Integrating Large Language Models and Knowledge Graphs for Next-level AGI

Linhao Luo
Monash University
Melbourne, Australia

Carl Yang
Emory University
Atlanta, USA

Evgeny Kharlamov
Bosch Center for Artificial Intelligence
Stuttgart, Germany

Shirui Pan
Griffith University
Gold Coast, Australia

Abstract

Large language models (LLMs), due to their emergent ability and generalizability, are making new waves in developing Artificial General Intelligence (AGI). However, LLMs are black-box models, which often fall short of capturing and accessing factual knowledge. In contrast, Knowledge Graphs (KGs) are structured knowledge models that explicitly store rich factual knowledge. KGs can enhance LLMs by providing external knowledge for inference and interpretability. Meanwhile, KGs are difficult to construct and evolve by nature, which challenges the existing methods in KGs to generate new facts and represent unseen knowledge. Therefore, it is complementary to integrating LLMs and KGs together and simultaneously leveraging their advantages to achieve AGI's ultimate goal: to reason, adapt, and synthesize knowledge with human-level nuance and factual accuracy.

This tutorial aims to bridge this gap by presenting a comprehensive overview of the unification of LLMs and KGs for next-level AGI. Specifically, we will cover three key frameworks: (1) KG-enhanced LLMs, which focus on augmenting LLMs with KGs for pre-training, fine-tuning, and inference, thereby enriching the LLMs' factual and contextual accuracy; (2) LLM-augmented KGs, which leverage LLMs to assist in tasks such as KG completion, construction, and question answering, ultimately facilitating KG scalability and adaptability; and (3) Synergized LLM-KG Systems and Applications, where LLMs and KGs function symbiotically to enable real-time, bidirectional reasoning, transforming static knowledge structures into dynamic, AGI-driven frameworks. Through this tutorial, participants will gain a structured understanding of the architectures, underlying methodologies, and key advancements in LLM-KG unification, alongside insights into current real-world applications and challenges. We will also explore future research directions, encouraging the development of AGI systems that are not only knowledgeable but also faithful in reasoning. This tutorial will empower researchers and practitioners to unlock the next level of AGI by integrating the strengths of LLMs and KGs into cohesive, intelligent systems.

WWW Companion '25, Sydney, NSW, Australia

ACM ISBN 979-8-4007-1331-6/2025/04
https://doi.org/10.1145/3701716.3715866

CCS Concepts

• **Computing methodologies → Artificial intelligence**.

Keywords

Large Language Models, Knowledge Graph, Artificial General Intelligence

ACM Reference Format:
Linhao Luo, Carl Yang, Evgeny Kharlamov, and Shirui Pan. 2025. Integrating Large Language Models and Knowledge Graphs for Next-level AGI. In *Companion Proceedings of the ACM Web Conference 2025 (WWW Companion '25), April 28-May 2, 2025, Sydney, NSW, Australia.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3701716.3715866

1 Contributors

The tutorial will be presented by the following rising/world-leading LLMs and KGs researchers from academia and industry.

Linhao Luo

- **Affiliation**: Monash University
- **Email**: Linhao.Luo@monash.edu
- **Short Bio**: Linhao Luo is a final-year Ph.D. candidate at Monash University, Australia. Prior to that, he received his B.S. degree from Harbin Institute of Technology, China. His research interests mainly focus on the areas of large language models, knowledge graphs, and graph neural networks. To date, he has published more than 10 research papers in top-tier journals and conferences, including ICLR, NeurIPS, ACL, EMNLP, IJCAI, WSDM, CIKM, and IEEE TKDE. He is the author of various leading research in the field LLMs and KGs, e.g., roadmap of unifying LLMs and KGs [14], reasoning-on-graph (RoG) [8], GCR [10], GFM-RAG [11], ChatRule [7], LLM-DA [18], assessing factual knowledge of LLMs with KGs [9], and evaluating CoT reasoning of LLMs with KGs [13].

Carl Yang

- **Affiliation**: Emory University
- **Email**: j.carlyang@emory.edu
- **Short Bio**: Carl Yang is an Assistant Professor at Emory University, Atlanta, GA. He earned a B.Eng. in computer science from Zhejiang University in 2014 and a Ph.D. in computer science from the University of Illinois in 2020. His research focuses on graph data mining, machine learning, knowledge graphs, and federated learning, with applications in recommender systems, social networks, neuroscience, and healthcare. He has published over 70 papers in top venues

such as IEEE Transactions on Knowledge and Data Engineering, KDD, and WWW. He is the author of various leading research in the field LLMs and KGs, e.g., knowledge graphs for healthcare [1], ClinGen [24], and GuardAgent [23].

Evgeny Kharlamov

- **Affiliation**: Bosch Center for Artificial Intelligence
- **Email**: evgeny.kharlamov@de.bosch.com
- **Short Bio**: Evgeny Kharlamov is a Senior Expert at the Bosch Centre for AI and an Associate Professor at the University of Oslo. His research interests encompass various AI topics, bridging Knowledge Representation and reasoning with learning, recently in LLM-centered and agentic scenarios. Evgeny's work focuses on both theoretical aspects and practical applications across diverse industrial sectors that require explainable, robust, and scalable AI solutions in manufacturing, energy, oil, and gas. He has published 190+ papers in prestigious international venues and journals such as WWW, IJCAI, AAAI, and SIGIR. He is the author of various papers in the field of KGs and LLMs, e.g., on Robust Knowledge Extraction from LLMs [15], OAG [26], P2TAG [27].

Shirui Pan

- **Affiliation**: Griffith University
- **Email**: s.pan@griffith.edu.au
- **Short Bio**: Shirui Pan is a Professor and an ARC Future Fellow with the School of Information and Communication Technology, Griffith University, Australia. He received his Ph.D degree in computer science from University of Technology Sydney (UTS), Australia. His research interests include artificial intelligence and machine learning. He has made contributions to advance graph machine learning methods for solving hard AI problems for real-life applications, including anomaly detection, recommender systems, and time series forecasting. He has published 150+ papers at top conferences and journals including NeurIPS, ICML, KDD, and WWW. He is the author of various leading research in the field LLMs and KGs, e.g., roadmap of unifying LLMs and KGs [14], reasoning-on-graph (RoG) [8], GCR [10], ChatRule [7], and LLM-DA [18].

2 Topic and Relevance

Tutorial Topic. This tutorial explores the unification of large language models (LLMs) and knowledge graphs (KGs) as a pathway toward advancing Artificial General Intelligence (AGI). The synergy of LLMs and KGs provides essential insights into the next level of AGI, benefiting various web applications, such as search engines, recommendation systems, and chatbots. Due to its significance, this topic has garnered considerable attention recently. The research community has recognized the potential synergy between LLMs and KGs, resulting in a surge of efforts in this area. Although their integration is still in its early stages, it shows promise for applications such as question answering, knowledge reasoning, and recommendation systems. This tutorial will provide a structured overview of recent attempts, architectures, and applications aimed at unifying LLMs and KGs to develop AGI systems that are both intelligent and grounded in factual data.

Scope and Depth. This tutorial will deliver an in-depth examination of frameworks and methodologies for LLM-KG integration, organized into three high-level frameworks: KG-enhanced LLMs, LLM-augmented KGs, and Synergized LLM-KG Systems. We will begin by covering the core challenges and limitations of both LLMs and KGs when used independently. For KG-enhanced LLMs, we will discuss methods for incorporating KG data into the training and inference stages of LLMs, enhancing factual accuracy and reducing hallucination. For LLM-augmented KGs, we will introduce techniques in which LLMs support KG tasks, including construction, and completion, improving KG flexibility and adaptability. Finally, we will examine bidirectional reasoning in Synergized LLM-KG Systems and their applications, where both technologies interact to create a more powerful, mutually reinforcing model. The tutorial will conclude with applications such as AGI-driven decision-making, question-answering systems, and recommendations, demonstrating how unified LLM-KG systems can transform various real-world web applications.

Importance and Timeliness. The synergy of LLMs and KGs is a timely topic due to the rapid advancements and increased reliance on LLMs in diverse fields. Despite the remarkable performance of LLMs in various tasks, they often lack factual reliability and are prone to hallucination, limiting their effectiveness in high-stakes applications. Meanwhile, KGs, although explicit in their structure, struggle to adapt quickly to new information and require extensive resources to maintain. As these limitations become more apparent, the demand for unifying these technologies to unlock their full potential is critical. This tutorial addresses this urgent need, providing the audience with a systematic and forward-looking perspective on this frontier area, and helping researchers, practitioners, and industry leaders understand and adopt these innovative approaches to advance AGI. Although there are a few tutorials on LLMs and KGs separately, there is a lack of comprehensive tutorials that bridge the gap between these two technologies. This tutorial aims to fill this gap by providing a structured overview of the integration of LLMs and KGs, offering insights into the latest advancements, applications, and challenges in this emerging field.

Relevance to WebConf. The topic of unifying LLMs and KGs aligns closely with several core tracks of WebConf, including "Graph Algorithms and Modeling", "Semantics and Knowledge" and "Search and retrieval-augmented AI". The integration of LLMs and KGs also has profound implications for "Responsible web" as a unified framework provides a more interpretable and reliable foundation for AGI, promoting transparency and trustfulness. This tutorial's focus on AGI-centered applications, such as personalized recommendations, question-answering systems, and knowledge-driven reasoning support, highlights the relevance and transformative potential of LLM-KG unification for web-based applications and services.

Qualification. The speakers are leading experts in the fields of large language models, knowledge graphs, and their integration. They have authored influential research in these areas, contributing to both the theoretical and practical aspects of LLM-KG systems. Their extensive experience in pioneering methods for enhanced factual reasoning and AGI applications provides them with a unique perspective to guide participants through this tutorial.

3 Style

It's a lecture-style tutorial. We will bring our own laptops for the tutorial. Attendees do not need to bring any equipment.

4 Schedule

The detailed tutorial schedule is as follows.

(1) **Introduction to LLMs and KGs**: We will provide an overview of LLMs and KGs, highlighting their strengths and limitations when used independently. We will discuss the importance of unifying LLMs and KGs to achieve next-level AGI and introduce the three key frameworks for integration: KG-enhanced LLMs, LLM-augmented KGs, and Synergized LLM-KG Systems.

(2) **KG-enhanced LLMs**

- Methodology I: KG-augmented LLM pre-training.
  In this part, we summarize how to incorporate KGs into the pre-training stage of LLMs to enhance their factual knowledge. We will highlight the three typical methods, including: integrating KGs into training objective [20], integrating KGs into LLM inputs [6], and KGs instruction-tuning [25]. We will also discuss the limitations and challenges of these methods and provide insights into future research directions.
- Methodology II: KG-augmented LLM inference.
  We present methods that inject KGs into the inference stage of LLMs to improve their facts awareness and reasoning capabilities. These methods need to retrieve the relevant KG information [3, 8, 12] and incorporate them into the LLM inference process [3, 10]. We will discuss how to design a graph retriever for effectively retrieving information from KGs, and then how to enable LLMs to conduct reasoning based on the structured knowledge.
- Methodology III: KG-augmented LLM analysis.
  Knowledge graphs store massive facts, which can be used to analyze the factual knowledge of LLMs to improve their transparency [9]. The structural nature of KGs can also be used to certify the reasoning process of LLMs [13], which is essential for the trustworthiness of LLMs. In this part, we will introduce the methods to assess the factual knowledge of LLMs with KGs and evaluate the reasoning process of LLMs with KGs.

(3) **LLM-augmented KGs**

- Methodology I: LLM-augmented KG construction.
  In this part, we present methods that leverage LLMs for constructing KGs. The process of knowledge graph construction typically involves multiple stages, including: (1) entity discovery [19], (2) coreference resolution [5], and (3) relation extraction [17]. The LLMs can be involved in each stage of KG construction. Recent approaches have explored (4) end-to-end knowledge graph construction, which involves constructing a complete knowledge graph in one step [28].
- Methodology II: LLM-augmented KG refinement.
  Knowledge graphs are evolving by nature, challenging the existing methods in KGs to generate new facts and represent unseen knowledge. In this part, we will introduce the methods to leverage LLMs to assist in tasks such as KG completion, ultimately facilitating KG scalability and adaptability. We will discuss how to utilize the general knowledge of LLMs to predict the missing facts in KGs [7], and how to leverage LLMs to model the temporal dynamics of KGs [18].

(4) **Applications of synergized KG-LLM systems**

- Application I: Synergy LLM-KG for healthcare.
  In this part, we will introduce the applications of LLM-KG systems in healthcare. We will discuss how to take advantage of the knowledge from the medical knowledge graph to enhance the generation and decision-making of LLMs. It would include the applications of LLM-KG systems in disease diagnosis, drug discovery, and personalized treatment [1, 22].
- Application II: Synergy LLM-KG for recommendation.
  In this part, we will introduce the applications of LLM-KG systems in recommendation systems. We will discuss how to combine KGs to capture the user-item interactions and leverage the powerful reasoning capabilities of LLMs to analyze the user preferences to improve the recommendation performance [21]. Besides, with the help of KGs and LLMs, we can also provide more interpretable and explainable recommendations to users [4].

(5) **Conclusions and future opportunities**

5 Intended Audience and Prerequisite Knowledge

This tutorial is designed for audiences at all levels, covering from basic concepts of large language modes and knowledge graphs, problems to advances, and industrial applications in unifying LLMs and KGs. While no specific knowledge is required from the audience, people who are familiar with large language models, knowledge graphs, graph neural networks, or work in industries related to LLMs and KGs will find it more beneficial in understanding the algorithms and case studies to be introduced in this tutorial.

6 Previous Editions and Relevant Tutorials

This tutorial has not been presented before but there is a list of relevant tutorials as follows.

- WSDM24's tutorial on "Bridging Text Data and Graph Data: Towards Semantics and Structure-aware Knowledge Discovery". This tutorial was organized between March 4th-8th, 2024. It provides a comprehensive overview of the latest advancements in graph mining techniques that utilize LLMs and enhance text mining methods by incorporating graph structure information. It mostly focuses on mining meaningful information from text data and graph data. However, in the proposed tutorial, we will focus on the board topic of unifying LLMs and KGs for next-level AGI, including reasoning, trustfulness, and transparency.
- KDD24's tutorial on "Automated Mining of Structured Knowledge from Text in the Era of Large Language Models" was organized on Aug 15th, 2024. It reviews the recent progress in mining structural knowledge from unstructured text. It mainly focuses on how to utilize LLMs to construct knowledge graphs, which is highly related to the LLM-augmented KG construction in the proposed tutorial. However, the proposed tutorial will cover more topics, including KG-enhanced LLMs, and applications of synergized LLM-KG systems.

Furthermore, the presenters have rich experience in tutorials in emerging research areas, which are highly related to this tutorial.

- ADC23's tutorial on "Towards Data-centric Graph Machine Learning" was organized by our contact tutor on Nov 1st, 2023. It introduces the recent works in data-enteric AI. It conducts an in-depth review of the recent advancements on the collection, management, and utilization of data to drive AI models and applications in graph machine learning. The tutorial is highly related to the proposed tutorial, as it also focuses on the utilizing the graph data to enhance the AI models.
- IJCNN23's tutorial on "Graph Self-Supervised Learning: Taxonomy, Frontiers, and Applications" is organized by our contact tutor on Jun 18th, 2023. This tutorial provides a comprehensive overview of the recent advancements in graph self-supervised learning. It covers the taxonomy, frontiers, and applications of graph self-supervised learning. The tutorial is highly related to the proposed tutorial, as it also focuses on using graph data for AI model training.

7 Tutorial Materials

This tutorial is primarily based on the latest roadmap of unifying LLMs and KGs [14], which has attracted over 500 citations in a couple of months since it was published in 2024. The tutorial also covers a series of follow-up research in this trending direction. This includes RoG [8], GCR [10], ChatRule [7], LLM-DA [18], assessing factual knowledge of LLMs with KGs [9], and evaluating CoT reasoning of LLMs with KGs [13]. Noticeably, the tutorial will not only cover research from the presenters' groups, but it will also present research from diverse research groups in the whole community. Some of these advancements include ToG [16], GNN-RAG [12], and KG-FIT [2]. The tutorial presentation will include toy examples, illustrative figures, and pseudocode to clarify the key concepts of the methods discussed. Each section will feature a Q&A session, and attendees are encouraged to ask questions at any time during the tutorial. There are no copyright issues.

8 Acknowledgments

Shirui Pan was supported in part by the Australian Research Council (ARC) under grants FT210100097 and DP240101547. Carl Yang was partially supported by the US National Science Foundation under Award Number 2319449 and Award Number 2312502, as well as the US National Institute of Diabetes and Digestive and Kidney Diseases of the US National Institutes of Health under Award Number K25DK135913. Evgeny Kharlamov was supported by EU Projects Graph Massivizer (GA 101093202), enRichMyData (GA 101070284), SMARTY (GA 101140087).

References

[1] Hejie Cui, Jiaying Lu, Shiyu Wang, Ran Xu, Wenjing Ma, Shaojun Yu, Yue Yu, Xuan Kan, Chen Ling, Tianfan Fu, et al. 2023. A Review on Knowledge Graphs for Healthcare: Resources, Applications, and Promises. *arXiv:2306.04802* (2023).
[2] Pengcheng Jiang, Lang Cao, Cao Xiao, Parminder Bhatia, Jimeng Sun, and Jiawei Han. 2024. KG-FIT: Knowledge Graph Fine-Tuning Upon Open-World Knowledge. *NeurIPS* (2024).
[3] Shiyang Li, Yifan Gao, Haoming Jiang, Qingyu Yin, Zheng Li, Xifeng Yan, Chao Zhang, and Bing Yin. 2023. Graph Reasoning for Question Answering with Triplet Retrieval. In *ACL*. 3366–3375.
[4] Yuhan Li, Xinni Zhang, Linhao Luo, Heng Chang, Yuxiang Ren, Irwin King, and Jia Li. 2025. G-Refer: Graph Retrieval-Augmented Large Language Model for Explainable Recommendation. In *WWW*.
[5] Qi Liu, Yongyi He, Tong Xu, Defu Lian, Che Liu, Zhi Zheng, and Enhong Chen. 2024. UniMEL: A Unified Framework for Multimodal Entity Linking with Large Language Models. In *CIKM*. 1909–1919.
[6] Weijie Liu, Peng Zhou, Zhe Zhao, Zhiruo Wang, Qi Ju, Haotang Deng, and Ping Wang. 2020. K-BERT: Enabling Language Representation with Knowledge Graph. In *AAAI*. 2901–2908.
[7] Linhao Luo, Jiaxin Ju, Bo Xiong, Yuan-Fang Li, Gholamreza Haffari, and Shirui Pan. 2023. Chatrule: Mining logical rules with large language models for knowledge graph reasoning. *arXiv preprint arXiv:2309.01538* (2023).
[8] LINHAO LUO, Yuan-Fang Li, Reza Haf, and Shirui Pan. [n.d.]. Reasoning on Graphs: Faithful and Interpretable Large Language Model Reasoning. In *ICLR*.
[9] Linhao Luo, Trang Vu, Dinh Phung, and Reza Haf. 2023. Systematic Assessment of Factual Knowledge in Large Language Models. In *EMNLP*. 13272–13286.
[10] Linhao Luo, Zicheng Zhao, Chen Gong, Gholamreza Haffari, and Shirui Pan. 2024. Graph-constrained Reasoning: Faithful Reasoning on Knowledge Graphs with Large Language Models. *arXiv preprint arXiv:2410.13080* (2024).
[11] Linhao Luo, Zicheng Zhao, Gholamreza Haffari, Dinh Phung, Chen Gong, and Shirui Pan. 2025. GFM-RAG: Graph Foundation Model for Retrieval Augmented Generation. *arXiv preprint arXiv:2502.01113* (2025).
[12] Costas Mavromatis and George Karypis. 2024. GNN-RAG: Graph Neural Retrieval for Large Language Model Reasoning. *arXiv preprint arXiv:2405.20139* (2024).
[13] Thi Nguyen, Linhao Luo, Fatemeh Shiri, Dinh Phung, Yuan-Fang Li, Thuy-Trang Vu, and Gholamreza Haffari. 2024. Direct Evaluation of Chain-of-Thought in Multi-hop Reasoning with Knowledge Graphs. In *ACL*. 2862–2883.
[14] Shirui Pan, Linhao Luo, Yufei Wang, Chen Chen, Jiapu Wang, and Xindong Wu. 2024. Unifying large language models and knowledge graphs: A roadmap. *TKDE* (2024).
[15] Nico Potyka, Yuqicheng Zhu, Yunjie He, Evgeny Kharlamov, and Steffen Staab. 2024. Robust Knowledge Extraction from Large Language Models using Social Choice Theory. In *AAMAS*. 1593–1601.
[16] Jiashuo Sun, Chengjin Xu, Lumingyuan Tang, Saizhuo Wang, Chen Lin, Yeyun Gong, Lionel Ni, Heung-Yeung Shum, and Jian Guo. 2024. Think-on-Graph: Deep and Responsible Reasoning of Large Language Model on Knowledge Graph. In *ICLR*.
[17] Somin Wadhwa, Silvio Amir, and Byron C Wallace. 2023. Revisiting relation extraction in the era of large language models. In *ACL*, Vol. 2023. 15566.
[18] Jiapu Wang, Kai Sun, Linhao Luo, Wei Wei, Yongli Hu, Alan Wee-Chung Liew, Shirui Pan, and Baocai Yin. 2024. Large Language Models-guided Dynamic Adaptation for Temporal Knowledge Graph Reasoning. *NeurIPS* (2024).
[19] Shuhe Wang, Xiaofei Sun, Xiaoya Li, Rongbin Ouyang, Fei Wu, Tianwei Zhang, Jiwei Li, and Guoyin Wang. 2023. Gpt-ner: Named entity recognition via large language models. *arXiv preprint arXiv:2304.10428* (2023).
[20] Xiaozhi Wang, Tianyu Gao, Zhaocheng Zhu, Zhengyan Zhang, Zhiyuan Liu, Juanzi Li, and Jian Tang. 2021. KEPLER: A Unified Model for Knowledge Embedding and Pre-trained Language Representation. *TACL* 9 (2021), 176–194.
[21] Wei Wei, Xubin Ren, Jiabin Tang, Qinyong Wang, Lixin Su, Suqi Cheng, Junfeng Wang, Dawei Yin, and Chao Huang. 2024. Llmrec: Large language models with graph augmentation for recommendation. In *WSDM*. 806–815.
[22] Junde Wu, Jiayuan Zhu, and Yunli Qi. 2024. Medical graph rag: Towards safe medical large language model via graph retrieval-augmented generation. *arXiv preprint arXiv:2408.04187* (2024).
[23] Zhen Xiang, Linzhi Zheng, Yanjie Li, Junyuan Hong, Qinbin Li, Han Xie, Jiawei Zhang, Zidi Xiong, Chulin Xie, Carl Yang, et al. 2024. GuardAgent: Safeguard LLM Agents by a Guard Agent via Knowledge-Enabled Reasoning. *arXiv preprint arXiv:2406.09187* (2024).
[24] Ran Xu, Hejie Cui, Yue Yu, Xuan Kan, Wenqi Shi, Yuchen Zhuang, Wei Jin, Joyce Ho, and Carl Yang. 2023. Knowledge-Infused Prompting Improves Clinical Text Generation with Large Language Models. In *NeurIPS 2023 Workshop on Synthetic Data Generation with Generative AI*.
[25] Hongbin Ye, Ningyu Zhang, Shumin Deng, Xiang Chen, Hui Chen, Feiyu Xiong, Xi Chen, and Huajun Chen. 2022. Ontology-enhanced Prompt-tuning for Few-shot Learning. In *WWW*. 778–787.
[26] Fanjin Zhang, Xiao Liu, Jie Tang, Yuxiao Dong, Peiran Yao, Jie Zhang, Xiaotao Gu, Yan Wang, Evgeny Kharlamov, Bin Shao, Rui Li, and Kuansan Wang. 2023. OAG: Linking Entities Across Large-Scale Heterogeneous Knowledge Graphs. *IEEE Trans. Knowl. Data Eng.* 35, 9 (2023), 9225–9239.
[27] Huanjing Zhao, Beining Yang, Yukuo Cen, Junyu Ren, Chenhui Zhang, Yuxiao Dong, Evgeny Kharlamov, Shu Zhao, and Jie Tang. 2024. Pre-Training and Prompting for Few-Shot Node Classification on Text-Attributed Graphs. In *KDD*. 4467–4478.
[28] Yuqi Zhu, Xiaohan Wang, Jing Chen, Shuofei Qiao, Yixin Ou, Yunzhi Yao, Shumin Deng, Huajun Chen, and Ningyu Zhang. 2024. Llms for knowledge graph construction and reasoning: Recent capabilities and future opportunities. *WWW* 27, 5 (2024), 58.