INNOVATIONS IN PUBLISHING, PRINTING AND MULTIMEDIA TECHNOLOGIES, 2026
https://doi.org/10.59476/ilpmt2026.v2i1.793

FROM SOCIAL INTERACTION GRAPHS TO KNOWLEDGE GRAPHS: ENABLING GRAPHRAG FOR LLM-BASED ANALYSIS OF ONLINE CONVERSATIONS

Arber Ceni, Alda Kika
University of Tirana, Albania

Abstract

Large language models (LLMs) offer great support for the analysis of social media data; however, they often struggle to provide grounded and reliable responses when reasoning over social network structures. This limitation mainly results from their dependency on unstructured text which can lead to hallucinations or incomplete interpretations of complex interaction patterns. Retrieval-augmented generation (RAG) has been shown to improve grounding by including external knowledge; however, the existing approaches mainly use collections of documents and fail to exploit the relational and temporal structure already present in social media data. The aim of this research is to propose a framework for transforming social interaction graphs into knowledge graphs that can serve as context graphs for GraphRAG, enabling more grounded and interpretable LLM-based reasoning over online conversations.

The proposed framework starts from a social interaction network constructed from X (formerly Twitter) data, where nodes represent users and edges represent interactions between users such as mentions, replies, reposts (retweets), and quotes. This interaction graph is then transformed into a heterogeneous knowledge graph which has posts (tweets) at its center and explicitly models authorship, conversational structure, and temporal dynamics. Furthermore, this schema contains entities such as posts(tweets), users, hashtags, URLs, and conversations, and defines relations including authored, mentions, replies to, reposts (retweets), and quotes. Additional semantic enrichment is incorporated through entity extraction and linking, as well as metadata such as timestamps and community membership. The resulting knowledge graph can be used as a context graph within a GraphRAG pipeline, where relevant subgraphs are retrieved and provided to an LLM for answering questions.

This research highlights the potential of integrating social network analysis with knowledge graph construction to enhance LLM-based analysis of social media data. By transforming interaction graphs into temporally enriched knowledge graphs, the proposed framework will enable structure-aware retrieval and reasoning, improving the quality and interpretability of generated responses.

Keywords: *GraphRAG, knowledge graph, large language models, social networks, X (formerly Twitter)*

Introduction

In recent years, large language models (LLMs) have experienced considerable advances in terms of architectural diversity, performance, usage and user adoption, leading to their integration in everyday tasks related to natural language processing and beyond (Brown et al., 2020). Social media platforms such as X (formerly Twitter), Facebook, Reddit produce large amounts of textual or visual content on which LLMs can reason and answer questions. However, users on these platforms do not only produce content, but also engage in different interactions between them including mentions, replies, reposts, quotes and shares. Such interactions form rich network structures which can be represented as social networks, where users are represented by nodes and the different interactions between them, form the edges of the network (Newman, 2003). These representations are often used in social network analysis but remain underutilized in LLM-based reasoning in contrast to unstructured textual content which continues to serve as the primary input of LLMs (Lewis et al., 2021).

Although they have strong language understanding capabilities, LLMs often struggle to provide reliable answers. When there is no clear path to the answer, LLMs often rely on assumptions, and this leads to incorrect or hallucinated responses (Brown et al., 2020). To alleviate hallucinations, retrieval-augmented generation (RAG) has been proposed. RAG works by incorporating external knowledge into the text generation process (Lewis et al., 2021). While RAG works well over unstructured text content, it fails to utilize the already present relational structure in social media interactions (Edge et al., 2025).

GraphRAG on the other hand, has demonstrated that modelling the knowledge as a graph, improves the grounding of LLMs by enabling multi-hop reasoning over nodes and edges of the graph (Lewis et al., 2021) (Peng et al., 2024). Current GraphRAG approaches are based on the textual data extracted from documents and the semantic relation between words to generate the knowledge graph (Peng et al., 2024). Social media data, on the other hand, naturally provide relational information which often is reduced into simple user-user interactions missing the underlying event context which can be used to reason upon (Ravi et al., 2024).

This paper explores the idea that social media interaction networks can be transformed into heterogenous knowledge graphs that serve as context graph for GraphRAG (Edge et al., 2025) (Peng et al., 2024). The proposed framework transforms a social network derived from X (formerly Twitter) posts, into a knowledge graph centered on tweets which models the conversational structure of the nodes in the original graph (users) and makes use of entities, hashtags and URLs. The constructed knowledge graph can support better reasoning by LLMs by preserving the structural and semantic dimensions of the original graph (Ravi et al., 2024).

Furthermore, a distinction between direct and indirect interactions such as mentions versus mentions through a retweet, reply-to or quote is made. This is done to maintain provenance and lower possible misleading interpretations. Temporal and community-level information also play an important role in capturing the dynamics of online conversations.

This work is positioned as a conceptual contribution, so our main focus is the design of the framework and the definition of evaluation criteria rather than empirical validation. A set of evaluation metrics for assessing grounding, interpretability, and reasoning capabilities in GraphRAG systems applied to social network data, and how these metrics can be used in future experimental studies is also discussed in this paper.

The contributions of this paper are threefold: (i) a framework for transforming social interaction graphs into knowledge graphs suitable for GraphRAG is introduced; (ii) the definition of a context graph representation that integrates structural, semantic, and temporal information; and (iii) the evaluation perspective for assessing LLM-based reasoning over social network data. This work contributes to the development of more reliable and explainable applications of LLMs in social computing.

Methodology

Social interaction graph construction

The initial dataset consists of several posts (tweets) related to COVID-19 vaccines downloaded from X (formerly Twitter) in a period between November 2021 and July 2022. For each term and for each month, a social graph was constructed based on the methodology described in (Ceni et al., 2023). Users who engaged in the conversation constitute the nodes in the graph and edges were built based on the different interactions those users had. These interactions include mentions, reply-to, retweets and quotes allowing the capture of relational structure of the network. In Figure 1 is shown a social interaction graph created by analyzing 5 sample posts in X (formerly Twitter).

This representation is very effective for traditional social network analysis (Newman, 2003), but it fails to deliver information about the event itself. In this way, conversational structure, temporal dynamics and semantic content, which are important aspects of a social network, are lost. To overcome these limitations, a transformation of the original graph into a knowledge graph which focuses on posts instead of users and rebuilds all the underlying interactions is proposed.

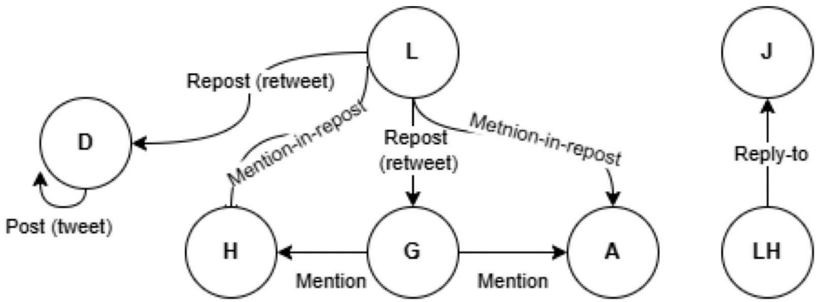

Fig. 1. Example of a social network created from X (formerly Twitter) data

Knowledge graph construction

The need to preserve structural and semantic dimensions, dictate the need to transform the social interaction graph into a knowledge graph (Peng et al., 2024). Edges which represent events around each individual post are now built from edges which represent user relationships. More specific, interactions between users are decomposed into one or more relations involving tweet nodes. If two users are connected by an edge in the original graph, those same two users are now connected through a set of intermediate edge and node entities that capture the content as well as the context of the interaction. This transformation represents the data more expressively because interactions are not anymore just edges between users, but structured events that can be used to reason upon.

The knowledge graph is now modelled as a heterogeneous graph consisting of different node and edge types.

The knowledge graph will contain these node types: (i) User – represents a user; (ii) Post – represents a post (tweet); (iii) Hashtag – represent a hashtag used in a post; (iv) URL and Domain – represent external references used in posts; (v) Entity – represents a real-world object extracted from post text; (vi) Conversation – grouping posts in a discussion thread; (vii) Community – represents a group of users identified through network structure

The knowledge graph will contain these edge types: (i) User → AUTHORED → Post; (ii) Post → MENTIONS → User; (iii) Post → REPLIES_TO → Post; (iv) Post → RETWEETS → Post; (v) Post → QUOTES → Post; (vi) Post → HAS_HASHTAG → Hashtag; (vii) Post → IN_CONVERSATION → Conversation; (viii) Post → HAS_MEDIA → Media; (ix) Post → MENTIONS_ENTITY → Entity; (x) Post → LINKS_TO → URL; (xi) URL → IN_DOMAIN → Domain; (xii) User → MEMBER_OF → Community

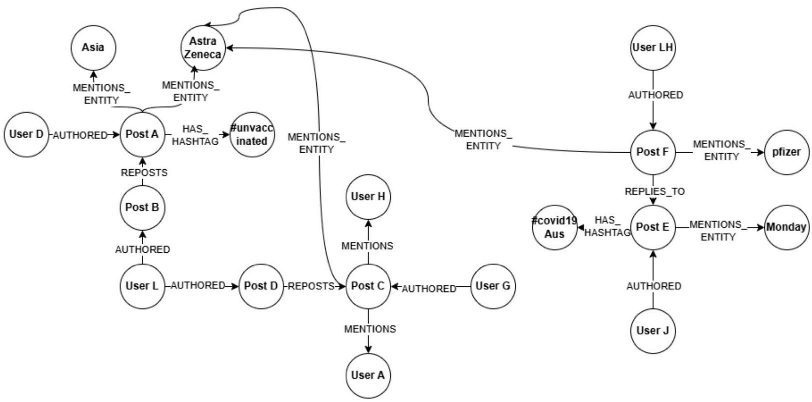

Fig. 2. The social network transformed into a knowledge graph following our proposed framework

Figure 2 shows the knowledge graph obtained by transforming the interaction graph of Figure 1 using the proposed framework.

This is a power framework to model different types of user interactions associated with their semantic. A particular aspect of this framework is the ability to distinguish between direct and indirect interactions. Direct interactions are considered interactions such as reply-to, mention, repost (retweet) and quote. Indirect interactions are considered mentions achieved through repost (retweet) or quote.

Instead of using a single relation type to represent all interactions and using attributes to distinguish them, the knowledge graph contains distinct relations (MENTIONS, REPLIES_TO, RETWEETS, QUOTES) thus preserving the semantic of the original interaction. Indirect interactions are stored in the structure of the graph itself and can be derived through multi-hop traversal thus always maintaining the origin of the interaction. The importance of this distinction is important because interaction through an intermediate step is not as important as a direct interaction. Users usually repost or quote without the primary intention of mentioning other users. We argue that by preserving this distinction, the graph will support better and more accurate reasoning.

Besides preserving the different users' interactions, the graph expressiveness is enriched by including additional semantic and contextual data generated from the post content and other metadata like: (i) Named entities – extracted from text using entity recognition techniques; (ii) Hashtags – representing different topics; (iii) URLS and domains – representing external references; (iv) Temporal information – timestamps when the post or the user profile was created; (v) Community information – derived from the graph's structure by using community detection algorithms.

The above elements allow the knowledge graph to not only capture interactions between users, but also *what* is being discussed, *when* interactions occur, and *how* the information travels across the network. So, this knowledge graph provides a richer context for downstream reasoning tasks (Ravi et al., 2024).

Context Graph for GraphRAG

One of the main differences between GraphRAG and classic RAG is that GraphRAG needs a context graph instead of context documents (Edge et al., 2025). The generated knowledge graph will serve to further generate context graphs for GraphRAG. When the user inputs a query, a subgraph is retrieved from the knowledge graph based on structural and semantic criteria which include user interactions, entities and the conversation context. The subgraph will serve as a structured context input for the language model, thus enabling it to generate responses based on relationships and evidence on the graph. By taking advantage of graph-based retrieval, this approach enables multi-hop reasoning on users, interactions and content and at the same time preserves interpretability given that nodes; and edges' connections are traceable.

Analysis

Analytical capabilities of the proposed framework

The proposed framework which transforms a social interaction graph into a knowledge graph enables us to perform richer analysis over social media content. The traditional user-user interaction network models relationships as aggregated edges while the transformed knowledge graph models interactions as event-level entities centered around posts. This way both structural relationships and their context is preserved.

As a result, more types of analytical queries are supported. For example, questions such as *"Which users interact most frequently"* can be answered using the structural aspect of the graph. At the same time, more complex queries like *"What topics are discussed between two communities over time?"* require the use of structural, semantic and temporal information, which are present in the entities of the knowledge graph. Answering these questions is not always possible and straightforward with the traditional graph representation.

Expected improvements in LLM-based reasoning

The use of the proposed knowledge graph as a context graph in the GraphRAG framework is expected to improve the quality of responses generated by LLMs in several ways. First, by limiting the model to operate on retrieved subgraphs, it will reduce the reliance of the LLM on assumptions, thus providing a better grounding (Lewis et al., 2021). Second, the model will be able to infer relationships that are not directly present in the graph or observable from pieces of text. This belief is based on the fact that the graph structure enables multi-hop reasoning across users, posts and other entities present in the graph (Edge et al., 2025). For example, the following query *"How are users A and B connected?"* may require traversing multiple nodes (conversation, mutual interaction, common hashtags/URLs). In a text-based RAG, such relationships may be difficult to reconstruct because they are not always present in the text entities.

Role of interaction semantics

In the proposed framework, the explicit modelling of interaction semantics like mentions, replies, reposts, and quotes, allow the knowledge graph to preserve the intent and context of user actions. This is an important aspect in social network analysis because different interaction types convey different meanings. Another important aspect of the design of the proposed framework is the separation between direct and indirect interactions. For example, a direct user mention may indicate intentional engagement, while a mention through a repost or quote may not reflect direct intentional communication. By keeping this distinction, the framework allows for more accurate analysis of user relationships by reducing the risk of misleading interpretations.

Temporal and community-level reasoning

User interactions in social media are dynamic, and they evolve over time and across communities, so the inclusion of temporal information and community structure is of the essence, and it allows the framework to capture these dynamics (Li et al., 2025). This way, questions about the evolution of the conversation can be asked, such as *"How has the interaction between two users changed over time?"* or *"Which communities are driving a particular discussion?"*. The addition of temporal information allows the model to distinguish early and late interactions, thus supporting causal and sequential reasoning. The presence of community-level information helps the identification of different interactions patterns within and across groups.

Comparison with traditional approaches

Traditional social network representation such as user-user interactions, lack the ability to capture content, context and temporal dynamics. Text-based approaches, on the other hand, focus more on content but ignore relational structure. Our proposed framework sits in between and bridges this gap by integrating structural, semantic and temporal dimensions. This way it supports both network-level analysis and content-aware reasoning, making it a good candidate for integration with LLM-based systems.

Evaluation perspective

For the purpose of evaluating the proposed framework, the focus is on four key dimensions: (i) grounding, (ii) reasoning capability, (iii) structural awareness and (iv) interpretability. The evaluation will be performed as a comparative experimental setup involving three approaches: (a) vanilla LLM – the model generates responses without access to external context, (b) text-based RAG – textual documents (posts/tweets) are retrieved and provided as context to the model (Lewis et al., 2021) and (c) GraphRAG – a query-specific subgraph is retrieved from the knowledge graph and provided as context to the model (Edge et al., 2025). The same dataset and query set will be used throughout the evaluation process.

Grounding – tries to measure how well the generated responses are supported by evidence in the retrieved subgraph and how well the model can avoid hallucinations (Peng et al., 2024). Two measures can be used to quantify grounding: (i) evidence-supported answer rate – the proportion of responses that can be attributed to nodes and edges in the retrieved subgraph, (ii) faithfulness score – the degree of consistency of the generated response compared to the supporting graph structure

Multi-hop reasoning – the ability of the model to identify indirect connections between users through shared conversations or common entities. Two metrics can be dsitinguished: (i) multi-hop accuracy – measures if responses to these queries are correct, (ii) path validity – determines if the identified relationships are valid paths in the graph.

Temporal reasoning – evaluates the dynamicity of the model by calculating the following measures: (i) temporal consistency – measures if responses follow the ordering and timing of interactions, (ii) temporal reasoning accuracy – evaluates if the model can answer time related questions regarding interactions (Li et al., 2025).

Structural awareness – tries to determine if the model correctly interprets the different interaction types present in the graph. The following metrics can be used to measure structural awareness: (i) relation correctness – does the models identify and use the correct interaction type? (ii) edge-type differentiation – is the model able to distinguish between direct and indirect interactions?

Interpretability – determines whether responses given by the model can be linked to the graph structure. (i) Traceability – if the response can be linked to specific entities on the graph.

Hallucination rate – can be regarded as the proportion of responses that are not supported by the retrieved subgraph. It provides an estimate of the model's reliability (Peng et al., 2024).

The following query categories are defined to evaluate the proposed framework:

- Relational queries – *"Which users interact more frequently?"*
- Multi-hop queries – *"How are users A and B connected?"*
- Semantic queries – *"What topics are discussed within a community?"*
- Temporal queries – *"How has the interaction between two users evolved over time?"*

Limitations and open challenges

The proposed approach shows great potential as a conceptual framework; however, it presents several challenges that need to be addressed when implementing it in future work. The construction of the knowledge graph relies on carefully extracting entities, interactions and metadata and often may be affected by noise and ambiguity. Second, the size and complexity of the graph may introduce challenges for retrieval and reasoning. It is known that the effectiveness of GraphRAG depends on the quality of the retrieved subgraph (Peng et al., 2024). Retrieving the optimal subgraph while maintaining completeness is not a trivial task. The evaluation of the reasoning over graph-structured data needs carefully designed benchmarks and metrics, aspects which are outlined in the following section.

Conclusions

Large language models have shown limitations when applied to social network data, especially in producing grounded and reliable responses. Although retrieval-augmented generation has improved the integration of external knowledge, it focuses mainly on text and does not take advantage of the inherent relation structure found in social media interactions.

In this paper, a framework for transforming social interactions graphs derived from social media into a heterogeneous knowledge graph which will then serve as a context graph within GraphRAG is proposed. A post-centered construction of the knowledge graph which explicitly models interaction semantics, temporal dynamics and community structure, thus preserving structural and semantic dimensions of social media data is further discussed. This will contribute to a richer analysis and more grounded and interpretable responses by LLMs.

Our primary contribution in this position paper is the conceptual design of the framework as well as the definition of an evaluation perspective focused on the graph-based reasoning over social media data. To this purpose, five evaluation dimensions, including grounding, multi-hop reasoning, temporal reasoning, structural awareness and interpretability are outlined.

The implementation of the proposed framework and the execution of the described experiments is the focus of our future work. This includes the creation of benchmark datasets, the design of retrieval strategies for constructing the context graph, and the comparison of GraphRAG approaches against traditional LLM and text-based RAG solutions. This line of research has the potential to significantly improve the reliability and explainability of LLM-based analysis in social computing.

List of references

1. Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., … Amodei, D. (2020). *Language Models are*

Few-Shot Learners (arXiv:2005.14165). arXiv. https://doi.org/10.48550/arXiv.2005.14165

2. Ceni, A., Kika, A., & Çollaku, D. X. (2023). *A Social Network Analysis of COVID-19 Vaccines Tweets.* 1–12.
3. Edge, D., Trinh, H., Cheng, N., Bradley, J., Chao, A., Mody, A., Truitt, S., Metropolitansky, D., Ness, R. O., & Larson, J. (2025). *From Local to Global: A Graph RAG Approach to Query-Focused Summarization* (arXiv:2404.16130). arXiv. https://doi.org/10.48550/arXiv.2404.16130
4. Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S., & Kiela, D. (2021). *Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks* (arXiv:2005.11401). arXiv. https://doi.org/10.48550/arXiv.2005.11401
5. Li, D., Niu, Y., Ai, Y., Zou, X., Qi, B., & Liu, J. (2025). T-GRAG: A Dynamic GraphRAG Framework for Resolving Temporal Conflicts and Redundancy in Knowledge Retrieval. *Proceedings of the 33rd ACM International Conference on Multimedia,* 11880–11889. https://doi.org/10.1145/3746027.3755628
6. Newman, M. E. J. (2003). The structure and function of complex networks. *SIAM Review*, 45(2), 167–256. https://doi.org/10.1137/S003614450342480
7. Peng, B., Zhu, Y., Liu, Y., Bo, X., Shi, H., Hong, C., Zhang, Y., & Tang, S. (2024). *Graph Retrieval-Augmented Generation: A Survey* (arXiv:2408.08921). arXiv. https://doi.org/10.48550/arXiv.2408.08921
8. Ravi, R., Ginde, G., & Rokne, J. (2024). *PRAGyan—Connecting the Dots in Tweets* (arXiv:2407.13909). arXiv. https://doi.org/10.48550/arXiv.2407.13909