The Second Workshop on Knowledge Graphs and Neurosymbolic AI

Shqiponja Ahmetaj[1], Fajar J. Ekaputra[2], Andreas Ekelhart[3] and Sebastian Neumaier[4]

[1]Institute of Logic and Computation, TU Wien, Austria
[2]Institute of Data, Process, and Knowledge Management, Vienna University of Economics and Business, Austria
[3]Research Group Security and Privacy, University of Vienna, Austria
[4]Institute of IT Security Research, St. Pölten University of Applied Sciences, Austria

Preface

Neurosymbolic AI has emerged as a promising paradigm that seeks to combine the robust reasoning capabilities of symbolic systems with the powerful pattern recognition and learning abilities of neural networks. At the same time, Knowledge Graphs (KGs) have proven to be highly effective for organizing and integrating complex, structured, and unstructured data, making them ideal candidates for synergistic integration with machine learning systems.

The KG-NeSy workshop[1] seeks to deepen understanding of how Knowledge Graphs and Neurosymbolic AI mutually benefit each other, fostering dialogue, collaboration, and progress in their integration. The 2025 edition marks the second iteration of the workshop, following the success of its first event co-located with the AIRoV symposium in 2024[2]. This year, KG-NeSy is hosted at the SEMANTiCS 2025 conference, further embedding the workshop in a vibrant community focused on semantics, knowledge engineering, and applied AI. The workshop encourages contributions that explore both theoretical foundations and practical applications of KG-NeSy integration—from reasoning with ontologies and symbolic representations, to enhancing the interpretability, explainability, and trustworthiness of AI systems using structured knowledge.

For KG-NeSy 2025, we received nine submissions. Each was evaluated through a single-blind review process, with three Program Committee members assigned to review each paper. Reviewers assessed submissions based on originality, technical soundness, clarity of presentation, and relevance to the workshop themes. After deliberation among the organizers, six submissions were accepted for presentation, consisting of one full paper, three short papers, and two extended abstracts. The selected works collectively span a diverse and timely set of topics, including OWL Lite reasoning in LLMs, semantic data augmentation for predictive modeling,

The Second Workshop on Knowledge Graphs and Neurosymbolic AI (KG-NeSy), co-located with SEMANTiCS'25: International Conference on Semantic Systems, September 3–5, 2025, Vienna, Austria
shqiponja.ahmetaj@tuwien.ac.at (S. Ahmetaj); fajar.ekaputra@wu.ac.at (F. J. Ekaputra); andreas.ekelhart@univie.ac.at (A. Ekelhart); sebastian.neumaier@fhstp.ac.at (S. Neumaier)

[1]KG-NeSy 2025 webpage: https://sites.google.com/view/kgnesy2025
[2]KG-NeSy 2024 proceedings: https://www.uibk.ac.at/iup/buecher/9783991061502.html

RDF mapping with lifecycle metadata, the role of intermediate representations in neurosymbolic reasoning, explainable visual question answering via logic-based inference, and AI risk assessment grounded in design and risk pattern taxonomies. The programme also features an invited keynote by Cogan Shimizu from the Wright State University entitled "*Accelerating Knowledge Engineering with Modularity*", whose work reflects the workshop's emphasis on integrating symbolic and neural methods in AI.

The accepted papers and their brief summaries are the following:

- Emanuele Damiano and Francesco Orciuoli. *Evaluating Large Language Models on OWL Lite Reasoning.* This paper evaluates the ability of large language models to perform OWL Lite reasoning by using a retrieval-augmented generation framework over embedded ontologies, revealing how well different models handle ontology-based inference tasks of varying complexity.
- Majlinda Llugiqi, Fajar J. Ekaputra and Marta Sabou. *Semantic-Driven Data Augmentation for Improved Machine Learning Predictions.* This extended abstract proposes a semantic-driven data augmentation method that integrates knowledge graph embeddings into tabular datasets to improve machine learning performance.
- Sarah Alzahrani and Declan O'Sullivan. *Guiding LLM Generated Mappings with Lifecycle-Based Metadata: An Early Evaluation.* This short paper explores how structured lifecycle metadata can guide large language models to generate more accurate, semantically rich, and reusable RDF mappings.
- Alexander Beiser, Nysret Musliu and David Penz. *Intermediate Languages Matter: Formal Languages and LLMs affect Neurosymbolic Reasoning.* This short paper investigates how the choice of formal language impacts the effectiveness of neurosymbolic reasoning with large language models, demonstrating that context-aware encodings enhance reasoning performance.
- Thomas Eiter, Jan Hadl, Nelson Higuera, Lukas Lange, Johannes Oetsch, Bileam Scheuvens and Jannik Strötgen. *Explainable Zero-Shot Visual Question Answering via Logic-Based Reasoning.* This extended abstract introduces a neurosymbolic system for zero-shot visual question answering that combines large language models, vision models, and logic-based inference over symbolic scene graphs to provide explainable answers with traceable reasoning.
- Muhammad Ikhsan, Elmar Kiesling, Salma Mahmoud, Alexander Prock, Artem Revenko and Fajar J. Ekaputra. *Pattern-based AI Risk Assessment: A Taxonomy Expansion Use Case.* This short paper proposes a pattern-based approach to AI risk assessment using semantic models of interlinked design and risk patterns, enabling scalable and context-adaptable evaluations across various domains.

We thank all authors for their high-quality submissions, our keynote speaker Cogan Shimizu for his thoughtful contributions, and the Programme Committee members for their careful and constructive reviews.