

Preface to the 1st Semantic Architectures for Governance and Explainability (SAGE 2026)

Amna Dridi¹, Edlira Vakaj¹, Cecilia Zanni-Merk² and Mohamed Ragab¹

¹Birmingham City University, Birmingham, United Kingdom

²INSA Rouen Normandie, France

Abstract

The 1st International Workshop on Semantic Architectures for Governance and Explainability (SAGE 2026) was held on 15 September 2026 in Ghent, Belgium, co-located with the 22nd International Conference on Semantic Systems (SEMANTiCS 2026). The workshop brought together researchers and practitioners interested in the role of semantic technologies in supporting trustworthy, explainable, auditable, and governable artificial intelligence systems.

Preface

The rapid adoption of Artificial Intelligence, and particularly Generative AI and Large Language Models, has created new challenges around transparency, accountability, provenance, regulatory compliance, and the reliability of automated decision-making. SAGE was established to provide a forum for exploring how semantic technologies—including knowledge graphs, ontologies, rule-based systems, provenance models, and semantic reasoning—can contribute to addressing these challenges.

A central motivation of the workshop is to move beyond explainability as an isolated model-level concern and towards semantic architectures in which governance, reasoning, evidence, provenance, and evaluation can be represented explicitly, inspected, and queried.

The contributions accepted to SAGE 2026 reflect this broad perspective. Several papers investigate the relationship between semantic technologies and explainable and trustworthy AI. One contribution presents a comparative evaluation of post-hoc explanation methods and explores the semantic enrichment of explanation qualification through knowledge graphs, ontology-based trust levels, provenance, and machine-interpretable validation. Another contribution introduces AEGIS, a modular ontology for Generative AI governance that connects models, prompts, inputs, generated outputs, risks, mitigation strategies, responsible actors, validation activities, and provenance within an integrated semantic representation.

A second group of contributions focuses on regulation, compliance, and audibility. One paper investigates GraphRAG architectures for legal question answering over the EU AI Act, comparing RDF knowledge graphs with labelled property graphs and examining how different graph representations affect retrieval quality and inspectability as regulatory questions become more complex. Another contribution reframes regulatory compliance as procedural knowledge, representing compliance decisions as explicit sequences of steps, conditions, decisions, and outcomes in order to support traceability and audibility.

The programme also reflects renewed interest in semantic reasoning and rule-based architectures in the age of Generative AI. Contributions investigate the construction and validation of interoperable semantic rule engines, including extensive Notation3 conformance testing for trustworthy and legally compliant data flows, as well as implementation experience with the emerging Shape Rules Language for explainable RDF reasoning. The workshop additionally includes work on auditable evaluation of knowledge-graph change, highlighting the need for evaluation frameworks that explain not only whether a system succeeds, but also the demands of the tasks being assessed.

1st Semantic Architectures for Governance and Explainability (SAGE), co-located with SEMANTiCS 2026: 22nd International Conference on Semantic Systems, September 15, 2026, Ghent, Belgium



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

SAGE 2026 received eight submissions, of which seven papers were accepted for presentation and inclusion in the workshop proceedings. Each submission underwent a double-blind peer-review process and was evaluated by three reviewers. The accepted papers demonstrate the diversity of approaches through which semantic technologies can contribute to AI governance and explainability, while sharing a common concern with making computational processes more transparent, inspectable, interoperable, and accountable.

We would like to express our sincere thanks to all authors for contributing their work to the first edition of SAGE. We are particularly grateful to the members of the Programme Committee for their careful reviews, constructive feedback, and support in maintaining the scientific quality of the workshop. We also thank the SEMANTiCS 2026 organisers and the joint workshop proceedings editors for their support in organising the event and preparing these proceedings.

We hope that SAGE 2026 provides a foundation for continued discussion and collaboration at the intersection of semantic technologies, Artificial Intelligence, governance, and explainability, and that this first edition contributes towards building a community around semantic foundations for trustworthy and accountable AI.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT in order to support language refinement and drafting of the workshop preface. The authors reviewed and edited the content as needed and take full responsibility for the publication's content.