Agents+Graph: 3rd International Workshop on Data Management Opportunities in Bringing Agents with Graph Data

Yixiang Fang
The Chinese University of Hong Kong, Shenzhen, China
fangyixiang@cuhk.edu.cn

Arijit Khan
Bowling Green State University
Ohio, USA
arijitk@bgsu.edu

Tianxing Wu
Southeast University
Nanjing, China
tianxingwu@seu.edu.cn

Da Yan
Indiana University Bloomington
Bloomington, USA
yanda@iu.edu

ABSTRACT

AI agents have become increasingly prominent, with LLM-based agents introducing stronger capabilities for autonomous reasoning, planning, and interaction. A growing foundation of these systems is their reliance on graph-structured data, including planning graphs, graph-of-thought structures, task dependency graphs, agent coordination graphs, tool calling graphs, and memory graphs, which collectively support planning, execution, tool invocation, and long-term memory management. At the same time, AI agents can enhance graph data systems through automated database tuning, adaptive query optimization, natural language interfaces such as text2Cypher translation, and agent-driven graph analytics. These developments create new opportunities for data management research, calling for algorithms, system architectures, and evaluation frameworks that jointly leverage agent capabilities and large-scale graph computation. The **"Agents+Graph"** workshop provides a timely venue for advancing this emerging area by uniting researchers focused on LLMs, AI agents, graph data management, and graph machine learning in real-world applications.

VLDB Workshop Reference Format:
Yixiang Fang, Arijit Khan, Tianxing Wu, and Da Yan. Agents+Graph: 3rd International Workshop on Data Management Opportunities in Bringing Agents with Graph Data. VLDB 2026 Workshop: Agent+Graph.

1 MOTIVATIONS, TOPICS, AND GOALS

Large language models (LLMs) have enabled a new class of AI agents, which are autonomous or semi-autonomous systems capable of reasoning, planning, tool use, and operating in complex environments [31]. By interpreting natural language instructions, decomposing tasks, coordinating diverse tools and data sources, and adapting to context, LLM-based agents are emerging as core enablers of general-purpose agency in data-intensive settings [9, 25]. As these agents mature, they are already reshaping fundamental data management tasks including data cleaning [23], query formulation and optimization [15], data exploration [2], and system tuning [18], underscoring the need for the database community to understand their computational foundations and their implications for future data systems.

Graph-structured data have become a central enabler of agentic behavior, supporting the reasoning, planning, and execution capabilities of modern LLM-based agents. These agents increasingly depend on diverse graph forms, such as planning graphs, graph-of-thought structures, task-dependency graphs, tool-calling graphs, agent-coordination graphs, and memory graphs to encode tasks, dependencies, tools, and long-term knowledge [3]. By leveraging these structures, agents can decompose complex objectives, orchestrate multi-step workflows, select and sequence tools effectively, and maintain persistent memory across interactions, positioning graphs as a critical framework for building reliable, interpretable, and scalable agentic systems.

Graph data has long been central to data management research, powering knowledge graphs, social networks, scientific workloads, and enterprise integration [6]. As AI agents become more capable and widely deployed, their dependence on–and influence over–graph-structured data continues to expand. Agents can construct, maintain, query, and optimize graph data systems [13, 24, 30], while graphs provide the structural grounding and relational context needed for agentic reasoning and decision-making [32]. This bidirectional reinforcement is driving renewed interest in unifying agent capabilities with graph-based representations, including integrating parametric and structured knowledge and advancing graph-aware modeling [1, 5, 20]. These developments motivate the **Agents+Graph** workshop, which brings together researchers and practitioners to rethink graph storage, indexing, querying, and analytics for agent-centric workloads and to chart the path toward next-generation data systems that tightly couple AI agents with graph-structured data.

1.1 Graph Data for Agents

Recent work shows that key elements of agent performance, such as sub-task structure, memory organization, and multi-agent communication, are often represented more effectively with graphs. This has sparked growing interest in designing agents that can interact with diverse graph types. The following section reviews recent progress in this direction, emphasizing how agent-graph interaction enhances four core agent capabilities: task planning, task execution, memory management, and multi-agent coordination.

• **Agent Interaction with Task Planning Graphs.** Planning enables AI agents to interpret user goals and generate action sequences. Graph structures enhance this capability by supporting task reasoning, task decomposition, and decision search. *Knowledge graphs* [26] and *Graph-of-Thought* [4] enrich reasoning with multi-hop context; *task-dependency graphs* [27] structure sub-task relationships for effective decomposition; and *state-space graphs* [17] guide sequential decision-making in complex environments.

• **Agent Interaction with Task Execution Graphs.** The execution phase centers on how agents use tools and interact with their

Proceedings of the VLDB Endowment. ISSN 2150-8097.

environments, and graph structures help organize both processes. *Tool-calling graphs* [19] model dependencies among functions to guide efficient tool sequencing, while *environment-interaction graphs* [10] capture relationships among agents and environmental entities, enabling more informed perception, reasoning, and action.

• **Agent Interaction with Memory Graphs.** Memory helps agents store and reuse past information. Graph-based representations strengthen this process by exposing relationships that improve pattern recognition, long-term reasoning, retrieval, and continual updating. Recent work leverages *hierarchical knowledge graphs* for structured long-term memory [8], *graph-based RAG* methods for accurate and efficient retrieval [7], and *dynamic graph* architectures that incrementally evolve as agents acquire new experiences [12].

• **Graphs for Multi-agent Interaction.** Multi-agent systems rely on efficient communication and coordinated decision making among specialized agents. An agent-coordination graph provides a structured view for such communication pathways [29].

• **KGs as Reasoning Guidelines, Refiners, & Validators.** Knowledge graphs can be incorporated into an agent's reasoning pipeline to provide structured knowledge that strengthens an LLM's decision making and problem solving capabilities [14]. They also improve output accuracy and reliability by filtering and validating candidate responses against verified, schema-grounded information [32].

1.2 Agents for Graph Data

As a powerful tool, agents can be used to solve many graph data management and mining problems.

• **Graph Reasoning.** Agents surpass stand-alone LLMs in graph reasoning by breaking problems into structured steps and invoking precise graph analysis tools, enabling reliable execution of tasks such as shortest-path search, cycle detection, triangle counting, maximum flow, PageRank, centrality analysis, community detection, and node similarity–areas where single LLMs fail due to weak structural recall and limited algorithmic precision [13, 24, 30].

• **Text2Cypher & Text2SPARQL.** Agentic architectures improve Text2Cypher and Text2SPARQL translations by retrieving schema elements, validating intermediate queries, and iteratively repairing errors–capabilities consistently emphasized in recent agent-based semantic-parsing research [11, 22].

• **KG Construction.** Agents improve knowledge graph construction by using domain-aware instructions and tool-based iterative refinement to extract diverse relations and handle complex reasoning more accurately than stand-alone LLMs [21].

1.3 Opportunities for Data Management

The combination of agents and graphs provides various exciting opportunities for data management and data science research.

• **Agent-centric Data and Input Modeling.** AI agents require principled graph-aware input modeling, spanning multimodal graph serialization, subgraph extraction, and structured task representation graphs, to support reliable planning, memory, and tool use.

• **Unified Vector & Graph Memory for Agent Retrieval.** Next-generation agent architectures demand integrated vector-graph indexes that support hybrid retrieval, semantic caching, and graph-augmented search to enable efficient, grounded decision making without repeated expensive inference.

• **Consistency & Safety via Graph-grounded Reasoning.** Graph-based external memory and verification offer agents a principled mechanism to reduce hallucinations, enforce factual consistency, and support trustworthy reasoning in high-stakes domains.

• **Explainability, Provenance, & Benchmarking.** Graph structures provide natural scaffolding for tracing agent decisions, attributing provenance, and developing standardized benchmarks that evaluate planning, execution, tool calling, memory, and multi-agent coordination.

1.4 Goals: Why Important? Why Now?

The global AI agents market is expanding rapidly from \$5.68B in 2024 to \$8.34B in 2025 (a 47% compound annual growth rate)[1]. This surge in adoption is exposing critical challenges in planning reliability, memory robustness, reasoning transparency, tool-use consistency, and security. Graph-structured data, including planning graphs, memory graphs, tool-calling graphs, and KG-grounded reasoning modules, are emerging as the essential foundations that enable agents to plan coherently, coordinate tools effectively, maintain long-term memory, and operate safely at scale. At the same time, agents are beginning to transform core data management tasks such as query formulation, KG construction, data cleaning, and graph reasoning, creating a mutually reinforcing research frontier that aligns directly with VLDB's strategic interests. With market momentum accelerating and graph-centric agent architectures becoming fundamental to real-world deployments, VLDB 2026 offers a timely and impactful venue to shape the next generation of **agent-first, graph-native data systems**.

1.5 Workshop Topics

We solicit unpublished papers discussing issues and successes of combining agents and graphs. Topics include but are not limited to:

- Agent-centric graph data modeling
- Graph-augmented planning
- Tool-calling & execution graphs
- Graph-based long-term memory
- Multi-agent coordination graphs
- Graph-grounded safety & verification
- Agents for graph reasoning & Text2Cypher/SPARQL
- Agent-oriented graph storage, indexing, & querying
- Next-generation agent-graph data systems

Papers must maintain a clear data management focus by addressing relevant challenges and offering contributions that advance data management solutions for the topics outlined in Section 1.3.

2 RELATED EVENTS AND PAST WORKSHOPS

Recent workshops at major database venues reflect the growing intersection of data management, AI, and graph technologies. At **VLDB 2025**, these included "Data Management Opportunities in Unifying Large Language Models and Graph Data (LLM+Graph)", "Data-Centric AI", and "Governance, Understanding and Integration of Data for Effective and Responsible AI". At **SIGMOD 2025**,

[1]https://www.researchandmarkets.com/reports/6103459/ai-agents-global-market-report?

relevant workshops featured "Exploiting Artificial Intelligence Techniques for Data Management (AIDM)", "Data Management for End-to-End Machine Learning (DEEM)", "Graph Data Management Experiences & Systems (GRADES) and Network Data Analytics (NDA)", "Novel Optimizations for Visionary AI Systems (NOVA)", and "Connecting Academia and Industry on Modern Integrated Database and AI Systems (MIDAS)". At **ICDE 2025**, related workshops included "Databases and Machine Learning (DBML)", "Coupling of Large Language Models with Vector Data Management (LLM+Vector Data)", "Data-AI Systems (DAIS)", and "Data Engineering Meets Large Language Models (DExLLM)". Among these, the most closely aligned was **LLM+Graph@VLDB 2025** [16], organized by the same team, which attracted one of the largest workshop audiences at VLDB 2025 (over 120 attendees). As the 3rd workshop in this growing series, the **Agents+Graph** workshop is well positioned to draw substantial interest, driven by the rapid adoption of AI agents and their emerging reliance on graph-structured data for planning, memory, and retrieval. To the best of our knowledge, *this represents one of the earliest efforts to connect the agentic AI revolution with graph-structured data in a database conference setting.*

3 WORKSHOP SCHEDULE

Agents+Graph will be a half-day workshop on Sep 4, 2026 at Boston, MA, USA, co-located with VLDB 2026, with the following tentative schedule.

Session 1 (13:45-15:15)
• **Keynote** - Yunyao Li (Adobe)
• **Keynote** - Chen Li (UC Irvine)
• **Keynote** - Eric Lo (The Chinese University of Hong Kong)

Session 2 (15:45-17:30)
• **Keynote** - Cheng Chen (ByteDance)
• **Panel** - *Databases, Data Science, and AI Futures*
• **Paper Presentation** - *How Much Structure Should Agentic Graph Memory Build for Text Retrieval?*: Satyananda Kashyap (IBM), Niharika Dsouza (IBM), Nandana Mihindukulasooriya (IBM), and Horst Samulowitz (IBM).
• **Paper Presentation** - *Knowledge Graphs as the Missing Data Layer for LLM-Based Industrial Asset Operations*: Madhulatha Mandarapu (Vaidhyamegha Pvt Ltd) and Sandeep Kunkunuru (Vaidhyamegha Pvt Ltd).
• **Paper Presentation** - *AIP: A Graph Representation for Learning and Governing Agent Skills*: Zachary Blumenfeld (Neo4j).
• **Paper Presentation** - *Reducing Hallucinations in Complex Question Answering using Simple Graph-based Retrieval-Augmented Generation*: Jacek Cala (Newcastle University), Christopher Wedge (Newcastle University), Danny Dixon (Newcastle University), and Joshua Stutter (Newcastle University).

Poster Session (17:30-18:15)

The workshop will be organized as an in-person event, along with a virtual component in case of travel restrictions or other exceptional situations. We expect about around 120 participants based on prior experience.

4 PROGRAM COMMITTEE

• A. Erdem Sariyuce (University at Buffalo)
• Chao Zhang (Lyon 1 University)
• Chenhao Ma (The Chinese University of Hong Kong, Shenzhen)
• Chuangtao Ma (Aalborg University)
• Dazhuo Qiu (Lyon 1 University)
• Fan Zhang (Guangzhou University)
• Guimu Guo (Rowan University)
• Hussein Abdallah (McGill University)
• Jianzhong Qi (The University of Melbourne)
• Jingbang Chen (The Chinese University of Hong Kong, Shenzhen)
• Lyuheng Yuan (Indiana University Bloomington)
• Meng Wang (Tongji University)
• Miao Qiao (The University of Auckland)
• Nandish Jayaram (Intuit)
• Qihao Cheng (Tsinghua University)
• Renata Borovica-Gajic (University of Melbourne)
• Tianhao Wu (Nanyang Technological University)
• Wensheng Luo (Hunan University)
• Wentao Li (University of Leicester)
• Xiang Lian (Kent State University)
• Xin Huang (Hong Kong Baptist University)
• Yingli Zhou (The Chinese University of Hong Kong, Shenzhen)
• Yiwen Zhu (Microsoft)
• Yubo Chen (Minzu University of China)
• Yuxiang Wang (Hangzhou Dianzi University)

5 WORKSHOP CO-CHAIRS

Yixiang Fang is an Associate Professor at the School of Data Science and School of Artificial Intelligence in the Chinese University of Hong Kong, Shenzhen, China. His research interests focus on data management, data mining, and artificial intelligence over big graph data. He has published extensively in the areas of database and data mining, including One of the Best Papers in SIGMOD 2020, and most of them are published in top-tier conferences and journals. He was awarded the 2021 ACM SIGMOD Research Highlight Award. He is an editorial board member of the journal of Information & Processing Management (IPM). He has also served as program committee members for several top conferences (e.g., PVLDB, ICDE, and KDD) and invited reviewers for top journals (e.g., TKDE and VLDBJ) in the areas of database and data mining. He is a member of ACM, IEEE, and CCF. More information at https://fangyixiang.github.io/

Arijit Khan is an Associate Professor at Bowling Green State University (Ohio, USA). His PhD is from University of California, Santa Barbara, USA, and he did a post-doc in the Systems group at ETH Zurich, Switzerland. He has been an assistant professor at Nanyang Technological University, Singapore and an associate professor at Aalborg University, Denmark. 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. Arijit is the recipient of the IBM Ph.D. Fellowship (2012-13), a VLDB Distinguished Reviewer award (2022), three SIGMOD Distinguished PC awards (2024, 2025, 2026), and a KDD outstanding PC award (2025). He is the author of a book on uncertain graphs and over 120 publications in top venues including SIGMOD, PVLDB, TKDE, ICDE, ICLR, KDD, EMNLP, SDM, USENIX ATC, EDBT, Web Conference (WWW), WSDM, CIKM, TKDD, VLDB Journal, and SIGMOD Record. Dr Khan served/is serving as an associate editor of TKDE and TKDD, co-Editor-in-Chief of Knowledge Engineering Review (KER), CIKM short paper track co-chair 2024, ICDE demonstration paper track program co-chair 2025, KDD 2025 PhD consortium track program co-chair, and Australasian Database Conference (ADC) 2025 PC co-chair. More information at https://www.cs.bgsu.edu/arijitk/index.html.

Tianxing Wu is an Associate Professor working at the School of Computer Science and Engineering of Southeast University, China. He is one of the main contributors to build Chinese large-scale encyclopedic knowledge graph: Zhishi.me and schema knowledge graph: Linked Open Schema. He has achieved SIGIR 2024 Best Short Paper Nominee, CCKS 2022 Best Paper Award, WISA 2024 Best Student Paper Award. His research interests include knowledge graph, knowledge representation and reasoning, and data mining. He has published over 70 papers in top-tier venues, e.g., SIGMOD, ICDE, NeurIPS, SIGIR, ACL, AAAI, IJCAI, and TKDE. He is an associate Editor-in-Chief of Knowledge Engineering Review, and an editorial board member of International Journal on Semantic Web and Information Systems, and Data Intelligence. He has served as the (senior) program committee member of AAAI, IJCAI, NeurIPS, ACL, KDD, WWW, ICML, and ICLR. More information at https://tianxing-wu.github.io/.

Da Yan is an Associate Professor in the Department of Computer Science of the Luddy School of Informatics, Computing, and Engineering at Indiana University Bloomington. He is a DOE Early Career Research Program (ECRP) awardee in 2023, the sole winner of the Hong Kong 2015 Young Scientist Award in Physical/Mathematical Science, and a senior member of ACM and IEEE. His research interests include parallel and distributed systems for big data analytics, data mining, and machine learning. He frequently publishes in top databases and AI conferences and journals, and he serves extensively in the major top databases and AI conferences and journals as reviewers. Dr. Yan leads the BIOKDD workshop with SIGKDD (now in its 23rd year), and co-organized a Dagstuhl seminar and a few top conferences. He also served as guest editors of journals such as IEEE/ACM TCBB, BMC Bioinformatics, and IEEE CG&A. More information at https://homes.luddy.indiana.edu/yanda/home.html.

The team possesses a diverse range of complementary expertise and a comprehensive scientific network to make the **"Agents+Graph"** workshop a great success. Specifically, Yixiang Fang and Arijit Khan are highly recognized for their contributions in designing efficient graph algorithms and graph AI systems. Da Yan is best known for developing systems for big graph analytics. Tianxing Wu is an expert in managing and mining knowledge graphs (KGs), and their AI applications in medicine, law, smart grid, etc. The team represents the geographic distribution (2xUS and 2xAsia). The team successfully organized the 1st edition **"LLM+KG"** workshop [16] at VLDB 2024 and the 2nd edition **"LLM+Graph"** workshop [28] at VLDB 2025, each attracting an audience over 120 participants.

REFERENCES

[1] L. Sun A. Khan and X. Huang. 2026. LLMs+Graphs: Toward Graph-Native, Synergistic AI Systems. In *Tutorials at PAKDD.*

[2] S. Amer-Yahia. 2024. Intelligent Agents for Data Exploration. *PVLDB* 17, 12 (2024), 4521–4530.

[3] Y. Bei, W. Zhang, S. Wang, W. Chen, S. Zhou, H. Chen, Y. Li, J. Bu, S. Pan, Y. Yu, I. King, F. Karray, and P. S. Yu. 2025. Graphs Meet AI Agents: Taxonomy, Progress, and Future Opportunities. *CoRR* abs/2506.18019 (2025).

[4] M. Besta, N. Blach, A. Kubicek, R. Gerstenberger, M. Podstawski, L. Gianinazzi, J. Gajda, T. Lehmann, H. Niewiadomski, P. Nyczyk, and T. Hoefler. 2024. Graph of Thoughts: Solving Elaborate Problems with Large Language Models. In *AAAI.*

[5] S. S. Bhowmick, B. Choi, and A. Khan. 2026. From Human-Graph to Agent-Graph Interaction: State of the Art and Future Directions. *PVLDB* (2026).

[6] A. Bonifati, M. T. Ozsu, Y. Tian, H. Voigt, W. Yu, and W. Zhang. 2025. A Roadmap to Graph Analytics. *SIGMOD Rec.* 53, 4 (2025), 43–51.

[7] B. Chen, Z. Guo, Z. Yang, Y. Chen, J. Chen, Z. Liu, C. Shi, and C. Yang. 2026. PathRAG: Pruning Graph-based Retrieval Augmented Generation with Relational Paths. In *AAAI.*

[8] W. Chen, T. Bai, J. Su, J. Luan, W. Liu, and C. Shi. 2025. KG-Retriever: Efficient Knowledge Indexing for Retrieval-Augmented Large Language Models. In *International Conference on Knowledge Graph.*

[9] Y. Fu, D. Wang, W. Ying, X. Zhang, H. Liu, and J. Pei. 2025. Autonomous Data Agents: A New Opportunity for Smart Data. *CoRR* abs/2509.18710 (2025).

[10] M. Gallici, M. Martin, and I. Masmitja. 2023. TransfQMix: Transformers for Leveraging the Graph Structure of Multi-Agent Reinforcement Learning Problems. In *AAMAS.*

[11] A. Gusarov, A. Volkova, V. Khrulkov, A. Kuznetsov, E. Maslov, and I. V. Oseledets. 2025. Multi-Agent GraphRAG: A Text-to-Cypher Framework for Labeled Property Graphs. *CoRR* abs/2511.08274 (2025).

[12] B. J. Gutierrez, Y. Shu, Y. Gu, M. Yasunaga, and Y. Su. 2024. HippoRAG: Neurobiologically Inspired Long-Term Memory for Large Language Models. In *NeurIPS.*

[13] Y. Hu, R. Lei, X. Huang, Z. Wei, and Y. Liu. 2024. Scalable and Accurate Graph Reasoning with LLM-based Multi-Agents. *CoRR* abs/2410.05130 (2024).

[14] J. Jiang, K. Zhou, X. Zhao, Y. Song, C. Zhu, H. Zhu, and J.-R. Wen. 2025. KG-Agent: An Efficient Autonomous Agent Framework for Complex Reasoning over Knowledge Graph. In *ACL.*

[15] Z. Kaoudi and I. Giurgiu. 2025. Query Optimization Beyond Data Systems: The Case for Multi-Agent Systems. *CoRR* abs/2512.11001 (2025).

[16] A. Khan, T. Wu, and X. Chen. 2025. LLM+KG@VLDB 24 Workshop Summary. *SIGMOD Rec.* 54, 2 (2025), 60–65.

[17] E. Leurent and O.-A. Maillard. 2020. Monte-Carlo Graph Search: the Value of Merging Similar States. In *ACML (PMLR)*, Vol. 129. 577–592.

[18] Y. Li, H. Li, J. Zhang, R. Borovica-Gajic, S. Wang, T. Zhang, J. Chen, R. Shi, C. Li, and H. Chen. 2025. AgentTune: An Agent-Based Large Language Model Framework for Database Knob Tuning. *Proc. ACM Manag. Data* 3, 6 (2025), 1–29.

[19] X. Liu, Z. Peng, X. Yi, X. Xie, L. Xiang, Y. Liu, and D. Xu. 2024. ToolNet: Connecting LLMs with Massive Tools via Tool Graph. *CoRR* abs/2403.00839 (2024).

[20] C. Ma, Y. Chen, T. Wu, A. Khan, and H. Wang. 2025. Unifying Large Language Models and Knowledge Graphs for Question Answering: Recent Advances and Opportunities. In *EDBT.*

[21] Y. Ning and H. Liu. 2024. UrbanKGent: A Unified Large Language Model Agent Framework for Urban Knowledge Graph Construction. In *NeurIPS.*

[22] A. Perevalov and A. Both. 2025. Text-to-SPARQL Goes Beyond English: Multilingual Question Answering Over Knowledge Graphs through Human-Inspired Reasoning. *CoRR* abs/2507.16971 (2025).

[23] D. Qi, Z. Miao, and J. Wang. 2025. CleanAgent: Automating Data Standardization with LLM-based Agents. *VLDB 2025 Workshop: Data Driven AI* (2025).

[24] B. Shi and I. Panagiotas. 2025. GDS Agent: A Graph Algorithmic Reasoning Agent. *CoRR* abs/2508.20637 (2025).

[25] Z. Sun, J. Wang, X. Zhao, J. Wang, and G. Li. 2025. Data Agent: A Holistic Architecture for Orchestrating Data+AI Ecosystems. *IEEE Data Eng. Bull.* 49, 3 (2025), 79–95.

[26] Y. Wen, Z. Wang, and J. Sun. 2024. MindMap: Knowledge Graph Prompting Sparks Graph of Thoughts in Large Language Models. In *ACL.*

[27] X. Wu, Y. Shen, C. Shan, K. Song, S. Wang, B. Zhang, J. Feng, H. Cheng, W. Chen, Y. Xiong, and D. Li. 2024. Can Graph Learning Improve Planning in LLM-based Agents?. In *NeurIPS.*

[28] T. Wu D. Yan Y. Fang, A. Khan and S. Wang. 2026. LLM+Graph@VLDB'2025 Workshop Summary. *SIGMOD Rec.* 55, 2 (2026), 58–64.

[29] J. Yu, Y. Ding, and H. Sato. 2025. DynTaskMAS: A Dynamic Task Graph-driven Framework for Asynchronous and Parallel LLM-based Multi-Agent Systems. In *International Conference on Automated Planning and Scheduling.*

[30] Z. Yuan, M. Liu, H. Wang, and B. Qin. 2025. MA-GTS: A Multi-Agent Framework for Solving Complex Graph Problems in Real-World Applications. In *EMNLP.*

[31] M. Zhang, Y. Yang, R. Xie, B. Dhingra, S. Zhou, and J. Pei. 2026. Generalizability of Large Language Model-Based Agents: A Comprehensive Survey. *ACM Comput. Surv.* 58, 10, Article 263 (2026).

[32] X. Zhao, M. Blum, F. Gao, Y. Chen, B. Yang, L. M.-Carpintero, M. P.-Navarro, Y. Fu, S. Morikawa, Y. Iwasawa, Y. Matsuo, C. Park, and I. Li. 2025. AGENTiGraph: A Multi-Agent Knowledge Graph Framework for Interactive, Domain-Specific LLM Chatbots. In *CIKM.*