Preface for the Third International Workshop on Retrieval-Augmented Generation Enabled by Knowledge Graphs (RAGE-KG 2026)

Daniel Dobriy[1,*], Marco Antonio Stranisci[2], Arianna Graciotti[3], Blerina Spahiu[4] and Sahar Vahdati[5]

[1]WU Vienna, Austria
[2]IT University of Copenhagen, Denmark
[3]University of Groningen, The Netherlands
[4]University of Milan–Bicocca, Italy
[5]Leibniz University of Hannover, Germany

Retrieval-Augmented Generation Enabled by Knowledge Graphs (RAGE-KG) is dedicated to advancing the combination of Retrieval-Augmented Generation (RAG) with Semantic Web infrastructure and exploring the synergies between Large Language Models (LLMs) and the Linked Open Data (LOD) ecosystem. The workshop promotes developing innovative RAG and GraphRAG architectures that leverage Semantic Web standards and novel methods to improve LOD usability for LLMs, enabling them to generate reliable, verifiable, context-aware responses from structured, decentralised, trustworthy data sources. The primary focus is on techniques for fine-tuning, prompting, and integrating LLMs with Knowledge Graphs (KGs), ultimately examining how to roll out RAG architectures across use cases and supplying a platform to further advance the field.

The first two iterations of the workshop, held in 2024[1] and 2025[2], demonstrated growing community interest, as evidenced by an increase in accepted publications from 8 to 15, four sessions, and substantial in-person participation (approximately 70 attendees). Although the workshop was initially motivated by the potential of RAG pipelines to address key limitations of LLMs, such as hallucinations and knowledge cut-off, it has remained relevant. A growing number of ISWC contributions now utilise RAG for knowledge and information extraction, entity disambiguation, knowledge graph construction and completion, SPARQL query generation, ontology learning, and other domain-specific applications [1]. Recent progress in GraphRAG [2], Agentic RAG [3], knowledge-supported reasoning, and the broad use of knowledge extraction with RAG indicate that the technology is growing beyond its original extent.

Building on previous iterations, RAGE-KG 2026[3] solicited submissions on innovative knowledge-integration methodologies, intending to foster a linear development of approaches that start with LLMs and extend toward longer-term AGI through key steps in between, including RAG, Knowledge-Augmented Generation (KAG), and agentic AI, as a progression from fluent generation toward grounded reasoning and then toward goal-directed learning and action. World models are introduced as a bridging concept for skill acquisition, enabling systems to develop internal models that facilitate prediction, planning, and iterative improvement [4]. Knowledge graphs are vital to this framework, as they can represent discrete symbolic layers of world knowledge and provide explicit semantics, constraints, and provenance [5]. World models can subsequently learn dynamics as well as uncertainty over this state

RAGE-KG 2026: The Third International Workshop on Retrieval-Augmented Generation Enabled by Knowledge Graphs, co-located with ISWC 2026, October 25–26, 2026, Bari, Italy
*Corresponding author.
daniel.dobriy@wu.ac.at (D. Dobriy); marcos@itu.dk (M. A. Stranisci); a.graciotti@rug.nl (A. Graciotti); blerina.spahiu@unimib.it (B. Spahiu); sahar.vahdati@tib.eu (S. Vahdati)
0000-0001-5242-302X (D. Dobriy); 0000-0001-9337-7250 (M. A. Stranisci); 0009-0004-7918-809X (A. Graciotti); 0000-0002-6958-8215 (B. Spahiu)
© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

[1]See: https://2024.rage-kg.org/
[2]See: https://2025.rage-kg.org/
[3]See: https://2026.rage-kg.org

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073
published 2026-11-04
CEUR-WS.org/Vol-4312/preface.pdf

and, when necessary, connect to continuous representations [4]. Furthermore, the 2026 workshop featured a special track on RAG systems for Human-AI Value Alignment to encourage contributions from the NLP community.

The workshop topics were:

- RAG Architectures
  - RAG architectures leveraging KGs, Semantic Web standards and LOD
  - RAG design patterns including GraphRAG and Agentic RAG
  - Evaluating RAG architectures with structured data
- LLMs and Structured Data in Monolingual and Multilingual Settings
  - Training and fine-tuning LLMs with structured data
  - Prompting Language Models with structured data
  - Language Model and ontology-supported SPARQL query generation
- Innovative Approaches
  - Neurosymbolic approaches for integrating Language Models with LOD
  - Use Cases, Work-In-Progress and Bold Proposals for RAG systems
- RAG-systems for Human-AI Value Alignment (special track)
  - RAG-based approaches for content moderation and harm reduction
  - Designing RAG systems to deal with human subjectivity

We implemented a single-blind peer review process. Submissions were reviewed based on their novelty, potential impact, technical quality, and relevance to the workshop's topics (by at least three members of the PC). The submission deadline was July 24, 2026, and the acceptance notification was sent on August 21, 2026.

The workshop received 35 submissions. 22 were accepted: 16 full papers and 6 short papers. 8 papers were presented in 15-minute slots (including Q&A) and 14 papers in a lightning-talk format. We thank the authors, the programme committee and all participants.

Keynote

The keynote was given by Filip Ilievski,[4] Assistant Professor of Commonsense AI at Vrije Universiteit Amsterdam.

Title: Abstracting Like an Expert: Memory, Representation Alignment, and Abduction in KG-RAG

Abstract: Abstraction, distilling the essence of a situation from its superficial trappings, is a hallmark of human expertise. Experts group physics problems by conservation laws rather than inclined planes, and they recognize that a marathon runner and a guitarist share the same "no pain, no gain" structure. Whether AI systems abstract in this expert-like way is an open question. On narrative analogy benchmarks, LLMs that excel at near analogies drop sharply on far analogies in zero-shot settings. Multimodal foundation models perform near chance on abstract visual puzzles, though GPT-4V reaches human level when given the right abstraction schema. This talk examines knowledge graph–based retrieval-augmented generation (KG-RAG) through this lens. Existing KG-RAG systems already use abstraction in two ways: long-term memory indexing (e.g., community summaries in GraphRAG) and working-memory compression within a fixed schema. A third form, reframing knowledge structure at query time to support generalization, appears to be underexplored. I map where abstraction currently sits in the KG-RAG pipeline, from conceptual generalization and graph summarization to query and path refinement. I then argue for a shift toward generative contextual abstraction, meaning small, bespoke, transient knowledge models built on the fly. I frame such abstraction as an abductive process, shaped by the interaction context, the ambiguity of the input, and the user's purpose. I close with a metacognitive view in which the KG serves not as an external database but as a working memory over which a system inspects, prunes, abstracts, and negotiates meaning with its human user.

[4] See: https://ilievski.info/

Programme Committee

- Hanna Abi Akl, Data ScienceTech Institute
- Reham Alharbi, Taibah University
- Hamed Babaei Giglou, TIB Leibniz Information Centre for Science and Technology
- Naga Sowjanya Barla, University of Liverpool
- Xiaohe Bo, Renmin University of China
- Luana Bulla, University of Bologna
- Francesco Compagno, University of Amsterdam
- Enrico Daga, The Open University
- Arunav Das, King's College London
- Jacopo de Berardinis, University of Liverpool
- Lia Draetta, University of Turin
- Beatrice Fiumanò, University of Bologna
- Bart Gajderowicz, University of Toronto
- Aidan Hogan, DCC, Universidad de Chile
- Pere-Lluís Huguet Cabot, VU
- Muhammad Ikhsan, Vienna University of Economics and Business
- Jan-Christoph Kalo, University of Amsterdam
- Kabul Kurniawan, University of Vienna
- Nicolas Lazzari, University of Bologna
- Anna Sofia Lippolis, CNR ISTC
- Yongchao Liu, Ant Group
- Gianfranco Lombardo, Università di Parma
- Ruhi Mahadeshwar, University of Groningen
- Antonello Meloni, University of Cagliari
- Arief Purnama Muharram, WU Vienna
- Gianmarco Pappacoda, Università di Bologna
- Marco Polignano, Università degli Studi di Bari Aldo Moro
- Célian Ringwald, University of Bologna
- Sergio José Rodríguez Méndez, School of Computing / College of Engineering & Computer Science. The Australian National University
- Angelo Salatino, The Open University
- Cristian Santini, University of Macerata
- Andrea Schimmenti, University of Bologna
- Sola Shirai, IBM Research
- Lumingyuan Tang, SailFar Intelligence
- Stella Verkijk, Vrije Universiteit Amsterdam
- Maria-Esther Vidal, Leibniz University Hannover
- Miguel Vázquez, WU Vienna
- Xiao Zhang, Center for Language and Cognition (CLCG), University of Groningen
- Hongye Zheng, Tencent

Acknowledgements

The work of M.A. Stranisci is supported by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Actions, Grant Agreement No. 101277473. Views and opinions expressed are, however, those of the author only and do not necessarily reflect those of the European Union or the European Research Executive Agency (REA). The work of D. Dobriy is funded in whole or in part by the Austrian Science Fund (FWF) 10.55776/COE12, and by the COST Action CA23147 GOBLIN - Global Network on Large-Scale, Cross-domain and Multilingual Open Knowledge Graphs, supported by COST (European Cooperation in Science and Technology, https://www.cost.eu). The work of A. Graciotti is funded by the ERC grant GOLEM (101040938), funded by the European Union. Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. The work of B. Spahiu is supported by the COST Action CA23147 GOBLIN - Global Network on Large-Scale, Cross-domain and Multilingual Open Knowledge Graphs.

Declaration on Generative AI

During the preparation of this work, the author(s) used Claude (Anthropic) for formatting assistance. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

[1] G. Demartini, K. Hose, M. Acosta, M. Palmonari, G. Cheng, H. Skaf-Molli, N. Ferranti, D. Hernández, A. Hogan (Eds.), The Semantic Web – ISWC 2024, Springer, 2024.

[2] B. Peng, Y. Zhu, Y. Liu, X. Bo, H. Shi, C. Hong, Y. Zhang, S. Tang, Graph retrieval-augmented generation: A survey, ACM Trans. Inf. Syst. 44 (2025). URL: https://doi.org/10.1145/3777378. doi:10.1145/3777378.

[3] A. Singh, A. Ehtesham, S. Kumar, T. T. Khoei, A. V. Vasilakos, Agentic retrieval-augmented generation: A survey on agentic rag, 2026. URL: https://arxiv.org/abs/2501.09136. arXiv:2501.09136.

[4] D. Ha, J. Schmidhuber, World models (2018). URL: https://zenodo.org/record/1207631. doi:10.5281/ZENODO.1207631.

[5] L. D. Costa, S. Namjoshi, M. A. Ansari, B. Schölkopf, Natural building blocks for structured world models: Theory, evidence, and scaling, 2025. URL: https://arxiv.org/abs/2511.02091. arXiv:2511.02091.