The Role Evolution of KGs in Synthesizing with LLMs: From Background Knowledge to Joint Reasoning

Chuangtao Ma

Department of Computer Science, 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.

Keywords

Knowledge Graphs, Large Language Models, KG-RAG, Knowledge Augmentation

1. Introduction

Recent years have witnessed significant achievements and wide applications of large language models (LLMs) in natural language understanding and generation, such as question answering (QA), context generation, text summarization, etc. Knowledge Graphs (KGs), as graph-based structured knowledge, maintain the rich relationships among the factual entity, causality event, and other entities, which provides factual and trackable knowledge. There are several differences between KGs and LLMs in terms of data structure, knowledge type, processing style, and use case, as summarized in Table 1, which poses their limitations and strengths. For instance, LLMs encounter the challenges of hallucinations

Table 1
Comparison of Data Structure, Knowledge Type, Processing Style, and Use Case between LLMs and KGs

Feature	LLMs	KGs
Data Structure	Unstructured text-based, sequential tokens	Structured, graph-based (triples)
Knowledge Type	Implicit, parametric, commonsense knowledge	Explicit, factual, domain-specific knowledge
Processing Style	Intuitive, implicit, next token prediction	Logical reasoning, graph query, path traversal
Primary Use Case	QA, content generation, text summarization	KGQA, recommendation, entity disambiguation

and poor explainability due to a lack of up-to-date domain knowledge and struggle with complex and multi-hop reasoning. In contrast, KGs offer reliable and factual knowledge from commonsense and domain-specific domains, which enables hallucination mitigation and explainable responses based on symbolic reasoning and graph traversal.

To take advantage of their strengths and mutually benefit from LLMs and KGs, a roadmap for unifying LLMs and KGs was designed [1] by bidirectionally enhancing and augmenting with each other. Motivated by this roadmap, an increasing number of works in synthesizing LLMs and KGs have been investigated to address the inherent limitations of LLMs [2]. The role of KGs in synthesizing with LLMs

The 2nd International Workshop on Knowledge Graphs for Responsible AI (KG-STAR 2025), June 02, 2025, Portorož, Slovenia, In conjunction with ESWC 2025
chuma@cs.aau.dk (C. Ma)
0000-0002-6492-2579 (C. Ma)

has evolved from background knowledge for unidirectional enhancement of LLMs to joint reasoning with LLMs for mutual collaboration. This extended abstract is based on the talk at the KG-STAR@ESWC2025 workshop, which aims to give a brief overview of KGs' role evolution when synthesizing with LLMs and summarize the strengths, limitations, and KG requirements of the approaches with different KG roles in augmenting LLMs with KGs, and their applications.

2. Role Evolution of KGs in Synthesizing with LLMs

The role of KGs in synthesizing with LLMs is evolving from passive background knowledge to active involvement in reasoning with LLMs to augment their capabilities.

2.1. Background Knowledge

The early synthesis paradigm of using KGs as background knowledge for augmenting LLMs, where the factual knowledge retrieved from KGs was directly incorporated into LLMs via pre-training, fine-tuning, and KG-based RAG for knowledge-intensive tasks.

2.1.1. Pre-traing and Fine-tuning

The initial phase of synthesizing LLMs with KGs, where KGs serve as the background knowledge, has been exploited in the paradigm of pretraining and fine-tuning LLMs with KGs. In the paradigm of pretraining, text-KG pairs were retrieved and created based on entity linking, and then a cross-modal encoder with the modality interaction token was introduced to bidirectionally fuse the text-KG pair for joint learning and reasoning [3]. Unlike the pretraining paradigm, where retraining is required when updating knowledge, KG fine-tuning aims to fine-tune LLMs with domain-specific knowledge for the specific knowledge-intensive task in a cost-effective manner [4]. To integrate KGs with LLMs at the parameter level, a parameter-efficient fine-tuning (PEFT) method was proposed [5] by introducing a KG adapter layer to bidirectional fusion and updating the token representations for joint reasoning.

2.1.2. KG-based RAG

With the help of prompt engineering, the well-crafted prompts can guide the LLMs to effectively utilize the external knowledge from KGs for hallucination mitigation [6]. Thereby, KG-based retrieval augmentation generation (KG-RAG) [7, 8, 9] is designed to initially retrieve the relevant knowledge from KGs and then feed the retrieved knowledge to LLMs in the form of a prompt. For instance, KG^{2}RAG [10] expands the textual chunks with the retrieved KG by leveraging BFS search over KGs and then incorporates the expanded chunks with the prompt for augmenting the generation. Similarly, a KG-RAG-based fake news detection method [8] was proposed to retrieve the evidence from constructed KGs for augmenting LLMs in veracity prediction. To mitigate the hallucination of LLMs, KG-Infused RAG [9] augments RAG with KGs by integrating retrieved relevant triples based on query-based entity retrieval and iterative triples expansion and providing reasoning paths for the generated response.

2.2. From Background Knowledge to Joint Reasoning

Although the hallucination of LLMs can be mitigated by synthesizing the relevant factual knowledge from KGs with LLMs via the early synthesis paradigm of LLMs and KGs, i.e., joint learning, fine-tuning, and KG-based RAG, it faces several challenges and limitations. The limitations of the early synthesis paradigm, where KGs serve as background knowledge, lie in: *(1) Re-training or fine-tuning is required when KGs are updated.* In the approaches of LLMs and KGs joint training and fine-tuning, KGs were passively integrated with LLMs, as an external knowledge context where the factual knowledge from KGs was injected into LLMs via modality interaction and fusion layer. *(2) Fail to fully leverage structural knowledge of KGs.* In the approaches of KG-RAG, the relevant subgraphs from KGs were retrieved and incorporated into LLMs as parts of the prompt, but they treated retrieved KGs as flat textual knowledge rather than as a graph structure. *(3) KGs passive involvement with LLMs for unidirectional integration.* The knowledge integration in KG-RAG-based approaches is unidirectional, where KGs are passively involved with LLMs. It fails to support the sophisticated reasoning that combines both implicit and explicit knowledge from LLMs and KGs. To address the above challenges and limitations, the paradigms of joint reasoning with LLMs and KGs have been primarily exploited in rule-based reasoning, agent-based reasoning, and collaborative reasoning etc, where KGs are actively involved in the reasoning with LLMs via explicitly incorporating the graph structure from KGs into their reasoning process.

2.2.1. Rule-based Reasoning

The rule-based reasoning aims to guide LLMs to reasoning over KGs by mining the logical rules [11] from KGs or designing instruction-based prompting [12], where the reasoning of the KGs and LLMs is synthesized. For instance, ChatRule [11] ranks the high-quality logical rules that were mined from KGs via LLMs and then incorporates the facts induced by logical rules from KGs to support the reasoning of LLMs. CoT (Chain-of-Thoughts) prompting was designed to fully leverage the reasoning capabilities of LLMs by employing instruction-based prompting to guide LLMs to think step-by-step and decompose the complex task into multiple intermediate steps. Inspired by this, a KG-CoT [12] was proposed by generating responsible knowledge chains over the reasoning paths from KGs based on CoT-based prompting with LLMs to enable knowledge-aware reasoning.

2.2.2. Agent-based Reasoning

By introducing an AI agent to the synthesis of LLMs and KGs, KGs can actively interact with LLMs for joint reasoning. For example, ToG [13] treats the LLM as an agent that interacts with KGs by iteratively executing beam search on KGs to explore the related entities and relations, and performing reasoning based on the retrieved knowledge. KG-Agent [14] integrates the reasoning capabilities of small LLMs with an agent-based KG toolbox to autonomously execute the tool selection and knowledge updating for solving complex reasoning tasks. ODA [15] observes and retrieves the relevant knowledge from the KG environment based on the AI agent and then incorporates the retrieved observations into LLM reasoning for synthesizing the reasoning of KG and LLMs.

2.2.3. Collaborative Reasoning

The collaborative reasoning between LLMs and KGs has been investigated by introducing an adaptive knowledge retrieval [16, 17] and iterative path exploration [18, 19]. To bridge structured knowledge in KGs with unstructured knowledge in LLMs, graph-constrained reasoning (GCR) [16] was proposed to incorporate the reasoning of LLMs and the reasoning over KG-Trie, a trie-based index encoding the reasoning paths retrieved from KGs. CRF [17] introduces a collaborative reasoning method by introducing reinforcement learning to a hierarchical agent that retrieves the path from KGs for supporting reward-based reasoning. Rather than viewing KGs as a static repository of facts that can be referenced by LLMs, PoG [19] treats KGs as a dynamic knowledge resource that can be explored and updated during reasoning. It initially decomposes questions into several sub-objectives and then repeats the process of adaptively exploring reasoning paths, updating memory, and reflecting on the need to self-correct erroneous reasoning paths until arriving at the correct answer. TOG-2 [18] leverages KGs as a navigational tool to guide knowledge retrieval and then iteratively utilizes LLMs to evaluate the retrieved clues from KGs to ensure logical coherence and completeness of factual evidence.

3. Comparison and Application

The various approaches and applications of augmenting LLMs with KGs are compared and summarized.

3.1. Comparison of Different Approaches

As previously discussed, the role evolution of KGs in synthesizing with LLMs demonstrates that the structural information contained in KGs is not merely factual background knowledge, but a rich and structured knowledge that can actively guide and get involved in complex reasoning for knowledge-intensive task scenarios. Table 2 gives a summary and comparison of approaches with role evolution of KGs from background knowledge to joint reasoning.

Table 2
Comparison of Approaches with Roles Evolution of KGs: From Background Knowledge to Joint Reasoning

Approaches	Key Techniques	Strengths	Limitations	KGs Requriements
Background Knowledge	Pre-training, fine-tuning, KG-based RAG	Hallucinations mitigation	Re-training is needed when updating KG	Up-to-date factual knowledge
Rule-based Reasoning	Rule-based reasoning, CoT-based path generation	Explainable results with less hallucinations	Costly rule mining and prompt token overhead	Rich logical and semantic knowledge
Agent-based Reasoning	Agent-based knowledge exploration and reasoning	Knowledge interaction, path exploration	Path explosion and reasoning latency	Flexible interaction interface
Collaborative Reasoning	Multi-hop path exploration, reinforcement learning	Iterative reasoning, knowledge updating	High reasoning complexity and computing cost	Dynamic knowledge adaptation

By demonstrating the benefits of incorporating factual knowledge into LLMs, the early approaches established the foundation for collaborative reasoning that would iteratively explore the reasoning paths from KGs to support multi-hop reasoning and reward-based reasoning for complex knowledge-intensive tasks. Although the collaborative reasoning of LLM and KG shows advantages in complex and multi-hop reasoning in comparison with the early approaches, it still meets several limitations and challenges, including the incompleteness of KG, path explosion, low reasoning efficiency, high reasoning complexity and computing overhead, and unreliable reasoning results.

3.2. Application

The application of augmenting LLMs with KGs are widely investigated in question answering (QA) [2], while the other applications, such as personalized recommendation, healthcare copilot, system diagnostics and detection, and data management, have been recently investigated.

3.2.1. Recommendation

The intention of the recommendation differs from QA, as personalized and relevant recommendations that align with users' preferences are expected for the recommendation system, while precise and reliable responses that answer the user's question are expected for the QA system. To facilitate the personalized recommendation, LLMRG [20] introduces the adaptive reasoning module with LLMs to construct the personalized reasoning graphs from the semantic relationships of user behaviors and profiles and further augments the final recommendations via the user's next item prediction based on reasoning over LLMs and the reasoning graph. To address the issue of lack of up-to-date knowledge in LLMs, K-RagRec [21] designs a popularity selective and similarity-based ranking pipeline to retrieve the relevant subgraphs from an item KGs, and then leverages GNN and projector to align the retrieved subgraphs into the semantic space of LLM for knowledge-augmented recommendation.

3.2.2. Fake News Detection

The current LLMs-based fake news detection suffers a lot from the semantic ambiguity in understanding news, interpretability, and explainability, etc. [22]. Therefore, KGs have been introduced to augment the LLMs-based fake news detection [23, 8]. DKFND [23] evaluates the relevance and authenticity of the given news with the help of the retrieved relevant knowledge and related news from KGs and guides LLMs in detecting and providing the explainable results of news veracity. To mitigate the poor explainability of existing evidence-based fake news detections, a KG-RAG-based fake news detection method [8] was proposed to retrieve the evidence from KGs for augmenting LLMs in veracity prediction.

3.2.3. Healthcare Copilot

To enhance the reliability and explainability of results that were generated by the healthcare copilot, the applications have been investigated in augmenting clinical decision-making with the joint reasoning of LLMs and medical KGs. In order to mitigate the hallucination of LLMs in knowledge-intensive tasks of medical chatbots, CMedRAGBot [24] was designed to enhance LLMs in generating accurate answers for the given medical questions with retrieved relevant knowledge from the medical KGs. MedRAG [25] generates the final answers and follow-up questions for precise diagnosis based on the reasoning elicited by KG that was retrieved from healthcare KGs via multi-level matching and upward traversal.

3.2.4. System Diagnostics and Detection

The applications of augmenting LLMs with KGs in industrial system diagnostics have been studied in nuclear power plants [26], cybersecurity compliance analysis [27], cybersecurity detection [28], etc. KG-DML [26] was proposed to enhance system diagnostics in high-reliability environments by integrating KGs to LLMs based on model interaction and cause-effect reasoning. To improve the accuracy of IoT security compliance analysis, KGs are introduced to the RAG pipeline based on vector-based document retrieval and graph query-based KG retrieval [27]. Similarly, a CyKG-RAG [28] was designed to enhance the reliability of cybersecurity detection by retrieving the relevant knowledge from cybersecurity KG.

3.2.5. Data Management

The several challenges of data management have been exploited [29], such as, schema matching [30], column type annotation (CTA) [31], vector search [32], etc. Given that the existing similarity-based and LLMs-based schema matching methods are incapable of resolving semantic ambiguities and conflicts in complex schema matching, KG-RAG4SM [30] investigated a KG-based RAG for schema matching and heterogeneous data integration. To mitigate the challenges of existing LLM-based methods for CTA in semantic label assignment, RACOON [31] augments LLMs with the retrieved subgraphs from external KGs for CTA. TigerVector [32] integrates vector search and graph query based on the massively parallel processing (MPP) index to accelerate the vector search for vector-based graph retrieval.

4. Open Challenges and Future Directions

The open challenges in synthesizing LLMs with KGs for explainable and reliable reasoning remain in reasoning path explosion and dynamic knowledge interaction.

4.1. Reasoning Path Explosion

The number of paths can grow exponentially by using an LLM-based agent for reasoning path exploration, while the challenge is to optimize the generation of reasoning paths and iterative path exploration over large and sparse KGs. To mitigate the challenges of path explosion and reasoning complexity over large KGs in joint reasoning with LLMs, hybrid neuro-symbolic reasoning over LLMs and KGs should be investigated for more efficient and trustworthy reasoning.

4.2. Dynamic Knowledge Interaction

Enabling the flexible knowledge interaction between KGs and LLMs and incremental updates to KGs are essential for ensuring the completeness of KGs in dynamic environments. Developing mechanisms to seamlessly integrate new knowledge into KGs without disrupting existing structures or reasoning presents significant challenges. In order to facilitate the deep interactions between LLMs and KGs and keep reasoning over the up-to-date factual knowledge from KG, building semantically rich, schema-aware, dynamic knowledge adaptation KGs and developing a KG system that supports a flexible interaction with LLMs and incremental updates will also be crucial.

5. Conclusions

Recent years have witnessed the evolution of KGs' role in synthesizing with LLMs from approaches where KGs serve as background knowledge to KGs actively involving the reasoning of LLMs. This work gives a brief overview of recent advances in augmenting LLMs with KGs, highlighting the evolving role of KGs through a comparative summary of different approaches and applications. Despite the advancements, future work should focus on developing a neuro-symbolic reasoning system for more efficient reasoning and a KG system that enables flexible interaction with LLMs and incremental update.

Acknowledgments

This work was supported by the Novo Nordisk Foundation grant (NNF22OC0072415) and the GOBLIN COST Action (CA23147). I would like to thank Prof. Arijit Khan for his inspiration and support.

Declaration on Generative AI

During the preparation of this work, the author used Grammarly to check spelling and grammar. After using the tool, the author reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] S. Pan, L. Luo, Y. Wang, C. Chen, J. Wang, X. Wu, Unifying large language models and knowledge graphs: A roadmap, IEEE Transactions on Knowledge and Data Engineering 36 (2024) 3580–3599.
[2] C. Ma, Y. Chen, T. Wu, A. Khan, H. Wang, Large language models meet knowledge graphs for question answering: Synthesis and opportunities, arXiv preprint arXiv:2505.20099 (2025).
[3] M. Yasunaga, A. Bosselut, H. Ren, X. Zhang, C. D. Manning, P. S. Liang, J. Leskovec, Deep bidirectional language-knowledge graph pretraining, in: Proceedings of the 36th Conference on Neural Information Processing Systems, 2022, pp. 37309–37323.
[4] P. Jiang, L. Cao, C. D. Xiao, P. Bhatia, J. Sun, J. Han, KG-FIT: Knowledge graph fine-tuning upon open-world knowledge, in: Proceedings of the 38th Conference on Neural Information Processing Systems, 2024, pp. 136220–136258.
[5] S. Tian, Y. Luo, T. Xu, C. Yuan, H. Jiang, C. Wei, X. Wang, KG-Adapter: Enabling knowledge graph integration in large language models through parameter-efficient fine-tuning, in: Findings of the Association for Computational Linguistics: ACL 2024, 2024, pp. 3813–3828.
[6] K. Soman, P. W. Rose, J. H. Morris, R. E. Akbas, B. Smith, B. Peetoom, C. Villouta-Reyes, G. Cerono, Y. Shi, A. Rizk-Jackson, et al., Biomedical knowledge graph-optimized prompt generation for large language models, Bioinformatics 40 (2024) btae560.
[7] Y. Zhou, X. Yan, H. Huang, Z. Hao, H. Zhu, F. Liu, Knowledge graph-based entity alignment with unified representation for auditing, Complex & Intelligent Systems 11 (2025) 1–12.
[8] J. J. Thomas, R.-C. Mihailescu, An explainable KG-RAG-based approach to evidence-based fake news detection using LLMs, in: Proceedings of the AAAI Symposium Series, 2025.
[9] D. Wu, Y. Yan, Z. Liu, Z. Liu, M. Sun, KG-Infused RAG: Augmenting corpus-based RAG with external knowledge graphs, arXiv preprint arXiv:2506.09542 (2025).
[10] X. Zhu, Y. Xie, Y. Liu, Y. Li, W. Hu, Knowledge graph-guided retrieval augmented generation, arXiv preprint arXiv:2502.06864 (2025).

[11] L. Luo, J. Ju, B. Xiong, Y.-F. Li, G. Haffari, S. Pan, ChatRule: Mining logical rules with large language models for knowledge graph reasoning, in: Proceedings of the 29th Pacific-Asia Conference on Knowledge Discovery and Data Mining, 2025.

[12] R. Zhao, F. Zhao, L. Wang, X. Wang, G. Xu, KG-CoT: Chain-of-thought prompting of large language models over knowledge graphs for knowledge-aware question answering, in: Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, 2024, pp. 6642–6650.

[13] J. Sun, C. Xu, L. Tang, S. Wang, C. Lin, Y. Gong, L. Ni, H.-Y. Shum, J. Guo, Think-on-Graph: Deep and responsible reasoning of large language model with knowledge graph, in: Proceedings of the Twelfth International Conference on Learning Representations, 2024.

[14] J. Jiang, K. Zhou, W. X. Zhao, Y. Song, C. Zhu, H. Zhu, J.-R. Wen, Kg-Agent: An efficient autonomous agent framework for complex reasoning over knowledge graph, arXiv:2402.11163 (2024).

[15] L. Sun, Z. Tao, Y. Li, H. Arakawa, ODA: Observation-driven agent for integrating LLMs and knowledge graphs, in: Findings of the Association for Computational Linguistics: ACL 2024, 2024, pp. 7417–7431.

[16] L. Luo, Z. Zhao, G. Haffari, Y.-F. Li, C. Gong, S. Pan, Graph-constrained reasoning: Faithful reasoning on knowledge graphs with large language models, in: Proceedings of the 42st International Conference on Machine Learning, 2025.

[17] Z. Zhang, W. Zhao, A collaborative reasoning framework powered by reinforcement learning and large language models for complex questions answering over knowledge graph, in: Proceedings of the 31st International Conference on Computational Linguistics, 2025, pp. 10672–10684.

[18] S. Ma, C. Xu, X. Jiang, M. Li, H. Qu, J. Guo, Think-on-Graph 2.0: Deep and interpretable large language model reasoning with knowledge graph-guided retrieval, in: Proceedings of the Thirteenth International Conference on Learning Representations, 2025.

[19] X. Tan, X. Wang, Q. Liu, X. Xu, X. Yuan, W. Zhang, Paths-over-graph: Knowledge graph empowered large language model reasoning, in: Proceedings of the ACM on Web Conference, 2025, pp. 3505–3522.

[20] Y. Wang, Z. Chu, X. Ouyang, S. Wang, H. Hao, Y. Shen, J. Gu, S. Xue, J. Y. Zhang, Q. Cui, et al., LLMRG: Improving recommendations through large language model reasoning graphs, in: Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence, 2024, pp. 19189–19196.

[21] S. Wang, W. Fan, Y. Feng, X. Ma, S. Wang, D. Yin, Knowledge graph retrieval-augmented generation for LLM-based recommendation, in: Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics, 2025.

[22] J. Yi, Z. Xu, T. Huang, P. Yu, Challenges and innovations in LLM-powered fake news detection: A synthesis of approaches and future directions, in: Proceedings of the 2025 2nd International Conference on Generative Artificial Intelligence and Information Security, 2025, pp. 87–93.

[23] Y. Liu, J. Zhu, K. Zhang, H. Tang, Y. Zhang, X. Liu, Q. Liu, E. Chen, Detect, investigate, judge and determine: A novel LLM-based framework for few-shot fake news detection, arXiv preprint arXiv:2407.08952 (2024).

[24] D. Zhang, H. Du, X. Wang, M. Zhu, X. Pang, D. Wei, X. Wang, CMedRAGBot: A chinese medical chatbot based on graph RAG and large language models, Interdisciplinary Sciences: Computational Life Sciences (2025) 1–16.

[25] X. Zhao, S. Liu, S.-Y. Yang, C. Miao, MedRAG: Enhancing retrieval-augmented generation with knowledge graph-elicited reasoning for healthcare copilot, in: Proceedings of the ACM on Web Conference 2025, 2025, pp. 4442–4457.

[26] S. Marandi, Y.-S. Hu, M. Modarres, Complex system diagnostics using a knowledge graph-informed and large language model-enhanced framework, arXiv preprint arXiv:2505.21291 (2025).

[27] M. Islam, L. Elluri, K. P. Joshi, et al., Integrating knowledge graphs with retrieval-augmented generation to automate iot device security compliance, in: 2025 IEEE International Conference on Intelligence and Security Informatics, 2025.

[28] K. Kurniawan, E. Kiesling, A. Ekelhart, CyKG-RAG: Towards knowledge-graph enhanced retrieval augmented generation for cybersecurity, in: Proceedings of the International Workshop on Retrieval-Augmented Generation Enabled by Knowledge Graphs, 2024.

[29] A. Khan, T. Wu, X. Chen, Data management opportunities in unifying large language models+knowledge graphs, in: Proceedings of the International Workshop on Data Management Opportunities in Unifying Large Language Models+Knowledge Graphs, 2024.
[30] C. Ma, S. Chakrabarti, A. Khan, B. Molnár, Knowledge graph-based retrieval-augmented generation for schema matching, arXiv preprint arXiv:2501.08686 (2025).
[31] L. Wei, G. Xiao, M. Balazinska, RACOON: An LLM-based Framework for Retrieval-Augmented Column Type Annotation with a Knowledge Graph, in: Proceedings of the Third Table Representation Learning Workshop, 2024.
[32] S. Liu, Z. Zeng, L. Chen, A. Ainihaer, A. Ramasami, S. Chen, Y. Xu, M. Wu, J. Wang, TigerVector: Supporting vector search in graph databases for advanced RAGs, arXiv preprint arXiv:2501.11216 (2025).