

Governing Multimodal Knowledge Graphs: Current Challenges and Possible Solutions

Anelia Kurteva^{1,*†}, Gabriela Kurteva^{2,†}

¹*School of Computer Science, University of Birmingham, The UK*

²*ADAPT Centre, Trinity College Dublin, College Green, Dublin, Ireland*

Abstract

Knowledge Graphs (KGs) have emerged as a popular technology for improving the reliability and relevance of generative AI's decisions by acting as grounding truth infrastructures that provide factual domain knowledge as a result of the integration of diverse personal and private data. While the applications of KGs are gaining momentum, an overlooked and, if not even more pressing, investigation into the governance of KGs through their lifetime is needed to ensure their responsible development, maintenance and use. This is, however, a challenge as there is no standard in building and using KGs in a responsible manner in the same way we approach AI. This issue of governance deepens when dealing with multimodal data subject to various licences and legal bases for data sharing under the GDPR and its KG representation. As AI increasingly automates data collection, ontology engineering and KG construction and maintenance, we are witnessing a soft shift towards a "black box" development model. By limiting human expert involvement and lacking rigorous findable, accessible, interoperable, reusable (FAIR) documentation of KG lifecycle events (e.g. maintenance, access and usage for various purposes by AI), we create a systemic blind spot. This contradicts principles of responsible AI development and poses threats to AI safety when KGs are utilised as main knowledge sources for AI improvement (e.g. in Graph Retrieval Augmented Generation (RAG) of Large Language Models (LLMs)). Without a verifiable and FAIR audit documentation of a KG's lifecycle, a KG can easily become a liability rather than an asset. This paper addresses the governance of multimodal KGs, highlighting that they fit neither solely under data nor entirely under AI governance. Motivated by this, we present an idea for a KG governance framework centred around FAIR documentation of events, their regulatory compliance and transparent communication to end-users.

Keywords

Knowledge Graphs, Multimodal Data, Data Governance, AI Governance, Responsible AI, Trust

1. Introduction

Knowledge Graphs (KGs) have become one of the key technologies for mitigating Generative AI hallucinations when used as grounding truth sources of structured factual expert knowledge [1][2][3] via Graph Retrieval Augmentation (Graph RAG) [4][5]. By providing a rich semantic framework for findable, accessible, interoperable, reusable (FAIR) data sharing and a shared understanding of the data itself, KGs help streamline the understanding and actions of both humans and autonomous AI agents in complex multiagent settings. Among other, KGs and their underlying semantic layer can also enable the implementation of ethical and legal principles for data sharing, processing and use in a machine-readable format that AI can understand and reason over [6][7]. As such, KG are becoming a key technology for data governance in AI [8] and aiding its responsible development from the ground up - a key principle of data-centric AI many consider can help unlock AI's full potential.

However, building and using KGs is increasingly shifting toward fully automated, opaque processes that lack both transparency and the integration of human expertise (e.g. from domain and knowledge engineering experts). This shift has its advantages and disadvantages. Recent work such as [9][10] has focused on using LLMs for ontology generation, others such as [11][12] have focuses on LLMs for KG

3rd NeXt-generation Data Governance workshop (NXDG 2026), co-located with the 22nd SEMANTiCS conference, 15 September 2026, Ghent, Belgium

*Corresponding author.

†These authors contributed equally.

✉ a.kurteva@bham.ac.uk (A. Kurteva)

ORCID 0000-0003-4512-5969 (A. Kurteva)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

construction. We have also already witness the use of industry-ready frameworks such as LangGraph¹ that can fully automate the construction of graphs for Graph RAG systems (i.e. documents are mined and algorithmically structured into complex knowledge networks, eliminating the need for manual orchestration and schema definition). While such work provides significant benefit to KG engineering in terms of saving time and effort, issues regarding ethical and responsible development arise, especially when it comes to the data and AI-based processing used in the KG engineering process and later on during KG utilisation. Zhong et al. [13] discuss the pressing need to protect cyber Intellectual Property (IP) and unauthorised access of LLMs for information retrieval tasks (e.g. data collection and harvesting). Others such as [14], highlight the need to remember data can be personal and subject to human rights and regulatory compliance when dealing with LLMs. Feretzakis et al. [15] also tie legal compliance directly to the responsible use of LLMs to handle and use data.

The absence of standardised FAIR documentation regarding key lifecycle events such as schema definition, update, data collection, processing, annotation and maintenance (e.g. when privacy-preservation and legal compliance need to be ensured), further undermines their verifiability, reliability and trustworthiness. This becomes a significant roadblock when KGs are used to augment AI in a reliable and responsible manner. These challenges are further magnified when KGs are used as infrastructures to integrate public and private multimodal data such as text, images, documents. Examples of two use cases utilising multimodal knowledge graphs (MMKGs) are presented in section 3. MMKGs have recently emerged as powerful tools for enhancing AI capabilities, allowing for richer data representation by linking contextual information across various modalities [9]. Similarly to KGs, MMKGs' engineering is also subject to various automation with the help of AI. One possible example of semi-automated MMKGs engineering involving human-AI collaboration is presented on Figure 1.




Figure 1: MMKGs' Development Example, Human and Agent Involvement, and Applicable Legislation. The process begins with defining a **Use Case (1)** and establishing an **Ontology (2)**, which feeds into **Data Source Identification (3)** and **Resource Acquisition (4)**. This data gathering is strictly governed by a legal framework covering informed consent, contracts, and intellectual property. The acquired multimodal data (datasets, documents, video, audio, images) undergo **Data Processing (5)**, **Ontology Alignment (6)**, are turned into a **MMKG (7)**. The **MMKG Deployment (8)** phase powers downstream AI applications such as Predictive Analytics, Graph RAG, and Automated Decision Making subject to the EU AI Act, all while undergoing continuous **Maintenance (9)** that needs to be carried out by authorised agents.

With current advances in AI, an agent can replace a human in many tasks such as multimodal

¹<https://www.langchain.com/langgraph>

data collection, integration and KG construction [16]. From a legal and regulatory standpoint, this automated collection and integration introduces significant risks regarding privacy, data processing and regulatory compliance. Under regulations like the GDPR, every piece of personal data must be tied to a specific legal basis (e.g., explicit consent, contractual necessity, or legitimate interest). Further, when data is already collected and made available as datasets, different licenses might be attached to them. When multiple data modalities are interconnected and integrated with a MMKG, tracking and adhering to such distinct legal bases to ensure compliant data sharing and use becomes a true challenge. There is a pressing need to establish an effective KG governance framework that can help us oversee diverse lifecycle events on both data and semantic levels, to ultimately ensure responsible technology development and use for AI. Without such operational governance over KGs' lifecycle, we risk KGs becoming sources of misinformation for AI and open the door to their unethical and irresponsible development and use. All of this contributes to a bottom-up erosion of the public's and institutions' trust in AI.

This paper aims to highlight the need for establishing clear KG lifecycle governance framework(s) that ensure KGs are verifiable and reliable sources of knowledge for improving and not misinforming AI. Motivated by this, we propose a framework for KG and MMKG governance focused on the development of mechanisms for documenting, monitoring, compliance verification and communication of KG lifecycle events. Having such framework, based on providing actionable FAIR KG lifecycle documentation can equip diverse users of KGs, such as KG engineers, software developers and regulatory compliance auditors, with information for verifying the legal compliance and operational transparency of diverse KGs. This will ultimately help build more trust in the KG we utilise and will motivate their responsible development and use in the future. The remainder of this paper is organised as follows. Section 2 reviews related literature. Section 3 introduces two motivating use cases for KG lifecycle governance, from which a series of challenges regarding KG and MMKG governance are derived in Section 4. An overview of our proposed framework is detailed in Section 5, and Section 6 concludes the paper.

2. Relevant Work for KG Governance

KG governance is bounded by standard data management principles. Conceptual frameworks established in literature emphasise that governance is far beyond a technical configuration, but a crucial design challenge that requires concrete decision-making rights and accountability for data processes, ensuring data is governed as a valuable asset. As defined in [17], governance determines *what* decisions are necessary to be done and *who* executes them, while management refers to the technical enforcement of those decisions. Significantly, failing to differentiate between these layers can result in data asset fragmentation, which impairs data-driven decisions and impedes organisational effectiveness. Based on this notion, Khatri and Brown et al. [17] outline five decision domains such as (i) data principles which dictate top-level standards and set boundary requirements, (ii) metadata which standardises taxonomies, (iii) data quality and lifecycle which determine validity, security, and the lifecycle parameters. Together, they form a structured data governance structure to assist organisations in managing data, nevertheless, the five decision domains provide no semantic or ontological abstractions for controlling accountability.

As relational databases evolve towards more diverse data environments, these decision domains must also evolve to accommodate Big Data factors. In [18], Big Data governance is defined as managing massive organisational data for decision-making using a data framework validated against the ISO 8000 standard, based on the volume, velocity, variety, value, and veracity of data. However, the emergence of heterogeneous settings shift the governance concern from database storage to semantic heterogeneity and structural complexity. With this, Big Data jeopardises the quality, security and integrity of data, indicating that conceptual frameworks fail to scale in heterogeneous settings. For Big Data governance to be successful in addressing challenges of data variety and heterogeneous environments, it must rely on frameworks that balance open sharing, specifically the FAIR principles, with human-centric sovereignty. For example, to follow the principles of collective benefit, authority to control, responsibility, ethics (CARE). In this work, they are lightly explored, but not incorporated as semantic techniques.

Moreover, integrating Big, Open, and Linked Data (BOLD) with algorithmic systems can demand for mechanisms that handle real-time data from various sources. Janssen et al. [19] presents a data governance framework for trustworthy Big Data Algorithmic Systems (BDAS). Across these systems, trustworthiness refers to how a system upholds the interests of the trusted entity (e.g., society). To integrate and support trust, essentially, the proposed framework is based on thirteen design principles, not limited to, mandates that handle data, processes, and algorithms, alongside trusted data sharing, risk-based governance, and shared ownership to maintain sovereign data control. Nevertheless, the framework offers no functionality in utilising KGs to capture data-to-decision lineage. In parallel, Jernite et al. [20] addresses human data management by proposing a Data Stewardship Organisation (DSO) framework built for language technology and training of large language models (LLM). Crucially, the DSO framework establishes a governance structure to foster the agency of data subjects by operating on (i) regulations, (ii) norms and values, and (iii) tools and implementation that apply technical policies. Similarly to the latter, the DSO framework relies primarily on manual administrative roles and standards, thus it is limited in the technical operationalisation of machine-readable KGs. Further discussions on governance [21] highlight the discrepancy between technological advancements of AI architectures and the slow pace of legislative or regulatory responses. In such contexts, to prevent risks and negative consequences of autonomous applications, there is a need for transition from '*principle to process*' through AI audit frameworks and technical certifications, which challenge traditional concepts of jurisdictional sovereignty, and require more adaptive regulatory environments.

Dafoe et al. [22] expands on this, by presenting a research agenda for AI governance that focuses on the norms and institutions shaping how AI is built and deployed. The focus is set on AI as (i) a general-purpose, (ii) an information, and (iii) an intelligence technology. Through this, the research agenda addresses the systemic failures of transformative AI (TAI) systems or AI-assisted human actors, which occur due to qualitative changes (e.g., private information about conflict costs, bargaining inefficiencies), and the need for new conceptual tools to prevent cooperative malfunctions. This perspective is significant in the era of agentic generative AI and Artificial General Intelligence (AGI) for which Joshi et al. [23] proposes a governance architecture to address post-deployment risks such as recursive self-improvement (RSI) and multi-agent coordination with minimal human oversight. This, however, offers only a review of static policies and conceptual proposals lacking concrete execution of semantics, leaving the proposed multi-agent coordination layer unable to monitor and verify autonomous agent actions in real-time.

Overall, there is limited research on the topic of governance of KGs themselves. Research such as [24][25] address the topic from different perspective related to implementing trust and accountability in KG-based AI, while in [26] the authors touch upon KG management principles, which can aid governance. Most of the work in the field has focused on broader data management [27][28] and on AI governance [29][30], is not open source (i.e. part of specific graph database management systems' functionalities requiring subscription plans), not FAIR (e.g. hidden in layers of system logs) and not operational and actionable for humans and AI agents. This creates a unique and timely opportunity to investigate the topic, especially when it comes to facilitating the integration and interpretation of diverse public and private multimodal data, such as text, images, and documents, with KGs in different cases as discussed in the next sections.

3. Use Cases motivating the Need for MMKG Governance

This section outlines two use cases in the financial risk management and digital product passport (DPP) domains, on which we have worked on previously in [31][32], where KGs have proven useful for multimodal data integration, interpretation and knowledge discovery. Use case 1 focuses on MMKGs for AI-driven digital twins of financial risk management processes, which can be seen as high risk AI case according to the AI Act due to the highly sensitive personal data used for risks evaluation of people. Use case 2, on the other hand, focuses on DPP-based data sharing enabled by MMKGs, which can also carry personal information about owners and their devices during their in-use lifetime. Both of these

use cases require multimodal data from various sources to be integrated and interpreted in a reliable and trustworthy manner by both humans and AI. Governing how these KGs are build, used, accessed, modified, utilised etc. becomes critical to enable their verifiability and reliability.

Use Case 1: MMKGs for Financial Operational Risks Management. Effective financial risk prediction is heavily affected by the fragmented and heterogeneous nature of financial data scattered across disparate departments and organisations. To mitigate these vulnerabilities, which saw global operational risk losses alone skyrocket past €17.8 billion in 2022 [33], the industry must move beyond isolated data. Overcoming this, requires integration of multimodal data (e.g. transaction logs, scanned pdf documents such as contracts and passports, biometrics, images from documents) and mapping of the complex, implicit relationships connecting institutional entities, financial products (e.g. loans and mortgages) and the risks associated with them. Financial Digital Twins (FDTs) have emerged as a promising solution by integrating heterogeneous data and dynamically modeling financial operations with the help of KGs and MMKGs as main data integration infrastructures. However, achieving this integration introduces significant hurdles regarding data governance and the management of the KGs used to represent it. Financial data is highly sensitive thus its sharing, processing, and utilisation (particularly when feeding KGs into downstream AI models) demands strict, robust governance frameworks to ensure security, regulatory compliance, and privacy. When data is transformed into a KGs/MMKGs, it is also accompanied by semantics that bring time-relevant context. The KGs/MMKGs should be built responsibly by respecting personal data privacy and securing explicit, informed consent. They must remain consistently reliable, ensuring that the underlying data and graph structures maintain high quality and are securely accessible, usable and editable on demand only by authorised and verified agents.

Use Case 2: MMKGs for Digital Product Passports-based Safety and Security monitoring of Products

Over the past decade, the exponential surge in global e-commerce has driven supply and demand to unprecedented levels. Legal entities such as Customs, are tasked with verifying the safety, origin, and valuation of millions of incoming products. This is a challenge due to lack of data and process digitisation and pressing legislation such as the Ecodesign for Sustainable Products Regulation (ESPR) [34]. The ESPR requires the implementation of Digital Product Passports (DPPs) to serve as digital identity cards tracking the components and sustainability profiles of goods. However, deploying DPPs at a scale that genuinely eases the enforcement bottleneck requires a robust, foundational digital infrastructure capable of processing, linking, and interpreting vast volumes of diverse multimodal data such as textual customs reports, structured product metadata, and visual imagery of goods and packaging [35][36]. KGs and specifically MMKGs can serve as an infrastructure that seamlessly unifies these distinct data modalities into a single, interconnected network of semantic relationships, enabling automated, intelligent cross-referencing and anomaly detection to flag non-compliant or fraudulent shipments instantly.

This approach introduces a variety of technical challenges alongside significant legal risks. From a data perspective, cross modality misalignment (e.g. discrepancies between visual data (like a product photo) and textual metadata) can result in incorrect data interpretation and its KG representation. This will also have an effect on decision making along the chain. Data sovereignty and privacy concerns present a high barrier to adoption, as private enterprises may resist sharing granular supply chain data due to the fear of exposing proprietary trade secrets. Personal privacy is also a factor as many ICT devices hold information that can be classified as personal. This issue has been further discussed in [35][37][38]. Adversarial risks pose a significant threat to compliance, as malicious actors could manipulate product images and metadata prior of after its KGs transformation. Mitigating this requires absolute trust in the underlying DPPs (stored as KGs/MMKGs). Strict verification protocols need to be implemented to guarantee that the data and its KG/MMKG representation(s) are managed by only authorised and trustworthy humans and AI agents.

4. Challenges in (Multimodal) KG Governance

Following our work in the fields of KGs and AI, data governance and regulatory compliance, and based on literature, this section, presents some of the existing challenges associated with establishing efficient and effective MMKGs governance. We have grouped these challenges in several categories as shown in the next subsections. While we focus on multimodal data and MMKGs, these challenges can also be abstracted to single data modality and traditional KGs.

4.1. Data-Oriented Challenges

- Lack of established multimodal data lineage tracking processes and documentation.
- Lack of structured frameworks to systematically verify data collection ethics, legal provenance, and licensing restrictions prior to graph integration.
- Enforcement of data processing restrictions based on dynamic user consent agreements during data ingestion for KG development and deployment in AI ecosystems.
- Navigating conflicting intellectual property (IP) laws, data licenses, and copyright when integrating multimodal data sources with KGs/MMKGs.

4.2. Knowledge Representation-Oriented Challenges

- Inconsistent adherence to best practices in ontology engineering (e.g. undocumented ontologies, lack of ontology validation and evaluation.)
- Difficulties in enforcing, validating, and maintaining standardised conceptual modeling practices and design patterns across complex, evolving knowledge schema.
- Semantic drift and misalignment in multimodal knowledge representations. This arises when unable to accurately capture, synchronise, and update the semantic abstractions of diverse data modalities (e.g., text, vision, structured tables) within a unified, high-fidelity graph schema over time.
- Lack of transparency and documentation of automated KG engineering pipelines and data extraction workflows, which hinders accountability and error verification.

4.3. Human-AI Collaboration Challenges

- Ambiguous definition of data access permissions, usage scope/boundaries, and legal and ethical responsibilities for autonomous agents operating on the KG.
- Formulating and enforcing distinct operational boundaries and obligations for humans and AI agents involved in KGs' lifecycle.
- Difficulties in seamlessly aligning, managing, and governing collaborative actions between human domain experts and autonomous AI agents within the broader KG ecosystem.
- Technical hurdles in securely verifying and authenticating the identities of both humans and AI agents.

4.4. Compliance-oriented challenges

- Unclear governance principles technologies such as KGs. Currently, most of the work in the research fields and industry has focused on either data or AI governance. KGs can be seen as a technology that stands in between. They represent and store data but also have been considered as a type of (symbolic) AI² [39].
- Lack of actionable documentation of how KGs are used by AI, what parts of them when and why during decision making.

²<https://www.lenovo.com/us/en/knowledgebase/symbolic-artificial-intelligence-a-comprehensive-guide/>




Figure 2: Knowledge Graph Lifecycle based on [40] and proposed Governance Mechanisms (in Blue)

- Lack of standardised data for regulatory compliance and audit. FAIR documentation frameworks capable of supporting complex, multi-jurisdictional compliance audits, such as ensuring GDPR compliance (e.g., handling the “right to be forgotten” within a MMKG are needed.).
- Lack of shared expert knowledge and structured, interdisciplinary collaboration between KG engineers and legal/technology governance experts.

5. Towards a Framework for MMKGs Governance

Effective MMKG governance requires a hybrid approach combining both data and AI governance. Traditional, standalone frameworks fall short because they cannot fully manage the complex evolution of graph data, automated AI processing, or the shifting roles of human and AI agents across the lifecycle. To build responsible AI applications, we must implement safeguards that ensure transparency and accountability at every stage. This section introduces an idea of a governance framework (see Figure 2) focused on documenting and monitoring lifecycle events, verifying their regulatory compliance, and keeping humans and AI agents informed via notifications.

5.1. How to document KG lifecycle events?

Proposal: We propose the development of a mechanism for the semi- and fully automated, machine-readable documentation of MMKG lifecycle events at both the data and semantic levels. Inspired by Data and AI Cards [41][42], this mechanism utilizes a dedicated governance ontology that reuses and extends standard vocabularies such as DCAT [43], PROV-O [44], DPV [45], Croissant ML Commons [46], and domain-specific baselines—to automatically generate FAIR, custom, and dynamic documentation templates. These templates capture granular metadata regarding event typologies (e.g., schema updates, new data modalities), underlying rationales, operational dynamics (status, changes, timestamps), and full agent provenance (human, software, and AI systems).

Expected Result: The successful implementation of this mechanism will yield machine-interoperable documentation providing a highly granular, transparent view of MMKG events (e.g. access to the KG, updates of specific data and semantics, annotation of new data), specifically detailing data sources, modality-specific processing, and semantic metadata. Furthermore, by integrating these templates with the Open Digital Rights Language (ODRL) [47], the framework will deliver an automated layer of governance that defines explicit usage policies (permissions and obligations) in direct compliance with the regulatory mandates of the EU AI Act [48].

5.2. How to enable the monitoring for KG lifecycle events?

Proposal: Drawing inspiration from code versioning and semantic conflict resolution research, we propose a monitoring mechanism designed to track the evolution of MMKGs and KG events by leveraging existing machine-readable documentation. The mechanism will use SPARQL queries to capture KG snapshots that establish authoritative baselines representing various KG lifecycle event outcomes. By comparing historical documentation versions against the active, latest KG, this monitoring mechanism will apply SHACL (Shapes Constraint Language) [49] to validate multimodal content and detect structural or schema-level shifts both before and after lifecycle events.

Expected Results: The deployment of this monitoring mechanism will enable the automated identification of previously unreported semantic or syntactic discrepancies across the knowledge graph. This will ensure continuous alignment between the physical state of the KG and its machine-readable documentation, significantly improving data integrity, validating multimodal content consistency, and providing a robust, automated layer for detecting structural regressions during KG evolution.

5.3. How to perform compliance-checking of KG lifecycle events?

Proposal: We propose a compliance verification mechanism designed to support the systematic verification and auditing of KG lifecycle events against regulatory frameworks such as the EU AI Act and GDPR. This mechanism operates by tightly integrating the previously defined machine-readable documentation with the active monitoring mechanisms. By doing so, it cross-references documented metadata against actual graph states to continuously verify data provenance, license adherence, and the absolute integrity of various KG lifecycle events. This capability is critical for multimodal KGs, where diverse, multi-source data types frequently introduce conflicting licenses and distinct legal bases for data processing and usage.

Expected Results: The implementation of this compliance verification mechanism will enable automated, systematic audits that identify and resolve legal or ethical conflicts before they result in compliance failures. It will successfully reconcile conflicting licenses across different data modalities and flags processing activities that lack a valid legal basis under the GDPR. Ultimately, this delivers a transparent, legally auditable trail for MMKGs, ensuring that complex data integrations and AI-driven lifecycle events remain strictly aligned with the ethical and regulatory mandates of the EU AI Act.

5.4. How to facilitate effective proactive communication with end-users of KGs?

Proposal: We propose a proactive communication mechanism designed to establish an automated notification layer within the MMKG ecosystem that bridges the gap between back-end monitoring and human-in-the-loop governance. This mechanism transforms raw system logs, validation states, and compliance status updates into actionable, user-centric alerts. Central to this mechanism is a dedicated User Interface (UI) featuring a dashboard and alerts tailored to visualise complex governance data in an accessible manner for both end-users and KG engineers. The system is engineered to proactively inform stakeholders of critical KG lifecycle events, including structural schema updates, system errors, and ongoing compliance audits (such as license and consent verification when collecting multimodal data to build or extend KGs).

Expected Results: The deployment of this proactive communication mechanism will deliver a transparent, accessible, and highly actionable view of the legal and technical "health" of the knowledge graph. By providing rich contextual insights through an intuitive dashboard, it will ensure that both human administrators and automated AI agents can immediately respond to validation failures, potential regulatory risks, or system anomalies. Ultimately, this mechanism fosters accountability and streamlined oversight across the entire MMKG lifecycle.

6. Conclusions

In conclusion, as KGs and MMKGs become indispensable pillars for grounding and improving the reliability of Generative AI, the current lack of dedicated governance guidelines poses significant risks to their integrity, safety, and regulatory compliance. This paper has highlighted the critical challenges inherent in the MMKGs lifecycle, emphasising that they exist at the complex intersection of data and AI governance. We have proposed a starting point for the development of a governance framework centered on FAIR KG lifecycle event documentation and compliance verification. By providing a more transparent, auditable lifecycle model, we can ensure that MMKGs serve as trusted, verifiable assets. Moving forward, the implementation of these proposed mechanisms is essential for fostering trust, ensuring adherence to legal frameworks such as the GDPR and the EU AI Act, and ultimately realising the potential of responsible, KG-driven AI.

Acknowledgments

Anelia Kurteva would like to thank Aad van Moorsel, Paolo Missier and Ruben Dedeker for the fruitful discussions on the topic. Gabriela Kurteva is funded by the European Union's Horizon Europe programme under the MarieSkłodowska-Curie Actions project HARNES (grant agreement No.101169409). The ADAPT Centre for Digital Media Technology is funded by Research Ireland through the Research Centres Programme and is co-funded under the European Regional Development Fund through Grant 13/RC/2106/P2. We acknowledge the assistance of Generative AI in grammar refinement and minor text paraphrasing.

References

- [1] S. Pan, L. Luo, Y. Wang, C. C. Chen, J. Wang, X. Wu, Unifying large language models and knowledge graphs: A roadmap, *IEEE Transactions on Knowledge and Data Engineering* 36 (2024) 3580–3599. doi:10.1109/TKDE.2024.3352100.
- [2] N. Matsumoto, J. Moran, H. C. Choi, M. E. Hernandez, M. Venkatesan, P. Wang, J. H. Moore, Kragen: a knowledge graph-enhanced rag framework for biomedical problem solving using large language models, *Bioinformatics* 40 (2024). doi:10.1093/bioinformatics/btae353.
- [3] X. Zhu, X. Guo, S. Cao, S. Li, J. Gong, Structgraphrag: Structured document-informed knowledge graphs for retrieval-augmented generation, in: *Proceedings of the AAAI Symposium Series*, volume 4, 2024, pp. 242–251. doi:10.1609/aaai.v4i1.31811.
- [4] S. Wang, W. Fan, Y. Feng, L. Shanru, X. Ma, S. Wang, D. Yin, Knowledge graph retrieval-augmented generation for LLM-based recommendation, in: *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2025, pp. 27152–27168.
- [5] T. T. Procko, O. Ochoa, Graph retrieval-augmented generation for large language models: A survey, in: *2024 Conference on AI, science, engineering, and technology (AIxSET)*, 2024, pp. 166–169.
- [6] T. R. Chhetri, A. Kurteva, R. J. DeLong, R. Hilscher, K. Korte, A. Fensel, Data protection by design tool for automated GDPR compliance verification based on semantically modeled informed consent, *Sensors* 22 (2022) 2763. doi:10.3390/s22072763.
- [7] D. Golpayegani, H. J. Pandit, D. O'Sullivan, D. Lewis, Semantic frameworks to support implementation of the EU AI Act, *Computer Law & Security Review* 61 (2026) 106315.
- [8] A. Merono-Penuela, E. Simperl, A. Kurteva, I. Reklós, Kg. gov: Knowledge graphs as the backbone of data governance in ai, *Journal of Web Semantics* 85 (2025) 100847.
- [9] J. Li, D. Garijo, M. Poveda-Villalón, Large language models for ontology engineering: a systematic literature review, *Semantic Web Journal* 20 (2025) 1–45.
- [10] Y. Teplyi, D. Dosyn, Prompt-guided LLM agent for end-to-end ontology learning, *Computer Engineering* 22 (2025) 35–47.
- [11] H. Huang, C. Chen, Z. Sheng, Y. Li, W. Zhang, Can LLMs be good graph judge for knowledge

- graph construction?, in: *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 2025, pp. 10940–10959.
- [12] V. K. Kommineni, B. König-Ries, S. Samuel, From human experts to machines: An LLM supported approach to ontology and knowledge graph construction, *arXiv preprint arXiv:2403.08345* (2024).
 - [13] Y. Zhong, Y. Wen, J. Guo, M. Kafai, H. Huang, H. Guo, Z. Zhu, Web intellectual property at risk: Preventing unauthorized real-time retrieval by large language models, in: *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 2025, pp. 17222–17235.
 - [14] W. H. Wong, Don't forget that there are people in the data: LLMs in the context of human rights: A response to samuel r. bowman, *Critical AI 2* (2024).
 - [15] G. Feretzakis, E. Vagena, K. Kalodanis, P. Peristera, D. Kalles, A. Anastasiou, GDPR and large language models: Technical and legal obstacles, *Future Internet 17* (2025) 151.
 - [16] L. Zhong, J. Wu, Q. Li, H. Peng, X. Wu, A comprehensive survey on automatic knowledge graph construction, *ACM Computing Surveys 56* (2023) 1–62.
 - [17] V. Khatri, C. V. Brown, Designing data governance, *Communications of the ACM 53* (2010) 148–152.
 - [18] A. Al-Badi, A. Tarhini, A. I. Khan, Exploring big data governance frameworks, *Procedia computer science 141* (2018) 271–277.
 - [19] M. Janssen, P. Brous, E. Estevez, L. S. Barbosa, T. Janowski, Data governance: Organizing data for trustworthy artificial intelligence, *Government information quarterly 37* (2020) 101493.
 - [20] Y. Jernite, H. Nguyen, S. Biderman, A. Rogers, M. Masoud, V. Danchev, S. Tan, A. S. Luccioni, N. Subramani, I. Johnson, et al., Data governance in the age of large-scale data-driven language technology, in: *Proceedings of the 2022 ACM conference on fairness, accountability, and transparency*, 2022, pp. 2206–2222.
 - [21] A. Taeihagh, Governance of artificial intelligence, *Policy and society 40* (2021) 137–157.
 - [22] A. Dafoe, AI governance: a research agenda, *Governance of AI Program*, Future of Humanity Institute, University of Oxford: Oxford, UK 1442 (2018) 1443.
 - [23] S. Joshi, AI governance in the era of agentic Generative AI and AGI: Frameworks, risks, and policy directions, *International Journal of Innovative Research in Computer Science and Technology* (2025).
 - [24] L.-D. Ibáñez, J. Domingue, S. Kirrane, O. Seneviratne, A. Third, M.-E. Vidal, Trust, accountability, and autonomy in knowledge graph-based ai for self-determination, *arXiv preprint arXiv:2310.19503* (2023).
 - [25] J. Domingue, L.-D. Ibáñez, S. Kirrane, M.-E. Vidal, P. D. Rohde, Trust and accountability in knowledge graph-based AI for self determination (Dagstuhl seminar 25051), *Dagstuhl Reports 15* (2025) 136–200.
 - [26] M.-E. Vidal, Y. Chudasama, H. Huang, D. Purohit, M. Torrente, Integrating knowledge graphs with symbolic ai: The path to interpretable hybrid ai systems in medicine, *Journal of Web Semantics 84* (2025) 100856.
 - [27] R. R. Pansara, S. A. Vaddadi, R. Vallabhaneni, N. Alam, B. Y. Khosla, P. Whig, Fortifying data integrity using holistic approach to master data management and cybersecurity safeguarding, in: *2024 11th International Conference on Computing for Sustainable Global Development (INDIA-Com)*, IEEE, 2024, pp. 1424–1428.
 - [28] I. M. Leghemo, C. Azubuike, O. D. Segun-Falade, C. S. Odionu, Data governance for emerging technologies: A conceptual framework for managing blockchain, IoT, and AI, *Journal of Engineering Research and Reports 27* (2025) 247–267. doi:10.9734/jerr/2025/v27i111100.
 - [29] M. A. Camilleri, Artificial intelligence governance: Ethical considerations and implications for social responsibility, *Expert Systems 41* (2024) e13406. doi:10.1111/exsy.13406.
 - [30] A. s. Batool, D. Zowghi, M. Bano, Ai governance: a systematic literature review, *AI and Ethics* (2025) 1–15. doi:10.1007/s43681-024-00612-4.
 - [31] N. Al Ahmad, Y. Zhang, A. Kurteva, A multimodal knowledge graph infrastructure for digital twins of financial operational risks: Mortgages and FX stocks, 2026.
 - [32] A. Kurteva, C. van der Valk, K. McMahon, A. Bozzon, R. Balkenende, Replanit ontology for

digital product passports of ICT: Laptops and data servers, *Semantic Web* 16 (2025). doi:10.1177/22104968251361274.

- [33] Operational Riskdata eXchange Association (ORX), Fraud drives frequency in operational risk losses, 2023. URL: <https://orx.org/blog/fraud-drives-frequency-in-operational-risk-losses>, Accessed on November 18, 2025.
- [34] European Parliament and Council of the European Union, Regulation (EU) 2024/1781 of the European Parliament and of the Council of 13 June 2024 establishing a framework for setting ecodesign requirements for sustainable products, amending Directive (EU) 2020/1828 and Regulation (EU) 2023/1542 and repealing Directive 2009/125/EC, *Official Journal of the European Union L* (2024). URL: <https://eur-lex.europa.eu/eli/reg/2024/1781/oj>.
- [35] A. Kurteva, K. McMahon, A. Bozzon, R. Balkenende, Semantic Web and its role in facilitating ICT data sharing for the circular economy: An ontology survey, *Semantic Web* 15 (2024) 2035–2067. doi:10.3233/SW-243586.
- [36] B. Rukanova, Y.-H. Tan, S. van Engelenburg, J. Ubacht, Cross-border e-commerce customs risk management: Exploring the potential of linking digital product passport data, X-ray scanned images, and AI, in: *Proceedings of the International Conference on Electronic Government (EGOV)*, Springer, 2026.
- [37] C. Ducuing, R. H. Reich, Data governance: Digital product passports as a case study, *Competition and Regulation in Network Industries* 24 (2023) 3–23.
- [38] K. Berger, M. Rusch, A. Pohlmann, M. Popowicz, B. C. Geiger, H. Gursch, J.-P. Schöggel, R. J. Baumgartner, Confidentiality-preserving data exchange to enable sustainable product management via digital product passports-a conceptualization., *Procedia CIRP* 116 (2023) 354–359.
- [39] S. Zhu, S. Sun, A short survey: Exploring knowledge graph-based neural-symbolic system from application perspective, *arXiv preprint arXiv:2405.03524* (2024).
- [40] D. Fensel, U. Simsek, K. Angele, E. Huaman, E. Kärle, O. Panasiuk, I. Toma, J. Umbrich, A. Wahler, *Knowledge Graphs*, Springer International Publishing, Switzerland, 2020. doi:10.1000/978-3-030-37439-6.
- [41] M. Pushkarna, A. Zaldivar, O. Kjartansson, Data cards: Purposeful and transparent dataset documentation for responsible ai, in: *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 2022, pp. 1776–1826. doi:10.1145/3531146.3534642.
- [42] D. Golpayegani, I. Hupont, C. Panigutti, H. J. Pandit, S. Schade, D. O’Sullivan, D. Lewis, AI cards: Towards an applied framework for machine-readable ai and risk documentation inspired by the EU AI act, in: *Annual Privacy Forum*, Springer Nature Switzerland, Cham, 2024, pp. 48–72. doi:10.1007/978-3-031-68194-3_3.
- [43] World Wide Web Consortium, Data Catalog Vocabulary (DCAT), W3C Recommendation, 2014. URL: <https://www.w3.org/TR/vocab-dcat/>.
- [44] T. Lebo, S. Sahoo, D. McGuinness, K. Belhajjame, J. Cheney, D. Corsar, D. Garijo, S. Soiland-Reyes, S. Zednik, J. Zhao, PROV-O: The PROV Ontology, W3C Recommendation, 2013. URL: <https://www.w3.org/TR/prov-o/>.
- [45] H. J. Pandit, B. Esteves, G. P. Krog, P. Ryan, D. Golpayegani, J. Flake, Data privacy vocabulary (dpv)–version 2.0, in: *International Semantic Web Conference*, Springer Nature Switzerland, Cham, 2024, pp. 171–193. doi:10.1007/978-3-031-77850-6_10.
- [46] M. Akhtar, O. Benjelloun, C. Conforti, L. Foschini, J. Giner-Miguel, P. Gijsbers, S. Goswami, N. Jain, M. Karamousadakis, M. Kuchnik, S. Krishna, Croissant: A metadata format for ml-ready datasets, *Advances in Neural Information Processing Systems* 37 (2024) 82133–82148.
- [47] R. Iannella, Open digital rights language (ODRL), in: *Open Content Licensing: Cultivating the Creative Commons*, Sydney University Press, 2007.
- [48] European Union, Artificial intelligence act, <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>, 2024. *Official Journal of the European Union*.
- [49] B. Bogaerts, M. Jakubowski, J. Van den Bussche, SHACL: A description logic in disguise, in: *International Conference on Logic Programming and Nonmonotonic Reasoning*, Springer International Publishing, Cham, 2022, pp. 75–88. doi:10.1007/978-3-031-15707-3_7.