

Preface

Manas Gaur, Efthymia Tsamoura, Edward Raff, Nikhita Vedula, and Srinivasan Parthasarathy

The advent of ChatGPT on November 30, 2022, marked a pivotal moment in AI, garnering both attention and criticism. This milestone catalyzed the development of other models like Google BARD, Google GEMINI, and Meta LLAMA, collectively called Large Language Models (LLMs). Despite their remarkable capabilities, these models have faced significant challenges, notably providing confidently asserted but factually inaccurate information, known as hallucinations. Hallucinations undermine their reliability and trustworthiness, compounded by their propensity to deliver inconsistent answers and irrelevant explanations. Addressing these issues is crucial for enhancing the applicability of LLMs. Researchers from the Natural Language Processing (NLP), the Artificial Intelligence (AI), and the database communities are actively developing diverse approaches to mitigate these pitfalls. Notable among these is the Retrieval Augmentation Generation exemplified by Facebook AI's Semantic Search, which aims to improve text generation accuracy. However, the widespread adoption of such solutions remains a formidable challenge. This workshop aims to accelerate efforts at the intersection of Symbolic Knowledge and Statistical Knowledge in LLMs.

The workshop will address key topics to improve the performance and trustworthiness of LLMs. The primary focus is enhancing LLMs' consistency, reliability, explainability, and safety, developing new methods and metrics for consistent outputs, and adapting safety certification frameworks for sensitive fields like health, cybersecurity, and law. We will explore general-purpose and domain-specific knowledge to enhance LLMs' decision-making, including procedural and declarative types. This involves comparing human and machine information processing, focusing on abstraction, linguistic structures, and generalization. Innovative interpretability and transparency methods will be discussed, especially in healthcare, to ensure comprehensive contextual understanding. Additionally, neuroscience and computer vision methodologies will also be applied to assess LLMs and foundational models. The workshop will explore using LLMs and generative AI for healthcare and well-being, developing open-source tools for analysis and visualization, and creating robust safety metrics to protect against adversarial attacks. Reliable metrics for hallucination and real-world deployment experiences will be shared for valuable insights. Our theme, neurosymbolic AI and knowledge-infused learning, involves using retrieval-augmented LLMs in health, legal, and crisis domains, advancing knowledge-infused reinforcement learning, and developing knowledge-injected foundational models. Automating rule learning and inference in various domains aims to enhance LLMs' decision-making capabilities. Promoting user-explainable machine learning ensures transparency and understandability. Ensuring safety in conversational systems and addressing bias within deep learning fosters fair AI systems. Enhancing user controllability through rules, constraints, and guidelines empowers the safe and effective use of LLMs.

By addressing diverse and interconnected topics, the workshop aims to develop technically advanced, reliable, safe LLMs that align with the needs and values of application domains. This holistic approach ensures that the advancements in LLMs will not only push the boundaries of what is possible but also foster trust and utility across different sectors, ultimately contributing to the broader goal of integrating AI responsibly and effectively into society.

The main program of KiL'24 consists of 14 papers, selected out of 25 submissions, covering topics such as knowledge graphs, knowledge-infused learning, neurosymbolic AI, LLMs, hallucination evaluation, interpretable and explainable AI, and high stakes decision-

making applications (e.g., autonomous driving). We sincerely thank the authors of the submissions as well as the attendees of the workshop. We wish to thank our program committee members for their help in selecting high-quality papers. Furthermore, we are grateful to Alex Jaimes, Amit Sheth, and Huzefa Rangawala for giving keynote presentations on their recent work and invited talks from Andrii Skomorokhov and Negar Foroutan Eghlidi. This year, the workshop is proud to be sponsored by ASIMOV by Haltia.inc.

August 2024

ACM SIGKDD Workshop on Knowledge-Infused Learning (KIL)

Organizers:

Manas Gaur, University of Maryland, Baltimore County, USA
Efthymia Tsamoura, Samsung Research, Cambridge, UK
Edward Raff, Booz Allen Hamilton, USA
Nikhita Vedula, Amazon, USA
Srinivasan Parthasarathy, Ohio State University, USA

Program Committee:

Amit Sheth, AI Institute, University of South Carolina, USA
Tim Finin, University of Maryland Baltimore County, USA
Varish Mulwad, GE Research, India
Prasad Yalamanchi, Lead Semantics, USA
Amélie Gyrard, Trialog, France
Aman Chaddha, Amazon AI, USA
Ponnurangam Kumaraguru, IIIT Hyderabad, India
Vikas Yadav, Samsung Research America, USA
Stella Biderman, Eleuther AI, USA
Iman Azimi, University of California, Irvine, USA
Kalpa Gunaratna, Samsung Research America, USA
Goonmeet Bajaj, Ohio State University, USA
Cory Henson, Bosch Research, USA
Swati Padhee, Ohio State University, USA
Qiwei Han, Nova School of Business and Economics, Lisbon, Portugal
Aritran Piplai, University of Texas at El Paso, USA
Tarmo Koppel, Tallinn Institute of Technology, Estonia
Abhishek Tiwari, Indian Institute of Technology, Patna, India
Prasoon Goyal, Amazon, New York, USA
Vignesh Narayanan, AI Institute, University of South Carolina, USA
Ziyu Yao, George Mason University, USA
Biplav Srivastava, AI Institute, University of South Carolina, USA
Koorosh Aslansefat, University of Hull
Ioannis Sorokos, Fraunhofer ISIE, Germany
Tayyaba Azim, Southampton University, UK
Suparna Bhattacharya, Hewlett Packard Enterprise, India
Roberto Yus, University of Maryland Baltimore County, USA
Amit Sheth, AI Institute, University of South Carolina, USA
Houbing Song, University of Maryland Baltimore County, USA
Lara Martins, University of Maryland Baltimore County, USA
Kaushik Roy, AI Institute, University of South Carolina, USA
Ruwan Wickramarachchi, AI Institute, University of South Carolina, USA
Harsha Kokel, IBM Research, Almaden, USA
Aamir Hamid, University of Maryland Baltimore County, USA
Jiajun Xu, Meta AI, USA
Mukuntha Narayanan, Pinterest Search, USA

Ramana Malladi, Meta AI, USA
Manas Paldhe, Infinitus, USA
Ankur Padia, Liberty Mutual Insurance, USA
Dibyendu Roy Chowdhury, Care Daily, Texas, USA
Vedant Palit, Indian Institute of Technology, Kharagpur, India
Anish Dubey, Meta AI, USA
Subhash Neupane, Mississippi State University, USA
Pengfei Zhang, University of California - Irvine, USA
Aseem Srivastava, IIIT Delhi, India
Akash Vartak, University of Maryland, Baltimore County, USA

Webmaster:

Yash Saxena, University of Maryland, Baltimore County, USA
Deepa Tilwani, AI Institute, University of South Carolina, USA