2nd International Workshop on Data Management for Knowledge Graphs (DMKG 2026)

Ruben Taelman[1], Olaf Hartig[2] and Katja Hose[3]

[1]Ghent University, IDLab, Department of Electronics and Information Systems, imec, 9052 Ghent, Belgium
[2]Linköping University, Department of Computer and Information Science, 581 83 Linköping, Sweden
[3]TU Wien, Institute of Logic and Computation, 1040 Vienna, Austria

Abstract

Knowledge graphs keep growing, and so does the gap between the data we publish and the infrastructure we have to manage it. Storing, indexing, validating, and querying knowledge graphs at scale remains hard, and it becomes harder still when the data is spread over multiple autonomous sources. The 2nd International Workshop on Data Management for Knowledge Graphs (DMKG 2026) brings together researchers, engineers, and practitioners to present and discuss work on these problems. This preface describes the scope of the workshop, its topics, the reviewing process, the keynote, and the people behind the event.

Keywords

Data management, Knowledge graphs, Semantic Web, Linked Data, Storage and indexing, Querying and benchmarking, Analytics and exploration

1. Preface

Knowledge graphs have been adopted widely over the past years, both in the open data domain and in industry. This growth is good news for the field, but it also means that data management solutions have to cope with ever larger amounts of data, and with data that no longer sits in a single place. Systems are expected to store and index graphs efficiently, to validate them, to keep track of how they change over time, and to answer queries over sources that are owned and operated by different parties. These requirements do not always line up with the tools we have, which is why the topic deserves a venue of its own.

The DMKG workshop invites novel research and advances in scalable data management solutions for large-scale knowledge graphs, covering centralized, distributed, federated, and decentralized settings. Alongside full research papers, we welcome vision papers, negative results, and system descriptions, because early ideas and honest reports of what did not work are exactly the kind of material that benefits from discussion in a room full of people working on the same problems. The main goal of the workshop is to bring together early-stage and established researchers as well as industrial partners, and to give them a place to compare notes.

DMKG 2026 is the second edition of the workshop. The first edition was held at the 20th Extended Semantic Web Conference (ESWC 2023) in Hersonissos, Greece [1]. This second edition is co-located with the 25th International Semantic Web Conference (ISWC 2026) in Bari, Italy, and takes place on October 25, 2026, as a full-day workshop. We chose to co-locate with ISWC in order to reach a wide audience of knowledge graph researchers, data scientists, and industry participants.

DMKG'26: 2nd International Workshop on Data Management for Knowledge Graphs, co-located with the 25th International Semantic Web Conference (ISWC 2026), October 25, 2026, Bari, Italy
ruben.taelman@ugent.be (R. Taelman); olaf.hartig@liu.se (O. Hartig); katja.hose@tuwien.ac.at (K. Hose)
https://www.rubensworks.net/ (R. Taelman); https://olafhartig.de/ (O. Hartig); http://katja-hose.de/ (K. Hose)
0000-0001-5118-256X (R. Taelman); 0000-0002-1741-2090 (O. Hartig); 0000-0001-7025-8099 (K. Hose)

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

CEUR-WS.org/Vol-4322/preface.pdf

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073
published 2026-11-09

2. Topics of Interest

The call for papers solicited contributions on all aspects of data management for knowledge graphs. We listed the following topics, without treating the list as exhaustive.

Storage and Management

- Storing and indexing knowledge graphs
- Partitioning knowledge graphs
- Decentralized, distributed and federated knowledge graph storage
- Graph databases and NoSQL
- Archiving and versioning
- Representing knowledge graphs using RDF 1.2

Analytics and Exploration

- Knowledge graph validation (SHACL/ShEx)
- Graph schema discovery and exploration
- Large-scale knowledge graph analytics (GraphX, Giraph, Pregel, …)

Querying and Benchmarking

- Efficient query processing
- Distributed and federated querying over knowledge graphs
- Querying over streaming graphs
- Benchmarking data management systems for knowledge graphs
- Querying knowledge graphs using SPARQL 1.2

3. Submissions and Reviewing

We accepted three categories of submissions, with page limits including references: research papers of up to 12 pages, which also covered vision papers; short papers of up to 6 pages, for position papers, negative results, and descriptions of systems, libraries, APIs, and datasets; and demo or poster papers of up to 4 pages. All submissions had to be original and could not be under simultaneous review elsewhere. Papers were submitted through EasyChair in the single-column CEURART style. We also allowed alternative formats such as self-hosted HTML pages.

The workshop schedule is as follows.

- Abstract deadline: July 17, 2026
- Paper submission deadline: July 24, 2026
- Notification of acceptance: August 21, 2026
- Camera-ready version due: September 11, 2026
- Workshop: October 25, 2026

We received 28 submissions in total, of which 24 were valid and went into review. Every valid submission was reviewed by members of the program committee, with most papers receiving three reviews and none receiving fewer than two. Reviews were not anonymous to the organizers, and submissions were not anonymized. Reviewer names were hidden from the authors. Papers were judged on significance, novelty, originality, technical contribution, clarity of the writing, and relevance to the topics of the workshop. Based on these reviews we accepted 16 papers, which amounts to an acceptance rate of 67% over the valid submissions. All accepted papers are included in this volume. Of the accepted papers, the majority were research papers, with the remainder consisting of short papers.

4. Keynote

The program for DMKG 2026 includes a keynote by Pascal Molli, Full Professor at the University of Nantes and head of the GDD team at LS2N, entitled *Can We Finally Query Decentralized Knowledge Graphs?*

The talk revisits what the Semantic Web set out to do and compares it with what we have built so far. Pascal claims that the aim was never just to put structured data online. Instead, the point was that machines should be able to work with knowledge that sits in many different places. According to him, Linked Data made this practical, and knowledge graphs are by now a common way to organize and combine data. What exists today, however, is mostly a large pile of datasets, identifiers, vocabularies, and endpoints rather than something that can be queried as a whole. Locating the relevant sources, running queries across them, deciding how much of what they return can be trusted, and respecting the conditions their owners attach are all hard once those sources are operated by different parties. He claims that language models can take some of this work off the hands of users and software agents, but they leave the underlying data management problems untouched. Pascal will discuss which of these obstacles weigh most heavily, and will argue that querying decentralized knowledge graphs reliably requires several elements at once, among them precise queries, approximate source discovery, federation, provenance, access control, and interfaces supported by language models.

5. Organizing Committee

- Ruben Taelman, Ghent University – imec, Belgium
- Olaf Hartig, Linköping University, Sweden
- Katja Hose, TU Wien, Austria

6. Program Committee

We thank the following people for serving on the program committee of DMKG 2026.

- Aidan Hogan, DCC, Universidad de Chile
- Andreas Harth, University of Erlangen-Nuremberg and Fraunhofer IIS
- Andreas Schwarte, metaphacts
- Anisa Rula, University of Brescia
- Beatriz Esteves, Universidad Politécnica de Madrid
- Bryan-Elliott Tam, Ghent University – imec
- Carlos Buil Aranda, Philip Morris
- Daniel Fernández Álvarez, University of Oviedo
- Dominik Tomaszuk, University of Białystok
- Hala Skaf-Molli, University of Nantes – LS2N
- Jesse Wright, University of Oxford
- Jose Emilio Labra Gayo, University of Oviedo
- Julián Arenas-Guerrero, Universidad Politécnica de Madrid
- Luis Galarraga, INRIA/IRISA
- Maria-Esther Vidal, Leibniz Information Centre for Science and Technology (TIB)
- Maribel Acosta, TU Munich
- Maxime Jakubowski, TU Wien
- Olivier Curé, Université Gustave Eiffel
- Oscar Corcho, Universidad Politécnica de Madrid
- Pascal Molli, University of Nantes – LS2N

- Philipp D. Rohde, Leibniz Universität Hannover
- Pierre-Antoine Champin, W3C
- Pieter Bonte, KU Leuven
- Robin Keskisärkkä, Linköping University
- Ruben Eschauzier, Ghent University – imec
- Sebastián Ferrada, DCC, Universidad de Chile

Acknowledgments

We thank the authors for submitting their work to DMKG 2026, and the members of the program committee for the reviews they wrote under a tight schedule. The workshop would not be possible without them. We also thank Pascal Molli for accepting our invitation to give the keynote, the ISWC 2026 workshop chairs and local organizers for hosting us in Bari, and everyone who takes part in the discussion at the workshop.

Declaration on Generative AI

During the preparation of this work, the authors used a generative AI assistant in order to draft and copy-edit parts of this preface. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of this publication.

References

[1] C. Aebeloe, A. Azzam, O. Hartig, K. Hose (Eds.), Proceedings of the 1st International Workshop on Data Management for Knowledge Graphs (DMKG 2023) co-located with the 20th Extended Semantic Web Conference (ESWC 2023), Hersonissos, Greece, May 29, 2023, Vol. 3443 of CEUR Workshop Proceedings, CEUR-WS.org, 2023. URL: https://ceur-ws.org/Vol-3443/.