A Survey on Hallucination-Aware LLMs for Social Network Analysis Using Dynamic Knowledge Graphs

Riju Bhattacharya[1], Sanjay Kumar Singh[2]
[1]Lincoln University College, 47301, Petaling Jaya, Selangor Darul Ehsan, Malaysia
[2]Department of Information Technology, Amity University Lucknow Campus, Lucknow, UP, India
E-mail ID: pdf.riju@lincoln.edu.my, sksingh1@amity.edu

Abstract: The widespread tendency of Large Language Models (LLMs) to "hallucinate"—that is, to produce network structures, relationships, and graph metrics that are fluent and syntactically correct but structurally or factually incorrect—is a major barrier to their employment in Social Network Analysis (SNA). Fixed mitigation methods are ineffective in social networks due to the constant evolution of data caused by real-time interactions, the construction of edges, and community migrations. Dynamic Knowledge Graphs (DKGs) became popular as a feasible solution for anchoring large language models (LLMs) with perpetually updated structured knowledge to tackle these difficulties. DKGs enable semantic relationships among entities, facilitating context-dependent reasoning, temporal modification, and factual consistency during inferences. This survey provides a comprehensive review of Hallucination-Aware LLMs optimized for SNA via Dynamic Knowledge Graphs (DKGs). Systematically categorize state-of-the-art methodologies into pre-generation grounding, in-generation graph-text alignment, and post-generation structural verification. Furthermore, analyze the core structural challenges unique to social graphs, provide a comparative evaluation of existing approaches, highlight key architectural contributions, and map out open research directions for building self-correcting, temporally-aware network intelligence.

Keywords: Large Language Models, Hallucination Detection, Social Network Analysis, Dynamic Knowledge Graphs, Graph Neural Networks, Explainable AI, Trustworthy AI.

Introduction

People now interact, share content, and build communities online in entirely new ways, all because of the meteoric rise of social media websites. Text, user interactions, multimedia, and ever-changing relationship structures are just a few examples of the many types of data produced by social networking sites. The importance of analyzing these intricate networks is growing in various fields, including cybersecurity, public sentiment monitoring, recommendation systems, opinion mining, community finding, and disinformation identification.

Studies into social network analysis has benefited greatly from new, potent approaches made possible by developments in transformer-based design. A number of methods have been developed to address language model hallucinations; these include retrieval-augmented generation, reinforcement learning, graph neural networks, and temporal graph embeddings. These methodologies enable more adaptable thinking in dynamic social contexts by establishing models on confirmed facts and maintaining chronological accuracy. Notwithstanding these developments, the area still necessitates thorough examinations of scalability, processing costs, data privacy, explainability of hallucinations, and the real-time synchronization of knowledge graphs. In response, we assess in detail LLMs that are resistant to hallucinations and designed for use in social network analysis on dynamic graphs.

Using the combined reasoning capacity of LLMs and Dynamic Knowledge Graphs (DKGs), current research has attempted to construct systems with context understanding and hallucination resistance. Here, event-driven data updates from social interactions are paired with structural conceptual frameworks. Directly tackling the factual reliability difficulties typical in standalone language models, the resultant hybrid models provide greater data dependability and contextual adaption.

One of the challenges in developing reliable and adaptable social intelligence systems is reducing the occurrence of LLM hallucinations. Approaches addressing this challenge include retrieval-augmented frameworks, temporal embeddings, graph neural networks, and reinforcement learning. To successfully traverse the dynamic landscape of social data, these tools ensure that models remain grounded in reality and consistent throughout time. However, in order to improve these hybrid designs, future studies should focus on addressing computational costs, scalability limits, user privacy, explainable error mechanisms, and dynamic graph synchronization. This research meets these objectives by providing a comprehensive evaluation of hallucination-aware LLMs using dynamic social network analysis as a framework. This allows us to reflect on current events, identify significant gaps in our understanding, and propose practical answers.

Challenges in Hallucination-Aware LLMs for Social Network Analysis

Methods for detecting and reducing hallucinations have also advanced in tandem with LLMs, with to their fast development and widening range of potential uses. Therefore, in order to stay current, it is necessary to do extensive analyses of hallucination causes, detection methods, and mitigation strategies. To better understand LLM hallucinations and find ways to prevent or detect them, it is important to investigate these potential causes. The combination of LLMs with growing knowledge graphs offers a lot of unrealized promise for dynamic knowledge synchronization.

Efficiently integrating large-scale language models with dynamic knowledge graphs is another major challenge. Knowledge graphs fluctuate through time in response to changes in user interactions. Maintaining real-time graph topologies, temporal embeddings, and semantic links is computationally extremely challenging. Current systems often struggle to simultaneously achieve quick inference, scalability, and factual foundation.

Hallucination detection is naturally difficult because hallucinated material is typically compelling in context and has good language skills. Since hallucinations aren't always obviously incorrect, automating verification in social network analysis is challenging. Problems with multi-hop reasoning, entity ambiguity, sarcasm, and multilingual material only add to the difficulty of accurate hallucination identification.

Privacy and ethical concerns are two more major roadblocks. On social media platforms, data regarding users' activities, interactions, and private information is kept. Incorporating user-level knowledge into dynamic graph topologies can provide a number of issues, including data protection, consent management, and information misuse. Federalized learning, privacy-preserving graph learning, and other related methods have not been fully integrated with LLM systems that are cognizant of hallucinations.

Related work

Classical machine learning techniques, statistical modeling, and graph mining comprised the backbone of the initial social network analysis literature. These systems attempted to investigate user interactions, community detection, and information transmission by means of static graph representations and features that were manually constructed. Conversely, they had little idea how social networks worked or how their semantic context changed over time.

Both the ability to grasp natural language and to reason contextually were greatly enhanced with the advent of transformer-based big language models. Opinion analysis, social behavior prediction, recommendation systems, and rumor detection were some of the first areas where LLMs were used by researchers. However, problems with hallucinations quickly became apparent as a significant restriction impacting dependability and trustworthiness.

The goal of the latest studies on retrieval-augmented generation (RAG) methods is to lessen hallucinations by tying LLM results to external data sources. Knowledge graphs are becoming more popular due to the fact that they provide ordered semantic relationships between entities. Mixing graph neural networks with LLMs is what several academics have proposed for social network applications to improve contextual reasoning and factual consistency.

Dynamic knowledge graphs substantially increase these options by allowing for time-based updates and altering entity relationships. Methods like temporal graph embedding, adaptive graph reasoning based on reinforcement learning, and interactive processing of graphs have shown promising results in maintaining current contextual awareness. By using these methods, computers that are conscious of hallucinations can compare newly generated data with existing knowledge systems.

Additionally, frameworks for explainable hallucination detection have been the subject of further research. Some have suggested that we can increase transparency and user trust by making use of graph-based interpretability approaches, confidence estimations, causal reasoning, but also attention visualization. More studies looking into privacy-preserving social intelligence systems that don't rely on centralized data sharing have focused on federated learning methodologies. These enhancements are not enough because the present situation is still plagued with too many problems. Graph reasoning pipelines have a high processing cost, many frameworks use static knowledge stores, and they don't handle language or temporal adaptation well. Explainability, efficient social network analysis, dynamic graph reasoning, and hallucination detection are further areas where there is a lack of cohesive structures.

Table 1. Summary Table

Ref.	Objective	Methodology	Key Limitations	Improvements Expected in Future Frameworks
[1]	Hallucination reduction in LLMs	Retrieval-Augmented Generation	Limited temporal adaptation	Dynamic graph-based real-time grounding
[2]	Knowledge-enhanced reasoning	Knowledge Graph + Transformer	Static knowledge representation	Dynamic and evolving graph integration
[3]	Social misinformation detection	Graph Neural Networks	Limited explainability	Explainable graph reasoning
[4]	Context-aware social analytics	Temporal Graph Embedding	High computational overhead	Lightweight adaptive embeddings
[5]	Trustworthy content generation	Reinforcement Learning	Limited scalability	Distributed collaborative learning
[6]	Fake news verification	Dynamic Graph Learning	Poor multilingual support	Cross-lingual graph reasoning
[7]	Privacy-aware social intelligence	Federated Learning	Weak hallucination handling	Federated hallucination-aware systems
[8]	Explainable AI for SNA	Attention-based reasoning	Black-box decision behavior	Transparent causal reasoning
[9]	Real-time event monitoring	Streaming Knowledge Graphs	Synchronization challenges	Efficient graph updating mechanisms
[10]	Adaptive social recommendation	LLM + GNN integration	Resource-intensive architecture	Efficient hybrid optimization

Key Contributions

The main contributions of this work are outlined as follows:

- Concentrated on current breakthroughs in merging graph neural networks, explainable AI approaches, and dynamic knowledge graphs, this article gives a comprehensive assessment of hallucination-aware LLM frameworks for social network research.
- The study extensively examines significant research issues such as hallucination identification, factual consistency, time adaptability, scalability, explainability, privacy protection, and dynamic knowledge synchronization.
- We can enhance the reliability of future social intelligence system advancements by exploring new avenues of research in causal reasoning, linked graph learning, reinforcement learning, and dynamic graph embeddings.
- In order to create hallucination-aware LLM systems that can adjust to new social network environments and give accurate, flexible, and repeatable analysis, this survey is necessary.

Conclusions and Future Works

- Significant advancements in natural language comprehension, semantic reasoning, and intelligent data generation have been made possible by Large Language Models, which have revolutionized the field of social network research. But hallucination is still a major problem that makes LLM-driven social intelligence systems untrustworthy and unreliable. In ever-changing social contexts, misleading or false results can cause the spread of false information, the making of untrustworthy suggestions, and the making of analytically flawed judgments.
- One potential approach to establishing the foundation of LLM reasoning with dynamically developing structured knowledge is the use of dynamic knowledge graphs. Factual consistency and contextual awareness are greatly enhanced by DKGs through the integration of temporal updates, graph-based contextual reasoning, and semantic linkages. Further enhancement of hallucination-aware analytical capacities is achieved through the use of explainable AI approaches, retrieval augmentation, reinforcement learning, graph neural networks, and other similar technologies.
- Although significant strides have been made, numerous obstacles continue to await resolution. Scalability issues, computational complexity, lack of explainability, and inadequate real-time adaption mechanisms are common in current systems. Additionally, there is a need for additional research on privacy protection and multilingual social reasoning. Without comprehensive frameworks that include features like explainable inference, dynamic graph reasoning, and hallucination detection, the need for continuous research becomes even more apparent.

References

1. Lewis, P., Perez, E., Piktus, A., et al. (2020). Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. *NeurIPS*.
2. Ji, Z., Lee, N., Frieske, R., et al. (2023). Survey of Hallucination in Natural Language Generation. *ACM Computing Surveys*.

3. Pan, S., Hu, R., Long, G., et al. (2024). Graph Neural Networks for Social Network Analysis: A Review. *IEEE Transactions on Neural Networks and Learning Systems*.
4. Wang, X., Zhang, Y., Shi, C., et al. (2023). Dynamic Knowledge Graph Learning: A Survey. *IEEE Transactions on Knowledge and Data Engineering*.
5. OpenAI. (2025). Hallucination Reduction Techniques in Large Language Models. *Technical Report*.
6. Zhao, Z., Wang, J., Li, Q., et al. (2024). Dynamic Graph Learning for Fake News Detection in Social Media. *Information Sciences*.
7. Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated Machine Learning: Concept and Applications. *ACM Transactions on Intelligent Systems and Technology*.
8. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why Should I Trust You?” Explaining the Predictions of Any Classifier. *KDD*.
9. Trivedi, R., Dai, H., Wang, Y., & Song, L. (2019). Know-Evolve: Deep Temporal Reasoning for Dynamic Knowledge Graphs. *ICML*.
10. Wu, Z., Pan, S., Chen, F., et al. (2021). A Comprehensive Survey on Graph Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems*.