The 5th International Workshop on Natural Language Processing for Knowledge Graph Construction (NLP4KGC 2026)

Edlira Vakaj[1], Mohamed Ragab[1], Youssef Senousy[1], Nandana Mihindukulasooriya[2], Rizou Matina[3] and Sanju Tiwari[4,5]

[1]Birmingham City University, Birmingham, United Kingdom
[2]IBM Research, New York, United States
[3]SingularLogic S.A., Greece
[4]Sharda University, India
[5]TIB, Hannover, Germany

Abstract

This preface introduces the proceedings of the 5th International Workshop on Natural Language Processing for Knowledge Graph Construction (NLP4KGC 2026), held on 26 October 2026 in Bari, Italy, co-located with the 25th International Semantic Web Conference (ISWC 2026). The workshop brought together researchers and practitioners exploring the integration of natural language processing (NLP), knowledge graphs (KGs), and artificial intelligence (AI). Its central themes included entity and relation extraction, entity linking, ontology learning, large language model (LLM)-based knowledge graph construction, retrieval-augmented generation (RAG), GraphRAG, and responsible AI. The proceedings feature 10 accepted contributions, reflecting recent advances, practical applications, and emerging challenges in constructing reliable, explainable, and trustworthy knowledge graphs from natural language and heterogeneous data sources.

Keywords

Natural Language Processing, Knowledge Graphs, Semantics, Large Language Models, Graph Neural Networks, Responsible AI

CEUR-WS.org/Vol-4335/preface.pdf

1. Introduction

The 5th International Workshop on Natural Language Processing for Knowledge Graph Construction (NLP4KGC 2026) was held on 26 October 2026 in Bari, Italy, as a half-day workshop co-located with the International Semantic Web Conference (ISWC 2026). The workshop brought together researchers and practitioners working at the intersection of natural language processing (NLP), knowledge graphs (KGs), and responsible artificial intelligence (AI). Building on four previous editions, NLP4KGC continued to provide a forum for sharing new methods, applications, and open research challenges in the construction of knowledge graphs from natural language and other heterogeneous sources.

Knowledge graphs support structured and semantically rich representations of information, but constructing them at scale remains challenging [1]. Key steps include entity recognition and resolution, relation extraction and linking, ontology learning, and the validation of extracted knowledge [2]. Recent progress in large language models (LLMs) and emerging agentic AI systems has opened new possibilities for automating these tasks while also introducing questions about reliability, interpretability, grounding, and the quality of constructed graphs. The workshop therefore examined how NLP, semantic technologies, and knowledge graphs can be combined to support more accurate, explainable, and trustworthy AI systems. Recent studies have also demonstrated the value of knowledge graphs in improving the accuracy of LLM-based question answering by providing structured semantic context [3].

5th International Workshop on Natural Language Processing for Knowledge Graph Construction (NLP4KGC 2026), co-located with the 25th International Semantic Web Conference (ISWC 2026), 26 October 2026, Bari, Italy

edlira.vakaj@bcu.ac.uk (E. Vakaj); mohamed.ragab@bcu.ac.uk (M. Ragab); youssef.senousy@mail.bcu.ac.uk (Y. Senousy); nandana@ibm.com (N. Mihindukulasooriya); srizou@singularlogic.eu (R. Matina); tiwarisanju18@ieee.org (S. Tiwari)



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

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073
published 2026-11-19

The workshop's scope encompassed three complementary areas: (i) natural language understanding, including entity extraction, disambiguation, and relation linking; (ii) knowledge graph construction, including ontology learning, neural approaches, LLM-based agents and tools, retrieval-augmented generation (RAG), and graph-based RAG (GraphRAG); and (iii) responsible AI, including the use of knowledge graphs for bias mitigation, interpretability, and explainability. The programme reflected these themes through sessions on LLM agents and knowledge graph construction, and on knowledge graph applications, entity linking, and reasoning.

2. Previous Editions of NLP4KGC

NLP4KGC 2026 builds upon four previous editions of the workshop, established in 2022 to foster collaboration between the natural language processing and knowledge graph research communities. The second edition was co-located with SEMANTiCS 2023 in Leipzig, Germany, with its proceedings published in CEUR Workshop Proceedings (Vol. 3510). The third edition took place at SEMANTiCS 2024 in Amsterdam, the Netherlands, with its proceedings published in CEUR Workshop Proceedings (Vol. 3874). The fourth edition was held in conjunction with The Web Conference 2025 in Sydney, Australia, with contributions published in the ACM Web Conference Companion Proceedings.

Across these editions, the workshop has explored advances in information extraction, entity and relation linking, ontology learning, and automated knowledge graph construction. More recently, its scope has expanded to encompass large language models, agentic AI, and responsible AI, reflecting the evolving relationship between language technologies and structured knowledge representation.

Previous proceedings are available at: https://ceur-ws.org/Vol-3510/ (2023), https://ceur-ws.org/Vol-3874/ (2024).

3. Submissions, review, and selection

The workshop invited original full research papers of up to 12 pages and short papers of up to 6 pages, including work in progress, position papers, datasets, and system demonstrations. The abstract registration deadline was 17 July 2026, and the paper submission deadline was 24 July 2026. Acceptance notifications were issued on 21 August 2026, with camera-ready versions due on 11 September 2026. Selecting qualified reviewers with relevant expertise while avoiding potential conflicts of interest is an important aspect of ensuring a rigorous and fair peer-review process [4].

NLP4KGC 2026 received 16 submissions, of which **10 papers** were accepted for presentation. All submissions underwent double-blind peer review, with each paper evaluated by at least two programme committee members. Acceptance decisions were based on technical soundness, novelty, relevance, clarity, and significance of the contributions. Selection considered technical soundness and novelty, relevance to the workshop, clarity of presentation, and significance of the contributions. The accepted contributions form the NLP4KGC 2026 part of these proceedings.

4. Accepted Papers and Workshop Programme

The NLP4KGC 2026 programme featured 10 accepted papers organised into two thematic sessions, reflecting both methodological advances and practical applications of natural language processing for knowledge graph construction.

1. **Session 1: LLM Agents and Knowledge Graph Construction**
   The first session comprised six papers investigating the integration of large language models and agentic AI into knowledge graph construction workflows. The accepted papers were:
   - *Governance Architecture for Knowledge Graph Construction Using LLM Agents.* John Strassner.
   - *AutoKGen: Autonomous Knowledge Graph Construction through Agentic AI-driven Code Generation and Ontology-based Validation.* Giorgos Anagnostou and Konstantinos Kotis.

- *SHACL-in-the-Loop: Enhancing Structural Conformance in Agentic Knowledge Graph Construction.* Vladislav Snytko, Rishabh Jakhar, and Remzi Celebi.
- *Where, How, and What to Ground: Evaluating Ontology-Guided LLM-based Knowledge Graph Construction.* Itisha Yadav, Roman Klinger, and Diana Peters.
- *When Should We Build the Graph? An Amortized Pareto View of Demand-Driven Knowledge Graph Construction.* Niharika Dsouza, Satyananda Kashyap, Nandana Mihindukulasooriya, Jui-Chien Lin, Oshani Seneviratne, and Horst Samulowitz.
- *Knowledge Graph Construction from Online Discussions: Modeling Context-Dependent Opinions and Their Evolution.* Ren Shimoide, Rikuto Sasaki, and Kazuhiro Takeuchi.

2. **Session 2: Knowledge Graph Applications, Entity Linking and Reasoning**
   The second session featured four papers demonstrating knowledge graph technologies across different domains, including healthcare, bias mitigation, e-commerce, fact verification, and financial document analysis. The accepted papers were:

   - *A Multilingual Workflow for Constructing Semantic Bridge Edges between Local Clinical Knowledge Graphs and Standard Clinical Ontologies.* Artid Boonrerng, Nouh Sabri Elmitwally, Mohamed Ragab, Edlira Vakaj, and Hadeel Saadany.
   - *Combining Knowledge Graphs and Node Embeddings for Gender Bias Detection and Mitigation in Greek Language.* Andreas Soularidis, Dimitris Doumanas, and Konstantinos Kotis.
   - *E4EL: CRF-Based Efficient Zero-Shot End-to-End Framework for E-Commerce Entity Linking.* Hsiao-Ying Lu, Zeyu Wang, Zhejun Shen, and Petar Ristoski.
   - *Agentic RAG with Self-Validation for Conversational Financial Document Analysis.* Raghavendranath Kandula.

5. Workshop chairs and Organizers

- Edlira Vakaj — Birmingham City University, United Kingdom.
- Nandana Mihindukulasooriya — IBM Research, United States.
- Rizou Matina — SingularLogic S.A., Greece.
- Sanju Tiwari — Sharda University, India, and TIB, Germany.
- Mohamed Ragab — Birmingham City University, United Kingdom.
- Youssef Senousy — Birmingham City University, United Kingdom.

6. Programme committee

Sima Iranmanesh (Birmingham City University, United Kingdom); Sule Yildirim (Norwegian University of Science and Technology, Norway); Veronika Heimsbakk (KPMG, Norway); Amna Dridi (Birmingham City University, United Kingdom); Hansi Hettiarachchi (Lancaster University, United Kingdom); Thomas Beach (Cardiff University, United Kingdom); Hadeel Sadanay (Birmingham City University, United Kingdom); Vladimir Alexiev (Ontotext, Bulgaria); Jagdev Bhogal (Birmingham City University, United Kingdom); Sergio Méndez (Australian National University, Australia); Arunav Das (King's College London, United Kingdom); Ellie Young (Common Action, United States); Hossein Ghomeshi (Shipamax, United Kingdom); Ryan Mcgranaghan (NASA, United States); Stamartina Rizou (SingularLogic, Greece); Gaetano Rossiello (IBM, United States); Katariina Kari (The Knowledge Graph Academy, Finland); He Tan (Jönköping University, Sweden); Diellza Elshani (University of Stuttgart, Germany); Ronakkumar Panchal (Gujarat Technology University, India); Anelia Kurteva (University of Birmingham, United Kingdom); and Bohui Zhang (King's College London, United Kingdom).

Acknowledgements

We thank all authors for their contributions to NLP4KGC 2026, the programme committee members for their careful reviews and constructive feedback, and all participants for their engagement and discussions. We also acknowledge the ISWC 2026 organisers for their support in hosting the workshop.

References

[1] A. Hogan et al., "Knowledge Graphs," *ACM Computing Surveys*, vol. 54, no. 4, Article 71, 2021. https://doi.org/10.1145/3447772

[2] L. Zhong, J. Wu, Q. Li, H. Peng, and X. Wu, "A Comprehensive Survey on Automatic Knowledge Graph Construction," *ACM Computing Surveys*, vol. 56, no. 4, Article 94, pp. 1–62, 2024. https://doi.org/10.1145/3618295

[3] J. Sequeda, D. Allemang, and B. Jacob, "A Benchmark to Understand the Role of Knowledge Graphs on Large Language Model's Accuracy for Question Answering on Enterprise SQL Databases," in *Proceedings of the 7th Joint Workshop on Graph Data Management Experiences & Systems (GRADES) and Network Data Analytics (NDA)*, Article 5, pp. 1–12, ACM, 2024. https://doi.org/10.1145/3661304.3661901

[4] M. R. Moawad, M. Maher, A. Awad, and S. Sakr, "MINARET: A Recommendation Framework for Scientific Reviewers," in *Proceedings of the 22nd International Conference on Extending Database Technology (EDBT)*, pp. 538–541, 2019. https://doi.org/10.5441/002/edbt.2019.50