Scilight

Transactions on Graph Intelligence and Network Applications
https://www.sciltp.com/journals/tgina

Editorial

From Graph Learning to Graph Intelligence

Jia Wu

School of Computing, Macquarie University, Sydney, NSW 2113, Australia; tgina.eic@gmail.com

Received: 4 September 2026; Accepted: 4 September 2026; Published: 8 September 2026

How To Cite: Wu, J. From Graph Learning to Graph Intelligence. *Transactions on Graph Intelligence and Network Applications* **2026**.

1. Introduction

The move toward what we call graph intelligence reflects the expanding role of graphs in machine learning. Traditional machine learning pipelines often model observations independently or as ordered sequences, focusing primarily on object attributes or local patterns. Graphs bring irregular and complex relationships among objects into the computation, allowing models to process relational structure directly [1,2].

Once relationships were represented explicitly, researchers began to explore how models could learn from the resulting structure. Early network embedding methods, such as DeepWalk [3] and node2vec [4], learned node representations from patterns of proximity captured through random walks. Graph neural networks took this further by aggregating information from neighboring nodes, thereby incorporating graph structure directly into end-to-end representation learning. GCN [5], GraphSAGE [6], and GAT [7] are representative examples. Together, these developments expanded the role of graphs beyond organizing relational data: models could use graph structure directly to learn representations and make predictions.

These advances changed what graph models could learn from a graph but, in most cases, not how that graph was constructed. Many widely used graph-learning architectures and benchmarks assume that the graph has already been constructed and is provided as input [8,9]. Graph models can learn from the supplied relations, but they typically do not seek new evidence during operation to determine whether relevant relations are missing or existing ones remain valid. Consequently, missing, outdated, uncertain, or erroneous relations may prevent the graph from reflecting the current state of the world and may affect downstream tasks. Therefore, improving a model's ability to learn from a given graph is not enough. The graph must also be checked against new information so that relations can be added, updated, or removed when necessary. Graph construction, verification, and correction can therefore no longer be confined to preprocessing; they must continue as new information arrives. Graph structure learning [10] can infer or refine graph connectivity from observed data, while dynamic, temporal, and continual graph learning address graphs that change over time [11]. These paradigms broaden how graph structure is learned or processed, but the setting considered here goes further: the system actively seeks external evidence and uses that evidence to determine whether and how specific relations should change. Such continual graph maintenance requires a system that can interpret newly arriving information, seek external evidence, and act on that evidence.

Large language models (LLMs) offer new ways to interpret unstructured information [12]. Recent studies have connected language models with graph structure [13,14] and used graphs to organize retrieval [15,16]. Moving from representation and retrieval to ongoing graph maintenance requires more than language understanding; it also requires the ability to obtain external evidence and act on that evidence to update the graph. An AI agent provides this bridge by combining language-model reasoning with access to external tools [17]. It can query external resources and use the resulting evidence to create, check, or modify relations while a task is under way [18]. Thus, graph construction and updating move beyond preprocessing and into ongoing operation, allowing the graph to change as new evidence becomes available.

2. Toward Graph Intelligence

This continual interaction changes the role of the graph itself. Rather than serving only as input to a learning model, the graph becomes a maintained relational state within an intelligent system. In this sense, graph learning primarily focuses on learning from supplied relations, whereas graph intelligence also concerns how those relations are maintained and used. In terms of graph maintenance, the agent maintains this relational state by incorporating

Copyright: © 2026 by the authors. This is an open access article under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).
Publisher's Note: Scilight stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

new evidence, validating existing relations, and revising the graph as the environment or task evolves. In terms of agent support, the maintained graph provides structured relational context that supports retrieval, consistency checking, reasoning, and subsequent actions. The resulting interaction is therefore bidirectional: the agent maintains the graph, while the graph, in turn, supports the agent. Graph intelligence thus extends beyond learning from a fixed relational structure to the continual maintenance and operational reuse of relational knowledge within an intelligent system.

For this cycle to be useful, relation updates should be reliable. An agent's ability to modify a graph does not guarantee that its updates are correct. The actions and evidence behind each update should be recorded so that its origin and downstream effects can be traced and audited [19]. Relations should carry appropriate temporal or validity information where relevant, and erroneous updates should be reversible. Clear update rules should specify the evidence required before a relation is added or changed, preventing unverified inferences from becoming persistent graph state. Traceable evidence, temporal validity, controlled updates, and error recovery are therefore essential to making this cycle dependable.

These requirements apply across domains, but their implementation must reflect what relations mean in each setting. Links between web pages, roads in transportation networks, and interactions in social networks represent fundamentally different types of relations, and errors in these relations can have different consequences. Accordingly, what is most likely to transfer across domains is not a single definition of entities and relations, but a set of general principles for representing uncertainty, checking evidence, controlling updates, and recovering from errors.

Evaluation must change accordingly and should consider two complementary dimensions. The first is the quality of graph maintenance: whether relations remain valid, updates are timely, evidence is traceable, and errors can be detected and repaired. The second is downstream utility: whether the maintained graph improves the agent's retrieval, reasoning, prediction, or task performance. The central question is not only whether an agent can use a graph to complete a task, but whether it can maintain the graph reliably during ongoing operation and use updated relations to support subsequent reasoning and action. Together, these questions define a broader interdisciplinary research agenda spanning graph learning, network science, knowledge representation, databases, and natural language processing. This agenda also aligns closely with the scope of *Transactions on Graph Intelligence and Network Applications*.

3. Conclusions

The move toward graph intelligence described here is not about replacing existing graph learning methods. It expands the research focus from using existing graphs for learning and prediction to understanding how graphs and agents work together through continuing interaction. Agents can seek and interpret new evidence to determine whether and how relations should be updated, while graphs provide structured support for subsequent retrieval, reasoning, and actions. For this bidirectional interaction to operate reliably, relations must be constructed carefully, updated when needed, linked to their evidence, and corrected when errors occur. Looking ahead, a central challenge is to develop intelligent systems in which agents can reliably maintain graphs under uncertainty, while updated graphs continue to support retrieval, reasoning, and action.

Conflicts of Interest

The author declares no conflict of interest.

Use of AI and AI-Assisted Technologies

During the preparation of this work, the author used Chatgpt to polish the language. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

1. Battaglia, P.W.; Hamrick, J.B.; Bapst, V.; et al. Relational Inductive Biases, Deep Learning, and Graph Networks. *arXiv* **2018**, arXiv:1806.01261.
2. Bronstein, M.M.; Bruna, J.; Cohen, T.; et al. Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges. *arXiv* **2021**, arXiv:2104.13478.
3. Perozzi, B.; Al-Rfou, R.; Skiena, S. DeepWalk: Online Learning of Social Representations. In Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, New York, NY, USA, 24–27 August 2014; pp. 701–710.

4. Grover, A.; Leskovec, J. node2vec: Scalable Feature Learning for Networks. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, 13–17 August 2016; pp. 855–864.
5. Kipf, T.N.; Welling, M. Semi-Supervised Classification with Graph Convolutional Networks. In Proceedings of the International Conference on Learning Representations, Toulon, France, 24–26 April 2017.
6. Hamilton, W.L.; Ying, R.; Leskovec, J. Inductive Representation Learning on Large Graphs. In Proceedings of the Advances in Neural Information Processing Systems, Long Beach, CA, USA, 4–9 December 2017; Volume 30, pp. 1024–1034.
7. Veličković, P.; Cucurull, G.; Casanova, A.; et al. Graph Attention Networks. In Proceedings of the International Conference on Learning Representations, Vancouver, BC, Canada, 30 April–3 May 2018.
8. Xu, K.; Hu, W.; Leskovec, J.; et al. How Powerful Are Graph Neural Networks? In Proceedings of the International Conference on Learning Representations, New Orleans, LA, USA, 6–9 May 2019.
9. Hu, W.; Fey, M.; Zitnik, M.; et al. Open Graph Benchmark: Datasets for Machine Learning on Graphs. In Proceedings of the Advances in Neural Information Processing Systems, Virtual, 6–12 December 2020; Volume 33, pp. 22118–22133.
10. Jin, W.; Ma, Y.; Liu, X.; et al. Graph structure learning for robust graph neural networks. In Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Virtual, 23–27 August 2020; pp. 66–74.
11. Kazemi, S.M.; Goel, R.; Jain, K.; et al. Representation learning for dynamic graphs: A survey. *J. Mach. Learn. Res.* **2020**, *21*, 1–73.
12. Bommasani, R.; Hudson, D.A.; Adeli, E.; et al. On the Opportunities and Risks of Foundation Models. *arXiv* **2021**, arXiv:2108.07258.
13. Tang, J.; Yang, Y.; Wei, W.; et al. GraphGPT: Graph Instruction Tuning for Large Language Models. In Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, Washington, DC, USA, 14–18 July 2024; pp. 491–500.
14. Chen, R.; Zhao, T.; Jaiswal, A.K.; et al. LLaGA: Large Language and Graph Assistant. In Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria, 21–27 July 2024; Volume 235, pp. 7809–7823.
15. He, X.; Tian, Y.; Sun, Y.; et al. G-Retriever: Retrieval-Augmented Generation for Textual Graph Understanding and Question Answering. In Proceedings of the Advances in Neural Information Processing Systems, Vancouver, BC, Canada, 10–15 December 2024; Volume 37, pp. 132876–132907.
16. Hu, Y.; Lei, Z.; Zhang, Z.; et al. GRAG: Graph Retrieval-Augmented Generation. In Proceedings of the Findings of the Association for Computational Linguistics: NAACL 2025, Albuquerque, NM, USA, 29 April–4 May 2025; pp. 4145–4157.
17. Yao, S.; Zhao, J.; Yu, D.; et al. ReAct: Synergizing Reasoning and Acting in Language Models. In Proceedings of the Eleventh International Conference on Learning Representations, Kigali, Rwanda, 1–5 May 2023.
18. Rezaei, M.R.; Fard, R.S.; Parker, J.L.; et al. Agentic Medical Knowledge Graphs Enhance Medical Question Answering: Bridging the Gap Between LLMs and Evolving Medical Knowledge. In Proceedings of the Findings of the Association for Computational Linguistics: EMNLP 2025, Suzhou, China, 4–9 November 2025; pp. 12682–12701.
19. Souza, R.; Gueroudji, A.; DeWitt, S.; et al. PROV-AGENT: Unified Provenance for Tracking AI Agent Interactions in Agentic Workflows. In Proceedings of the 2025 IEEE 21st International Conference on e-Science (e-Science), Chicago, IL, USA, 15–18 September 2025; pp. 467–473.