LLMs+Graphs: Toward Graph-Native, Synergistic AI Systems

Arijit Khan[1[0000−0002−7312−6312]], Longxu Sun[2[0000−0002−7465−9153]], and Xin Huang[2[0000−0002−3650−0301]]

Bowling Green State University, Ohio, USA
arijitk@bgsu.edu
Hong Kong Baptist University, Hong Kong, China
{cslxsun,xinhuang}@comp.hkbu.edu.hk

Abstract. Large Language Models (LLMs) have advanced rapidly, but their limitations in structured and multi-hop reasoning underscore the need for graph-native, synergistic artificial intelligence (AI) systems. Graph-structured data underpins critical applications across social, biological, financial, transportation, web, and knowledge domains, making it essential to understand how LLMs can leverage graph computation for grounded, context-rich inference. Three complementary synergies are emerging: LLMs augmented with graph computation for retrieval and reasoning; bidirectional integration between LLMs and knowledge graphs (KGs), where LLMs support KG construction and curation while KGs enforce semantic constraints and factual consistency; and AI agents strengthened by graph algorithms for planning, decision making, and multi-step reasoning. In parallel, LLMs introduce new capabilities for graph data management and graph machine learning (ML) through natural language interfaces and hybrid LLM-graph neural network (GNN) pipelines. This tutorial synthesizes the algorithms, systems, and design principles driving these converging directions, offering data science and data mining researchers a unified perspective on integrating LLMs, graph data management, graph mining, graph ML, and agentic computation into next-generation graph-native AI systems.

1 Goals and Objectives

LLMs have quickly become a primary interface for data-intensive applications, reshaping how users query, explore, and manage information [11,17,58,83]. Their emergence coincides with a growing need to reason over graph-structured data, e.g., knowledge graphs, heterogeneous networks, and text-attributed graphs that encode rich semantics and relational structure. At the same time, existing graph neural networks (GNNs) and other deep graph models rely heavily on task-specific supervision and struggle to generalize across the diverse and rapidly evolving landscape of graph applications [53,67]. This has motivated the development of graph foundation models, which aim to unify LLMs' semantic and reasoning capabilities with graph-native computation to support broad, transferable graph intelligence [34,68]. As modern applications increasingly depend on complex graph data, understanding how LLMs can construct, curate, and leverage these structures–and how graphs can in turn ground and enhance LLM reasoning–has become a central challenge for the data mining community.

Concurrent advances in LLM-based AI agents have heightened the importance of graph-centric architectures. Agents rely on planning graphs, memory graphs, tool-calling graphs, and KG-grounded reasoning to achieve reliable planning, multi-step decision making, and long-term coherence [3]. At the same time, agents are beginning to reshape core data management tasks–including KG construction, graph reasoning, query formulation and optimization–creating a mutually reinforcing frontier between agentic computation and graph data systems [40,54,82]. This tutorial is motivated by the convergence of these trends: adapting LLMs' semantic and generalization capabilities to graph tasks, advancing graph foundation models, and leveraging graph structures to enable robust, interpretable, and scalable agentic behavior. For the PAKDD community, this convergence presents a timely opportunity to shape the next generation of graph-native, synergistic AI systems.

Previous Offering and Related Tutorials. We have not presented this tutorial at any prior venue, and it differs substantially from existing offerings. While prior tutorials focus narrowly on LLM–GNN alignment for graph learning [20], LLM+RAG [10], KG reasoning [35], or LLM+KG integration [41], our tutorial is the first to provide a unified, end-to-end treatment of the full synergy between LLMs and graph data. It covers both directions–LLMs for graphs and graphs for LLMs–while simultaneously integrating knowledge graphs, graph-structured data, and LLM-based agents into a single coherent framework. This breadth and unification constitute the core novelty of our proposal.

2 Target Audience

This tutorial is intended for participants working in the broader area of large language models, graph foundation models, graph learning, graph data mining, knowledge-augmented models, agentic AI, and graph AI systems from both academia and industry. Familiarity with basic LLM techniques would be helpful.

3 Outline

```
1 Introduction (30 minutes)
  - Large Language Models, Foundation Models, AI Agents,
    Retrieval-augmented Generation, Knowledge Graphs
2 LLMs for Graphs (20 minutes)
3 Graphs for LLMs (20 minutes)
4 Knowledge Graphs for LLMs (20 minutes)
5 LLMs for Knowledge Graphs (20 minutes)
6 Graphs for AI Agents (20 minutes)
7 AI Agents for Graphs (20 minutes)
8 Future Directions (30 minutes)
```

4 Description of Topics

4.1 LLMs for Graphs

LLMs are increasingly being used in many graph data management, mining, and ML problems [31,53,25,50]. • **Graph Querying.** Acting as *predictors*, the language understanding capacity of LLMs makes them suitable for processing natural language questions (NLQs) over structured graphs [13,45,38]; they enable natural language interfaces for graph querying, where they translate NLQs into executable GraphQL or Cypher queries [13,45,38] and power systems such as Neo4j's NLQ2Cypher pipeline [45]. • **Graph Mining.** In graph mining, LLMs again serve as *predictors*, using their reasoning and code generation abilities to extract graph properties and perform tasks such as graph classification, shortest path computation, cycle detection, and subgraph matching [33,80,22], as demonstrated by GraphWiz [6]. • **Graph Learning.** LLMs enhance graph learning by serving as *enhancers* and supporting *GNN-LLM alignment*. They enable zero-shot reasoning over text-attributed graphs and motivate graph foundation models [33,80]. Unified architectures such as Liu et al.'s model [33] handle node classification, link prediction, and related tasks without task-specific designs.

4.2 Graphs for LLMs

Graph-based retrieval-augmented generation (Graph RAG) improves LLM accuracy by supplying structured, relationship-rich context instead of relying only on flat text retrieval. Microsoft's GraphRAG builds a document-derived graph, organizes it into communities, and produces community-level summaries that offer coherent context for downstream reasoning [9]. ArchRAG advances this paradigm by enriching user queries with attributed subgraph communities from an external corpus and introducing an index that improves retrieval relevance and efficiency [65]. Zhou et al. provide an extensive empirical comparison of recent graph-based RAG methods, showing how graph structure enables more faithful and contextually grounded LLM outputs [84].

4.3 Knowledge Graphs for LLMs

We categorize the role of KGs in enhancing LLMs as follows [47,41]. • **Background Knowledge.** KGs provide structured facts that improve LLM reasoning by aligning subgraphs with text for joint training [75]. InfuserKI selectively integrates KG facts to reduce forgetting [62], while KG-Adapter and GAIL inject KG structure through efficient fine-tuning [59,81]. GRAG retrieves the top-k relevant subgraphs and aligns graph and text embeddings [19], and KG-RAG retrieves curated KG triples for fact-grounded reasoning across QA, recommendations, and data management [43,64,42,44]. • **Reasoning Guidelines.** KGs guide LLM reasoning by supplying candidate subgraphs or shaping each reasoning step. EtD extracts fine-grained KG facts for knowledge-enhanced prompts [32], and GCR encodes KGs as tries for graph-constrained decoding [39]. Online methods such as LLM-ARK and ToG support sequential KG-guided decisions [23,55], while agent-based systems like KG-Agent and ODA integrate KG tools and memory for iterative reasoning [24,56]. • **Refiners & Validators.** KGs refine and validate LLM outputs by filtering incorrect answers and grounding responses in factual structure. ACT-Selection and KG-Rank re-rank candidates using KG types and medical KGs [51,73], and KGR verifies factual statements in generated text [14]. EFSUM summarizes KG evidence for zero-shot QA [28], InteractiveKBQA enables iterative KG-guided correction [72], and LPKG fine-tunes LLMs with KG-based planning data to improve complex reasoning [63].

4.4 LLMs for Knowledge Graphs

LLMs augment KGs through knowledge extraction, completion, embedding, querying, analytics, and domain applications. • **KG Creation.** Multi-modal LLMs extract entities, relations, and facts from heterogeneous sources such as text, images, and tables [8,60]. They also support discovery, typing, linking, and end-to-end KG construction [61,27]. • **KG Completion.** LLMs improve link prediction by combining textual signals with KG facts [74]. Recent models directly generate missing entities in KG triples [52]. • **KG Embedding.** LLMs enrich KG embeddings by integrating textual semantics, as in KEPLER and K-BERT [66,36]. Multi-modal encoders further extend KG embeddings with image and graph information [21]. • **KG Querying.** LLMs interpret natural language questions, extract entities and relations, and support KG-grounded reasoning [76]. They also translate NLQs to SPARQL [1] and integrate KG facts into retrieval-augmented QA [2,71]. • **KG Analytics.** LLMs assist with graph reasoning tasks such as computing sizes, degrees, and connectivity, supported by prompting-based methods for natural language graph problems [77]. • **Domain-specific KG Applications.** LLM-KG synergy benefits healthcare, biomedical science [57], education [30], e-commerce [49], and spatio-temporal analysis [26].

4.5 Graphs for AI Agents

Recent work shows that key elements of agent performance are often represented more effectively with graphs, motivating agents that interact with diverse graph types. • **Task Planning Graphs.** Graphs support task reasoning, decomposition, and decision search. Knowledge graphs and Graph-of-Thought provide multi-hop context [69,4], task-dependency graphs structure sub-task relations [70], and state-space graphs guide sequential decisions [29]. • **Task Execution Graphs.** Execution improves when agents use graphs to organize tools and environments. Tool-calling graphs model function dependencies for efficient sequencing [37], and environment-interaction graphs capture relationships among agents and entities [12]. • **Memory Graphs.** Memory benefits from graph structures that expose relationships for retrieval and long-term reasoning. Recent work uses hierarchical knowledge graphs for structured memory [7], graph-based RAG for accurate retrieval [5], and dynamic graphs that evolve with new experiences [16]. • **Multi-agent Interaction Graphs.** Multi-agent systems rely on structured communication, and coordination graphs provide clear pathways for agent interaction [78].

4.6 AI Agents for Graphs

Agents can address many graph data management and mining tasks. • **Graph Reasoning.** Agents outperform stand-alone LLMs by decomposing problems into structured steps and invoking graph tools, enabling reliable execution of tasks such as shortest paths, cycles, triangle counting, maximum flow, PageRank, centrality, community detection, and node similarity [18,54,79]. • **Text2Cypher & Text2SPARQL.** Agentic systems improve semantic parsing by retrieving schema elements, validating intermediate queries, and iteratively repairing errors [15,48]. • **KG Construction.** Agents enhance KG construction by using domain-aware instructions and iterative tool-based refinement to extract relations and handle complex reasoning more accurately than stand-alone LLMs [46].

4.7 Future Directions

We conclude by summarizing open challenges and opportunities for data mining. • **Unification of LLM+KG+Vector DB & NeuroSymbolic AI.** Future systems must tightly integrate symbolic KGs, neural LLMs, vector retrieval, and rule-based reasoning into unified architectures that support consistent cross-modal inference and robust neuro-symbolic decision making. • **KG-based Agentic Memory.** Agents require scalable KG-structured memory that can store, update, and retrieve long-term knowledge to support reliable multi-step reasoning. • **Unifying Long Context with RAG.** Combining long-context models with KG-guided RAG demands mechanisms that balance extended attention with structured retrieval to prevent drift and hallucination. • **Explainability.** LLM+KG+agent systems need principled methods that expose graph-grounded evidence and trace decision pathways for transparent and trustworthy reasoning. • **Security and Privacy.** Protecting KG content, agent memory, and retrieval pipelines requires defenses against leakage, poisoning, unauthorized inference, and adversarial manipulation.

5 Biography

Arijit Khan is an Associate Professor at Bowling Green State University (Ohio, USA). His research is on data management and AI for the emerging problems in large graphs. He is an IEEE senior member and an ACM distinguished speaker. More info at `https://www.cs.bgsu.edu/arijitk/index.html`.
Longxu Sun is a postdoctoral research fellow in the Department of Computer Science, Hong Kong Baptist University (HKBU). Her research interests include graph mining, community search, and knowledge graph-based community search for Graph RAG. More info at `https://www.comp.hkbu.edu.hk/~cslxsun/`.
Xin Huang is an Associate Professor at the Department of Computer Science, Hong Kong Baptist University. His research interests include graph data management, graph mining and visualization, social network analysis, and privacy-aware computing. More info at `https://www.comp.hkbu.edu.hk/~xinhuang/`.

References

1. Avila, C.V.S., Vidal, V.M.P., Franco, W., Casanova, M.A.: Experiments with Text-to-SPARQL based on ChatGPT. In: ICSC (2024)
2. Baek, J., Aji, A.F., Saffari, A.: Knowledge-augmented language model prompting for zero-shot knowledge graph question answering. In: MATCHING (2023)
3. Bei, Y., Zhang, W., Wang, S., Chen, W., Zhou, S., Chen, H., Li, Y., Bu, J., Pan, S., Yu, Y., King, I., Karray, F., Yu, P.S.: Graphs meet AI agents: Taxonomy, progress, and future opportunities. CoRR **abs/2506.18019** (2025)
4. Besta, M., Blach, N., Kubicek, A., Gerstenberger, R., Podstawski, M., Gianinazzi, L., Gajda, J., Lehmann, T., Niewiadomski, H., Nyczyk, P., Hoefler, T.: Graph of thoughts: Solving elaborate problems with large language models. In: AAAI (2024)
5. Chen, B., Guo, Z., Yang, Z., Chen, Y., Chen, J., Liu, Z., Shi, C., Yang, C.: Pathrag: Pruning graph-based retrieval augmented generation with relational paths. CoRR **abs/2502.14902** (2025)
6. Chen, N., Li, Y., Tang, J., Li, J.: Graphwiz: An instruction-following language model for graph computational problems. In: KDD (2024)
7. Chen, W., Bai, T., Su, J., Luan, J., Liu, W., Shi, C.: Kg-retriever: Efficient knowledge indexing for retrieval-augmented large language models. CoRR **abs/2412.05547** (2024)
8. Deng, X., Sun, H., Lees, A., Wu, Y., Yu, C.: Turl: Table understanding through representation learning. SIGMOD Rec. **51**(1), 33–40 (2022)
9. Edge, D., Trinh, H., Cheng, N., Bradley, J., Chao, A., Mody, A., Truitt, S., Larson, J.: From local to global: A graph rag approach to query-focused summarization. CoRR **abs/2404.16130** (2024)
10. Fan, W., Ding, Y., Wang, S., Ning, L., Li, H., Yin, D., Chua, T.S., Li, Q.: Rag meets llms: Towards retrieval-augmented large language models. KDD Tutorial (2024)
11. Fernandez, R.C., Elmore, A.J., Franklin, M.J., Krishnan, S., Tan, C.: How large language models will disrupt data management. Proc. VLDB Endow. **16**(11), 3302–3309 (2023)
12. Gallici, M., Martin, M., Masmitja, I.: Transfqmix: Transformers for leveraging the graph structure of multi-agent reinforcement learning problems. In: AAMAS (2023)
13. Ganesan, B., Ghosh, S., Gupta, N., Kesarwani, M., Mehta, S., Sindhgatta, R.: Llm-powered graphql generator for data retrieval. In: IJCAI (2024)
14. Guan, X., Liu, Y., Lin, H., Lu, Y., He, B., Han, X., Sun, L.: Mitigating large language model hallucinations via autonomous knowledge graph-based retrofitting. In: AAAI. pp. 18126–18134 (2024)
15. Gusarov, A., Volkova, A., Khrulkov, V., Kuznetsov, A., Maslov, E., Oseledets, I.V.: Multi-agent graphrag: A text-to-cypher framework for labeled property graphs. CoRR **abs/2511.08274** (2025)
16. Gutierrez, B.J., Shu, Y., Gu, Y., Yasunaga, M., Su, Y.: Hipporag: Neurobiologically inspired long-term memory for large language models. In: NeurIPS (2024)
17. Halevy, A.Y., Choi, Y., Floratou, A., Franklin, M.J., Noy, N.F., Wang, H.: Will llms reshape, supercharge, or kill data science? Proc. VLDB Endow. **16**(12), 4114–4115 (2023)
18. Hu, Y., Lei, R., Huang, X., Wei, Z., Liu, Y.: Scalable and accurate graph reasoning with llm-based multi-agents. CoRR **abs/2410.05130** (2024)
19. Hu, Y., Lei, Z., Zhang, Z., Pan, B., Ling, C., Zhao, L.: GRAG: Graph retrieval-augmented generation. CoRR **abs/2405.16506** (2024)

20. Huang, C., Ren, X., Tang, J., Yin, D., Chawla, N.V.: Large language models for graphs: Progresses and directions. In: WWW. pp. 1284–1287 (2024)
21. Huang, N., Deshpande, Y.R., Liu, Y., Alberts, H., Cho, K., Vania, C., Calixto, I.: Endowing language models with multimodal knowledge graph representations. CoRR **abs/2206.13163** (2022)
22. Huang, X., Han, K., Yang, Y., Bao, D., Tao, Q., Chai, Z., Zhu, Q.: Can gnn be good adapter for llms? In: WWW (2024)
23. Huang, Y.: Evaluating and enhancing large language models for conversational reasoning on knowledge graphs. CoRR **abs/2312.11282** (2023)
24. Jiang, J., Zhou, K., Zhao, X., Song, Y., Zhu, C., Zhu, H., Wen, J.: Kg-agent: An efficient autonomous agent framework for complex reasoning over knowledge graph. In: ACL (2025)
25. Jin, B., Liu, G., Han, C., Jiang, M., Ji, H., Han, J.: Large language models on graphs: A comprehensive survey. IEEE Trans. Knowl. Data Eng. **36**(12), 8622–8642 (2024)
26. Jin, M., Wen, Q., Liang, Y., Zhang, C., Xue, S., Wang, X., Zhang, J., Wang, Y., Chen, H., Li, X., Pan, S., Tseng, V.S., Zheng, Y., Chen, L., Xiong, H.: Large models for time series and spatio-temporal data: A survey and outlook. CoRR **abs/2310.10196** (2023)
27. Joshi, M., Levy, O., Zettlemoyer, L., Weld, D.S.: Bert for coreference resolution: Baselines and analysis. In: EMNLP-IJCNLP (2019)
28. Ko, S., Cho, H., Chae, H., Yeo, J., Lee, D.: Evidence-focused fact summarization for knowledge-augmented zero-shot question answering. In: EMNLP. pp. 10636–10651 (2024)
29. Leurent, E., Maillard, O.: Monte-carlo graph search: the value of merging similar states. In: ACML. PMLR, vol. 129, pp. 577–592 (2020)
30. Li, X., Henriksson, A., Duneld, M., Nouri, J., Wu, Y.: Evaluating embeddings from pre-trained language models and knowledge graphs for educational content recommendation. Future Internet **16**(1) (2024)
31. Li, Y., Li, Z., Wang, P., Li, J., Sun, X., Cheng, H., Yu, J.X.: A survey of graph meets large language model: Progress and future directions. In: IJCAI. pp. 8123–8131 (2024)
32. Liu, G., Zhang, Y., Li, Y., Yao, Q.: Explore then determine: A GNN-LLM synergy framework for reasoning over knowledge graph. arXiv **abs/2406.01145** (2024)
33. Liu, H., Feng, J., Kong, L., Liang, N., Tao, D., Chen, Y., Zhang, M.: One for all: Towards training one graph model for all classification tasks. In: ICLR (2024)
34. Liu, J., Yang, C., Lu, Z., Chen, J., Li, Y., Zhang, M., Bai, T., Fang, Y., Sun, L., Yu, P.S., Shi, C.: Graph foundation models: Concepts, opportunities and challenges. IEEE Trans. Pattern Anal. Mach. Intell. **47**(6), 5023–5044 (2025)
35. Liu, L., Wang, Z., Bai, J., Song, Y., Tong, H.: New frontiers of knowledge graph reasoning: Recent advances and future trends. In: WWW. pp. 1294–1297 (2024)
36. Liu, W., Zhou, P., Zhao, Z., Wang, Z., Ju, Q., Deng, H., Wang, P.: K-bert: Enabling language representation with knowledge graph. In: AAAI (2020)
37. Liu, X., Peng, Z., Yi, X., Xie, X., Xiang, L., Liu, Y., Xu, D.: Toolnet: Connecting llms with massive tools via tool graph. CoRR **abs/2403.00839** (2024)
38. Liu, Y., Wang, X., Ge, J., Wang, H., Xu, D., Jia, Y.: Text to graph query using filter condition attributes. In: SGDA Workshop@VLDB (2024)
39. Luo, L., Zhao, Z., Gong, C., Haffari, G., Pan, S.: Graph-constrained reasoning: Faithful reasoning on knowledge graphs with large language models. arXiv:2410.13080 (2024)

40. Luo, Y., Li, G., Fan, J., Tang, N.: Data agents: Levels, state of the art, and open problems. In: SIGMOD (2026)
41. Ma, C., Chen, Y., Wu, T., Khan, A., Wang, H.: Unifying large language models and knowledge graphs for question answering: Recent advances and opportunities. In: EDBT. pp. 1174–1177 (2025)
42. Ma, C., Chakrabarti, S., Khan, A., Molnár, B.: Knowledge graph-based retrieval-augmented generation for schema matching. CoRR **abs/2501.08686** (2025)
43. Ma, C., Chen, Y., Wu, T., Khan, A., Wang, H.: Large language models meet knowledge graphs for question answering: Synthesis and opportunities. In: EMNLP. pp. 24578–24597 (2025)
44. Ma, C., Zhang, Z., Khan, A., Schelter, S., Groth, P.: Cost-efficient rag for entity matching with llms: A blocking-based exploration. CoRR **abs/2602.05708** (2026)
45. neo4j: neo4j. `https://neo4j.com/labs/neodash/2.4/user-guide/extensions/natural-language-queries/` (2024)
46. Ning, Y., Liu, H.: Urbankgent: A unified large language model agent framework for urban knowledge graph construction. In: NeurIPS (2024)
47. Pan, S., Luo, L., Wang, Y., Chen, C., Wang, J., Wu, X.: Unifying large language models and knowledge graphs: A roadmap. IEEE Trans. Knowl. Data Eng. **36**(7), 3580–3599 (2024)
48. Perevalov, A., Both, A.: Text-to-sparql goes beyond english: Multilingual question answering over knowledge graphs through human-inspired reasoning. CoRR **abs/2507.16971** (2025)
49. Regino, A.G., Caus, R.O., Hochgreb, V., dos Reis, J.C.: From natural language texts to rdf triples: A novel approach to generating e-commerce knowledge graphs. In: CIKM (2023)
50. Ren, X., Tang, J., Yin, D., Chawla, N.V., Huang, C.: A survey of large language models for graphs. In: KDD. pp. 6616–6626 (2024)
51. Salnikov, M., Lysyuk, M., Braslavski, P., Razzhigaev, A., Malykh, V.A., Panchenko, A.: Answer candidate type selection: Text-to-text language model for closed book question answering meets knowledge graphs. In: KONVENS. pp. 155–164 (2023)
52. Saxena, A., Kochsiek, A., Gemulla, R.: Sequence-to-sequence knowledge graph completion and question answering. In: ACL (2022)
53. Shang, W., Huang, X.: A survey of large language models on generative graph analytics: Query, learning, and applications. IEEE Trans. Knowl. Data Eng. **37**(12), 6799–6819 (2025)
54. Shi, B., Panagiotas, I.: GDS agent: A graph algorithmic reasoning agent. CoRR **abs/2508.20637** (2025)
55. Sun, J., Xu, C., Tang, L., Wang, S., Lin, C., Gong, Y., Ni, L., Shum, H.Y., Guo, J.: Think-on-Graph: Deep and responsible reasoning of large language model with knowledge graph. In: ICLR (2024)
56. Sun, L., Tao, Z., Li, Y., Arakawa, H.: ODA: Observation-driven agent for integrating LLMs and knowledge graphs. In: ACL. pp. 7417–7431 (2024)
57. Sung, M., Lee, J., Yi, S.S., Jeon, M., Kim, S., Kang, J.: Can language models be biomedical knowledge bases? In: EMNLP (2021)
58. Tan, W.: Unstructured and structured data: Can we have the best of both worlds with large language models? IEEE Data Eng. Bull. **47**(2), 5–11 (2023)
59. Tian, S., Luo, Y., Xu, T., Yuan, C., Jiang, H., Wei, C., Wang, X.: KG-Adapter: Enabling knowledge graph integration in large language models through parameter-efficient fine-tuning. In: ACL. pp. 3813–3828 (2024)

60. Vogel, L., Hilprecht, B., Binnig, C.: Towards foundation models for relational databases [vision paper]. In: Table Representation Learning Workshop@NeurIPS (2022)
61. Wadden, D., Wennberg, U., Luan, Y., Hajishirzi, H.: Entity, relation, and event extraction with contextualized span representations. In: EMNLP-IJCNLP (2019)
62. Wang, F., Bao, R., Wang, S., Yu, W., Liu, Y., Cheng, W., Chen, H.: Infuserki: Enhancing large language models with knowledge graphs via infuser-guided knowledge integration. In: EMNLP. pp. 3675–3688 (2024)
63. Wang, J., Chen, M., Hu, B., Yang, D., Liu, Z., Shen, Y., Wei, P., Zhang, Z., Gu, J., Zhou, J., Pan, J.Z., Zhang, W., Chen, H.: Learning to plan for retrieval-augmented large language models from knowledge graphs. In: EMNLP. pp. 7813–7835 (2024)
64. Wang, S., Fan, W., Feng, Y., Lin, S., Ma, X., Wang, S., Yin, D.: Knowledge graph retrieval-augmented generation for llm-based recommendation. In: ACL (1). pp. 27152–27168. ACL (2025)
65. Wang, S., Fang, Y., Zhou, Y., Liu, X., Ma, Y.: Archrag: Attributed community-based hierarchical retrieval-augmented generation. In: AAAI (2026)
66. Wang, X., Gao, T., Zhu, Z., Zhang, Z., Liu, Z., Li, J., Tang, J.: Kepler: A unified model for knowledge embedding and pre-trained language representation. Trans. Assoc. Comput. Linguistics **9**, 176–194 (2021)
67. Wang, Z., Liu, Z., Ma, T., Li, J., Zhang, Z., Fu, X., Li, Y., Yuan, Z., Song, W., Ma, Y., Zeng, Q., Chen, X., Zhao, J., Li, J., Jiang, M., Lio, P., Chawla, N.V., Zhang, C., Ye, Y.: Graph foundation models: A comprehensive survey. arXiv:2505.15116 (2025)
68. Wang, Z., Zhang, C., Li, J., Chawla, N.V., Ye, Y.: Graph foundation models: Challenges, methods, and open questions. In: KDD. pp. 6184–6194 (2025)
69. Wen, Y., Wang, Z., Sun, J.: Mindmap: Knowledge graph prompting sparks graph of thoughts in large language models. In: ACL (2024)
70. Wu, X., Shen, Y., Shan, C., Song, K., Wang, S., Zhang, B., Feng, J., Cheng, H., Chen, W., Xiong, Y., Li, D.: Can graph learning improve planning in llm-based agents? In: NeurIPS (2024)
71. Wu, Y., Hu, N., Bi, S., Qi, G., Ren, J., Xie, A., Song, W.: Retrieve-rewrite-answer: A kg-to-text enhanced llms framework for knowledge graph question answering. In: IJCKG (2023)
72. Xiong, G., Bao, J., Zhao, W.: Interactive-KBQA: Multi-turn interactions for knowledge base question answering with large language models. In: ACL. pp. 10561–10582 (2024)
73. Yang, R., Liu, H., Marrese-Taylor, E., Zeng, Q., Ke, Y., Li, W., Cheng, L., Chen, Q., Caverlee, J., Matsuo, Y., Li, I.: KG-Rank: Enhancing large language models for medical QA with knowledge graphs and ranking techniques. In: BioNLP Workshop@ACL. pp. 155–166 (2024)
74. Yao, L., Mao, C., Luo, Y.: Kg-bert: Bert for knowledge graph completion. CoRR **abs/1909.03193** (2019)
75. Yasunaga, M., Bosselut, A., Ren, H., Zhang, X., Manning, C.D., Liang, P.S., Leskovec, J.: Deep bidirectional language-knowledge graph pretraining. In: NeurIPS. pp. 37309–37323 (2022)
76. Yasunaga, M., Ren, H., Bosselut, A., Liang, P., Leskovec, J.: Qa-gnn: Reasoning with language models and knowledge graphs for question answering. In: NAACL-HLT (2021)
77. Ye, R., Zhang, C., Wang, R., Xu, S., Zhang, Y.: Language is all a graph needs. In: Findings of the Association for Computational Linguistics: EACL (2024)

78. Yu, J., Ding, Y., Sato, H.: Dyntaskmas: A dynamic task graph-driven framework for asynchronous and parallel llm-based multi-agent systems. In: ICAPS (2025)
79. Yuan, Z., Liu, M., Wang, H., Qin, B.: MA-GTS: A multi-agent framework for solving complex graph problems in real-world applications. In: EMNLP (2025)
80. Zhang, M., Sun, M., Wang, P., Fan, S., Mo, Y., Xu, X., Liu, H., Yang, C., Shi, C.: Graphtranslator: Aligning graph model to large language model for open-ended tasks. In: WWW (2024)
81. Zhang, Z., Wen, L., Zhao, W.: A GAIL fine-tuned LLM enhanced framework for low-resource knowledge graph question answering. In: CIKM. pp. 3300–3309 (2024)
82. Zhao, X., Blum, M., Gao, F., Chen, Y., Yang, B., Marquez-Carpintero, L., Pina-Navarro, M., Fu, Y., Morikawa, S., Iwasawa, Y., Matsuo, Y., Park, C., Li, I.: Agentigraph: A multi-agent knowledge graph framework for interactive, domain-specific LLM chatbots. In: CIKM. pp. 6757–6761 (2025)
83. Zhou, X., He, J., Zhou, W., Chen, H., Tang, Z., Zhao, H., Tong, X., Li, G., Chen, Y., Zhou, J., Sun, Z., Hui, B., Wang, S., He, C., Liu, Z., Zhou, J., Wu, F.: A survey of llm $\times$ data. arXiv:2505.18458 (2025)
84. Zhou, Y., Su, Y., Sun, Y., Wang, S., Wang, T., He, R., Zhang, Y., Liang, S., Liu, X., Ma, Y., Fang, Y.: In-depth analysis of graph-based RAG in a unified framework. Proc. VLDB Endow. **18**(13), 5623–5637 (2025)