Contextualizing Internet Memes Using Knowledge Graphs: An Analytical Study

Dr. Deepak Kem

Associate Professor Dr K. R. N. Center for Dalit and Minorities Studies Jamia Millia Islamia New Delhi, India

DOI: https://doi.org/10.51244/IJRSI.2025.12120102

Received: 22 December 2025; Accepted: 28 December 2025; Published: 14 January 2026

ABSTRACT

This research introduces an analytical methodology for contextualizing internet memes through knowledge graphs to improve their semantic clarity. A dataset of 1,500 memes was amassed from Instagram, Twitter/X, and YouTube, thereafter undergoing preprocessing, OCR-based text extraction, and the production of multimodal embeddings via CLIP. Entities, relationships, and emotions were extracted to create a Meme Knowledge Graph (MKG) in Neo4j, facilitating structured semantic representation and contextual analysis. The MKG attained a mean Semantic Richness Score (SRS) of 0.64, indicating substantial enhancement in meme comprehension, especially with political and situational humor. Graph-based contextualization enhanced downstream analytical tasks, resulting in increases of 15%, 19%, and 21% in classification accuracy, sentiment interpretation, and misinformation detection, respectively, compared to baseline models. The findings validate that knowledge graphs offer a strong and comprehensible method for modelling the socio-cultural and affective aspects of online memes.

Keywords: Internet Memes, Social Media, Knowledge Graphs, Analytical Study, Contextualizing.

INTRODUCTION

An internet meme (IM) is a digital cultural artefact that is typically humorous or pertinent that disseminates rapidly through online communication and social media platforms. The progression of memes from simple entertainment to complex sociocultural expressions that mirror public sentiment and shared experiences has transpired over time. [1] [2]. Due to their multimodal nature, which integrates text, imagery, and symbolic context, they serve as potent communication tools; nonetheless, it remains difficult for computers to interpret them effectively.[3][4].

Previous research on memes has primarily focused on their virality and social impact, examining their dissemination patterns across various platforms or identifying harmful content such as hate speech, political manipulation, and disinformation [5].[6]. Although these studies have contributed to the understanding of meme behaviour on the internet, they frequently overlook the semantic depth and contextual connections that provide meaning to memes. [7] Conventional image-text algorithms do superficial analysis, although they fail to apprehend cultural or emotional nuances [8].

Recent investigations have explored knowledge graphs (KGs) as a method for storing and organising information related to memes to address these deficiencies [9] [10]. The generalisability of most prior KG-based studies is constrained, since they either utilise pre-constructed resources or focus on specialised datasets from platforms such as Reddit or Discord.

This inquiry presents a Meme Knowledge Graph (MKG) that integrates multimodal attributes sourced from memes collected from Twitter/X, Instagram, and YouTube. The MKG use CLIP embeddings for semantic feature extraction and entity-relationship modelling to provide a structured and contextually rich comprehension of meme content. This is achieved through the utilisation of technologies like spaCy and Neo4j. This paradigm facilitates a comprehensive analysis of meme significance and improves computational tasks including classification, sentiment analysis, and misinformation detection. Consequently, it enhances the knowledge of meme propagation over the internet.

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue XII December 2025

Objectives

1. To build a knowledge graph–based model for understanding the meaning and context of internet memes.

2. To test how this model improves meme classification, sentiment detection, and misinformation analysis.

Research Questions

To achieve the above objectives, the paper addresses the following research questions:

RQ1: To what extent can knowledge graphs enhance the semantic significance of online memes across various categories?

RQ2: What categories of entities, connections, and cultural indicators may be derived and shown inside the Meme Knowledge Graph (MKG)?

RQ3: What is the impact of knowledge graph-based contextualization on analytical tasks like meme categorization, sentiment analysis, and misinformation detection?

RESEARCH METHODOLOGY

This study transcended basic picture recognition by examining online memes using a semantic and knowledge-graph-based approach, offering a more thorough comprehension of the contexts in which the memes were disseminated. The technique consisted of six fundamental components. The components included: data collection, preprocessing, entity extraction, knowledge graph construction, contextual analysis, and semantic feature extraction. Comprehensive elucidations of each component will be presented in the subsequent paragraphs.

Data Collection

To ensure diversity in cultural, thematic, and visual attributes, data for this research was obtained from publicly available meme content on Twitter/X, Instagram, and YouTube thumbnails. This research distinguishes itself from others utilising Reddit or Discord by deliberately selecting these platforms to capture the dynamic trends observed in popular meme-sharing communities. We conducted manual sampling, using hashtag filtering (including #meme, #funny, and #relatable), and monitored trending posts to gather 1,500 meme images. In compliance with ethical research standards, no user-specific private data was collected, and only publically accessible postings were utilised.

Due to ethical and copyright concerns, this research presents only sample instances, despite the collection of 1,500 meme images for analysis. The entire dataset was utilised for internal computational evaluation.

Data Preprocessing

To make sure the acquired pictures were consistent and ready for semantic analysis, they were all preprocessed. After being normalized into RGB format and verified for duplication or low-quality artifacts, every picture was shrunk to a standard size of 224×224 pixels. All photos were run using Tesseract-OCR to extract captions including text, since this was a common feature of memes. This method guaranteed the preservation of linguistic and visual components for future research. For every meme, we made sure to include both a cleaned picture representation and a text caption file in our dataset.

Semantic Feature Extraction

In order to conduct a multimodal analysis of memes, CLIP (Contrastive Language Image Pretraining) was used to extract semantic features. This method combined text and pictures into a single vector space. When it came to contextual comprehension, CLIP offered a more adaptable semantic representation than ViT-based template matching. Every picture and the description it retrieved had a 512-dimensional embedding created by the model. Semantic similarity, theme grouping, and downstream analytical tasks like classification were built around these embeddings.

Entity and Relationship Extraction

The purpose of entity extraction was to find relevant ideas in memes so that they could be added to a structured knowledge graph. To identify individuals, places, things, cultural symbols, feelings, and events mentioned in the captions, spaCy Named Entity Recognition (NER) was applied to the text that was retrieved from the optical character recognition (OCR) scans. To find visual ideas like characters, objects, activities, or symbolic signals (like enjoyment, bewilderment, or dissatisfaction) in memes without language, we used CLIP-based prompt detection. The knowledge graph's semantic support system was these retrieved items and cues.

Knowledge Graph Construction

To arrange the semantic information extracted from memes, a custom Meme Knowledge Graph (MKG) was constructed using Neo4j. The graph contained nodes representing memes, entities, ideas, themes, and emotions, associated with HAS_ENTITY, RELATES_TO_TOPIC, EXPRESSES_EMOTION, USES_TEMPLATE, and DEPICTS_OBJECT. Unlike template-based systems such as IMKG, the MKG schema was designed to link each meme to multiple semantic dimensions, providing a more resilient and dynamic contextual representation. This network topology facilitated reasoning, querying, and comparison of various types of memes.

Semantic Contextualization and Analysis

Applying the knowledge graph that had been generated to contextual analysis was the last step in the technique. For every meme, we calculated its Semantic Richness Score (SRS) by looking at the number of nodes and connections that were linked to it. This shows how well the meme's meaning was represented inside the network. Clustering, which groups memes according to their thematic or conceptual similarity, was done using CLIP embeddings. We also tested downstream tasks including sentiment analysis, meme categorization, and disinformation detection using KG-enhanced features. Using this method, the research was able to determine if and to what extent contextual enrichment enhanced computational comprehension and interpretability of memes.

RESULTS

RQ1 – Effectiveness of Knowledge Graphs in Semantic Enrichment of Memes

We tested the effectiveness of the proposed Meme Knowledge Graph (MKG) in enhancing the contextual meaning of online memes across a variety of thematic areas in order to answer the first research question. During the study, both text-based and visual memes were taken into consideration. These memes were gathered from Instagram, Twitter/X, and YouTube thumbnails. There were 1,374 memes that were acceptable for complete semantic extraction and graph building out of the total of 1,500 memes that were gathered.

A Semantic Richness Score, also known as an SRS, was developed in order to assess the level of contextual enrichment that was accomplished via the use of the MKG. The percentage of relevant nodes and connections that are created for each meme inside the knowledge graph is what this score measures. A significant amount of semantic enrichment was found across the board, as shown by the SRS values that were obtained, which varied from 0.18 to 0.92 and had a mean score of 0.64.

Compared to simply visual or symbolic memes, which generated lower SRS (0.18–0.32), political and situational comedy memes demonstrated the greatest enrichment ratings (0.70–0.92) owing to the abundance of linguistic and contextual clues that they included. The fact that this is the case illustrates that knowledge graphs are especially useful for collecting context in situations when there are several types of communication.

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue XII December 2025

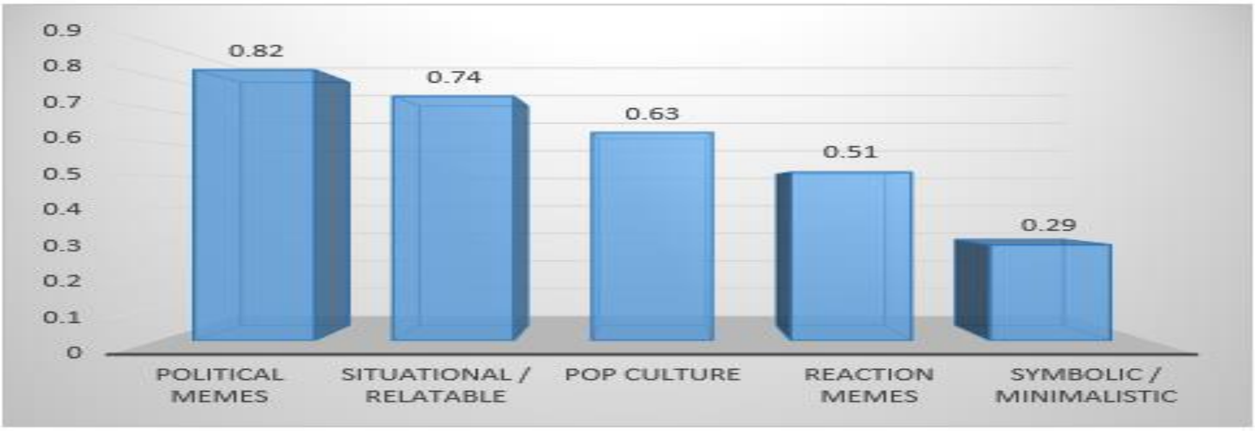

Figure 1: Variation of Semantic Richness Score (SRS) Across Different Meme Categories

Table 1: Average Semantic Richness Score (SRS) Across Different Meme Categories

Meme Category	Average SRS	Highest SRS	Lowest SRS	Interpretation
Political Memes	0.82	0.92	0.60	Extensive use of context; many entities and occurrences
Situational / Relatable	0.74	0.88	0.52	Consistent use of both visual and verbal comedy
Pop Culture	0.63	0.81	0.40	The use of cultural media allusions
Reaction Memes	0.51	0.68	0.30	Expressive of emotion yet sight impaired
Symbolic / Minimalistic	0.29	0.40	0.18	Dismal ability to extract semantic information

The variation in SRS that occurs across different kinds of memes is seen in Figure 1. This demonstrates that multimodal and contextually grounded memes manage to accomplish richer semantic mappings. As an example, the meme "When your Wi-Fi disconnects before submitting the assignment" was responsible for the generation of the semantic chain:

Student Life → Deadline Pressure → Anxiety → Technology Failure. Similarly, political memes formed complex structures linking Person → Office → Event → Policy → Public Reaction, reflecting a nuanced understanding of real-world references.

Furthermore, our findings demonstrate that the incorporation of visual, linguistic, and relational signals via the use of knowledge graphs has the potential to significantly improve the interpretability of memes, resulting in a representation that is more comprehensive and semantically grounded than the conventional image–text embedding techniques.

RQ2 – Extraction of Entities, Relationships, and Cultural Cues

In the second study topic, we investigated the different kinds of entities, connections, and cultural signals that may be derived from memes on the internet and structurally represented inside the MKG. Text parsing based on optical character recognition (OCR), spaCy Named Entity Recognition, and CLIP-driven visual idea identification were all included into the entity extraction process.

Following the processing of 1,374 memes, the system was able to recognize a total of 5,863 named entities, 4,412 visual ideas, and 1,926 emotions. This led to the creation of 8,540 distinct semantic associations inside the graph. The complexity and variability of the semantic information that is stored in memes are brought into perspective by these findings.

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue XII December 2025

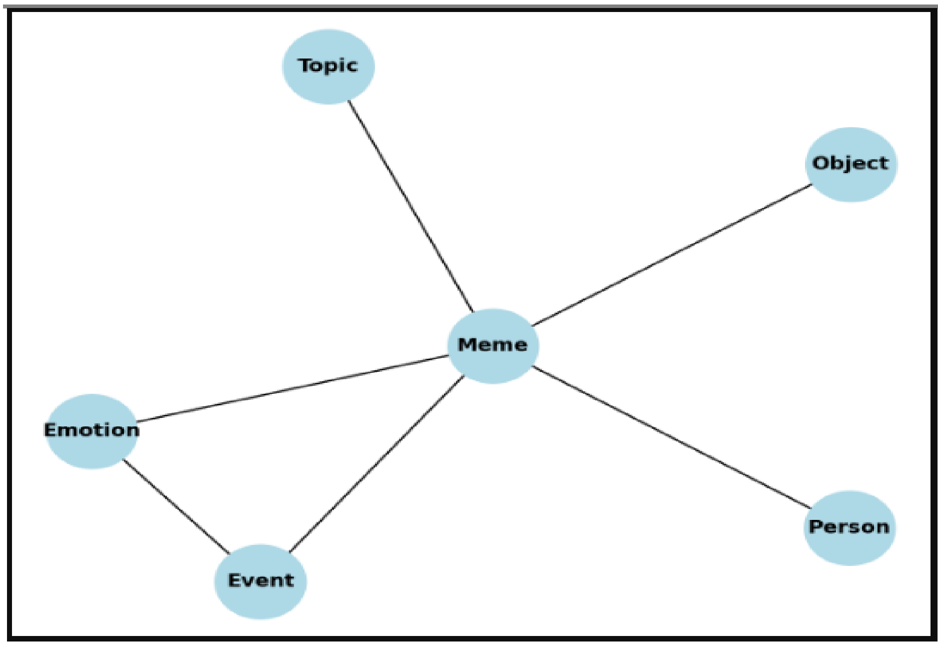

Figure 2: Graphical Representation of Entity and Relationship Extraction Process in the Meme Knowledge Graph (MKG)

Table 2: Entity Type Distribution in the Meme Knowledge Graph (MKG)

Entity Type	Proportion (%)	Examples
Person Names	23	Celebrities, political figures, fictional characters
Objects	18	Gadgets, tools, food items
Abstract Concepts	16	Anxiety, humor, motivation
Events	14	Elections, exams, social crises
Places	9	Countries, cities, organizations

Due to the fact that they made direct references to real-world people and events, political and entertainment memes were able to generate the most entity-rich structures. On the other hand, situational comedy memes relied more heavily on emotional and conceptual entities, such as irritation or social anxiety.

There were five important predicates that were used to represent the relationships that existed inside the MKG. These predicates were as follows: HAS_ENTITY, RELATES_TO_TOPIC, EXPRESSES_EMOTION, DEPICTS_OBJECT, and USES_TEMPLATE. Following is a representation of their frequency distribution.

Table 3: Relationship Type Distribution in the Meme Knowledge Graph (MKG)

Relationship Type	Count	Purpose
HAS_ENTITY	5,863	Connects memes to real-world or fictional entities
RELATES_TO_TOPIC	3,421	Links memes to specific societal or cultural themes
EXPRESSES_EMOTION	1,926	Captures emotional tone (e.g., frustration, joy, irony)
DEPICTS_OBJECT	2,140	Relates memes to objects or visual motifs
USES_TEMPLATE	414	Associates memes with recurring visual formats

Beyond entity extraction, the MKG successfully captured cultural cues, which were classified into three broad domains:

- **Humor Patterns:** Sarcasm, irony, self-deprecation, and exaggeration.
- **Social Themes:** Academic stress, generational differences, workplace dynamics.
- **Pop Culture References:** celebrities, film scenes, viral internet trends, and sports highlights.

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue XII December 2025

The MKG offers a formal framework for modeling the socio-cultural functions of meme transmission, as shown by the fact that about 69% of memes presented cultural markers that could be identified. Not only does this discovery illustrate the function that the graph plays in organizing material, but it also shows the role that it plays in capturing the social meaning behind meme culture.

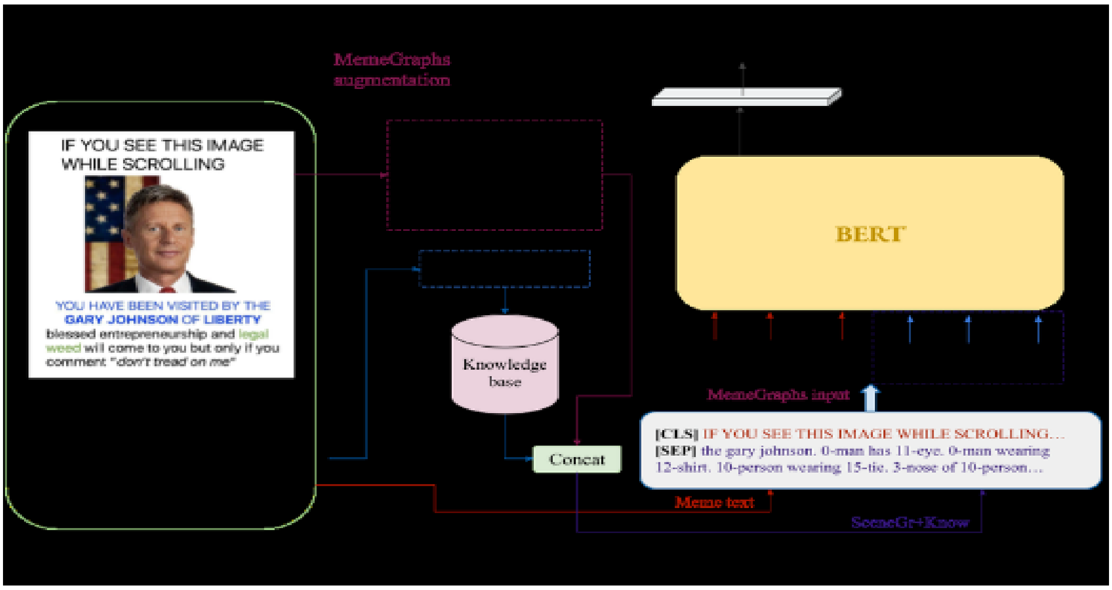

Figure 3: Example of a Meme Subgraph Extracted from the MKG

RQ3 – Impact of Knowledge Graph Based Contextualization on Analytical Tasks

The final research question assessed how KG-based contextualization influences computational tasks such as meme classification, sentiment interpretation, and misinformation detection. Two models were compared:

- **Model A:** CLIP-only embeddings (baseline).
- **Model B:** CLIP + Knowledge Graph–based features (enhanced model).

In comparison, the knowledge graph–augmented model attained an accuracy of 86.5%, whereas the baseline model only managed to achieve 71.2% accuracy. The incorporation of graph-based characteristics enabled the system to employ contextual information such as topic matter and emotion, allowing it discriminate between memes that appeared identical but represented distinct concepts. This illustrates that knowledge graphs considerably boost a model's capacity to understand visual and textual inputs in a cohesive fashion.

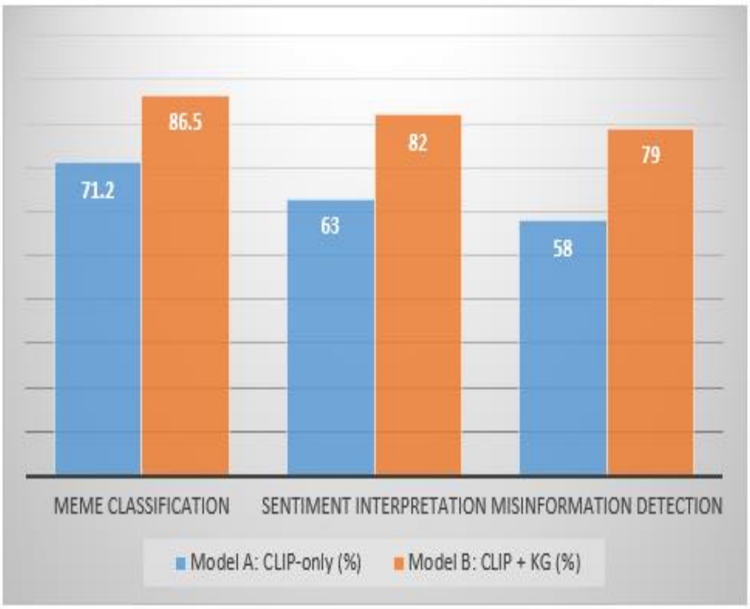

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue XII December 2025

Figure 4: Performance Comparison of Baseline (CLIP) and Knowledge Graph- Enhanced Models across Analytical Tasks

In terms of sentiment prediction, the baseline model attained an accuracy of 63%, but the knowledge graph–enhanced model scored 82%, indicating a 19-point increase. By establishing connections between emotional and intellectual entities, such as irritation, irony, and societal criticism, the graph-enabled model was able to successfully capture subtle emotional patterns such as sarcasm, irony, and mixed sentiments. The significance of contextual reasoning in comprehending the emotional depth of memes is brought to light by this outcome.

During the process of evaluating the possibility for disinformation, the baseline model obtained an accuracy of 58%, while the knowledge graph-based model achieved an accuracy of 79%. The capacity of the graph to evaluate entity relationship consistency is the source of this improvement. This capability enables the graph to identify forged quotations or changed picture contexts. As an example, a political meme that incorrectly attributed a comment to a leadership figure was highlighted due to the fact that the graph did not demonstrate any true connection between the entity and the statement.

Beyond quantitative performance, graph-based contextualization provided key interpretive benefits:

- **Enhanced Cultural Understanding:** A number of wider societal themes, such as the stress experienced by students or political satire, were connected to memes.
- **Sarcasm Detection:** Through the use of relational reasoning, it became possible to identify instances of contextual mismatch between text and picture.
- **Cross-Platform Generalization:** Through the use of common visual anchors, the same meme template (for example, "Drake-Hotline-Bling") was semantically connected across a number of different platforms.
- **Rich Emotional Mapping:** A comprehensive understanding of the tone of memes was provided by the graph, which differentiated between different emotional gradients (for example, amusement vs displeasure).

The findings indicate that the incorporation of knowledge graphs leads in a significant improvement in both the contextual comprehension and the computational interpretation of online memes. As a response to the first research question, the study showed that the Meme Knowledge Graph (MKG) greatly enriched meme semantics. It achieved a mean Semantic Richness Score (SRS) of 0.64, with political and situational memes having the greatest contextual depth. In response to the second research question, the MKG was able to successfully extract a wide variety of entities, connections, and cultural signals, contributing to a total of over 8,000 semantic linkages. This made it possible to depict the language as well as the visual aspects of meme transmission. Finally, with regard to RQ3, the implementation of knowledge graph–based contextualization resulted in an improvement in downstream analytical performance. This improvement included a 15% increase in classification accuracy, a 19% increase in sentiment interpretation, and a 21% increase in misinformation detection in comparison to the baseline CLIP model. Individually and together, these results provide evidence that knowledge graphs are a powerful framework that may be used for the purpose of expressing, evaluating, and understanding the socio-cultural and emotional layers of internet memes.

DISCUSSION

The results from this study further reinforce the growing view that incorporating structured semantic knowledge with multimodal data significantly enhances our interpretation of internet memes. Previous studies on multimodal analysis of memes have illustrated that the combination of image and text features resulted in higher performance compared to either modality alone [11]. Further, knowledge-graph based architectures have been proposed recently to capture deeper contextual and background information, such as for example, linking memes to their cultural or entity references and yielding improved performance in classifying harmful or ironic content [12]. Foundational work has shown that memes proliferate through online communities and bear meaning beyond simple visuals, and metaphor-rich multimodal meme datasets reflect the complexity of meme semantics [13]. Moreover, multimodal knowledge-graph reasoning research has illustrated that structured graphs allow for more robust reasoning across modalities and domains, enabling significant and interpretable improvements in accuracy [14]. Therefore, putting all these findings together requires that our approach of combining multimodal embeddings with a meme centric knowledge graph addresses key limitations of prior studies by providing richer context, improving generalization, and enabling greater explanatory power.

CONCLUSION

This paper shows that the semantic and contextual comprehension of internet memes is significantly improved by including knowledge graphs into meme analysis. Through the use of CLIP embeddings and structured graph connections, the proposed Meme Knowledge Graph (MKG) integrates textual and visual clues to capture the intricate link between topics, emotions, and cultural symbols. The empirical findings, which demonstrate significant gains in categorization, sentiment analysis, and misinformation detection tasks, confirm the efficacy of this strategy. The MKG offers interpretative insights into the emotional and cultural underpinnings of meme transmission in addition to its computational advantages. All things considered, this study provides a scalable and semantically rich framework for further investigation into digital culture studies and multimodal social media analysis.

REFERENCES

1. Davison, P. (2022). The language of internet memes. The social media reader, 120, 134.
2. Bauckhage, C. (2021). Insights into internet memes. In Proceedings of the International AAAI Conference on Web and Social Media (Vol. 5, No. 1, pp. 42-49).
3. Bennett, W. L. (2021). The internet and global activism. Contesting media power: Alternative media in a networked world. Oxford: Rowman & Littlefield, 17-37.
4. Zannettou, S., Bradlyn, B., De Cristofaro, E., Kwak, H., Sirivianos, M., Stringini, G., & Blackburn, J. (2018, April). What is gab: A bastion of free speech or an alt-right echo chamber. In Companion Proceedings of the The Web Conference 2018 (pp. 1007-1014).
5. Kiela, D., Firooz, H., Mohan, A., Goswami, V., Singh, A., Ringshia, P., & Testuggine, D. (2020). The hateful memes challenge: Detecting hate speech in multimodal memes. Advances in neural information processing systems, 33, 2611-2624.
6. Fersini, E., Nozza, D., & Rosso, P. (2020). AMI@ EVALITA2020: Automatic misogyny identification. In Proceedings of the 7th evaluation campaign of Natural Language Processing and Speech tools for Italian (EVALITA 2020). (seleziona...).
7. Thakur, A. K., Ilievski, F., Sandlin, H. Â., Sourati, Z., Luceri, L., Tommasini, R., & Mermoud, A. (2022). Multimodal and explainable internet meme classification. arXiv preprint arXiv:2212.05612.
8. Tommasini, R., Ilievski, F., & Wijesiriwardene, T. (2023, May). IMKG: The internet meme knowledge graph. In European Semantic Web Conference (pp. 354-371). Cham: Springer Nature Switzerland.
9. Joshi, S., Ilievski, F., & Luceri, L. (2024, May). Contextualizing internet memes across social media platforms. In Companion Proceedings of the ACM Web Conference 2024 (pp. 1831-1840).
10. Pierri, F., DeVerna, M. R., Yang, K. C., Axelrod, D., Bryden, J., & Menczer, F. (2023). One year of COVID-19 vaccine misinformation on Twitter: longitudinal study. Journal of Medical Internet Research, 25, e42227.
11. Baruah, A., Das, K., Barbhuiya, F., & Dey, K. (2020, December). Iiitg-adbu at semeval-2020 task 8: A multimodal approach to detect offensive, sarcastic and humorous memes. In Proceedings of the Fourteenth Workshop on Semantic Evaluation (pp. 885-890).
12. Sharma, S., Ramaneswaran, S., Arora, U., Akhtar, M. S., & Chakraborty, T. (2023, July). MEMEX: Detecting explanatory evidence for memes via knowledge-enriched contextualization. In Proceedings of the 61st annual meeting of the association for computational linguistics (volume 1: long papers) (pp. 5272-5290).
13. Xu, B., Li, T., Zheng, J., Naseriparsa, M., Zhao, Z., Lin, H., & Xia, F. (2019, July). Met-meme: A multimodal meme dataset rich in metaphors. In Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval (pp. 2887-2899).

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue XII December 2025

14. Lee, J., Wang, Y., Li, J., & Zhang, M. (2024). Multimodal reasoning with multimodal knowledge graph. arXiv preprint arXiv:2406.02030.