILAI Solver”
- (“BILAI Knowledge Graph”?)

