



LREC 2026

**Knowledge Graphs and Large Language Models
Workshop 2026 (KG-LLM) @ LREC 2026**

Workshop Proceedings

Editors

**Gilles Sérasset, Katerina Gkirtzou, Michael Cochez and
Jan-Christoph Kalo**

16 May 2026

©ELRA Language Resources Association (ELRA), 2026
These proceedings are licensed under a Creative Commons Attribution-
NonCommercial 4.0 International License (CC BY-NC 4.0)

ISBN 978-2-493814-49-4

Preface

Welcome to the KG & LLM Workshop, taking place during the post-conference day (16 May 2026) in Palma de Mallorca, Spain, co-located with the LREC 2026 conference.

Large Language Models (LLMs) have become foundational in Natural Language Processing (NLP), yet they continue to face challenges related to bias, hallucination, explainability, environmental impact, and the cost of training. Knowledge Graphs (KGs), in contrast, provide high-quality, interpretable, and reusable ontological and linguistic structures that support reasoning, fact-checking, and knowledge preservation.

This workshop brings together original research that leverages both Knowledge Graphs and Large Language Models in various domains of NLP and language resource development, aiming to better understand how both paradigms may be joined to take advantage of their complementary strengths.

The workshop did attract 36 submissions, which is very good for a one day workshop. This further attests that the workshop intention was clearly aligned with a lively trend in Natural Language Processing research.

After thorough reviews, the Programme Committee accepted 21 of them for publication. Nine of them will be presented orally and twelve as posters. We hope each of the accepted papers will reach a broad audience and lead to fruitful discussions, new ideas, and further collaborative research.

We tried to select a representative panel of approaches in our selection, and organise the oral sessions around three topics: "Knowledge Graphs and Evaluation", "Knowledge Graphs and Retrieval" and "Knowledge Graphs and Generation". The diversity of approaches shows a very lively research area, with diverse approaches leading to many open research questions.

We hope that you will all enjoy this workshop and gain many insights from it.

April 2026

Gilles Sérasset,
Katerina Gkirtzou,
Michael Cochez, and
Jan-Christoph Kalo

Acknowledgments: The KG & LLM workshop is proudly endorsed by 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>).

Organizing Committee

- Gilles Sérasset, Université Grenoble Alpes, France
- Katerina Gkirtzou, Athena Research Center, Greece
- Michael Cochez, Ellis Institute Finland & Åbo Akademi, Finland
- Jan-Christoph Kalo, University of Amsterdam, Netherlands

Program Committee

- Hanna Abi Akl, Inria Sophia-Antipolis, France
- Bradley Allen, University of Amsterdam, The Netherlands
- Russa Biswas, Aalborg University, Denmark
- Francesco Compagno, University of Amsterdam, The Netherlands
- Milan Dojchinovski, Czech Technical University in Prague and InfAI, Germany
- Radovan Garabik, L. Štúr Institute of Linguistics, Slovak Academy of Sciences, Slovakia
- Daniela Gifu, Romanian Academy, Iasi Branch, Romania
- Hugo Gonalo Oliveira, University of Coimbra, Portugal
- Jorge Gracia, University of Zaragoza, Spain
- Paul Groth, University of Amsterdam, The Netherlands
- Ilan Kernerman, Lexicala, Israel
- Aditi Khullar, Two Sigma, USA
- Xue Li, CWI, Amsterdam
- Chaya Liebeskind, Jerusalem College of Technology, Israel
- Chuangtao Ma, Aalborg University, Denmark
- Marco Passarotti, Università Cattolica del Sacro Cuore, Milan, Italy
- Heiko Paulheim, University of Mannheim, Germany
- Emrick Poncet, Université Grenoble Alpes, France
- Harald Sack, Karlsruhe Institute of Technology, Karlsruhe, Germany
- Didier Schwab, Université Grenoble Alpes, France
- Ranka Stanković, University of Belgrade, Serbia
- Riste Stojanov, Ss. Cyril and Methodius University – Skopje, Macedonia
- Andon Tchechmedjiev, \Ecole des Mines, Alès, France
- Dimitar Trajanov, Ss. Cyril and Methodius University – Skopje, Macedonia
- Nakanyseth Vuth, Université Grenoble Alpes, France
- Krzysztof Weceł, Poznań University of Economics and Business, Poland
- Slavko Zitnik, University of Ljubljana, Slovenia

Table of Contents

<i>Linguistic Initialization for Inductive Reasoning in Heterogeneous Knowledge Graphs</i> Daniele Pasquini, Danilo Croce and Roberto Basili	1
<i>OntoBook: Ontology-Grounded Synthetic Textbooks for Medical Encoder Pretraining</i> Rian Touchent and Éric de la Clergerie	11
<i>Conversational Control with Ontologies for Large Language Models: A Lightweight Framework for Constrained Generation</i> Barbara Gendron, Gael Guibon and Mathieu d'Aquin	20
<i>Is One Token All It Takes? Graph Pooling Tokens for LLM-based GraphQA</i> Ankit Grover, Lodovico Giarretta, Remi Bourgerie and Sarunas Girdzijauskas	36
<i>Improving Text2Cypher with Confidence-Based Test-Time Strategies</i> Rima Dessi and Makbule Gulcin Ozsoy	45
<i>Integrating Knowledge Graph and Large Language Models for Defining Business Strategies</i> Eleonora Ghizzota, Alex Jordan, Alessandro Petruzzelli, Lucia Siciliani, Giuseppe Spillo, Pierpaolo Basile, Davide Sola, Giovanni Scarso Borioli and Giovanni Semeraro	53
<i>GROUNDEDKG-RAG: Grounded Knowledge Graph Index for Long-document Question Answering</i> Tianyi Zhang and Andreas Marfurt	63
<i>Ontology-Guided Synthetic Data Generation for Low-Resource Information Extraction: A Case Study in IT Heritage Domain</i> Nakanyseth Vuth, Emrick Poncet, Gilles Sérasset, Didier Schwab, Caroline Djambian and Benjamin Lecouteux	73
<i>A Clinical SKOS Ontology and Evaluation Benchmark for LLM Query Generation over ICU Knowledge Graphs</i> Khurram Ali	82
<i>ReX-GG: A LLM Ensemble Pipeline for Relation-extraction and Graph Generation</i> Giacomo Magnifico and Eduard Barbu	93
<i>Graph Fusion across Languages Using Large Language Models</i> Kaung Myat Kyaw, Khush Agarwal and Jonathan Chan	102
<i>Efficient KG-Augmented RAG with Reusable Graph Community Summaries</i> Maha Karkout, Maria Andreevna Khodorchenko, Nikolay Alekseevich Butakov and Denis Nasonov	110
<i>Stack2Graph: A Structured Knowledge Representation of Stack Overflow Data for Retrieval-based Question Answering</i> Lukas Amadeus Kleybolte, Viviana Ventura and Alessandra Zarcone	120

<i>LLM-based Atomic Propositions Help Weak Extractors: Evaluation of a Propositioner for Triplet Extraction</i>	
Luc Pommeret, Thomas Gerald, Christophe Servan, Sahar Ghannay, Patrick Paroubek and Sophie Rosset	134
<i>Towards Knowledge Graph-Grounded Evaluation of Agentic LLMs on Cybersecurity Capture-the-Flag Challenges</i>	
Daniel Schlör, Marius Bohn, Maximilian Wolf, Kevin Bergner, Christian Goldschmied and Andreas Hotho	144
<i>End-to-End Graph Retrieval Pipeline for Specialized Domains</i>	
Haraldur Davidsson and Hazar Harmouch	155
<i>The Structure-Content Trade-off in Knowledge Graph Retrieval: A Diagnostic Study of Question Decomposition</i>	
Valentin Six, Gaël de Chalendar and Evan Dufraisie	166
<i>Quantifying Retrieval Quality in GraphRAG: A Schema-Agnostic Approach</i>	
Thibaud Vanmechelen, Alexandre Achten, Zaineb Gabsi and Sabri Skhiri	176
<i>A Wikidata-Based Framework to Measure Cross-Lingual Bias in Multilingual Large Language Models</i>	
Mouloud Iferroudjene, Lisa Poggel, Andrea Schimmenti, Duo Yang, Kanchan Shivashankar, Jan-Christoph Kalo and Marta Boscariol	190
<i>Evaluating Large Language Models for Strategic Knowledge Extraction in Capability-Based Planning</i>	
Hein C. Kolk, Julia García-Fernández, Julia Bronkhorst and Roos M. Bakker	210
<i>Large Language Models for Knowledge Graph Extraction: A Schema-Constrained Evaluation Framework</i>	
Markus Ilves, Eduard Barbu and Jaan Übi	221

Conference Program

Saturday, May 16, 2026

09:00–10:30 Session Oral 1: KG, LLMs and Generation

Chair: Michael Cochez

Linguistic Initialization for Inductive Reasoning in Heterogeneous Knowledge Graphs

Daniele Pasquini, Danilo Croce and Roberto Basili

OntoBook: Ontology-Grounded Synthetic Textbooks for Medical Encoder Pre-training

Rian Touchent and Éric de la Clergerie

Conversational Control with Ontologies for Large Language Models: A Lightweight Framework for Constrained Generation

Barbara Gendron, Gael Guibon and Mathieu d'Aquin

10:30–12:00 Session Poster: Poster Session

Chair: Gilles Sérasset

Is One Token All It Takes? Graph Pooling Tokens for LLM-based GraphQA

Ankit Grover, Lodovico Giarretta, Remi Bourgerie and Sarunas Girdzijauskas

Improving Text2Cypher with Confidence-Based Test-Time Strategies

Rima Dessi and Makbule Gulcin Ozsoy

Integrating Knowledge Graph and Large Language Models for Defining Business Strategies

Eleonora Ghizzota, Alex Jordan, Alessandro Petruzzelli, Lucia Siciliani, Giuseppe Spillo, Pierpaolo Basile, Davide Sola, Giovanni Scarso Borioli and Giovanni Semeraro

GROUNDEDKG-RAG: Grounded Knowledge Graph Index for Long-document Question Answering

Tianyi Zhang and Andreas Marfurt

Ontology-Guided Synthetic Data Generation for Low-Resource Information Extraction: A Case Study in IT Heritage Domain

Nakanyseth Vuth, Emrick Poncet, Gilles Sérasset, Didier Schwab, Caroline Djambian and Benjamin Lecouteux

A Clinical SKOS Ontology and Evaluation Benchmark for LLM Query Generation over ICU Knowledge Graphs

Khurrum Ali

Saturday, May 16, 2026 (continued)

ReX-GG: A LLM Ensemble Pipeline for Relation-extraction and Graph Generation

Giacomo Magnifico and Eduard Barbu

Graph Fusion across Languages Using Large Language Models

Kaung Myat Kyaw, Khush Agarwal and Jonathan Chan

Efficient KG-Augmented RAG with Reusable Graph Community Summaries

Maha Karkout, Maria Andreevna Khodorchenko, Nikolay Alekseevich Butakov and Denis Nasonov

Stack2Graph: A Structured Knowledge Representation of Stack Overflow Data for Retrieval-based Question Answering

Lukas Amadeus Kleybolte, Viviana Ventura and Alessandra Zarccone

LLM-based Atomic Propositions Help Weak Extractors: Evaluation of a Propositioner for Triplet Extraction

Luc Pommeret, Thomas Gerald, Christophe Servan, Sahar Ghannay, Patrick Paroubek and Sophie Rosset

Towards Knowledge Graph-Grounded Evaluation of Agentic LLMs on Cybersecurity Capture-the-Flag Challenges

Daniel Schlör, Marius Bohn, Maximilian Wolf, Kevin Bergner, Christian Goldschmied and Andreas Hotho

12:00–13:00

KG, LLMs and Retrieval

Chair: Jan-Christoph Kalo

End-to-End Graph Retrieval Pipeline for Specialized Domains

Haraldur Davidsson and Hazar Harmouch

The Structure-Content Trade-off in Knowledge Graph Retrieval: A Diagnostic Study of Question Decomposition

Valentin Six, Gaël de Chalendar and Evan Dufraisse

Saturday, May 16, 2026 (continued)

14:00–16:00 KG, LLMs and Evaluation

Chair: Katerina Gkirtzou

Quantifying Retrieval Quality in GraphRAG: A Schema-Agnostic Approach

Thibaud Vanmechelen, Alexandre Achten, Zaineb Gabsi and Sabri Skhiri

A Wikidata-Based Framework to Measure Cross-Lingual Bias in Multilingual Large Language Models

Mouloud Iferroudjene, Lisa Poggel, Andrea Schimmenti, Duo Yang, Kanchan Shivashankar, Jan-Christoph Kalo and Marta Boscarol

Evaluating Large Language Models for Strategic Knowledge Extraction in Capability-Based Planning

Hein C. Kolk, Julia García-Fernández, Julia Bronkhorst and Roos M. Bakker

Large Language Models for Knowledge Graph Extraction: A Schema-Constrained Evaluation Framework

Markus Ilves, Eduard Barbu and Jaan Übi

