Preface for the 2nd International Workshop on Knowledge Graphs for Responsible AI (KG-STAR 2025)

Edlira Vakaj[1], Nandana Mihindukulasooriya[2], Manas Gaur[3] and Arijit Khan[4]

[1]Birmingham City University, UK
[2]IBM Research, New York, USA
[3]University of Maryland Baltimore County, USA
[4]Aalborg University, Denmark

Abstract

Responsible AI hinges on formalizing fairness, transparency, accountability, and inclusivity throughout the AI lifecycle–a need that grows ever more urgent as generative models scale in capability and complexity. Knowledge Graphs (KGs) offer a structured semantic backbone that enriches generative AI by injecting contextual priors, elucidating model inferences, and curbing bias propagation. By encoding entities and relations, KGs enable interpretable reasoning paths–allowing practitioners to audit decision logic–and unify diverse data sources to ensure comprehensive, equitable coverage. This semantic scaffolding thus underpins responsible AI by making generative outputs more explainable, traceable, and aligned with ethical safeguards. The 2nd International Workshop on Knowledge Graphs for Responsible AI (KG-STAR 2025) focused on the role of KGs in promoting Responsible AI principles and creating a cooperative space for researchers, practitioners, and policymakers to exchange insights and enhance their understanding of KGs' impact on achieving Responsible AI solutions. It aimed to facilitate collaboration and idea-sharing to advance the understanding of how KGs can contribute to Responsible AI. The workshop featured two thought-provoking keynote talks and four insightful research presentations exploring the intersection of Knowledge Graphs and Responsible AI.

Keywords

Knowledge Graphs, Responsible AI, Explainable AI, Bias and Fairness

1. Introduction

Industry leaders like QinetiQ are paving the way for safe and ethical AI use in government agencies. They're doing this by developing an AI Assurance Framework, which helps ensure responsible deployment of AI and large language models (LLMs) within organizations like UK Defence. Knowledge Graphs (KGs) have been identified as key enablers for explainability in AI [1, 2, 3, 4, 5]. Knowledge Graphs and their byproducts, such as KG embeddings, can have their own implicit biases that have to be taken into account [6, 7, 8], e.g., the KG-BIAS workshop [9] at Automated Knowledge Base Construction (AKBC) focuses on identifying biases in automatic KG construction. A deeper integration of Knowledge graphs into AI is seen through NeuroSymbolic AI. This innovative approach merges the power of statistical machine learning (black-box neural network technologies, e.g., Deep Learning and LLMs), known for its data-driven predictions, with structured symbolic systems such as Knowledge Graphs [10, 11]. The Knowledge-infused Learning workshop [1] and the NeuroSymbolic AI [2] Workshop have been central to building a community for transparent and trustworthy AI [12]. Our workshop takes this vision a step further by inviting researchers to explore how to use knowledge graphs to detect hallucinations, falsehoods, contradictions, and knowledge gaps in LLM outputs and contribute to Responsible AI principles (e.g., Fairness, Bias, Consistency, Explainability [13, 14, 15]). We believe that ISWC (The

KG-STAR 2025: 2nd International Workshop on Knowledge Graphs for Responsible AI, June 1–5, 2025, Portorož, Slovenia.
edlira.vakaj@bcu.ac.uk (E. Vakaj); nandana@ibm.com (N. Mihindukulasooriya); manas@umbc.edu (M. Gaur); arijitk@cs.aau.dk (A. Khan)
0000-0002-0712-0959 (E. Vakaj); 0000-0003-1707-4842 (N. Mihindukulasooriya); 0000-0002-5411-2230 (M. Gaur); 0000-0002-7312-6312 (A. Khan)

[1]https://aiisc.ai/kiml2023
[2]https://sites.google.com/view/nesy2023

International Semantic Web Conference) can be an excellent venue to foster the discussion about using Knowledge Graphs for responsible AI.

2. Topics of Interest

We invited submissions of original research, case studies, and position papers on topics related to Knowledge Graphs and their applications in advancing Responsible AI. The workshop explores the intersection of Knowledge Graphs and ethical considerations in AI development. Submissions may include, but are not limited to, the following topics:

Knowledge Graphs for Bias Mitigation:

- Techniques and methodologies for using Knowledge Graphs to identify and mitigate biases in AI models.
- Case studies demonstrating the successful application of Knowledge Graphs in addressing bias challenges.

Interpretability and Explainability:

- Approaches to enhancing the interpretability and explainability of black-box AI models through integrating Knowledge Graphs.
- Evaluating the effectiveness of Knowledge Graphs in making AI decision-making processes more transparent.

Privacy-Preserving Knowledge Graphs:

- Methods for constructing Knowledge Graphs that prioritize privacy and comply with data protection regulations.
- Applications of Knowledge Graphs in privacy-preserving AI systems.

Fairness in AI with Knowledge Graphs:

- How Knowledge Graphs contribute to ensuring fairness in AI applications.
- Techniques for using Knowledge Graphs and their embeddings to identify and rectify unfair biases in AI models.

Ethical Considerations in Knowledge Graph Construction:

- Ethical challenges in the creation and maintenance of Knowledge Graphs.
- Best practices for ensuring responsible and ethical Knowledge Graph development.
- Real-world applications of Knowledge Graphs in Responsible AI.

Integration of Large Language Models (LLMs) and Knowledge Graphs (KGs):

- Enhancing LLMs' accuracy, consistency, reducing hallucinations and harmful contents generation, fake news detection, fact checking, etc. with knowledge-grounded techniques.
- Enhancing the interoperability of KG downstream tasks through LLMs' natural language interfaces, transferability, and generalization capacity.

3. Keynotes

- **Keynote 1**: Responsible AI with LLMs: Why We Need Knowledge Graphs.
  Speaker: Dr. Sven Hertling, University of Mannheim, Mannheim, Germany
  Abstract: This keynote provided an overview of cases where knowledge graphs are relevant in Responsible AI. It furthermore highlighted what still needs to be solved in the knowledge graph community to increase the adoption of KGs in the industry. The talk was structured around four key topics: reliability, privacy, fairness, and explainability.

- **Keynote 2**: The Role Evolution of KGs in Synthesizing with LLMs: From Background Knowledge to Joint Reasoning
  Speaker: Dr. Chuangtao Ma, Aalborg University, Aalborg, Denmark.
  Abstract: Knowledge Graphs (KGs), as graph-based structured knowledge, maintain the rich relationships among the trackable and verifiable facts and evidence, which have been investigated to address the inherent limitations of large language models (LLMs), such as hallucinations, limited reasoning capabilities, and interoperability. Recent years have witnessed the role of KGs in synthesizing with LLMs evolving from background knowledge to joint reasoning. This work aims to give a brief introduction to the recent works in augmenting LLMs with KGs and highlights the evolving role of KGs, i.e., from KGs serving as passive background knowledge to actively getting involved in joint reasoning processes with LLMs. It summarizes the key techniques, strengths, limitations, and KG requirements of the approaches with different KG roles in augmenting LLMs with KGs, and their applications in several downstream tasks. It also discusses the open challenges and future directions for developing more efficient and trustworthy reasoning over LLMs and KGs.

4. Accepted Papers

- **Title:** ELSA Knowledge Graphs for Animal Treatment Recommendations
  Authors: Varsha Kalidas, André Gomes Regino, Anderson Rossanez, Julio Cesar dos Reis, Tarek Alskaif, and Ricardo da Silva Torres

- **Title:** Understanding Vulnerable Road User Behavior using Spatio-Temporal Knowledge Graphs
  Authors: He Tan and Erick Escandon Bailon

- **Title:** Towards Supporting AI System Engineering with an Extended Boxology Notation
  Authors: Fajar J. Ekaputra, Alexander Prock, and Elmar Kiesling

- **Title:** GOSt-MT: A Knowledge Graph for Occupation-related Gender Biases in Machine Translation
  Authors: Orfeas Menis Mastromichalakis, Giorgos Filandrianos, Eva Tsouparopoulou, Dimitris Parsanoglou, Maria Symeonaki, and Giorgos Stamou

5. Program Committee Members

- Ernesto Jiménez-Ruiz, City St George's, University of London, UK
- Simon Razniewski, TU Dresden, Germany
- Ahmed Zalouk, Birmingham City University, UK
- He Tan, Jönköping University, Sweden
- Hansi Hettiarachchi, Lancaster University, UK
- Muhammad Afzal, Birmingham City University, UK

- Jagdev Bhogal, Birmingham City University, UK
- Gaetano Rossiello, IBM Research, USA
- Edlira Vakaj, Birmingham City University, UK
- Nandana Mihindukulasooriya, IBM Research, USA
- Manas Gaur, University of Maryland Baltimore County, USA
- Arijit Khan, Aalborg University, Denmark

6. Workshop Organizers

- **Edlira Vakaj** is leading the Natural Language Processing AI Lab and is an Associate Professor of Neuro-Symbolic AI at Birmingham City University, UK. She conducts research in multidisciplinary projects focusing on semantic web, knowledge graphs, AI, and semantic data spaces. Edlira is the principal investigator of the ACCORD Horizon project and is engaged in several European and UK-funded projects of various domains where Semantic Web Technologies are applied, such as Renewable Energy (FP7 RENESENG), Industrialised Construction and Industry 4.0 (Innovate UK DfMA Knowledge Transfer Partnership), Higher Education and Youth (Erasmus + Capacity Building, Learning mobility of individuals, Cooperation for innovation and the exchange of good practices action). Dr. Vakaj is an active member of various communities such as The Alan Turing Knowledge Graph Community, Linked Building Data, Knowledge Graph Creation and Common Action. She organized the 1st NLP4KGC workshop in conjunction with the WWW 2023 conference and the 2nd NLP4KGc workshop in the SEMANTICs 2023 conference.
- **Nandana Mihindukulasooriya** is a senior research scientist at IBM Research, New York, USA. Prior to that, he was a postdoctoral fellow at MIT-IBM Watson AI Lab. He holds a PhD in artificial intelligence from Universidad Politecnica De Madrid, Spain. His research interests include relation extraction and linking, information extraction, knowledge representation and reasoning, and neuro-symbolic AI. Dr. Mihindukulasooriya is also an inventor of several patents related to the same area. He has won several prestigious international awards for the impact of his research contributions to IBM. Nandana was co-organizer of several workshops including the International Workshop on Knowledge Graph Generation from Text at ESWC 2022, International Workshop on Knowledge Graph Summarization at ISWC 2022, SMART Semantic Web Challenge at ISWC 2020-2022, International Workshop on Knowledge Graphs on Travel and Tourism at IWCE 2018, International Workshop on Quality Engineering Meets Knowledge Graph at K-CAP 2017. As a PC member, he has contributed to more than 35 international conferences, including AAAI, IJCAI, ISWC, ESWC, ACL, EMNLP, K-CAP, SEMANTiCS, and others.
- **Manas Gaur** is an assistant professor at the University of Maryland Baltimore County (UMBC) in the Department of Computer Science and Electrical Engineering. He directs the Knowledge-infused AI and Inference (KAI2) lab. Prior to his academic pursuits, Dr. Gaur served as the lead research scientist in Natural Language Processing at the AI Centre in Samsung Research America and he also contributed as a visiting researcher at the prestigious Alan Turing Institute. Dr. Gaur is acknowledged with AAAI New Faculty for 2023 and IEEE Intelligent Systems and Internet Computing best paper awards. He is serving as Guest editor for the IEEE Special Issue on Knowledge-infused Learning and ACM Health Special Issue on Large Language Models, Conversational Systems, and Generative AI in Health.
- **Arijit Khan** is an IEEE senior member, an ACM distinguished speaker, and an associate professor in the Department of Computer Science, Aalborg University, Denmark. He earned his PhD from the Department of Computer Science, University of California, Santa Barbara, USA, and did a post-doc in the Systems group at ETH Zurich, Switzerland. Arijit is the recipient of the prestigious IBM PhD Fellowship in 2012-13, a VLDB Distinguished Reviewer award (2022), and two SIGMOD Distinguished PC awards (2024, 2025). He published more than 100 papers in premier databases and data mining conferences and journals including ACM SIGMOD, VLDB, ICLR, ACM KDD, IEEE TKDE, IEEE ICDE, SIAM SDM, USENIX ATC, EDBT, The Web Conference

(WWW), ACM WSDM, ACM CIKM, ACM TKDD, and ACM SIGMOD Record. He served in the program committee of SIGMOD, VLDB, ICDE, ICDM, EDBT, CIKM, and in the senior program committee of KDD and WWW. Dr. Khan served as the co-chair of Big-O(Q) workshop co-located with VLDB 2015 and LLM+KG workshop co-located with VLDB 2024.

References

[1] F. Lécué, On the role of knowledge graphs in explainable AI, Semantic Web 11 (2020) 41–51. URL: https://doi.org/10.3233/SW-190374. doi:10.3233/SW-190374.

[2] I. Tiddi, S. Schlobach, Knowledge graphs as tools for explainable machine learning: A survey, Artificial Intelligence 302 (2022) 103627. URL: https://doi.org/10.1016/j.artint.2021.103627. doi:10.1016/j.artint.2021.103627.

[3] E. Rajabi, K. Etminani, Knowledge-graph-based explainable ai: A systematic review, Journal of Information Science (2022).

[4] M. Gaur, K. Faldu, A. P. Sheth, Semantics of the black-box: Can knowledge graphs help make deep learning systems more interpretable and explainable?, IEEE Internet Computing 25 (2021) 51–59.

[5] C. Ma, Y. Chen, T. Wu, A. Khan, H. Wang, Unifying large language models and knowledge graphs for question answering: Recent advances and opportunities, in: International Conference on Extending Database Technology, (EDBT), 2025, pp. 1174–1177.

[6] G. Demartini, Implicit bias in crowdsourced knowledge graphs, in: The World Wide Web Conference (WWW), 2019, pp. 624–630.

[7] S. Bourli, E. Pitoura, Bias in knowledge graph embeddings, in: IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), 2020, pp. 6–10.

[8] R. Wickramarachchi, C. A. Henson, A. P. Sheth, An evaluation of knowledge graph embeddings for autonomous driving data: Experience and practice, in: AAAI Symposium on Combining Machine Learning and Knowledge Engineering in Practice (AAAI-MAKE), Volume I, volume 2600 of *CEUR Workshop Proceedings*, 2020.

[9] T. Safavi, E. Meij, F. Özcan, M. Redi, G. Demartini, C. Xiong, Report on the first workshop on bias in automatic knowledge graph construction at akbc 2020, SIGIR Forum 54 (2021). URL: https://doi.org/10.1145/3483382.3483393. doi:10.1145/3483382.3483393.

[10] A. P. Sheth, K. Roy, M. Gaur, Neurosymbolic artificial intelligence (why, what, and how), IEEE Intelligent Systems 38 (2023) 56–62.

[11] Y. Wang, A. Khan, T. Wu, J. Jin, H. Yan, Semantic guided and response times bounded top-k similarity search over knowledge graphs, in: IEEE International Conference on Data Engineering (ICDE), 2020.

[12] M. Gaur, E. Tsamoura, S. Sreedharan, S. Mittal, Kil 2023: 3rd international workshop on knowledge-infused learning, in: Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2023, pp. 5857–5858.

[13] M. Gaur, A. Sheth, Building trustworthy neurosymbolic ai systems: Consistency, reliability, explainability, and safety, AI Magazine (2024).

[14] A. Khan, E. B. Mobaraki, Interpretability methods for graph neural networks, in: IEEE International Conference on Data Science and Advanced Analytics (DSAA), 2023.

[15] T. Chen, D. Qiu, Y. Wu, A. Khan, X. Ke, Y. Gao, View-based explanations for graph neural networks, Proc. ACM Manag. Data 2 (2024) 40:1–40:27.