

KaLLM 2024

**The First Workshop on Knowledge Graphs and Large
Language Models**

Proceedings of the Workshop

August 15, 2024

The KaLLM organizers gratefully acknowledge the support from the following sponsors.

Best Paper Award by

Bloomberg

Engineering

©2024 Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)
317 Sidney Baker St. S
Suite 400 - 134
Kerrville, TX 78028
USA
Tel: +1-855-225-1962
acl@aclweb.org

ISBN 979-8-89176-147-6

Introduction

Welcome to KaLLM 2024, the first workshop on Knowledge Graphs and Large Language Models.

Large Language Models (LLMs) have revolutionised the field of Natural Language Processing (NLP) by improving contextual understanding, zero-shot and few-shot learning, text generation, and question answering. However, LLMs have limitations such as accuracy, bias, interpretability, and context. They often produce factually inaccurate information, struggle to understand complex contexts, and may unintentionally produce biased or stereotyped information. KGs, on the other hand, have high-quality explicit knowledge, symbolic reasoning potential, and the ability to evolve with new knowledge, making them essential for various applications. The intersection of LLMs and KGs has sparked significant interest and innovation in NLP. LLM-enhanced KGs can be introduced into pre-training and inference stages to provide external knowledge and assess interpretability. LLM-augmented KGs are designed to improve coverage and ease the use of KGs in various tasks, including embedding learning, completion, construction, KG-to-text generation, and KGQA. Combining the benefits of LLMs and KGs can also improve performance in knowledge representation.

KaLLM 2024 intends to provide a platform for researchers, practitioners, and industry professionals to explore the synergies between LLMs and KGs. We aim to provide a space for the LLM community and the community of KG researchers to interact and explore how these two communities could collaborate and support one another. The goal of the workshop is to seize on the tremendous opportunities arising from investigating cutting-edge approaches, addressing challenges and limitations, and applications in different domains.

We received a total of 18 submissions; 1 non-archival and 17 archival. 1 archival submission was withdrawn as the topic did not fit the workshop. We accepted the non-archival submission and 13 out of the 16 archival submissions. We used reviewers' recommendations and scores to shortlist a set of three papers nominated for the Best Paper Award.

The program will feature oral presentations of the three papers nominated for best papers, and poster presentations of all accepted papers. We are also excited to have invited talks by four speakers: Xin Luna Dong (Meta Reality Labs), Marko Grobelnik (Jozef Stefan Institute), Heng Ji (University of Illinois Urbana-Champaign) and Ivan Titov (University of Edinburgh).

Organizing Committee

General Chairs

Russa Biswas, Aalborg University, Denmark
Lucie-Aimée Kaffee, HuggingFace, Germany
Oshin Agarwal, Bloomberg LP, USA
Pasquale Minervini, University of Edinburgh, UK
Sameer Singh, University of California Irvine, USA
Gerard de Melo, University of Potsdam, Germany

Program Committee

Reviewers

Nikita Bhutani, Megagon Labs Inc
Margarita Bugueño, Hasso Plattner Institute
Yiyi Chen, Aalborg University
Daniel Garijo, Universidad Politécnica de Madrid
Shrestha Ghosh, Max-Planck Institute for Informatics
Paul Groth, University of Amsterdam
Stefan Heindorf, Paderborn University
Fabian Hoppe, Vrije Universiteit Amsterdam
Filip Ilievski, Vrije Universiteit Amsterdam
Jan-Christoph Kalo, University of Amsterdam
Mayank Kejriwal, University of Southern California
Maria Koutraki, Leibniz University of Hannover
Pengwei Li, Meta
Matteo Lissandrini, University of Verona
Daniele Metilli, University College London
Swati Padhee, Wright State University
Heiko Paulheim, University of Mannheim
Simon Razniewski, Technische Universität Dresden
Sneha Singhanian, Max-Planck Institute for Informatics
Mary Ann Tan, Karlsruhe Institute of Technology
Rosni Vasu, University of Zurich

Keynote Talk

The Journey to A Knowledgeable Assistant with Retrieval-Augmented Generation (RAG)

Xin Luna Dong
Meta Reality Labs

Abstract: For decades, multiple communities (Database, Information Retrieval, Natural Language Processing, Data Mining, AI) have pursued the mission of providing the right information at the right time. Efforts span web search, data integration, knowledge graphs, question answering. Recent advancements in Large Language Models (LLMs) have demonstrated remarkable capabilities in comprehending and generating human language, revolutionizing techniques in every front. However, their inherent limitations such as factual inaccuracies and hallucinations make LLMs less suitable for creating knowledgeable and trustworthy assistants.

This talk describes our journey in building a knowledgeable AI assistant by harnessing LLM techniques. We start with our findings from a comprehensive set of experiments to assess LLM reliability in answering factual questions and analyze performance variations across different knowledge types. Next, we describe our federated Retrieval-Augmented Generation (RAG) system that integrates external information from both the web and knowledge graphs for trustworthy text generation on real-time topics like stocks and sports, as well as on torso-to-tail entities like local restaurants. Additionally, we brief our explorations on extending our techniques towards multi-modal, contextualized, and personalized Q&A. We will share our techniques, our findings, and the path forward, highlighting how we are leveraging and advancing the decades of work in this area.

Bio: Xin Luna Dong is a Principal Scientist at Meta Reality Labs, leading the ML efforts in building an intelligent personal assistant. She has spent more than a decade building knowledge graphs, such as the Amazon Product Graph and the Google Knowledge Graph. She has co-authored books Machine Knowledge: Creation and Curation of Comprehensive Knowledge Bases and “Big Data Integration”. She was named an ACM Fellow and an IEEE Fellow for significant contributions to knowledge graph construction and data integration, awarded the VLDB Women in Database Research Award and VLDB Early Career Research Contribution Award. She serves in the PVLDB advisory committee, was a member of the VLDB endowment, a PC co-chair for KDD’2022 ADS track, WSDM’2022, VLDB’2021, and Sigmod’2018.

Keynote Talk

Extracting Common Sense World Models from LLMs

Marko Grobelnik

AiLab, Jozef Stefan Institute

Abstract: LLMs are often criticized for not operating with a notion of world models, which could provide robustness, explainability, and multi-hop reasoning. In this keynote, we will show the methodology and concrete examples of how to extract non-trivial symbolic world models from LLMs for an arbitrary domain. The extracted world models will be represented in an operational first-order logic; concretely in the Prolog programming language, in its basic and probabilistic versions. In the second step, the extracted world models will be used to ground textual data into the semantics of world models to enable reasoning, explanation, and possibly efficient agent communication operating with explicit representations. The approach aims to integrate black-box LLM representations with transparent symbolic representations close to humans without significant loss of information for practical applications.

Bio: Marko Grobelnik is a researcher in the field of Artificial Intelligence (AI). Focused areas of expertise are Machine Learning, Data/Text/Web Mining, Network Analysis, Semantic Technologies, Deep Text Understanding, and Data Visualization. Marko co-leads Artificial Intelligence Lab at Jozef Stefan Institute, cofounded UNESCO International Research Center on AI (IRCAI), and is the CEO of Quintelligence.com specialized in solving complex AI tasks for the commercial world. He collaborates with major European academic institutions and major industries such as Bloomberg, British Telecom, European Commission, Microsoft Research, New York Times, OECD. Marko is co-author of several books, co-founder of several start-ups and is/was involved into over 100 EU funded research projects in various fields of Artificial Intelligence. Significant organisational activities include Marko being general chair of LREC2016 and TheWebConf2021 conferences. Marko represents Slovenia in OECD AI Committee (AIGO/ONEAI), in Council of Europe Committee on AI (CAHAI/CAI), NATO (DARB), and Global Partnership on AI (GPAI). In 2016 Marko became Digital Champion of Slovenia at European Commission.

Keynote Talk

Making Large Language Model’s Knowledge More Accurate, Organized, Up-to-date and Fair

Heng Ji

University of Illinois Urbana-Champaign

Abstract: Large language models (LLMs) have demonstrated remarkable performance on knowledge reasoning tasks, owing to their implicit knowledge derived from extensive pretraining data. However, their inherent knowledge bases often suffer from disorganization and illusion, bias towards common entities, and rapid obsolescence. Consequently, LLMs frequently make up untruthful information, exhibit resistance to updating outdated knowledge, or struggle with generalizing across multiple languages. In this talk I will discuss our recent research efforts at tackling these challenges. I will begin by presenting theoretical and empirical analyses that illuminate when and why LLMs frequently produce factual errors and struggle to determine knowledge updating boundary in order to reach “ripple effect”. Our investigations reveal several underlying causes. First, LLMs acquire implicit knowledge primarily through attention-weighted associations between words, rather than explicit understanding of concepts, entities, attributes, relations, events, semantic roles, and logics. Second, frequent word associations overshadow uncommon ones due to training data imbalance and wide context, particularly in contexts involving dynamic events. Third, counter-intuitive updating behaviors are elucidated through a novel gradient similarity metric. Fourth, LLMs are often unaware of real-world events occurring after their pretraining phase, complicating the anchoring of related knowledge updates. While existing methods focus largely on updating entity attributes, our research underscores the necessity of updating factual knowledge—such as participants, semantic roles, time, and location—based on real-world events. We propose a novel framework for knowledge updating in LLMs that leverages event-driven signals to identify factual errors preemptively and introduce a training-free self-contrastive decoding approach to mitigate inference errors.

Bio: Heng Ji is a professor at Computer Science Department, and an affiliated faculty member at Electrical and Computer Engineering Department and Coordinated Science Laboratory of University of Illinois Urbana-Champaign. She is an Amazon Scholar. She is the Founding Director of Amazon-Illinois Center on AI for Interactive Conversational Experiences (AICE). She received her B.A. and M. A. in Computational Linguistics from Tsinghua University, and her M.S. and Ph.D. in Computer Science from New York University. Her research interests focus on Natural Language Processing, especially on Multimedia Multilingual Information Extraction, Knowledge-enhanced Large Language Models and Vision-Language Models. She was selected as a Young Scientist by the World Laureates Association in 2023 and 2024. She was selected as Young Scientist and a member of the Global Future Council on the Future of Computing by the World Economic Forum in 2016 and 2017. She was named as part of Women Leaders of Conversational AI (Class of 2023) by Project Voice. The other awards she received include two Outstanding Paper Awards at NAACL2024, AI’s 10 to Watch Award by IEEE Intelligent Systems in 2013, NSF CAREER award in 2009, PACLIC2012 Best paper runner-up, Best of ICDM2013 paper award, Best of SDM2013 paper award, ACL2018 Best Demo paper nomination, ACL2020 Best Demo Paper Award, NAACL2021 Best Demo Paper Award, Google Research Award in 2009 and 2014, IBM Watson Faculty Award in 2012 and 2014 and Bosch Research Award in 2014-2018. She was invited to testify to the U.S. House Cybersecurity, Data Analytics, IT Committee as an AI expert in 2023. She was selected to participate in DARPA AI Forward in 2023. She was invited by the Secretary of the U.S. Air Force and AFRL to join Air Force Data Analytics Expert Panel to inform the Air Force Strategy 2030, and invited to speak at the Federal Information Integrity RD Interagency Working Group (IIRD IWG) briefing in 2023. She is the lead of many multi-institution projects and tasks, including the U.S. ARL projects on information

fusion and knowledge networks construction, DARPA ECOLE MIRACLE team, DARPA KAIROS RESIN team and DARPA DEFT Tinker Bell team. She has coordinated the NIST TAC Knowledge Base Population task 2010-2020. She is the Chief Editor of Data Intelligence Journal. She served as the associate editor for IEEE/ACM Transaction on Audio, Speech, and Language Processing, and the Program Committee Co-Chair of many conferences including NAACL-HLT2018 and ACL-IJCNLP2022. She was elected as the North American Chapter of the Association for Computational Linguistics (NAACL) secretary 2020-2023. Her research has been widely supported by the U.S. government agencies (DARPA, NSF, DoE, ARL, IARPA, AFRL, DHS) and industry (Amazon, Google, Bosch, IBM, Disney).

Keynote Talk

Understanding and Navigating Human Control and Transparency in LLMs

Ivan Titov
University of Edinburgh

Abstract: Language models are an exciting technology that has transformed our field and used by millions of people daily. However, both users and researchers often find themselves puzzled by LLM’s responses and struggle to understand the underlying decision processes or attribute their responses to specific data sources. I will talk about our work which tries to enhance the transparency of these models for human users, ensure their behavior is systematic, and uncover the sources of their decisions. This transparency should enable finer control of these models, including model editing, the unlearning of undesirable behaviors or data sources, integration of extra information (e.g., in the form of knowledge bases).

In this talk, I will discuss the approaches my group (as well as colleagues) have been developing, highlighting not only methods but also some cautious lessons learned along the way. This includes pitfalls in data attribution and the challenges of guiding model responses with human rationale. Although progress in these areas may seem slow and sometimes illusory, it is a crucial direction, given the growing reliance on collaboration between humans and large language models. I also hope to convince you that this area holds a diverse range of intriguing open problems for us, researchers, to explore.

Bio: Ivan Titov is a Full Professor at the University of Edinburgh, UK, and also a faculty member at the University of Amsterdam, Netherlands. Ivan’s current interests lie in making deep learning models interpretable, robust, and controllable, or more generally in machine learning for NLP. He has received awards at leading NLP conferences. Ivan has been a program co-chair of ICLR 2021 and CoNLL 2018, and has served on the editorial boards of the Transactions of the ACL, Journal of Artificial Intelligence Research, and Journal of Machine Learning Research, and on the advisory board of the European chapter of ACL. Ivan is an ELLIS fellow and co-directs the ELLIS NLP program and Edinburgh ELLIS unit. Ivan’s research group has been supported by personal fellowships (e.g., ERC, Dutch Vici, and Vidi grants) as well as industrial funding (e.g., Google, SAP, Booking.com and Amazon).

Table of Contents

<i>Multi-hop Database Reasoning with Virtual Knowledge Graph</i> Juhee Son, Yeon Seonwoo, Alice Oh, James Thorne and Seunghyun Yoon	1
<i>Zero- and Few-Shots Knowledge Graph Triplet Extraction with Large Language Models</i> Andrea Papaluca, Daniel Krefl, Sergio José Rodríguez Méndez, Artem Lensky and Hanna Suominen	12
<i>Analysis of LLM’s “Spurious” Correct Answers Using Evidence Information of Multi-hop QA Datasets</i> Ai Ishii, Naoya Inoue, Hisami Suzuki and Satoshi Sekine	24
<i>Application of Generative AI as an Enterprise Wikibase Knowledge Graph Q&A System</i> Renê De Ávila Mendes, Dimas Jackson De Oliveira and Victor Hugo Fiuza Garcia	35
<i>KGAST: From Knowledge Graphs to Annotated Synthetic Texts</i> Nakanyseth Vuth, Gilles Sérasset and Didier Schwab	43
<i>HRGraph: Leveraging LLMs for HR Data Knowledge Graphs with Information Propagation-based Job Recommendation</i> Azmine Touseh Wasi	56
<i>Adapting Multilingual LLMs to Low-Resource Languages with Knowledge Graphs via Adapters</i> Daniil Gurgurov, Mareike Hartmann and Simon Ostermann	63
<i>Ontology-guided Knowledge Graph Construction from Maintenance Short Texts</i> Zeno Van Cauter and Nikolay Yakovets	75
<i>Educational Material to Knowledge Graph Conversion: A Methodology to Enhance Digital Education</i> Miquel Canal-Esteve and Yoan Gutierrez	85
<i>STAGE: Simplified Text-Attributed Graph Embeddings using Pre-trained LLMs</i> Aaron Zolnai-Lucas, Jack Boylan, Chris Hokamp and Parsa Ghaffari	92
<i>Zero-Shot Fact-Checking with Semantic Triples and Knowledge Graphs</i> Moy Yuan and Andreas Vlachos	105
<i>Fine-tuning Language Models for Triple Extraction with Data Augmentation</i> Yujia Zhang, Tyler Sadler, Mohammad Reza Taesiri, Wenjie XU and Marek Reformat	116
<i>Improving LLM-based KGQA for multi-hop Question Answering with implicit reasoning in few-shot examples</i> Mili Shah and Jing Tian	125

Program

Thursday, August 15, 2024

09:00 - 09:15 *Opening Remarks*

09:15 - 10:00 *Invited Talk by Xin Luna Dong*

10:00 - 10:30 *Spotlight Paper Presentations*

Multi-hop Database Reasoning with Virtual Knowledge Graph

Juhee Son, Yeon Seonwoo, Alice Oh, James Thorne and Seunghyun Yoon

Zero- and Few-Shots Knowledge Graph Triplet Extraction with Large Language Models

Andrea Papaluca, Daniel Krefl, Sergio José Rodríguez Méndez, Artem Lensky and Hanna Suominen

KGAST: From Knowledge Graphs to Annotated Synthetic Texts

Nakanyseth Vuth, Gilles Sérasset and Didier Schwab

10:30 - 11:00 *Break*

11:00 - 11:45 *Invited Talk by Marko Grobelnik*

11:45 - 13:00 *Poster Session I*

13:00 - 14:00 *Break*

14:00 - 14:45 *Invited Talk by Heng Ji*

14:45 - 15:30 *Invited Talk by Ivan Titov*

15:30 - 16:00 *Break*

16:00 - 16:45 *Poster Session II*

16:45 - 17:00 *Closing Remarks*

Thursday, August 15, 2024 (continued)