The Second Workshop on Graph-enhanced LLMs for trustwOrthy Web data management (GLOW'26)

Gianluca Bonifazi[1], Stefano Cirillo[2], Eliana Pastor[3] and Luca Virgili[1]

[1]Department of Information Engineering, Marche Polytechnic University, Italy
[2]Department of Computer Science, University of Salerno, Italy
[3]Department of Control and Computer Engineering, Politecnico di Torino, Italy

Abstract

The recent growth of Large Language Models (LLMs) has expanded possibilities in data management, enabling powerful natural language access, reasoning, and decision support. However, reliability and trustworthiness remain major challenges when deploying LLMs in sensitive domains. Graph-based representations of knowledge and data, e.g., knowledge and property graphs, provide a promising avenue to address these challenges. LLMs generate fluent responses but often struggle with factuality, bias, hallucinations, and a lack of explainability. Conversely, graphs provide structured, interconnected representations that can serve as grounding and validation layers for LLM-based systems. In this context, Semantic Web technologies offer standardized graph-based frameworks, e.g., RDF and ontologies, to formally model, link, and reason over knowledge.

The second edition of the *Graph-enhanced LLMs for trustwOrthy Web data management* (GLOW) workshop explores how Semantic Web and graph technologies can be used to enhance the reliability of LLM-driven data management systems. It investigates retrieval and reasoning techniques, as well as the use of ontologies, constraints, and provenance metadata to assess and improve the trustworthiness of LLM outputs. The goal of the workshop is to foster discussion on how established Semantic Web principles can guide the design of reliable LLM-based web applications.

Keywords

Web Data Management, Graph-based Representations, Large Language Models, Semantic Web

CEUR-WS.org/Vol-4313/glow-preface.pdf

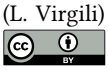
1. Introduction

The increasing adoption of Large Language Models (LLMs) is reshaping how users access and reason over heterogeneous data, while knowledge graphs, ontologies, and Semantic Web technologies continue to provide structured foundations for representation, integration, querying, and reasoning. Their combination creates new opportunities for data management, but also raises important questions about the reliability, provenance, and accountability of LLM-driven systems. These challenges are particularly relevant when LLMs interact with structured and evolving knowledge, where outputs may depend on multiple sources, complex relationships, and domain-specific constraints. Graph-based representations can help by providing structured evidence for retrieval and reasoning, while supporting verification, validation, provenance tracking, and traceability of generated outputs.

The relationship between graphs and LLMs is also bidirectional. Structured knowledge can improve the reliability of LLMs, while LLMs are increasingly used to construct, augment, query, and maintain knowledge graphs. This introduces challenges related to ontology alignment, schema mapping, semantic consistency, and validation of machine-generated knowledge. These issues are especially relevant to the Semantic Web community, where provenance, interoperability, and machine-processable semantics are long-standing concerns. Semantic Web technologies therefore provide a natural foundation for developing LLM-based systems whose outputs can be grounded, inspected, validated, and ultimately trusted.

Graph-enhanced LLMs for trustwOrthy Web data management (GLOW), in ISWC 2026 Workshops Joint Proceedings, October 25–26, 2026, Bari, Italy

g.bonifazi@univpm.it (G. Bonifazi); scirillo@unisa.it (S. Cirillo); eliana.pastor@polito.it (E. Pastor); luca.virgili@univpm.it (L. Virgili)
0000-0002-1947-8667 (G. Bonifazi); 0000-0003-0201-2753 (S. Cirillo); 0000-0002-3664-4137 (E. Pastor); 0000-0003-1509-783X (L. Virgili)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073
published 2026-11-05

Following its first edition, held in conjunction with The ACM Web Conference (WWW'26), the second edition of the *Graph-enhanced LLMs for trustwOrthy Web data management* (GLOW) workshop, co-located with ISWC 2026, continues to investigate the emerging intersection of graph-based data management and LLMs. GLOW provides a forum for discussing methods and applications in which graphs contribute to trustworthy LLM-based systems, as well as approaches in which LLMs support the creation and exploitation of structured knowledge. Particular attention is given to grounding and retrieval, reasoning and verification, provenance and explainability, semantic validation, knowledge graph construction and querying, and the evaluation of graph-enhanced LLM systems. Through these perspectives, the workshop aims to advance the design of more reliable, transparent, and accountable approaches to knowledge-intensive data management.

2. Workshop Statistics

The second edition of GLOW received contributions from authors across 19 countries. A total of 30 papers were submitted to the workshop. The submissions underwent a peer-review process in which 61 reviews were completed by the members of the Program Committee. Based on this evaluation, 13 papers were accepted for presentation and inclusion in the proceedings, resulting in an acceptance rate of 43%.

3. Workshop Program

The workshop program brought together contributions exploring how graph-based representations and Semantic Web technologies can enhance the reliability, transparency, and effectiveness of Large Language Models for Web data management. The accepted papers reflected the broad scope of GLOW, spanning graph-grounded retrieval and reasoning, knowledge graph construction and access, provenance and verification, as well as applications in scientific and industrial domains.

A significant part of the program focused on grounding LLMs in structured evidence. The presented approaches investigated graph-enhanced retrieval and knowledge-augmented generation for question answering, fact verification, recommendation, and multimodal reasoning, with particular attention to evidence selection, multi-hop reasoning, and the separation between generation and verification. Several contributions addressed hallucination and factual reliability directly, using knowledge graphs not only as sources of contextual information, but also as mechanisms for probing model knowledge, validating generated claims, and supporting explicit abstention or verification when the available evidence is insufficient.

A second major theme concerned the trustworthy construction, integration, and querying of structured knowledge. Contributions explored LLM-assisted schema mapping, knowledge graph augmentation with provenance and semantic validation, and natural-language interfaces to RDF graphs. Related work also considered the reliability of evaluation itself, including the adaptation of established benchmarks to evolving knowledge graphs. Across these contributions, provenance, ontology alignment, constraint validation, and graph-grounded correctness emerged as key mechanisms for making LLM-based data management processes more transparent and auditable.

Finally, the workshop highlighted the applicability of graph-enhanced LLMs across heterogeneous and knowledge-intensive settings, including biodiversity research infrastructures, biomedical question answering, industrial maintenance, power-line operations, and data-driven recommendation. These applications illustrated how structured domain knowledge can connect heterogeneous evidence sources while preserving their semantics, provenance, and operational context. Taken together, the program highlighted the multiple roles that graphs can play in trustworthy LLM systems: as sources of grounding evidence, structures for retrieval and reasoning, interfaces to heterogeneous data, and explicit substrates for validation, provenance tracking, and verification.

4. Program Committee

We would like to thank the members of the GLOW Program Committee for their valuable contribution to the review process:

- Alessia Antelmi, University of Turin
- Sajid Anwar, Institute of Management Sciences Peshawar
- Marco Arazzi, University of Pavia
- Francesca Bugiotti, Paris-Saclay University
- Francesco Cauteruccio, University of Salerno
- Gaetano Cimino, University of Salerno
- Domenico Desiato, University of Bari Aldo Moro
- Stefania Dumbrava, Télécom SudParis
- Lisa Ehrlinger, Hasso Plattner Institute
- Alberto García S., Universitat Politècnica de València
- Flavio Giobergia, Politecnico di Torino
- Gianpaolo Iuliano, University of Salerno
- Weiwei Jiang, Beijing University of Posts and Telecommunications
- Andrea Mauri, University of Lyon I
- Vincenzo Moscato, University of Naples "Federico II"
- Matteo Paganelli, University of Modena and Reggio Emilia
- Vincenzo Pasquadibisceglie, University of Bari Aldo Moro
- Maria Angela Pellegrino, University of Salerno
- Francesco Pugnaloni, Hasso Plattner Institute
- Georgia Troullinou, CNRS, Univ. Grenoble Alpes
- Gerardo Vitagliano, MIT CSAIL
- Luca Zecchini, BIFOLD & TU Berlin
- Zhiqiang Zhong, University of Luxembourg
- Andrea Zingoni, University of Tuscia

Acknowledgments

We would like to thank all authors for submitting their work and the members of the Program Committee for their careful and constructive reviews. We also thank the organizers of ISWC 2026 for hosting GLOW and supporting its organization.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT in order to: grammar and spelling check, paraphrase and reword. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.