Preface of NORA 2025: First International Workshop on KNOwledge GRaphs & Agentic Systems Interplay

Btissam Er-Rahmadi[1], Sébastien Montella[1], Damien Graux[2] and Hajira Jabeen[3]

[1]Huawei Technologies Research and Development UK Ltd, Edinburgh, United Kingdom
[2]EcoVadis Ltd., London, United Kingdom
[3]Institute of Biomedical Informatics (BI-K), Uniklinik Köln, Germany

Abstract

Agents have experienced significant growth in recent years, largely due to the rapid technological advancements of Large Language Models (LLMs). Although these agents benefit from LLMs' advanced generation proficiency, they still suffer from catastrophic forgetting and a limited context window size compared to the agents' needs in terms of contextual information. Knowledge Graphs (KGs) are a powerful paradigm for structuring and managing connected pieces of information while unlocking deeper insights than traditional methods. Their value is immense for tasks that require context, integration, inter-linking, and reasoning. However, this power comes at the cost of significant upfront and ongoing investment in construction, curation, and specialised expertise. The NORA workshop aims at analysing and discussing emerging and novel practices, ongoing research efforts and validated or deployed innovative solutions that showcase the growing synergy between LLMs agents and KGs.

Webpage: https://nora-workshop.github.io/2025/

Keywords

Agentic AI, LLM, Knowledge Graphs, Workshop Series

Knowledge Graphs & Agentic Systems Interplay

The recent proliferation of large language models (LLMs) has opened the doors for new paradigms that benefit many applications like intelligent assistants, content creation & summarisation, code generation & debugging, and knowledge discovery, to name a few. Such applications are achieved through prompt engineering & in-context learning, retrieval augmented generation, fine-tuning & alignment, and function calling & tool usage. These families of techniques can be used on their own or combined for better results.

Thanks to the constantly improving reasoning and function-calling capabilities of LLMs, LLM-based agents have attracted more attention [14, 15, 16, 17, 18]. While performing their allocated tasks, these agents usually need to accumulate memory and feedbacks from tool calls and maintain a long run of these tasks. Consequently, they can easily exceed the context window size, explode costs, and degrade both latency and performance, due to their growing usage of tokens.

Depending on their tasks, agents usually need access to minimal portions of *semantic memory* (*i.e.* facts)[19, 20], *episodic memory* (*i.e.* events) [20, 21, 22, 19], and *procedural memory* (*i.e.* instructions) [23, 24]. However, it remains challenging for agents to select relevant examples from different memories, especially in large-scale applications (*e.g.*, personal memories for personal assistance).

Knowledge Graphs (KGs) [25, 26] model data and knowledge in a structured and explicit format known as *graphs*. Thanks to this native structure, they have demonstrated great capabilities in capturing

NORA'25: 1st Workshop on Knowledge Graphs & Agentic Systems Interplay co-located with NeurIPS, Dec.1, 2025, Mexico City, Mexico

btissam.errahmadi@gmail.com (B. Er-Rahmadi); sebastien.montella@huawei.com (S. Montella); dgraux@ecovadis.com (D. Graux); hajira.jabeen@uk-koeln.de (H. Jabeen)
https://scholar.google.com/citations?user=COfbBvAAAAAJ (B. Er-Rahmadi);
https://montellasebastien.github.io/index.html (S. Montella); https://dgraux.github.io/ (D. Graux);
https://hajirajabeen.github.io/ (H. Jabeen)
0000-0003-0526-661X (B. Er-Rahmadi); 0000-0002-2605-7238 (S. Montella); 0000-0003-3392-3162 (D. Graux);
0000-0003-1476-2121 (H. Jabeen)

rich semantics and connections between entities and concepts in both closed and open domains [25, 27]. This feature has enabled both 1) complex logical reasoning, which is needed for multi-hop queries [28] and deriving new implicit knowledge from explicit facts; and 2) graph-based learning through richer features of the structured data. However, curating knowledge can be challenging, especially from heterogeneous data sources and formats (*e.g.*, personal assistants). As a consequence, large-scale and industrial applications' scenarios are even more impacted by this bottleneck, which thereby lower the adoption of pure KG-based solutions in some Industrial use-cases.

Therefore, this **first** edition of the workshop aims to unveil the emerging yet growing interplay between two key paradigms of recent AI systems: Agents and Knowledge Graphs. On the one hand, the efficiency and performance of agentic systems can benefit greatly from KGs as a structured data model and reasoning foundation, especially in designing and implementing their various memories [19, 20]. On the other hand, KGs can leverage the advanced linguistic capabilities of LLM agents in extracting, computing and engineering knowledge from unstructured, multi-modal & multi-lingual data sources [29, 30, 31, 32].

Relevant Topics This is a non-exhaustive list of relevant topics to the workshop (alphabetically):

- Agentic and Knowledgeable Systems with Small Language Models
- Agentic Information Extraction and Retrieval
- Agentic KG Construction & Enrichment
- Agents for Complex Reasoning over KGs
- Agents and KGs for private and proactive personal assistants & Personalisation
- Augmenting Agents with External Knowledge
- Collaborative Agents for Knowledge Computing and Serving
- Context Engineering enhanced by KGs
- Efficient Reinforcement Learning for better performance
- Graph Retrieval Augmented Generation in Agentic systems
- KGs serving agents' memories: episodic (experiences, events, etc.), semantic (facts, concepts, etc.), and procedural (skills, tasks, etc.)
- Multi-Lingual & Multi-modal integrations
- On-Device or Hybrid (Device-Cloud) systems combining Agents and KGs
- Personalisation via Agents and KGs
- Personas and digital twins enabled by Agents and KGs
- Theoretical and experimental analysis of close and open Domain applications scenarios

NORA Scientific Program

Overall, the workshop spanned one full day and provided the audience with 3 keynotes and a panel discussion. The rest of the program was completed by oral presentations of accepted articles and poster sessions were organized during the breaks.

Invited speakers

The keynote presentations articulated a unifying perspective on how structured knowledge and explicit world models can address fundamental limitations of current large-scale learning systems.

Sebastián Ferrada [1] focused on the problem of grounding and long-horizon coherence in agentic AI, arguing that autonomy and adaptivity at scale require explicit representations of memory, semantics, and relational structure. He positioned knowledge graphs as a form of inductive bias that complements large language models by supporting persistent state, compositional reasoning, and controlled interaction with external knowledge. Through examples including GraphRAG, large-scale multimodal knowledge graphs, and socio-political simulation systems, the talk highlighted both the representational advantages

of graph-based memory and the systems-level challenges that arise when integrating symbolic and vector-based components, including scalable graph–vector architectures, evaluation methodologies for grounded agents, and mechanisms for agents to incrementally read from, write to, and revise structured knowledge.

Mustafa Jarrar's keynote [2] addressed the upstream problem of constructing reliable knowledge representations from unstructured data. Focusing on recent advances in named entity recognition and relation extraction, with particular emphasis on Arabic language resources, the talk surveyed state-of-the-art models, datasets, and extraction pipelines, and examined how their outputs can be composed into semantically coherent knowledge graphs. A central contribution was the introduction of an Information Extraction Ontology that provides a unifying semantic layer over heterogeneous extraction systems while maintaining alignment with widely adopted schemas such as schema.org and Wikidata. By embedding this ontology directly into large language model prompts, the approach enables more controlled, portable, and sample-efficient knowledge graph construction, and illustrates how symbolic constraints can be integrated into neural workflows without sacrificing flexibility.

Finally, Tiffany Callahan's keynote [3] examined the limitations of learning and retrieval-based approaches in scientific and high-stakes domains where critical data is rare, missing, or counterfactual in nature. She introduced the concept of agentic simulators, in which agents reason over knowledge graphs encoding domain constraints and causal structure, propose interventions, and invoke mechanistic and large quantitative models as causal simulation engines. The simulated outcomes, together with their causal assumptions and provenance, are written back into the knowledge graph, transforming it from a static repository into a computable causal world model. Drawing on examples from patient modeling and agentic chemistry, the talk argued that simulation-augmented, causally grounded reasoning over structured world models provides a principled extension beyond retrieval-augmented generation, enabling AI systems to reason about rare events, latent mechanisms, and the unmeasured world.

Accepted articles

Collectively, the accepted papers advance a unifying agenda: moving from static retrieval and pattern matching toward agentic, structured, and adaptive reasoning systems grounded in explicit knowledge representations, causal structure, and principled evaluation. Several contributions investigate foundational principles for reasoning over knowledge graphs (KGs) as world models. Graph Distance as Surprise [4] connects KG traversal to the Free Energy Principle, framing graph distance as a proxy for epistemic surprise and positioning KGs as generative models that guide agent behavior through structure-induced inductive bias. Complementing this theoretical perspective, Validation-Gated Hebbian Learning for Adaptive Agent Memory [5] introduces neuro-inspired plasticity mechanisms that allow agent memory graphs to evolve over time, while guarding against hallucination through validation-gated consolidation. At the representation-learning level, Elastic Weight Consolidation for Knowledge Graph Continual Learning [6] empirically studies catastrophic forgetting in KG embeddings, demonstrating that continual learning techniques can stabilize long-term knowledge acquisition while highlighting the sensitivity of results to task partitioning and evaluation protocol design.

A second cluster of papers focuses on agentic systems that construct, validate, and exploit structured knowledge in complex real-world domains. In the biomedical setting, Agentic Knowledge Computing for Automated Biomarker Validation [7] presents a large-scale, multi-model NLP pipeline that constructs weighted causal knowledge graphs from ALS literature, introduces a triangulated validation score to ensure reliability, and demonstrates counterfactual reasoning over the resulting graph. Biomedical Evidence Retrieval with Agentic RAG and Dual Text Encoders [8] similarly adopts an agentic perspective, using iterative query refinement and domain-specific encoders to improve evidence retrieval across heterogeneous biomedical corpora. Together, these works illustrate how agentic orchestration, multi-model fusion, and explicit causal structure can substantially reduce manual curation while preserving expert-level accuracy.

Several contributions demonstrate how knowledge-augmented and schema-constrained agents enable robust decision support in applied domains. AgentTravel [9] integrates domain-adapted language

models, structured itinerary memory, and real-time data retrieval to address the spatial and temporal constraints of urban travel planning, introducing a benchmark that jointly evaluates grounding and plan feasibility. RAPTOR: Reasoned Agentic Portfolio Trading [10] extends agentic design to financial decision-making, proposing a multi-agent, blackboard-based architecture coupled with Bayesian portfolio optimization, where structured debate and checkable financial indicators enable interpretable, risk-aware portfolio construction. In the legal domain, RuleSum [11] injects rulesets and KG-based representations into LLM summarization pipelines, using the IRAC method (Issue, Rule, Application, Conclusion) as a reasoning scaffold to improve factual fidelity, interpretability, and accessibility—demonstrating how symbolic structure can guide generation without sacrificing readability.

Finally, the workshop highlights the growing importance of evaluation and benchmarking for structured and agentic reasoning. ATLAS [12] introduces Harmonized Tariff Schedule (HTS) code classification as a challenging, hierarchically structured benchmark for global trade compliance, showing that fine-tuned LLMs can achieve significant gains in accuracy and cost efficiency while exposing the demands of legally grounded reasoning. Measuring What Matters [13] argues that standard accuracy metrics are insufficient for safety-critical reasoning tasks, proposing a transit-domain benchmark that probes consistency, robustness, and multi-step reasoning through perturbation and coherence-based evaluations. Together, these works underscore that progress in agentic AI requires not only richer world models, but also evaluation frameworks that can meaningfully assess grounding, causality, and long-horizon consistency.

Panel Discussion

The last session of the workshop kicked off preliminary discussions on the symbiosis between KGs and LLM Agents in the form of a panel session. This session involved both the invited keynote speakers [1, 2, 3] and the audience. The discussions pivoted around the trendiest architectures and application scenarios of combining KGs and agents, the practical methods to leverage agents in the construction of Knowledge graphs/bases domain-specific scenarios and modalities, and the potential of KG-based reasoning for improving agents' performance. The participants also discussed safety issues that might arise from delegating tasks to autonomous agents and the role of the ground truth knowledge stored in KGs in reducing such risks. The discussions closed up on the conclusion that the interplay between KGs and LLM Agents is an emergent topic that needs to be divided and explored deeply from different perspectives for at least the next five years.

Taken as a whole, the contributions reflect a shift toward AI systems that reason with, learn from, and act upon structured knowledge: treating knowledge graphs not merely as static retrieval indices, but as adaptive memories, causal world models, and coordination substrates for multi-agent intelligence. This convergence of symbolic structure, learning dynamics, and principled evaluation aligns closely with the workshop's goal of advancing grounded, reliable, and agentic AI beyond today's retrieval-centric paradigms.

Organisation

Organising Committee

- Btissam Er-Rahmadi *[Huawei Technologies R&D UK Ltd.]* is a Senior Scientist at the Knowledge Graph Lab. She has been developing innovative and novel solutions designed for Knowledge Computing and Serving in many application scenarios belonging to different domains (*e.g.*, personal assistants, e-commerce, geospatial search, etc.). She also has experience in applying operations research methods, mainly mathematical optimisation and simulations, to enhance the performance of distributed systems. She is interested in leveraging Deep Learning and NLP with knowledge engineering to construct advanced, seamless and practical approaches. Previously, she has organized the first N2Women Meeting at WiMob 2015 in UAE.

- Damien Graux *[EcoVadis]* leads a team of research scientists at EvoVadis that is specialised in AI/ML. He has been contributing to research efforts in Knowledge Computing technologies: focusing *inter alia* on Semantic Web, designing complex pipelines for heterogeneous Big Data and LLM-based knowledge management. Prior to this, he had research positions at Huawei R&D (UK), at Inria (France), Trinity College Dublin (Ireland) and Fraunhofer IAIS (Germany). He has been involved in the organisations of many international workshops at major conferences such as the LASCAR (co-located with ESWC) or the MEPDaW (co-located with ISWC) series, and more recently PromptEng (co-located with the ACM WebConf).
- Sébastien Montella *[Huawei Technologies R&D UK Ltd.]* is a Senior Research Scientist at the Knowledge Graph Lab. During his Ph.D., he specialized in Natural Language Generation and Knowledge Graph Embeddings research areas. Additionally, he has a keen interest in statistical learning, geometric deep learning, natural language processing, and computer vision. In the past, Sebastien has co-organized the 18th Workshop on Spoken Dialogue Systems for PhDs, PostDocs & New Researchers (YRRSDS) in Edinburgh, Scotland (2022), but also the PromptEng workshop series at TheWebConf (ex-WWW). Currently, he is one of the Program Chairs of the EMNLP 2025 Industry Track.
- Hajira Jabeen *[UniKlinik]* leads the 'AI in Research Data Management' team within the Institute for Biomedical Informatics. Her team focuses on leveraging artificial intelligence and LLMs to improve research data management practices, particularly in the biomedical field. The team works on developing scalable AI-driven tools and workflows that enhance data organization, integration, and analysis, developing innovative data-driven solutions. She has a background in both research and teaching, with previous affiliations at the University of Bonn, the University of Cologne, and ITU Copenhagen, and has organized multiple workshops and conferences in data science and informatics.

Program Committee (Alphabetical)

Aashka Trivedi (International Business Machines), Abdulaziz Alhamadani (Florida Polytechnic University), Abhay Dutt Paroha (SLB), Ahmed Abdelali (Humain), Aldrian Obaja Muis (Singapore Polytechnic), Alexandra Lavrentovich (University of Florida), Ali Pesaranghader (LG Electronics), Alsu Sagirova (DeepPavlov), Ana Kotarcic (University of Zurich), Andrei Kucharavy (University of Applied Sciences Western Switzerland, Sierre (HES-SO Valais)), Ankani Chattoraj (NVIDIA), Ankit Arun (Facebook), Ankur Padia (Philips Research North America), Anmol Goel (Technische Universität Darmstadt), Anup Kalia (Millenium), Arpit Sharma (Walmart Inc.), Arun LN (University of Pittsburgh), Ashish Shenoy (Meta), ASWARTH ABHILASH DARA (School of Computer Science, Carnegie Mellon University), Baban Gain (Indian Institute of Technology, Patna), Baohang Zhou (Tiangong University), Bin Dong (Ricoh Software Research Center Beijing Co., Ltd.), Bonaventura Coppola (SAP Security Research), Brian Riordan (Cisco), Brian Ulicny (RTX BBN Technologies), Carlos Bobed Lisbona (University of Zaragoza), Cheng Yu (Technische Universität München), Cheoneum Park (Hanbat National University), Chong Li (Institute of automation, Chinese Academy of Sciences), Chun-Nam Yu (Nokia Bell Labs), Daniel Bauer (Columbia University), Daniel Dickinson (American Family Insurance), Daniel Varab (German Research Center for AI), Daryna Dementieva (Technische Universität München), David Elson (Google / Google DeepMind), Debmalya Biswas (UBS Group AG), Deborah Dahl (LF AI & Data Foundation), Derrick Higgins (Illinois Institute of Technology), Deven Santosh Shah (Microsoft), Diane Napolitano (The Washington Post), Diman Ghazi (Google), Dong Zhou (Guangdong University of Foreign Studies), Duygu Altinok (Independent Researcher), Elio Querze (Bose Corporation), Elnaz Nouri (Research, Microsoft), Emir Munoz (Genesys Cloud Services Inc.), Emmanuel Ngue Um (University of Yaoundé 1), Ethan Selfridge (LivePerson), Fabio Casati (ServiceNow Inc), Fuxiang Chen (University of Leicester), Georgios Alexandridis (University of Athens), Gilbert Lim (EyRIS), Giorgos Stoilos (Huawei Technologies Ltd.), Giuliano Tortoreto (University of Trento), Guimin Hu (University of Copenhagen), Hai Wang (Amazon), Hanna Abi Akl (INRIA), Hemant Misra (Simpl), Hideya Mino (NHK), Hyun-Je Song (Chonbuk National University), Ian Stewart (Pacific Northwest National Laboratory), Issei Yoshida

(Hosei University), Jiahe Huang (University of California, San Diego), Jiangning Chen (Cisco), Jiaying Gong (eBay Inc.), Jinseok Nam (Amazon), Jinyeong Yim (University of Michigan), Jiyue Jiang (The Chinese University of Hong Kong), John Hudzina (Thomson Reuters), Juanyong Duan (Microsoft), Kaige Xie (Georgia Institute of Technology), Katya Artemova (Toloka AI), Keith Trnka (Independent Research), Kemal Kurniawan (University of Melbourne), Keyi Li (Rutgers University), Kushagr Arora (Bloomberg), Lawrence Moss (Indiana University at Bloomington), Lei ZHANG (Meta), Leslie Barrett (Bloomberg, LP), Liang Ma (Thomson Reuters), Lisheng Fu (Meta), Long Bai (Institute of Computing Technology, Chinese Academy of Sciences), Long Qin (Alibaba Group), Lorenzo Malandri (University of Milan - Bicocca), Lori Moon (MoonWorks, Inc.), Lucas Pavanelli (aiXplain), Maeda Hanafi (International Business Machines), Mahdi Zakizadeh (State University of New York at Stony Brook), Mahnoosh Mehrabani (Interactions Corp.), Manabu Torii (Kaiser Permanente), Manali Sharma (Samsung Semiconductor), Marcin Namysl (Ringler Informatik AG), Marek Kubis (Adam Mickiewicz University of Poznan), Marek Suppa (Comenius University in Bratislava), Mark Steedman (University of Edinburgh), Masaaki Tsuchida (Tokyo University of Science), Matīss Rikters (National Institute of Advanced Industrial Science and Technology (AIST)), Matthew Dunn (New York University), Matthew Mulholland (Lattice), Mihaela Bornea (IBM, International Business Machines), Minoru Sasaki (Ibaraki University), Mithun Balakrishna (Morgan Stanley), Mohammad Yeghaneh Abkenar (Universität Potsdam), Mohsen Mesgar (Bosch) Mounir Ghogho (University Mohammed VI Polytechnic), Mukul Singh (Microsoft), Munira Syed (The Procter & Gamble Company), Nadjet Bouayad-Agha (Independent Researcher), Naoki Otani (Megagon Labs), Natalia Loukachevitch (Lomonosov Moscow State University), Ningyu Zhang (Zhejiang University), Oleg Okun (Writer and translator), Pengyu Hong (Brandeis University), Pierre-Henri Paris (Université Paris-Saclay), Pradyot Prakash (Meta), Prajit Dhar (Universität Potsdam), Quentin Brabant (Orange-labs), Rafael Anchiêta (Federal Institute of Maranhão), Rahul Divekar (Bentley University), Rajasekar Krishnamurthy (Adobe Systems), Runze Wang (alibaba), Ryan Wang (University of Illinois Urbana-Champaign), Sachin Agarwal (Apple), Sachin Pawar (Tata Consultancy Services Limited, India), Sallam Abualhaija (University of Luxemburg), Sanjeev Kumar (Quark Inc), Sarasi Lalithsena (Wright State University), Sashank Santhanam (Apple), Shailza Jolly (Amazon Alexa AI), Shamil Chollampatt (Zoom Video Communications), Sherrie Shen (University of Edinburgh), Shihao Ran (University of Houston), Shubhashis Sengupta (Accenture), Sidharth Mudgal (Google), Simona Frenda (Heriot-Watt University), Sourav Dutta (Huawei Research Center), Souvik Das (J.P. Morgan Chase), Srideepika Jayaraman (IBM TJ Watson Research Center), Srijani Mukherjee (Texas A&M University - College Station), Stefano Pacifico (Jozef Stefan Institute), Sucheta Ghosh (Heidelberg University, Ruprecht-Karls-Universität Heidelberg), Sudarshan Rangarajan (International Business Machines), Tanay Kumar Saha (Purdue University), Tianhao Shen (Tianjin University), Tianlin Zhang (University of Chinese Academy of Sciences), Tianxing Wu (Southeast University), Tong Guo (Meituan), Tracy King (Adobe Systems), Traian Rebedea (NVIDIA), Vaishali Mishra (Expedia Group), Vera Pavlova (burevestnik.ai), Veronica Liesaputra (University of Otago), Vinod Goje (IEEE), Voula Giouli (Aristotle University of Thessaloniki), Wang Xu (Tsinghua University), Wei Hu (Nanjing University), Weixu Zhang (McGill University), Wenjie Zhou (Baidu), Wolfgang Maier (Mercedes Benz Research & Development), Won Ik Cho (Samsung Advanced Institute of Technology), Xiaolei Lu (University of California, San Diego), Xiliang Zhu (Dialpad Inc.), Xin Ying Qiu (Guangdong University of Foreign Studies), Xuan Zhu (University of California, Berkeley), Xuemei Tang (Hong Kong Polytechnic University), Xueting Pan (Oracle), Xu Jinan (Beijing Jiaotong University), Yekun Chai (Baidu), Ye Liu (Tencent AI Lab), Yifan Deng (University of the Chinese Academy of Sciences), Yifan Zhou (Shanghai Jiao Tong University), Yihao Fang (Huawei Technologies Ltd.), Yinghui Li (Tsinghua University, Tsinghua University), Yingya Li (Harvard University), Yin Zhang (Research, Google), Yixin Ji (Soochow University), Yonghao Liu (Jilin University), Young-Suk Lee (IBM, International Business Machines), Yuanliang Meng (Tsinghua University, Tsinghua University), Yubo Chen (Zhongguancun Laboratory), Yuhang Yao (Carnegie Mellon University), Yuqicheng Zhu (Universität Stuttgart), Yuwei Bao (Microsoft), Yuwei Yin (University of British Columbia), Yuxia Wu (Singapore Management University), Zac Yu (Duolingo), Zhengzhe Yang (Google), Zhixin Ma (Singapore Management University), Zhuoxuan Jiang (Shanghai Business School), Zihao Wang (TSY Capital).

Acknowledgments

We would like to thank the authors, reviewers, committee members and speakers for their contributions, support and commitment. Same also goes to the people attending the workshop who made the event successful through their fruitful discussions.

Declaration on Generative AI

The authors have not employed any Generative AI tools to prepare this Preface.

Articles presented at NORA 2025

[1] S. Ferrada, Memory, Meaning, and Machines: Building the Knowledge Scaffolds of Agentic AI, in: Proceedings of the 1st International Workshop on Knowledge Graphs & Agentic Systems Interplay (NORA) co-located with the NeurIPS, 2025. Invited Keynote.

[2] M. Jarrar, Advances in Information Extraction and Knowledge Graphs, in: Proceedings of the 1st International Workshop on Knowledge Graphs & Agentic Systems Interplay (NORA) co-located with the NeurIPS, 2025. Invited Keynote.

[3] T. Callahan, Simulators for the Unmeasured World, in: Proceedings of the 1st International Workshop on Knowledge Graphs & Agentic Systems Interplay (NORA) co-located with the NeurIPS, 2025. Invited Keynote.

[4] G. Jhajj, F. Lin, Graph Distance as Surprise: Free Energy Minimization in Knowledge Graph Reasoning, in: Proceedings of the 1st International Workshop on Knowledge Graphs & Agentic Systems Interplay (NORA) co-located with the NeurIPS, 2025.

[5] P. Singh, S. Yu, Validation-Gated Hebbian Learning for Adaptive Agent Memory, in: Proceedings of the 1st International Workshop on Knowledge Graphs & Agentic Systems Interplay (NORA) co-located with the NeurIPS, 2025.

[6] G. Jhajj, F. Lin, Elastic Weight Consolidation for Knowledge Graph Continual Learning: An Empirical Evaluation, in: Proceedings of the 1st International Workshop on Knowledge Graphs & Agentic Systems Interplay (NORA) co-located with the NeurIPS, 2025.

[7] K. Nidamarthi, K. Zhu, Agentic Knowledge Computing for Automated Biomarker Validation: Triangulated Causal Graph Construction in ALS Research, in: Proceedings of the 1st International Workshop on Knowledge Graphs & Agentic Systems Interplay (NORA) co-located with the NeurIPS, 2025.

[8] D. Goyal, E. Seibert, R. Ding, M. Migliarini, K. Zhu, Biomedical Evidence Retrieval with Agentic RAG and Dual Text Encoders, in: Proceedings of the 1st International Workshop on Knowledge Graphs & Agentic Systems Interplay (NORA) co-located with the NeurIPS, 2025.

[9] J. Zhao, J. Feng, Y. Li, AgentTravel: Knowledge-Augmented LLM Agent Framework for Urban Travel Planning, in: Proceedings of the 1st International Workshop on Knowledge Graphs & Agentic Systems Interplay (NORA) co-located with the NeurIPS, 2025.

[10] B. Almon, M. Caliboso, A. Kim, R. Dutta, R. Raman, M. Srungarapu, V. Sharma, K. Zhu, S. Dev, RAPTOR: Reasoned Agentic Portfolio Trading with Orchestrated Rebalancing, in: Proceedings of the 1st International Workshop on Knowledge Graphs & Agentic Systems Interplay (NORA) co-located with the NeurIPS, 2025.

[11] S. Chauhan, A. Chandrashekar, RuleSum: Injecting Rulesets into Knowledge Graphs for Accurate and Accessible Legal Summarization, in: Proceedings of the 1st International Workshop on Knowledge Graphs & Agentic Systems Interplay (NORA) co-located with the NeurIPS, 2025.

[12] P. Yuvraj, S. Devarakonda, ATLAS: Benchmarking and Adapting LLMs for Global Trade via Harmonized Tariff Code Classification, in: Proceedings of the 1st International Workshop on Knowledge Graphs & Agentic Systems Interplay (NORA) co-located with the NeurIPS, 2025.

[13] H. Veeramani, S. Thapa, U. Naseem, Measuring What Matters: Probing Transit Reasoning Consistency in Large Language Models, in: Proceedings of the 1st International Workshop on Knowledge Graphs & Agentic Systems Interplay (NORA) co-located with the NeurIPS, 2025.

References

[14] X. Li, S. Wang, S. Zeng, Y. Wu, Y. Yang, A survey on LLM-based multi-agent systems: workflow, infrastructure, and challenges, Vicinagearth 1 (2024) 9. URL: https://doi.org/10.1007/s44336-024-00009-2. doi:10.1007/s44336-024-00009-2.

[15] X. Huang, W. Liu, X. Chen, X. Wang, H. Wang, D. Lian, Y. Wang, R. Tang, E. Chen, Understanding the planning of LLM agents: A survey, 2024. URL: http://arxiv.org/abs/2402.02716. doi:10.48550/arXiv.2402.02716, arXiv:2402.02716 [cs].

[16] M. A. Ferrag, N. Tihanyi, M. Debbah, From LLM Reasoning to Autonomous AI Agents: A Comprehensive Review, 2025. URL: http://arxiv.org/abs/2504.19678. doi:10.48550/arXiv.2504.19678, arXiv:2504.19678 [cs].

[17] H. Jin, L. Huang, H. Cai, J. Yan, B. Li, H. Chen, From LLMs to LLM-based Agents for Software Engineering: A Survey of Current, Challenges and Future, 2025. URL: http://arxiv.org/abs/2408.02479. doi:10.48550/arXiv.2408.02479, arXiv:2408.02479 [cs].

[18] M. Mohammadi, Y. Li, J. Lo, W. Yip, Evaluation and Benchmarking of LLM Agents: A Survey, in: Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2, ACM, Toronto ON Canada, 2025, pp. 6129–6139. URL: https://dl.acm.org/doi/10.1145/3711896.3736570. doi:10.1145/3711896.3736570.

[19] T. Kim, M. Cochez, V. Francois-Lavet, M. Neerincx, P. Vossen, A Machine with Short-Term, Episodic, and Semantic Memory Systems, Proceedings of the AAAI Conference on Artificial Intelligence 37 (2023) 48–56. URL: https://ojs.aaai.org/index.php/AAAI/article/view/25075. doi:10.1609/aaai.v37i1.25075.

[20] T. Kim, V. François-Lavet, M. Cochez, Leveraging Knowledge Graph-Based Human-Like Memory Systems to Solve Partially Observable Markov Decision Processes, 2024. URL: http://arxiv.org/abs/2408.05861. doi:10.48550/arXiv.2408.05861, arXiv:2408.05861 [cs].

[21] P. Anokhin, N. Semenov, A. Sorokin, D. Evseev, A. Kravchenko, M. Burtsev, E. Burnaev, AriGraph: Learning Knowledge Graph World Models with Episodic Memory for LLM Agents, 2025. URL: http://arxiv.org/abs/2407.04363. doi:10.48550/arXiv.2407.04363, arXiv:2407.04363 [cs].

[22] P. Rasmussen, P. Paliychuk, T. Beauvais, J. Ryan, D. Chalef, Zep: A Temporal Knowledge Graph Architecture for Agent Memory, 2025. URL: http://arxiv.org/abs/2501.13956. doi:10.48550/arXiv.2501.13956, arXiv:2501.13956 [cs].

[23] H. Qian, Z. Liu, MetaAgent: Toward Self-Evolving Agent via Tool Meta-Learning, 2025. URL: http://arxiv.org/abs/2508.00271. doi:10.48550/arXiv.2508.00271, arXiv:2508.00271 [cs].

[24] K. Roth, R. Gupta, S. Halle, B. Liu, Pairing Analogy-Augmented Generation with Procedural Memory for Procedural Q&A, volume 15874, 2025, pp. 260–271. URL: http://arxiv.org/abs/2409.01344. doi:10.1007/978-981-96-8186-0_21, arXiv:2409.01344 [cs].

[25] A. Hogan, E. Blomqvist, M. Cochez, C. D'amato, G. D. Melo, C. Gutierrez, S. Kirrane, J. E. L. Gayo, R. Navigli, S. Neumaier, A.-C. N. Ngomo, A. Polleres, S. M. Rashid, A. Rula, L. Schmelzeisen, J. Sequeda, S. Staab, A. Zimmermann, Knowledge Graphs, ACM Comput. Surv. 54 (2022) 1–37. URL: https://dl.acm.org/doi/10.1145/3447772. doi:10.1145/3447772.

[26] J. Z. Pan, S. Razniewski, J.-C. Kalo, S. Singhania, J. Chen, S. Dietze, H. Jabeen, J. Omeliyanenko, W. Zhang, M. Lissandrini, R. Biswas, G. de Melo, A. Bonifati, E. Vakaj, M. Dragoni, D. Graux, Large Language Models and Knowledge Graphs: Opportunities and Challenges, Transactions on Graph Data and Knowledge (TGDK) 1 (2023) 2:1–2:38. URL: https://drops.dagstuhl.de/entities/document/10.4230/TGDK.1.1.2. doi:10.4230/TGDK.1.1.2.

[27] K. Liang, L. Meng, M. Liu, Y. Liu, W. Tu, S. Wang, S. Zhou, X. Liu, F. Sun, K. He, A Survey of Knowledge Graph Reasoning on Graph Types: Static, Dynamic, and Multi-Modal, IEEE

Transactions on Pattern Analysis and Machine Intelligence 46 (2024) 9456–9478. URL: https://ieeexplore.ieee.org/abstract/document/10577554. doi:10.1109/TPAMI.2024.3417451.

[28] Z. Shen, C. Diao, P. Vougiouklis, P. Merita, S. Piramanayagam, E. Chen, D. Graux, A. Melo, R. Lai, Z. Jiang, Z. Li, Y. Qi, Y. Ren, D. Tu, J. Z. Pan, GeAR: Graph-enhanced agent for retrieval-augmented generation, in: W. Che, J. Nabende, E. Shutova, M. T. Pilehvar (Eds.), Findings of the Association for Computational Linguistics: ACL 2025, Association for Computational Linguistics, Vienna, Austria, 2025, pp. 12049–12072. URL: https://aclanthology.org/2025.findings-acl.624/. doi:10.18653/v1/2025.findings-acl.624.

[29] L. Kwan, P. G. Omran, K. Taylor, Using Knowledge Graphs and Agentic LLMs for Factuality Text Assessment and Improvement, CEUR Workshop Proceedings 3828 (2024). URL: http://www.scopus.com/inward/record.url?scp=85210233104&partnerID=8YFLogxK.

[30] R. Zhao, S. Conia, E. Peng, M. Li, S. Potdar, AgREE: Agentic Reasoning for Knowledge Graph Completion on Emerging Entities, 2025. URL: http://arxiv.org/abs/2508.04118. doi:10.48550/arXiv.2508.04118, arXiv:2508.04118 [cs].

[31] A. Arun, F. Dimino, T. P. Agarwal, B. Sarmah, S. Pasquali, FinReflectKG: Agentic Construction and Evaluation of Financial Knowledge Graphs, 2025. URL: http://arxiv.org/abs/2508.17906. doi:10.48550/arXiv.2508.17906, arXiv:2508.17906 [q-fin].

[32] D. W. Jo, AKA: Agentic Self-Knowledge Augmentation Framework, in: M. Iklé, A. Kolonin, M. Bennett (Eds.), Artificial General Intelligence, Springer Nature Switzerland, Cham, 2026, pp. 304–313. doi:10.1007/978-3-032-00686-8_27.