Unveiling Misinformation Spread with Graph Neural Networks: a Model-Driven Explainable Approach

Liliana Martirano[1 [0000−0001−6231−9373]] and Carmela Comito[1 [0000−0001−9116−4323]]

Institute of High Performance Computing and Networking (ICAR-CNR), Via P. Bucci, 87036 Rende (CS), Italy
{liliana.martirano,carmela.comito}@icar.cnr.it

Abstract. The widespread dissemination of misinformation on social media demands advanced detection strategies beyond traditional content-based approaches. This study introduces a model-driven, explainable framework for multimodal fake news detection, leveraging Graph Neural Networks (GNNs) to jointly capture textual, relational, and social context. The proposed method classifies news articles as real or fake while simultaneously identifying key misinformation spreaders within the network. Using GNN-Explainer, we enhance model transparency by identifying the most influential nodes, edges, and features that drive classification decisions. Experimental results on real-world datasets show that combining structural and content-based signals improves detection accuracy. Further analysis reveals that misinformation spreaders consistently rank high in centrality and act as amplifiers of false narratives; it also highlights the need for dataset-specific interventions — disrupting coordinated groups in dense, polarized networks and targeting key spreaders in sparse, less polarized ones.

Keywords: Graph Neural Networks · GNN-Explainer · Misinformation Spreading · Influential Nodes.

1 Introduction

Misinformation in online environments contributes to serious societal issues such as vaccine hesitancy, election interference, and social polarization. Its rapid spread is facilitated by the structure of social networks, where central actors, such as influencers, amplify both real and false information [5]. Despite the growing recognition of this issue, research gaps persist in understanding how to strategically disrupt the flow of misinformation and mitigate its impact; moreover, evaluation of targeted interventions against key nodes remains limited [1].

This study adopts a network-centric approach to analyze the structural and behavioral dynamics of misinformation dissemination. We investigate three key questions: (i) the role of central nodes in amplifying or mitigating fake news; (ii) the social interactions most impactful in spreading misinformation; and (iii) the potential of targeting key nodes to curb its spread. These inquiries aim to inform more effective and actionable mitigation strategies.

Our proposed framework leverages Graph Neural Networks (GNNs), specifically Graph Attention Networks (GATs), to integrate multimodal information — textual content, social interactions, and relational structures — into the fake news detection process. Unlike traditional GNN-based approaches that typically model homogeneous networks, our method explicitly accounts for heterogeneous networks, comprising multiple types of nodes (e.g., news, tweets, users, hashtags), edges (e.g., retweets, replies, friendship) and *meta-paths* (e.g., co-occurring hashtags, tweets discussing the same news - cf. Section 3). Modeling this heterogeneity allows for a more realistic and fine-grained representation of misinformation dynamics in complex online environments. The GAT architecture enables the model to weigh the relative importance of neighboring nodes, allowing for nuanced representation learning that captures both local and global patterns. To address the interpretability challenge commonly associated with GNNs, we incorporate GNN-Explainer [12], which provides fine-grained insights into model decisions by highlighting the most influential features and relationships that drive classification outcomes. To further analyze misinformation spread, we incorporate network centrality measures to pinpoint the most influential users in the network. By combining these components, our framework not only improves classification accuracy but also enhances transparency and supports actionable analysis of misinformation dynamics in complex, heterogeneous social systems.

The structure of the paper is as follows. Section 2 discusses major related studies. Section 3 introduces the proposed fake news detection approach. Section 4 describes the methods used to understand the role of nodes and edges in fake news propagation. Section 5 illustrates the experimental investigation. Finally, Section 6 concludes the work and provides pointers for future research.

2 Related Work

Deep learning approaches address the limitations of traditional methods, based solely on centrality metrics, by handling high-dimensional data and complex network structures, but they are mainly suited for homogeneous graphs, failing in processing multi-type information. Methods designed for heterogeneous networks primarily focus on meta-path influence while overlooking variations in individual node importance within meta-paths. Below we discuss some of the leading work based on machine/deep learning techniques for complex networks.

Yu et al. [13] transformed node influence detection into a regression task with Convolutional Neural Networks (CNNs). Kou et al. [7] proposed a multi-head attentional regression model to refine influence aggregation based on Graph Convolutional Networks (GCNs). Zhao et al. [15] combined structural and node characteristics in a GCN-based classification model, while Keikha et al. [6] used DeepWalk embeddings to filter influential nodes. Ahmad et al. [2] proposed a data-driven approach for significant node identification, further advancing centralized feature selection methods. The approach most similar to ours is MEGA

[11], which identifies influential author nodes in academic networks by leveraging meta-paths, subgraph construction via random walks, and GAT-based influence aggregation. In contrast, our method integrates centrality measures and GNN-Explainer to refine node ranking based on learned representations, rather than relying solely on predefined meta-paths. Additionally, it offers interpretability by revealing which specific nodes are influential, supporting more informed intervention strategies. Finally, MEGA is domain-specific and cannot be directly adapted to broader real-world scenarios like misinformation detection.

3 The Multi-modal Fake News Detection Framework

Our approach leverages two key modalities, textual content and social network structure, to classify fake news effectively. We exploit an early fusion approach of modalities as proposed in [8], employing a GAT-based architecture enhanced by meta-paths that operates on the heterogeneous graph at node-level, with additional modalities encoded as initial node features. The model architecture includes an embedding layer to encode the initial features of each node, followed by attention layers to refine the embeddings based on the influence of neighboring nodes. The final output layer classifies nodes by their likelihood of being fake news, relying thus on both the network topology and the other modalities.

For textual data, we exploit both news and related tweets using BERT-like architectures. Specifically, we utilize TwHIN-BERT [14] for tweet encoding, a multilingual model trained on Twitter data, and a sentence-transformer trained in a Siamese architecture for news embeddings [4]. The social network structure is modeled as a heterogeneous graph and processed in a type-aware manner, with distinct attention mechanisms applied to different edge types. We employ a GATv2 architecture [3] to capture the relational dependencies between nodes while dynamically adjusting the importance of neighboring connections. Formally, given the heterogeneous graph $G = \langle \mathcal{V}, \mathcal{E}, A, R, \phi, \varphi, \rangle$, where $\mathcal{V}$ and $\mathcal{E}$ are the sets of nodes and edges, A and R are the sets of node and relation types, with $|A| + |R| > 2$, $\phi : \mathcal{V} \to A$ and $\varphi : \mathcal{E} \to R$ are the node- and edge-type mapping functions, resp., the updated node representation $\mathbf{h}_i^{(l+1)}$ for each node $i \in \mathcal{V}$ at each layer l is computed by aggregating over all relation types:

$$\mathbf{h}_i^{(l+1)} = \sigma \left(\sum_{r \in R} \sum_{j \in \mathcal{N}_r(i)} \alpha_{ij,r}^{(l)} \mathbf{W}_r^{(l)} \mathbf{h}_j^{(l)} \right)$$

where σ is the ELU activation, $\mathcal{N}_r(i)$ denotes the neighborhood of node i under relation r, i.e., all nodes connected to i via an edge of type r; $\mathbf{W}_r^{(l)}$ is the learnable weight matrix for layer l specific to relation type r; $\alpha_{ij,r}^{(l)}$ is the relation-specific attention coefficient between i and j at layer l, calculated as follows: $\alpha_{ij,r}^{(l)} = \text{softmax}_{j \in \mathcal{N}_r(i)} \left(\text{LeakyReLU} \left(\mathbf{a}_r^{(l)T} \left[\mathbf{W}^{(l)} \mathbf{h}_i^{(l)} \odot \mathbf{W}^{(l)} \mathbf{h}_j^{(l)} \right] \right) \right)$, where $\mathbf{a}_r^{(l)}$ is the relation-specific learnable attention vector, and $\odot$ denotes element-wise multiplication. $h_i^{(0)}$ corresponds to the initial feature vector.

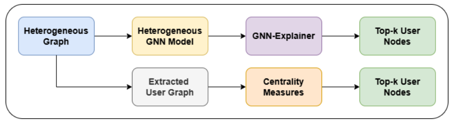

Fig. 1: Logical schema to detect top-k user nodes in a given heterogeneous graph.

To capture more nuanced patterns and indirect relationships, we extended the GAT architecture to include *meta-paths* as additional relations. A *meta-path type* or simply *meta-path* in a heterogeneous network is a composite relation modeling high-order proximity induced by a path $a_1 \xrightarrow{r_1} a_2 \xrightarrow{r_2} \ldots \xrightarrow{r_x} a_{x+1}$ between two node types, $a_1 \in A$ and $a_{x+1} \in A$, which are expected to share information – see Tables 1 and 2 for concrete meta-path examples. A *meta-path instance* is a sequence of connected nodes matching the node and edge types in the meta-path, able to make two distant nodes in the network reachable.

Please note that, as well as to tell fake and real news apart, the learned embeddings can be used for multiple downstream tasks, at node, edge and (sub)graph level, such as identifying critical nodes, predict the user interactions or detect coordinated misinformation campaigns.

4 Unraveling key actors in Fake News Classification

Identifying key nodes and relationships, such as influencers and spreading patterns, is crucial for understanding their role in either amplifying or suppressing misinformation within social media. To develop effective intervention strategies, we combine traditional graph theory measures with explainability of the GNN, leveraging its output in the multimodal fake news detection task (see Fig. 1).

4.1 Centrality-Based Approach

Centrality measures are traditionally employed to identify key nodes within the information diffusion process. We leverage the following metrics for user nodes:

- *Degree Centrality*. Reflects the number of direct connections a user has, capturing their immediate influence;
- *Betweenness Centrality*. Measures how often a user acts as a bridge between other nodes, highlighting their role in spreading information by facilitating communication between otherwise disconnected groups;
- *Closeness Centrality*. Indicates how quickly a user can reach others in the network, capturing their ability to spread information across the network;
- *Pagerank*. Quantifies a user's overall importance within the graph, based on the connectivity with other influential nodes;
- *Voterank*. Identifies influential nodes in a network based on an iterative voting process by selecting influential nodes based on their ability to influence others while avoiding redundancy.

Each metric captures different aspects of user influence and network dynamics. However, these methods may overlook contextual factors in misinformation spread and can be computationally intensive, especially in large or dense networks. To overcome these limitations, our framework incorporates model-driven analysis using the trained GAT and advanced explainability techniques.

4.2 GNN Explainer Approach

GNN-Explainer [12] is a post-hoc method designed to uncover the most important structural and feature-related factors influencing GNN's predictions. We thus identify which nodes have the highest contribution to classification predictions, revealing whether structural centrality aligns with model-relevant nodes. We compute a node-level score according to two complementary strategies:

- *feature-based score* (for short, *Expl. Feats.*), computed by summing along the feature dimension, i.e., by aggregating the contributions of all its features;
- *relation-based score* (for short, *Expl. Rel.*), computed by summing along the incident edges, i.e., by aggregating the contributions of all its neighbors.

For user nodes, node features include engagement metrics, credibility scores, and content characteristics, while relations between users include mentions or interactions on the same posts.

Combining centrality metrics with GNN-Explainer enables a comprehensive assessment of both structural and contextual influence in misinformation dynamics. Both approaches applied on the heterogeneous GAT deal with the multiplicity of types in the graph, allowing any node type to be chosen as target.

5 Experimental Evaluation

The evaluation aims to demonstrate the effectiveness of our fake news classification framework and validate its ability to identify influential users while providing interpretable insights into misinformation spread.

5.1 Datasets

Experiments have been conducted on two real-world datasets, i.e., MuMiN and PolitiFact datasets. Both are modeled as heterogeneous information networks, with multiple node and edge types and external information associated with nodes available as a set of attributes. The former is multi-topic; a complete description is provided in [9]. The latter is extracted from the FakeNewsNet data repository [10], which fact-checks news pertaining to the US political system.

In the following, we use the terms Claim and News interchangeably, as the MuMiN dataset refers to news items as Claims, while the PolitiFact dataset uses the term News. Classification is performed on nodes of type Claim (C) in the MuMiN dataset and nodes of type News (N) in the PolitiFact dataset.

Table 1: MuMiN statistics, in terms of no. of nodes, edges and meta-paths. In bold are highlighted the types involving users.

# Nodes	Claim (C)	2168	# Edges	T *discusses* C (C-T)	5081
	Tweet (T)	4340		R *reply_to* T (T-R_r)	90196
	Reply (R)	195459		R *quote_of* T (T-R_q)	101216
	User (U)	153168		T *has_hashtag* H (T-H)	2289
	Hashtag (H)	28091		T *has_article* A (T-A)	1898
	Image (I)	1020		T *has_image* I (T-I)	1028
	Article (A)	1453		T *mentions* U (U-T_m)	1119
# Meta-paths	C-T-U-T-C	28867		U *posted* T (U-T_p)	4091
	C-T-H-T-C	21577		U *posted* R (U-R)	179247
	C-T-R-T-C_r	2859		U *retweeted* T (U-T_r)	13402
	C-T-R-T-C_q	3042		**U *follows* U** (U-U_f)	18379
	U-T-U	11412		**U *mentions* U** (U-U_m)	2797
	U-R-U	146056		U *has_hashtag* H (U-H)	50451

Table 2: PolitiFact statistics, in terms of no. of nodes, edges and meta-paths. In bold are highlighted the types involving users.

# Nodes	News (N)	696	# Edges	T *discusses* N (N-T)	276676
	Tweet (T)	268306		T *has_hashtag* H (T-H)	59782
	User (U)	169106		U *posted* T (U-T_p)	285124
	Hashtag (H)	18631		U *retweeted* T (U-T_r)	539
# Meta-paths	N-T-H-T-N	44682		**U *mentions* U** (U-U)	84093
	N-T-U-T-N	46025			
	N-T-U-U-T-N	20056			
	U-T-U	533			

To enhance classification performance by capturing richer structural and semantic relationships, we build different meta-paths toward the Claim/News node type. Two meta-path types are common to both datasets: including connections between pairs of claims (news) discussed in tweets by the same user or associated with the same hashtag. Additionally, dataset-specific meta-paths are defined. For MuMiN, we identify pairs of claims belonging to the same conversation thread through reply or quote relationships, respectively. For PolitiFact, we include pairs of news discussed in tweets posted by users who mention each other. Tables 1 and 2 show the statistics in terms of number of nodes, number of edges, and number of meta-path instances for each type, for the two datasets.

Centrality measures are computed with respected to the user graph, a homogeneous weighted graph constructed from the heterogeneous one by keeping all nodes of type User. Edges include both direct relationships between users in the heterogeneous graph (e.g., mentions and friendship) and meta-paths connecting pairs of users based on their interactions on posts (e.g., retweet and reply).

Users are further categorized into *fake*, *real* and *mixed* based on the news they discuss and spread with their tweets. *Mixed* users are those who are not dominated by the spread of either real or fake news because the difference between the two on the total news discussed is less than 25%.

Table 3: Results of the Fake News Detection task averaged over 5 runs.

Dataset	**F1-micro**	**F1-macro**	**Precision T**	**Recall T**	**Precision F**	**Recall F**
MuMiN	0.954±0.007	0.788±0.117	0.721±0.139	0.647±0.191	0.967±0.014	0.985±0.009
Politifact	0.859±0.039	0.845±0.054	0.848±0.066	0.943±0.044	0.909±0.063	0.736±0.153

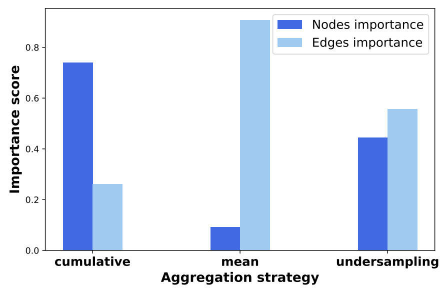

(a) PolitiFact dataset

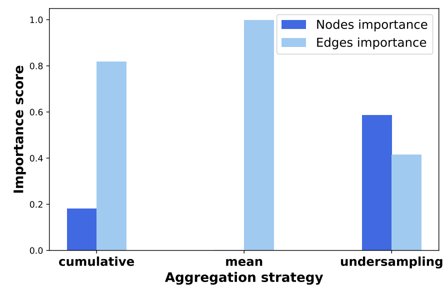

(b) MuMiN dataset

Fig. 2: Importance scores for nodes and edges w.r.t. diverse aggregation strategies

5.2 Results

Fake News Detection Performance We first assess the performance of the proposed GAT-based architecture in classifying multi-modal fake news. Table 3 reports accuracy metrics for both datasets, highlighting the model's ability to generalize across different misinformation domains. More details are provided in [8]. The strong performance justifies the explainability study.

Here, we report the main parameters used in the experimentation, selected via a focused grid search based on validation performance. As regards the graph neural network model, we employed a l-layer GAT [3] architecture with dropout set to 0.4, hidden channels dimension set to 64 and out channel dimension set to 2 as the number of classes. The number of layers l was set to 2 for the PolitiFact dataset and to 3 for MuMiN. We employed a weighted cross entropy loss in a fully supervised setting, with weights inversely proportional to the class frequency. We trained the model over 200 epochs and used the Adam optimization algorithm. The learning rate was set to 0.005 while the weight decay to 0.001.

The results validate the effectiveness of our approach in leveraging both textual and structural information for fake news detection, and enable our subsequent analysis on influential node identification and explainability.

Explainability Study This section delves into an in-depth evaluation of the classifier's decision process, helping to assess the contribution of the two modalities. We employed a Captum-based GNNExplainer [12] using the Integrated Gradients algorithm for multi-instance explanations, able to quantify each input feature and relation's impact on the model output via gradient analysis.

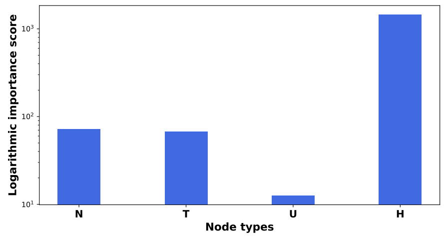

(a) PolitiFact dataset - Features

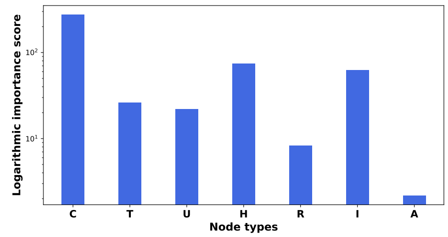

(b) MuMiN dataset - Features

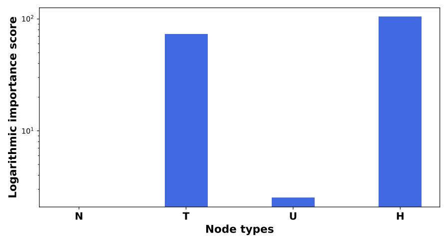

(c) PolitiFact dataset - Relations

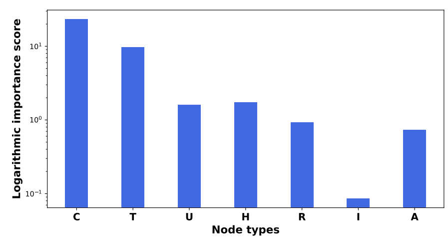

(d) MuMiN dataset - Relations

Fig. 3: Node types importance based on features (a-b) and on relations (c-d).

Features (nodes) vs relations (edges) importance. We evaluate component contributions in the social graph by computing importance scores. For nodes, scores are based on individual features, while for relations, scores are computed for edge and meta-path instances. We define three aggregation strategies to assess node/edge contributions: *cumulative* (sum of all feature or instance scores), *mean* (type-normalized scores), and *undersampling* (top-k instances only). Figure 2(a) shows that for PolitiFact, nodes dominate cumulative importance (overall behavior), while edges are more influential in mean importance (per-instance decisions). For MuMiN (Fig. 2(b)), cumulative and undersampling yield opposite trends, though mean importance again highlights edge dominance in misinformation propagation. The cumulative strategy particularly reflects the underlying network structure, where denser graphs naturally yield higher aggregated scores for edges. Top-k analysis reveals edges slightly outweigh nodes in PolitiFact, while the opposite accounts for MuMiN, confirming both components' importance for fake news detection and the value of multi-modal approaches.

Individual node and relation type importance. We analyze the impact of individual node and relation types using the undersampling strategy (evaluating only the top-k instances per type) to assess their structural and semantic contributions. To ensure clarity in presenting the results, we adopt the following notation. Node types are denoted by their initial letter (e.g., U for users), while edges and meta-paths are represented by the initials of connected node types (e.g., N-T for News-Tweet), plus any distinctive suffix if multiple relations between two node

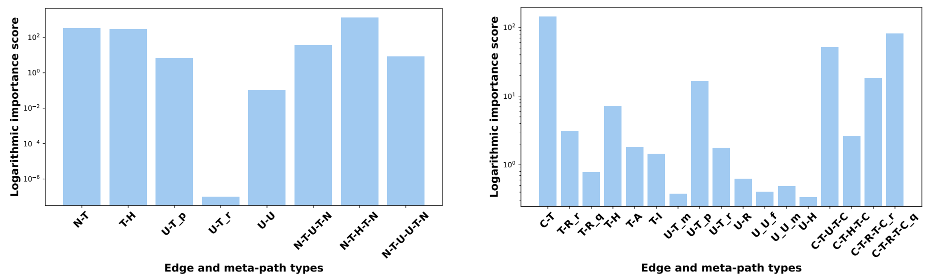

(a) PolitiFact dataset (b) MuMiN dataset

Fig. 4: Edge and meta-path types importance.

types exists. Reciprocal relationships, such as bidirectional interactions between news and tweets, are aggregated to ensure comprehensive analysis (e.g., N-T encompasses both "*news is discussed by a tweet*" and "*tweet discusses a news*").

A deeper analysis of individual node types (see Figure 3) reveals that hashtags and tweets contributed the most to model predictions across all categories, more than news node target type for the PolitiFact dataset. In both datasets, tweet connections drive the classification (Figure 3(c)-(d)) and meta-paths emerge as influential structural components for misinformation detection (Figure 4).

In the PolitiFact dataset, hashtags emerge as the most influential node type, both in terms of features and relationships. This prominence can be attributed to several factors. First, hashtags provide semantic coherence by grouping topic-specific discussions, which are often exploited in misinformation campaigns. Second, they foster network influence through dense communities that amplify message propagation. Third, their propagation patterns can reveal coordinated activity, such as bot-driven amplification. Lastly, hashtags act as cross-modal signals, bridging textual content and network structure.

Conversely, the MuMiN dataset exhibits distinct patterns: claims, hashtags, and images are the most impactful nodes, while claim-tweet interactions and retweet relations dominate edge importance. These differences arise from structural disparities between the datasets. Specifically, MuMiN is denser, with abundant relationships that enhance the role of multi-hop connections and meta-paths. In contrast, sparser graphs like PolitiFact rely more on node-level features due to fewer relational patterns. Overall, the results underscore the interplay between graph density and model explainability. Dense networks prioritize structural importance, whereas sparser ones emphasize feature-level contributions, highlighting the need for adaptable interpretation frameworks depending on network topology.

Identifying Key Influential Users We then focus on user nodes, who play a critical role in fake news propagation, by combining network centrality metrics with GNN-Explainer outcomes. These techniques generate independent node

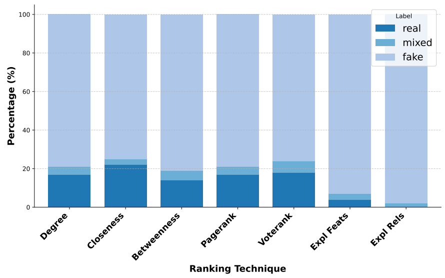

(a) PolitiFact dataset

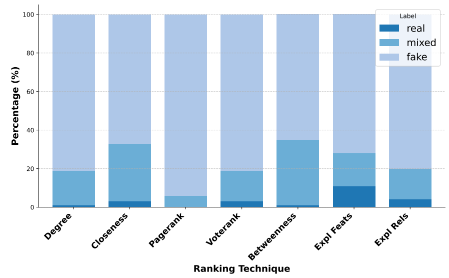

(b) MuMiN dataset

Fig. 5: Label distribution of top-k (k=100) most influential users. Stacked bar chart for different ranking techniques.

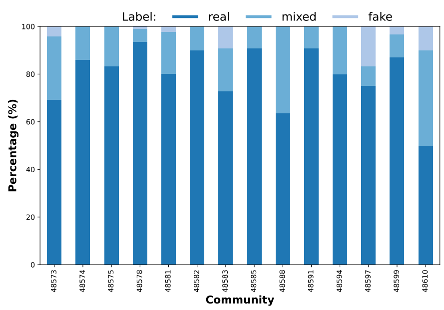

(a) PolitiFact dataset

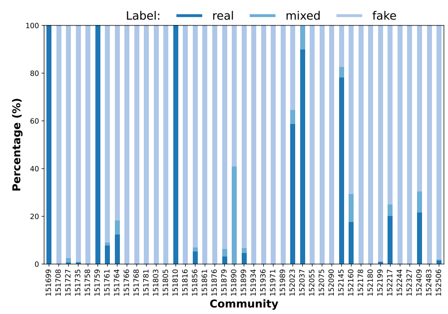

(b) MuMiN dataset

Fig. 6: User label distribution in communities with at least 10 members. Stacked bar chart for largest community IDs.

rankings, capturing both structural influence (via centrality measures) and model-learned significance (via GNN-Explainer). The results demonstrate consistent alignment between these methods, with top-ranked users predominantly being fake news spreaders based on each technique and across both datasets (Figure 5). This insight indicates that both structurally significant users within the network and those identified as critical by GNN-Explainer are predominantly misinformation actors, reinforcing the importance of focusing not only on network connectivity but also on content-driven influence. Their impact extends beyond simple connectivity; they shape narratives and engagement patterns that reinforce the persistence of fake news. Thus, effective mitigation strategies must target not just high-degree hubs but also clusters of mid-level influencers who operate within misinformation ecosystems, amplifying false information across different sub-networks.

Further investigation through community detection reveals distinct propagation patterns between datasets. The Louvain algorithm identifies power-lawdistributed community sizes, with most communities being small but a few significantly larger (e.g., 95 users in the largest PolitiFact community and 1,730 users in the largest MuMiN community). Figure 6 plots the distribution of users spreading fake and real news within larger communities (at least 10 members) in the two datasets. The community analysis reveals strong polarization patterns across datasets. In MuMiN, these larger communities show strong polarization, functioning as misinformation echo chambers where false narratives are reinforced through coordinated clusters of mid-level influencers rather than centralized hubs. This decentralized structure mirrors real-world misinformation campaigns that rely on distributed networks of actors. In contrast, PolitiFact exhibits more isolated misinformation spreaders within predominantly real-news communities, suggesting simpler propagation dynamics centered around individual superspreaders. These structural differences have important implications for intervention strategies. For MuMiN-like networks with dense, polarized communities, effective mitigation requires disrupting coordinated clusters and mid-level influencers. For PolitiFact-like structures with isolated spreaders, targeting key individual actors may be adequate.

These findings collectively demonstrate the necessity of adopting an integrated analytical framework that simultaneously considers both content characteristics and network structure. Such a dual perspective not only provides a more comprehensive understanding of misinformation propagation patterns but also enables the development of more robust and generalizable intervention strategies, capturing the nuanced mechanisms through which false narratives emerge, persist, and spread within different network ecosystems.

6 Conclusion

This study introduced a graph-based, explainable framework for fake news detection, combining Graph Attention Networks with GNN-Explainer and centrality-based influence analysis. Our multimodal approach effectively integrates textual and social signals to achieve high classification accuracy while offering interpretability. Experiments on real-world datasets highlight how graph density shapes misinformation dynamics: structural relationships dominate in denser networks, whereas node-level features play a larger role in sparser ones. Exploring centrality measures further reveals that misinformation spreaders consistently occupy top-ranked positions, acting as both influential nodes and amplifiers of false narratives. These insights underscore the importance of dataset-specific interventions: disrupting coordinated groups in highly connected, polarized networks, and targeting key individuals in sparser environments.

Future work will explore adaptive mitigation strategies through reinforcement learning and temporal modeling, aiming to enhance responsiveness to the evolving nature of misinformation. By advancing the understanding of misinformation patterns, our research lays the groundwork for more targeted and effective countermeasures in digital ecosystems.

Disclosure of Interests. The authors have no competing interests to declare.

Test reproducibility. The code necessary to replicate our experiments is available at: `https://github.com/lilymart/M3DUSA`.

Acknowledgments. This work has been partially supported by: (i) project SERICS (PE00000014) under the NRRP MUR program funded by the EU - NGEU; (ii) Italian MUR, PRIN PNRR 2022 Project "Limiting MIsinformation spRead in online environments through multi-modal and cross-domain FAKe news detection (MIRFAK)", Prot.: P2022C23K9, funded by the EU - NGEU.

References

1. Aïmeur, E., Amri, S., Brassard, G.: Fake news, disinformation and misinformation in social media: a review. Social Network Analysis and Mining **13**(1), 30 (2023)
2. Asgharian Rezaei, A., Munoz, J., Jalili, M., Khayyam, H.: A machine learning-based approach for vital node identification in complex networks. Expert Systems with Applications **214**, 119086 (2023)
3. Brody, S., Alon, U., Yahav, E.: How attentive are graph attention networks? CoRR **abs/2105.14491** (2021)
4. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In: NAACL-HLT. pp. 4171–4186 (2019)
5. Islam, M.R., Liu, S., Wang, X., Xu, G.: Deep learning for misinformation detection on online social networks: a survey and new perspectives. Social Network Analysis and Mining **10**(1), 82 (2020)
6. Keikha, M.M., Rahgozar, M., Asadpour, M., Abdollahi, M.F.: Influence maximization across heterogeneous interconnected networks based on deep learning. Expert Systems with Applications **140**, 112905 (2020)
7. Kou, J., Jia, P., Liu, J., Dai, J., Luo, H.: Identify influential nodes in social networks with graph multi-head attention regression model. Neurocomp. **530**, 23–36 (2023)
8. Martirano, L., Comito, C., Guarascio, M., Pisani, F.S., Zicari, P.: M3dusa: A modular multi-modal deep fusion architecture for fake news detection on social media. Social Network Analysis and Mining **15**(1), 1–17 (2025)
9. Nielsen, D.S., McConville, R.: Mumin: A large-scale multilingual multimodal fact-checked misinformation social network dataset. In: Proc. of the 45th Intl. ACM SIGIR Conf. on research and development in inf. retr. pp. 3141–3153 (2022)
10. Shu, K., Sliva, A., Wang, S., Tang, J., Liu, H.: Fake News Detection on Social Media: A Data Mining Perspective. ACM SIGKDD Explorations Newsletter **19**(1), 22–36 (2017)
11. Xie, J., Yu, J., Chen, Z.: Mega: identifying influential nodes in heterogeneous networks based on meta-path and attention. Journal of Statistical Mechanics: Theory and Experiment **2025**(2), 023401 (2025)
12. Ying, Z., Bourgeois, D., You, J., Zitnik, M., Leskovec, J.: Gnnexplainer: Generating explanations for graph neural networks. Advances in neural information processing systems **32** (2019)
13. Yu, E.Y., Wang, Y.P., Fu, Y., Chen, D.B., Xie, M.: Identifying critical nodes in complex networks via graph convolutional networks. Knowledge-Based Systems **198**, 105893 (2020)
14. Zhang, X., Malkov, Y., Florez, O., Park, S., McWilliams, B., Han, J., El-Kishky, A.: Twhin-bert: A socially-enriched pre-trained language model for multilingual tweet representations. arXiv preprint arXiv:2209.07562 (2022)
15. Zhao, G., Jia, P., Zhou, A., Zhang, B.: Infgcn: Identifying influential nodes in complex networks with graph convolutional networks. Neurocomp. **414**, 18–26 (2020)