European Commission

Large Language Models to support semantic interoperability

Workshop – 10 April 2025

Written jointly by the SEMIC and BDVA teams
May 2025

interoperable europe

BDV BIG DATA VALUE ASSOCIATION

EUROPEAN COMMISSION

Directorate-General for Digital Services
Directorate B — Digital Enablers & Innovation
Unit B.2 — Interoperability and Digital Government

Contact: The SEMIC Team

E-mail: DIGIT-SEMIC-TEAM@ec.europa.eu

European Commission
B-1049 Brussels

Large Language Models to support semantic interoperability

Workshop – 10 April 2025

Directorate-General for Digital Services
Semantic Interoperability Initiative

2025

Manuscript completed in May 2025

Contents

1. Introduction and motivation

On 10 April 2025, a joint workshop titled “Large Language Models to Support Semantic Interoperability” was organised by the Big Data Value Association (BDVA) and the SEMIC (Semantic Interoperability Community) initiative of the European Commission’s DG DIGIT.

This workshop aimed to explore the state of the art in techniques, tools, and applications leveraging Large Language Models (LLMs) to enhance semantic interoperability, including real-world use cases where LLMs facilitate data models alignment, semantic enrichment, and cross-domain interoperability. Beyond mapping advancements in the field, the workshop seeks to identify main challenges and issues, as well as key actors and communities working on them, fostering collaboration and establishing a knowledge-sharing group dedicated to LLM-driven semantic interoperability

The workshop was built around a call for contributions to researchers, practitioners, and industry experts to contribute with short presentations showcasing how Large Language Models (LLMs) enhance semantic interoperability, ensuring that data exchanged between systems is correctly interpreted and effectively used. The workshop explored techniques, tools, and real-world applications that improve the accuracy, automation, and scalability of semantic interoperability processes. Topics of interest included, but were not limited to:

- Advanced semantic matching and alignment
- Intelligent metadata generation and enrichment
- Multilingual semantic mapping and cross-language understanding
- Automated data annotation and structuring
- Integration of LLMs and knowledge graphs for enhanced reasoning
- Ontology enhancement and adaptive knowledge representation

This workshop aims to be just the starting point of a broader initiative to build a community around the topic, bringing together researchers, experts and industry to identify challenges, opportunities and next steps.

2. Workshop organisation

2.1. Organising entities

Big Data Value Association (BDVA): A European industry-driven organisation focused on boosting Data and AI innovation and data value creation for business, citizens and the environment, developing an innovation ecosystem that enables and accelerates data and AI economy according to European values.

SEMIC: A programme led by the European Commission's Directorate-General for Digital Services (DIGIT), fostering semantic interoperability through specifications, guidelines, and community efforts.

Table 1 – Organising team members

Name	Organisation
Daniel Alonso	BDVA
Marcello Gritta	SEMIC
Claudio Baldassarre	SEMIC
Juan Alvarez	SEMIC
Ine Weyts	SEMIC
Emilien Caudron	SEMIC

2.2. Selection process for workshop speakers

For the workshop, the organising team invited researchers, practitioners, and industry experts to submit short presentations (10 minutes each) demonstrating how Large Language Models (LLMs) can enhance semantic interoperability—ensuring that data exchanged between systems is accurately interpreted and effectively utilised.

A call for contributions was open from 25 February to 25 March. Interested contributors were asked to provide their name, affiliation, and a brief abstract outlining their proposed presentation via a designated survey that received 19 answers.

The SEMIC and BDVA teams reviewed all submissions by 31 March. A total of 10 presentations were selected based on their relevance to the workshop theme and the significance of the content. The workshop aimed to explore techniques, tools, and real-world applications that improve the accuracy, automation, and scalability of semantic interoperability processes. Presentations not selected are being considered for future events related to this topic.

2.3. Format and logistics

The workshop was held online via Webex on 10 April 2025 from 13:00 CET until 16:00 CET. Participation was open to experts from industry, academia, EU institutions, and various data-related communities, with over 160 participants involved.

The workshop combined keynote presentations, short technical talks, and interactive discussions. After every presentation or talk, the respective speaker presented the audience with two statements that participants could respond to on a scale from strongly disagree to strongly agree. The presentations were divided into two main segments, each concluding with time allocated for discussion of the statements and the ratings. During these discussions, participants were encouraged to elaborate on their opinions, providing an opportunity for clarification and deeper exploration of varying viewpoints. The detailed agenda of the workshop can be found in Table 22 – Agenda of the .

Table 22 – Agenda of the workshop

Time stamp	Name	Organisation	Presentation title
13:00 – 13:10	Daniel Alonso and Georges Lobo	BDVA and SEMIC	Welcome and logistics
13:10 – 13:20	Edward Curry	BDVA, Insight, University of Galway	The Symbiotic Relationship between Data Spaces and AI. Challenges and Trends
13:20 – 13:35	Emilien Caudron and Lorenzo Gabrielli	SEMIC and JRC	AI 4 Interoperability
13:35 – 13:45	Vladimir Alexiev	Graphwise	Using LLM to Generate In-Depth Column Descriptors in the UNDERPIN Manufacturing/Maintenance Dataspace
13:45 – 13:55	Suna Akbayir & Guven Fidan	Artech International	Architecting LLM-Powered Digital Product Passports for Semantic Interoperability
13:55 – 14:05	Tobias Jacobs	NEC Laboratories Europe GmbH	Semantic Interoperability with FIWARE and LLM
14:05 – 14:25	Discussion on the speakers' statements – part 1 – moderated by Claudio Baldassarre (SEMIC)		
14:25 – 14:30	Break		
14:30 – 14:40	Arunav Das	King's College London	Telecom Question Answering Systems: LLMs and the Future of Semantic Interoperability
14:40 – 14:50	Orfeas Menis Mastromichalakis	National Technical University of Athens	Detecting and Contextualizing Harmful Language in Cultural Heritage Collections
14:50 – 15:00	Rafiqul Haque	Insight – University of Galway	Towards LLM-based Semantic Mediation Engine
15:00 – 15:10	Harri Ketamo	Headai	Making unstructured data interoperable
15:10 – 15:20	Ronald Kok	Municipality of Rotterdam	Apply LLM's to find relevant data for societal challenges
15:20 – 15:30	Frédéric Bellaiche	DAWEX	AI-powered semantic hub for data spaces
15:30 – 15:50	Discussion on the speakers' statements – part 2 – moderated by Claudio Baldassarre (SEMIC)		
15:50 – 16:00	Wrap-up and next steps		

3. Workshop presentations summary

1. **Daniel Alonso (BDVA) and Georges Lobo (SEMIC)** kicked off the workshop by underscoring why data interoperability, particularly semantic interoperability, remains a major challenge, despite the availability of powerful tools like Large Language Models. They explained that this event marks the start of a longer process to explore where LLMs truly stand regarding readiness, industry adoption, and research outcomes.
2. **Keynote presentations**
    - **Professor Edward Curry (University of Galway, Insight, BDVA vice-president)** highlighted in his presentation that data spaces benefit from "good enough" interoperability—i.e., letting data coexist rather than forcing a fully unified schema. He presented a "unified data and AI lifecycle," showing that large language models can integrate seamlessly with data spaces to perform tasks such as schema alignment and metadata generation.
    - **Emilien Caudron (SEMIC)** demonstrated a prototype chatbot during his presentation that helps data modellers locate and reuse standard classes, properties, and relationships. Additionally, he emphasised how an LLM can map new concepts to existing standards for faster development of interoperable data models.
    - **Lorenzo Gabrielli (Joint Research Centre – JRC)** showed in his presentation how the JRC combines semantic search with LLM-based summarisation to improve data retrieval and interpretation. Furthermore, he stressed that "AI-ready" data—clean metadata, consistent structures—improves the reliability of LLM outputs while reducing hallucinations.
3. **First round of industry presentations**
    - **Vladimir Alexiev (Graphwise)** discussed the UNDERPIN project for critical sectors (e.g., wind farms) in his presentation, where CSV on the Web (CSVW) ontologies plus LLM-assisted column descriptor generation help reduce manual labelling of sensor data. He also showed how deeper, per-column metadata (instead of basic dataset-level descriptions) leads to richer discoverability and analytics.
    - **Suna Akbayır & Güven Fidan (Artech)** addressed the EU push for digital product passports (DPPs), referencing upcoming regulations (e.g., Eco-Design for Sustainable Products) in their presentation. They proposed using a "hybrid" LLM + knowledge graph approach, layering zero-knowledge proofs for data protection, ensuring DPP scalability across multiple industries.
    - **Tobias Jacobs (NEC)** discussed a presentation on LLM-based approaches to translate both structured and unstructured data into a standard called NGSI-LD (used in FIWARE contexts). He also explained that LLMs can automate the creation of mapping functions between arbitrary data schemas and common smart data models—lowering the barriers to interoperability.
4. **Presentations from academia**
    - **Arunav Das (King's College London)** explored during the presentation how LLMs can improve enterprise question answering systems in the

field of telecom by using SPARQL-based prompts to better interpret customer intent without relying on rule-based systems. They also proposed combining vector search with knowledge graphs to enable seamless access across structured and unstructured knowledge sources.
    - **Orfeas Menis Mastromichalakis (National Technical University of Athens)** showed a hybrid tool (DE-BIAS) during his presentation that uses an LLM plus a specialised vocabulary to detect harmful language in cultural-heritage text collections. It was also mentioned how LLMs can reproduce bias, so combining them with curated knowledge sources and manual oversight is crucial.
    - In his presentation, **Rafiqul Haque (Insight – University of Galway)** aimed to outline a concept for an LLM-based semantic mediation engine, which can interpret and reconcile differences between multiple data standards in real time. Rafiqul Haque had connectivity issues, therefore, only the two statements were presented at the end of the workshop.
5. **Second round of industry presentations**
    - **Harri Ketamo (Headai)** showed during his presentation how LLMs can speed up skills taxonomies and workforce-related data linking—e.g., generating standardised "competence definitions" for HR systems. Furthermore, he stressed that these AI-driven approaches only work if data owners label or "tag" data consistently.
    - **Ronald Kok (Municipality of Rotterdam)** talked in his presentation about how local governments can harness LLM-based integration to unify administrative records, IoT sensors, and citizen feedback. Privacy and local regulation constraints on city data were mentioned, implying that the city sees LLMs as an added "semantic layer" rather than a replacement for structured reference models.
    - During his presentation, **Frédéric Bellaiche (Dawex)** discussed federated data exchange solutions where each participant retains control of their data. It can be argued that LLMs can facilitate more flexible negotiation of data-sharing "contracts" by summarising legal terms or proposing standard metadata formats.
6. **Daniel Alonso (BDVA)** concluded the webinar by elaborating on the next steps (Conclusions and next steps). **Georges Lobo (SEMIC)** encouraged workshop participants to consider registering their AI-related and semantic-interoperability projects in the "Public Sector Tech Watch".

4. Overall workshop discussion summary

Insights and reflections raised during the discussion:

- Balancing “perfect” and “pragmatic”: several participants echoed the idea that strictly standardised or fully integrated data solutions are costly and often unfeasible across many domains. “good enough” solutions (a combination of approximate text-based matching, incremental mapping, and conversational interactions) can unlock immediate value.
- Unified lifecycle & governance: there was considerable emphasis on how governance frameworks (licensing, provenance, data quality) must blend with the ai lifecycle. Many in the chat stressed that LLM-based models need clear training data provenance to build trust.
- Concerns about hallucinations & reliability: questions arose about how to evaluate LLM outputs, e.g., verifying suggestions for data model classes or semantic links. The role of human oversight in “co-creating” solutions with an LLM, and thereby mitigating hallucinations, was widely underscored.
- Tools & implementation: while participants reported existing prototypes (e.g. Data modelling chatbots, summarisation tools), they also noted ongoing difficulties with scaling them. The prevailing sense was that many tools remain “experimental” rather than fully “plug and play.”
- Sector-specific use cases: presentations on digital product passports (DPPS) for sustainability, manufacturing data integration, and public-sector data emphasised that LLMs need domain knowledge and hybrid approaches (e.g. Knowledge graphs + rule-based validations + AI) to ensure robust interoperability.

Throughout the workshop, participants also actively contributed questions and observations in the chat, highlighting:

- Human-in-the-loop approaches: Many participants recognised that LLM-generated mappings or semantic annotations still require expert validation – especially in domains with legal or safety-critical stakes.
- Benchmarks and best practices: Attendees requested consolidated benchmarks for evaluating LLM’s performance in semantic tasks and more guidelines on making data “AI-ready.”
- Data readiness: The conversation touched on the distinction between truly machine-readable or structured data (e.g. well-modelled JSON-LD) versus semi-structured data (e.g. PDFs that require extensive extraction).
- Future synergy: Some emphasised potential synergies between data spaces and LLM-based solutions (e.g. microservices or “agents” that mediate or orchestrate tasks automatically).

In short, the chat underscored that LLM-based solutions can reduce friction, but success depends on coherent standards, robust data governance, and continued experimentation with domain-specific validations.

5. Conclusions and next steps

This workshop demonstrated that LLMs are widely viewed as powerful enablers of semantic interoperability. In particular:

- Growing consensus: There is broad alignment on the value of LLMs for “approximate” or “best effort” semantic tasks, especially when used with knowledge graphs and domain ontologies.
- Persistent gaps: Despite widespread recognition of interoperability’s importance, satisfaction with current semantic interoperability and discovery solutions remains low.
- Importance of governance: Implementing robust oversight, clarifying data provenance, and merging data governance with AI lifecycle management is seen as essential for scale.
- Next-level innovations: Tools for automated data model generation, question answering, and domain-specific “co-pilot” chatbots show promise. However, they still need reliability checks and standard ways to incorporate domain context.

Overall, the workshop validated that LLM-based approaches can indeed accelerate the adoption of interoperable data ecosystems, but also that careful measures for the technical, organisational, and legal aspects are needed to embed trust and avoid pitfalls like hallucinations or bias.

Next steps related to this topic include:

1. Follow-up workshop during the Data Week (27-28 May 2025 in Athens): Attendees are invited to join a session focused on deeper technical demonstrations and policy updates.
2. Working groups and additional activities after summer 2025: BDVA and the SEMIC team plan to host smaller, targeted meetings and a concluding workshop in Q4 2025. These will aim to:
    - Collect success stories of LLM-based semantic solutions in data spaces.
    - Identify obstacles and develop recommended guidelines for public administrations and industry.
3. Community building and tool consolidation: A number of participants proposed a shared resource repository (open-source code snippets, best-practice documents, curated semantic models) to ensure synergy in the broader community.

Appendix 1: Statement results

After each presentation during the workshop, the respective speaker proposed one or two statements to the audience, which were anonymously rated by the participants within a range from strongly disagree to strongly agree. This appendix summarises the statements and results from the workshop.

Table 3 – Statements and results

Speaker	Statement	1 (Strongly disagree)	2	3	4	5 (Strongly agree)	Average rating	Total votes
Edward Curry	1. Data Spaces need a data co-existence approach for “Good Enough” Interoperability	1	3	6	30	27	4.2	67
Edward Curry	2. Data Spaces will require a unified Data and AI Lifecycle if we are to maximise the potential of Generative AI	1	1	9	20	37	4.3	68
Emilien Caudron	3. AI tools that facilitate the reuse of existing information (standards, data models, ...) are critical in advancing interoperability by ensuring consistency and reducing redundancy.	1	2	6	18	36	4.4	63
Lorenzo Gabrielli	4. Data interoperability isn’t a bonus feature—it’s a prerequisite. Designing for interoperability by default is the only way to ensure that AI solutions can work across domains, datasets, and institutions without becoming brittle or biased.	2	3	9	13	35	4.2	62
Vladimir Alexiev	5. I am satisfied with the level of semantic interoperability in dataspace(s) that I have practical experience with.	17	13	12	3	1	2.1	46
Vladimir Alexiev	6. I am satisfied with the dataspace discovery tools in dataspace(s) that I have practical experience with.	17	12	16	3	2	2.2	50
Suna Akbayir	7. LLMs will be indispensable for generating Digital Product Passports at scale, especially within the ESPR's 2025-2026 implementation timeline (before 2027 mandatory period).	2	7	10	14	4	3.3	37

Speaker	Statement	1 (Strongly disagree)	2	3	4	5 (Strongly agree)	Average rating	Total votes
Suna Akbayir	8. Combining LLMs with knowledge graphs, zero-knowledge proofs, and rule-based validation is necessary to ensure trustworthy DPPs.	1	4	18	10	9	3.5	42
Tobias Jacobs	9. Generative AI is a powerful translator between formal semantics and everyday human language – making the benefits of semantic interoperability accessible to anyone.	1	0	1	18	30	4.5	50
Tobias Jacobs	10. No matter if it is AI agents or humans that communicate - the challenges of semantic interoperability remain.	2	8	9	12	25	3.9	56
Arunav Das	11. Managing heterogeneous datasets presents a greater challenge for achieving semantic interoperability compared to disambiguating user intent.	3	5	6	7	6	3.3	27
Arunav Das	12. Large Language Models (LLMs) could have the most significant impact on semantic interoperability in Question Answering (QA) systems by enhancing the interoperability between questions and knowledge representations, facilitating more accurate and contextually relevant answers.	2	1	3	16	6	3.8	28
Orfeas Menis Mastromichalakis	13. Large language models trained on web data inevitably reproduce societal biases, making them unsuitable for tackling sensitive matters.	2	0	7	12	12	4	33
Orfeas Menis Mastromichalakis	14. Grounding LLMs in knowledge graphs or other external sources of knowledge eliminates the need for human oversight.	15	13	10	3	0	2	41
Rafiqul Haque	15. The integration of LLM-based semantic mediation mechanisms is a critical enabler for achieving scalable and context-aware semantic interoperability across heterogeneous participants in federated data spaces	0	2	2	5	5	3.9	14
Rafiqul Haque	16. Large Language Model driven semantic mediation can substantially enhance the effectiveness of standards like DCAT and NGSI-LD by enabling dynamic interpretation, alignment of heterogeneous vocabularies, and context-sensitive mapping across domains within data spaces.	0	1	3	6	8	4.2	18
Harri Ketamo	17. Semantic interoperability can be sorted out without AI.	5	9	5	4	6	2.9	29
Harri Ketamo	18. Hallucination and re-generated-concepts are going to be a major challenge in terms of semantic interoperability.	2	4	2	8	10	3.8	26

Speaker	Statement	1 (Strongly disagree)	2	3	4	5 (Strongly agree)	Average rating	Total votes
Ronald Kok	19. One should not use LLM's to process data resulting from a prompt as this will deliver unreliable outcomes.	1	5	7	3	3	3.1	19
Ronald Kok	20. A knowledge graph / triple store is a great source for LLM's.	1	1	5	8	10	4	25
Frédéric Bellaiche	21. Without an LLM-powered semantic hub, the scalability of data conformity assessment across data spaces becomes practically unmanageable.	0	4	6	6	3	3.4	19
Frédéric Bellaiche	22. The true enabler of data value in shared ecosystems is not access—but semantic understanding.	1	2	4	6	9	3.9	22

Appendix 2: Speaker reflections on audience responses to the statements

As part of the workshop, each speaker concluded their presentation by offering one or two statements for the audience to assess, using a scale from 1 (strongly disagree) to 5 (strongly agree). These statements prompted attendees to consider the assumptions, implications and priorities underpinning LLM applications for semantic interoperability. The responses offer a snapshot of community perspectives, and in several cases, speakers provided further interpretation of the results. This section summarises those insights.

Growing consensus on the need for “good enough” approaches

Edward Curry's statements received strong support. His call for a "data co-existence" strategy rather than rigid standardisation reflects a pragmatic view of interoperability. Respondents endorsed the notion that LLMs enable flexible integration across data spaces without enforcing a single schema. Edward Curry noted surprise at the high level of agreement, interpreting it as recognition of the potential of LLMs to reshape data practices.

In his second statement, on the necessity of a unified AI and data lifecycle, the audience showed even greater agreement. Edward Curry emphasised that generative AI cannot be treated in isolation from the data that powers it, reinforcing calls for tool support that bridges data governance and AI development lifecycles.

AI tools for reusing standards gain approval

Emilien Caudron highlighted the need for AI-driven reuse of standards and data models, citing growing complexity in the semantic modelling landscape. The audience's response demonstrated strong support. Emilien Caudron interpreted this as validation of the prototype he presented during the session, suggesting a clear demand from practitioners for tools that reduce duplication and streamline semantic modelling efforts.

Lorenzo Gabrielli took a firmer stance, arguing that interoperability should be designed by default rather than added later. His statement resonated with many participants. Lorenzo Gabrielli saw this as proof that interoperability is no longer viewed as optional in AI deployment, even if some divergent views remained.

Dissatisfaction with the status quo in dataspace tools

By contrast, Vladimir Alexiev’s two statements received some of the lowest scores. He presented a critical view of current dataspaces, arguing that both semantic interoperability and discovery tooling fall short of expectations. The audience appeared to share this frustration. In his reflection, Vladimir Alexiev noted widespread underuse of semantic technologies and inadequate metadata standards. He called for stronger enforcement and incentives to improve interoperability at both technical and governance levels.

Hybrid architectures for digital product passports

Suna Akbayir introduced statements on the role of LLMs in generating Digital Product Passports (DPPs) under upcoming EU regulation. The first statement, which positioned LLMs as essential for scaling DPPs, received a mixed but positive response. Suna Akbayir attributed this to ongoing concerns around integration and maturity of LLMs. The second statement, advocating a combination of LLMs, knowledge graphs, zero-knowledge proofs and validation logic, received more support. This was seen as a confirmation that hybrid architectures are gaining traction among practitioners.

LLMs as translators and collaborators

Tobias Jacobs proposed that generative AI can bridge the gap between formal semantics and human understanding. This was one of the most positively rated statements, suggesting that the community sees LLMs as tools for accessibility and adoption. His follow-up statement—that semantic interoperability remains a challenge regardless of whether agents or humans communicate—was met with general agreement, pointing to ongoing uncertainty about the extent to which LLMs can resolve these challenges on their own.

Enterprise search and QA applications

Arunav Das explored the comparative difficulty of managing heterogeneous datasets versus disambiguating user queries in enterprise settings. While there was no overwhelming consensus, responses suggested that both remain pressing challenges. Arunav Das interpreted the feedback as evidence that LLMs are useful but insufficient on their own, and that further work is needed to bridge data silos through semantic representation.

Nuanced thinking about bias and oversight

Orfeas Menis Mastromichalakis presented two intentionally provocative statements to prompt reflection on bias and oversight in LLMs. The first, suggesting that societal bias in LLMs makes them unsuitable for sensitive matters, received general agreement from the audience. In his reflection, Orfeas clarified that the goal was not to argue against using such models, but to challenge a binary view and highlight the potential value of LLMs when paired with clear task boundaries and proper oversight. The second statement, claiming that grounding in external knowledge removes the need for human oversight, was met with strong disagreement. He interpreted this as healthy skepticism, while noting that human oversight also has limitations. He encouraged more flexible oversight models that consider the strengths and weaknesses of both humans and machines.

Critical voices on reliability and human oversight

Harri Ketamo raised concerns around over-reliance on AI, suggesting semantic interoperability may not require AI at all. The audience response was highly mixed, which Harri Ketamo had anticipated. His second statement, on hallucination and concept regeneration as core risks for LLM-based systems, was received more favourably, reflecting a widely shared caution around AI reliability.

Trust and knowledge foundation

Ronald Kok reflected on two statements addressing LLM reliability and the value of structured knowledge. The first, warning against using LLMs to process prompt-generated data due to unreliable outcomes, received mixed responses. Ronald noted that the question could be interpreted in several ways and might have been clearer if it focused directly on hallucination risks. He took the responses as a sign of broader concerns around trust and emphasised that trust should remain a guiding principle. The second statement, which endorsed knowledge graphs and triple stores as valuable sources for LLMs, received strong agreement. Ronald viewed this as confirmation of his belief in hybrid approaches and concluded that investment in structured knowledge is key to building trustworthy systems.


Publications Office
of the European Union


interoperable
europe
BDV
BIG DATA VALUE
ASSOCIATION