International Journal of
Computational Linguistics Research

DLINE JOURNALS

Print ISSN: 0976-416X
Online ISSN: 0976-4178

IJCLR 2026: 17 (2)
https://doi.org/10.6025/ijclr/2026/17/2/59-77

Temporal Bursts and Structural Persistence in Online Misinformation Networks

Pit Pichappan
Digital Information Research Labs
Chennai. 600017. India
pichappan@dirf.org

ABSTRACT

The rapid proliferation of online misinformation poses a critical challenge to digital information ecosystems, driven by coordinated communities and dynamic network interactions. While existing detection models often treat social networks as static or focus narrowly on content, this study investigates the complex interplay between temporal bursts, structural persistence, and community dynamics. Utilizing a novel Misinformation Dynamics Testbed, we analyze a longitudinal Reddit dataset (2015–2024, n=1,455) across multiple analytical dimensions. Statistical and network analyses reveal that misinformation engagement follows a heavily skewed, heavy-tailed distribution (Gini ≈ 0.71), with a pronounced visibility credibility mismatch (Spearman's p ≈ 0.32). Temporal patterns are non stationary and episodic, characterized by event driven activity surges rather than linear trends. Structurally, the network exhibits scale free properties and densification driven growth, in which stable core communities show highly volatile interconnections. Cross community analysis uncovers a hybrid propagation mechanism: localized amplification within tightly knit echo chambers, coupled with controlled diffusion through strategic bridge nodes. These findings demonstrate that misinformation operates as a self reinforcing, temporally adaptive process rather than a purely content-driven phenomenon. Consequently, effective mitigation requires a paradigm shift from static, post-level moderation to dynamic, system level interventions. Prioritizing high centrality hubs, monitoring bridge nodes, and deploying real time burst detection are essential for disrupting propagation pathways. This research underscores the necessity of integrated, network aware frameworks to address the evolving landscape of digital misinformation.

Keywords: Misinformation propagation, Temporal dynamics, Graph neural networks, Community dynamics, Echo chambers, Burst detection, Network densification, Engagement credibility mismatch

Received: 26 August 2025, Revised 9 December 2025, Accepted 5 January 2026

1. Introduction

With the increasing reliance on social networks as primary sources of information, the proliferation of misinformation has emerged as a critical challenge to the modern information ecosystem. Misinformation spreads rapidly across social media platforms [1] influencing public opinion, electoral processes, and public health decisions [2]. Notably, this spread is often driven by coordinated communities that reinforce misleading narratives rather than isolated individuals [3].

Despite extensive research on misinformation detection, existing approaches often treat propagation either as a content classification problem or as a static network phenomenon. Such perspectives overlook the inherently dynamic and community-driven nature of misinformation spread, where temporal fluctuations and structural evolution play a critical role. Consequently, there is a need for a more integrated analytical perspective that captures not only what misinformation contains, but how it evolves, amplifies, and persists within complex social network structures over time.

2. Graph-Based Approaches to Misinformation Detection

Social network structures provide a natural framework for modeling misinformation dissemination. Relationships among users, posts, and interactions can be represented as homogeneous or heterogeneous graphs, enabling effective fake news detection [4, 5, 6]. Early approaches often treated these relationships as static; however, real-world social networks are inherently dynamic.

To address this limitation, temporal network models have been introduced. For instance, Wei Wang [7] proposed a theoretical framework for misinformation propagation in temporal networks. Similarly, Dou et al. [8] (2021) incorporated both user credibility (internal factors) and propagation patterns (external factors) to enhance detection performance. Shu, Wang, and Liu [9] (2019) modeled publisher–post–user relationships using heterogeneous information networks and improved representation learning via matrix factorization. More recently, Park and Chai [10] (2023) integrated user behavior, content features, and social network structures using a social capital framework to better capture propagation characteristics.

2.1 Multimodal Learning for Fake News Detection

News content on social media is inherently multimodal, consisting of text, images, videos, and audio. To leverage this richness, researchers have developed multimodal approaches that significantly enhance detection accuracy [11, 12, 13, 14]. Zhu et al. [13] introduced a reinforcement-driven subgraph selection mechanism that retrieves entity-level knowledge and captures cross-modal correlations through heterogeneous graph learning. Luvembe et al. [15] proposed a complementary attention fusion mechanism that combines image captions and textual features, thereby reducing semantic noise and improving robustness. Similarly, Peng et al. [16] emphasized contextual semantic learning by integrating global and local representations, addressing the challenge that fake news often lacks consistent semantic patterns.

Despite these advancements, many multimodal models fail to incorporate temporal dynamics, which are crucial for understanding how misinformation evolves over time. Real-time monitoring and dynamic information integration remain essential for addressing the complexities of misinformation spread [17].

2.2 Temporal and Dynamic Modeling of Misinformation

Recent research has increasingly focused on incorporating temporal dynamics into misinformation detection. Xue et al. [18] highlighted the importance of multimodal consistency in capturing social media characteristics, while Yadav and Gupta [19] leveraged emotional cues and vision transformers to enhance classification performance. Zhang [20] addressed early fake news detection by modeling propagation paths, enabling the use of rich social context for timely identification.

Graph-based dynamic approaches have also gained prominence. Sivasankari [21] demonstrated the effectiveness of social network graph representations for modeling fake news propagation. Jing [22] proposed a multilevel attention network combined with social situation analytics to analyze disinformation diffusion trends.

Temporal modeling techniques such as Temporal Point Processes (TPPs) have been applied to capture event sequences in misinformation spread. Zhang, Q. [23] modeled misinformation propagation as a dynamic graph and extracted temporal evolution patterns using TPPs, which are well-suited for modeling stochastic event sequences in user engagement data.

2.3 Temporal Graph Neural Networks and Advanced Models

Advancements in Temporal Graph Neural Networks (TGNNs) have enabled more sophisticated modeling of time-evolving interactions. [24]. Caleb Stam [25] demonstrated that temporal patterns of user interactions and semantic similarities can independently indicate the spread of misinformation and proposed a dynamic graph-based detection framework. Ghosh [26] combined Temporal Graph Networks (TGN) with Recurrent Neural Networks (RNNs) to jointly capture structural and temporal dynamics.

Victor Ethan [27] applied TGNNs to track the evolution of communities involved in coordinated misinformation campaigns. Modern architectures such as DySAT [28] and TGAT [29] utilize self-attention and inductive message passing to learn temporal embeddings, while TGN[30] introduces memory modules to model evolving node states. These approaches enable fine-grained temporal representation learning and have shown strong performance in prediction tasks [31].

However, despite their effectiveness, these models are often limited to node-level predictions and link forecasting. Community-level analysis remains challenging due to sensitivity to noise and a lack of temporal stability mechanisms. To address scalability and real-time monitoring needs, streaming community detection methods have been proposed, allowing continuous updates of community structures as new interactions occur [32].

2.4 Research Gap

Taken together, the reviewed studies demonstrate significant progress in modeling misinformation using graph structures, multimodal data, and temporal learning frameworks. However, these approaches remain fragmented, often emphasizing isolated dimensions of the problem without fully capturing the interdependence between temporal dynamics, structural evolution, and community-level behavior. This fragmentation limits the ability to explain how misinformation persists and adapts within real-world ecosystems. Although significant progress has been made in graph-based, multimodal, and temporal approaches, a key limitation persists: the lack of unified frameworks that effectively integrate content, network structure, and temporal evolution. Existing models often focus on individual aspects, such as static structure or node-level dynamics, while overlooking the complex interplay between evolving communities and misinformation propagation.

Addressing this gap requires a controlled and extensible framework capable of jointly modeling temporal bursts, structural persistence, and cross-community interactions. To this end, this study proposes a Misinformation Dynamics Testbed, designed to simulate and analyze the interplay between these factors in a unified environment. The following section introduces the architecture and components of this testbed.

3. Testbed

To systematically investigate the complex dynamics of misinformation propagation, a simulation-driven analytical framework is required. Real-world datasets, while rich, often lack the controllability and reproducibility needed to isolate key factors such as burst frequency or network volatility. The proposed Misinformation Dynamics Testbed addresses this limitation by providing a semi-synthetic environment in which observed data characteristics can be replicated and experimentally manipulated.

The testbed framed for this work is outlined as follows.

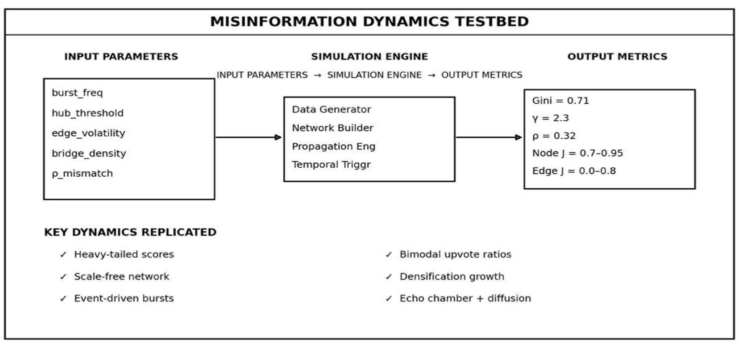

Figure 1. Testbed

The testbed presents a conceptual pipeline, the "Misinformation Dynamics Testbed", organised inside a framed layout with a bold title header. Across the top, it shows a left to right flow: Input Parameters → Simulation Engine → Output Metrics, emphasizing that controllable knobs drive a synthetic simulation whose behavior is then evaluated with quantitative diagnostics.

On the left, the Input Parameters box lists five tunable levers that shape the simulated scenario: burst_freq, hub_threshold, edge_volatility, bridge_density, and ρ _mismatch. These represent, respectively, how often activity spikes occur, how influential hubs emerge, how unstable connections are over time, how strongly communities are connected through bridges, and how mismatched correlation (ρ) might be between modeled vs observed conditions.

In the center, the Simulation Engine box summarizes four internal modules: Data Generator, Network Builder, Propagation Engine, and Temporal Trigger. Together, they imply the system first creates content/activity, constructs the underlying social graph, simulates how misinformation spreads, and then injects time-based events that create bursts.

On the right, the Output Metrics box lists example measurements (e.g., Gini, γ, ρ, and node/edge Jaccard ranges) that quantify inequality, scaling behavior, correlation, and structural similarity over time.

At the bottom, a Key Dynamics Replicated section highlights that the testbed is intended to replicate real-world phenomena such as heavy-tailed attention, bimodal feedback, scale-free structure, densification, event bursts, and echo-chamber diffusion.

4. Dataset

This study utilizes a dataset comprising 1,455 textual submissions collected from Reddit, spanning the period 2015–2024. The data were sourced from multiple subreddits associated with news dissemination, satire, and online discourse, including communities such as *neutralnews*, *nottheonion*, and *savedyouaclick* . Each instance in the dataset contains the following structured attributes: *title*, *content*, *timestamp*, *category*, *subreddit*, *score*, and *upvote_ratio*.

The dataset is annotated into four primary categories: *True*, *misleading content*, *satire*, and *imposter content*, reflecting the nature of the information rather than independently verified factual accuracy. The corpus exhibits substantial stylistic diversity, ranging from objective journalistic reporting to ironic and fabricated narratives. Additionally, the dataset is characterized by class imbalance, with satire and imposter content being more prevalent than factual and misleading entries.

Given its temporal coverage and thematic diversity, the dataset is well-suited for tasks such as misinformation detection, satire classification, and longitudinal analysis of online narratives. However, limitations include missing values in engagement metrics and potential label noise due to reliance on source-based annotations.

4.1 Overview of Dataset and Analytical Scope

The dataset consists of Reddit posts annotated for misinformation-related analysis, incorporating attributes such as title, content, timestamp, subreddit, score, and upvote ratio. These features enable a multi-dimensional investigation of misinformation dynamics, including engagement behavior, temporal evolution, and network-based propagation.

Building on these characteristics, the dataset serves as both an empirical foundation and a calibration reference for the proposed testbed. By aligning simulation parameters with observed data properties, the subsequent analysis ensures both realism and interpretability. The following section presents a multidimensional examination of misinformation dynamics based on this integrated framework.

5. Analysis

The statistical formulations used throughout this section are presented to formally validate observed patterns; however, emphasis is placed on their interpretive implications rather than mathematical derivation. Each test is used to confirm distributional properties, group-level differences, or temporal trends relevant to misinformation dynamics.

The analytical results presented in this section are derived using the proposed *Misinformation Dynamics*

Testbed, which enables systematic examination of engagement behavior, temporal evolution, and network-driven propagation dynamics. The analysis follows a progressive structure, moving from distributional properties to network structure and cross-community diffusion, thereby providing a unified understanding of misinformation behavior.

5.1 Distributional Characteristics of Engagement

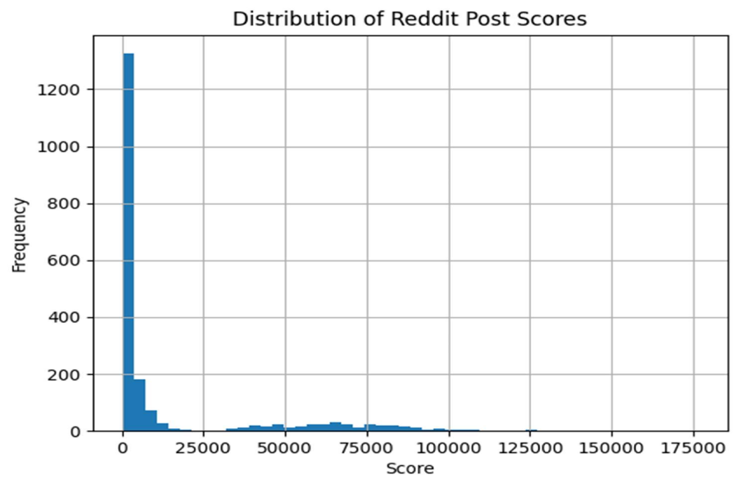

Figure 2. Distribution of post scores (heavy-tailed behavior)

The distribution of post scores exhibits a strongly right skewed, heavy tailed pattern, indicating that a small subset of posts accounts for a disproportionately large share of user engagement. This reflects unequal visibility dynamics, where only a limited number of posts achieve high amplification.
To statistically validate this observation, the Shapiro Wilk test is applied,

$$W = \frac{\left(\sum_{i=1}^{n} a_i\, x_{(i)}\right)^2}{\sum_{i=1}^{n} (x_i - \bar{x})^2} \qquad (1)$$

Where:

- $x_{(i)}$: ordered sample values
- a_i: constants from the covariance matrix
- $\bar{x}$: sample mean

yielding $p < 0.001$, which confirms a significant deviation from normality. Consequently, non-parametric statistical methods are employed throughout the analysis.

Additionally, missing values in the upvote ratio suggest incomplete engagement metadata, while the wide dispersion of scores highlights heterogeneous interaction patterns across subreddits. These findings collectively emphasize that attention allocation is inherently unequal, a key factor underlying misinformation amplification.

5.2 Subreddit-Level Concentration of Misinformation

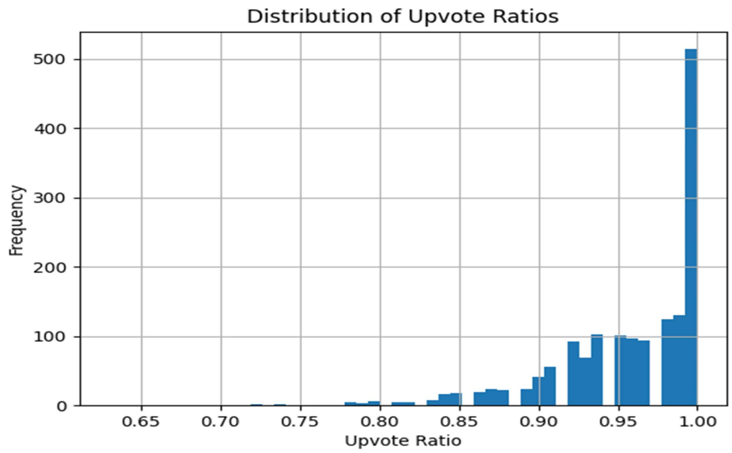

Figure 3. Distribution of upvote ratios across posts

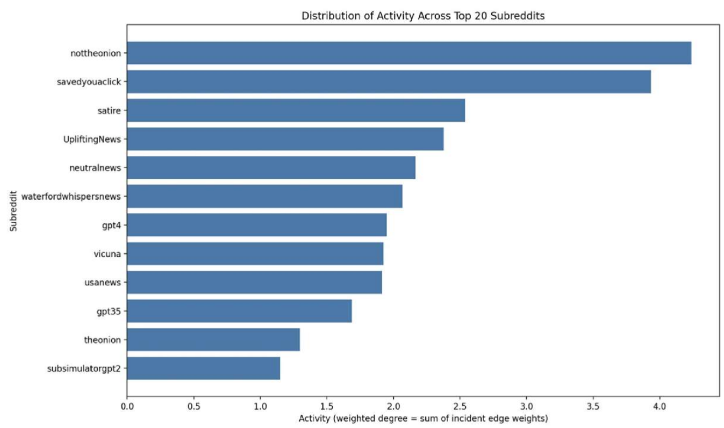

Figure 4. Distribution of activity across top subreddits

The most "active" nodes (by total incident edge weight) are *nottheonion* and *savedyouaclick*, followed by satire, *UpliftingNews*, and *neutralnews*.

Misinformation is not uniformly distributed; it is highly concentrated in a limited number of subreddits. This inequality is quantified using the Gini coefficient,

$$G = \frac{\sum_{i=1}^{n} \sum_{j=1}^{n} | x_i - x_j |}{2n^2 \bar{x}} \quad (2)$$

which yields a value of approximately 0.71, indicating strong concentration.

A Kruskal–Wallis test

$$H = \frac{12}{N(N+1)} \sum_{i=1}^{k} n_i R_i^2 - 3(N+1) \quad (3)$$

Where:

- n_k: size of group K
- $\bar{R}_k$: average rank of group
- N: total observations

(p < 0.001), further confirms statistically significant differences in engagement patterns across subreddits. High-activity communities display substantial variability, whereas smaller communities tend to exhibit more consistent thematic behavior.

This concentration fosters information silos, where misinformation is produced and reinforced within tightly interconnected communities, limiting exposure to corrective information.

5.3 Engagement Dynamics: Visibility versus Consensus

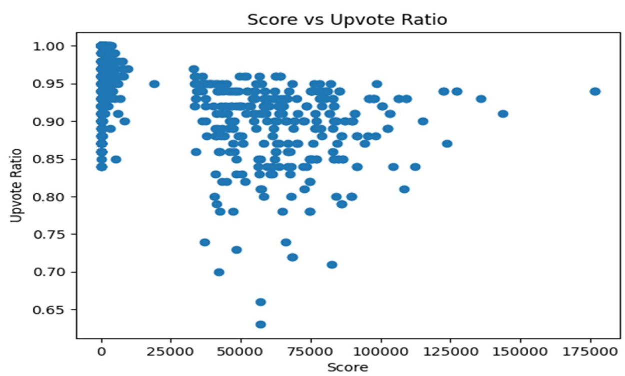

Figure 5. Relationship between post score and upvote ratio

To investigate whether visibility corresponds to credibility, the Spearman rank correlation coefficient is computed.

$$\rho = 1 - \frac{6\sum d_i^2}{n(n^2 - 1)} \quad (4)$$

di: difference between ranks

n: number of observations

The observed correlation ($\rho \approx 0.32$, $p < 0.001$) indicates a weak to moderate positive relationship.

This suggests that highly visible posts are not necessarily widely endorsed. The bimodal distribution of upvote ratios further reveals two distinct regimes:

- highly endorsed content (upvote ratio ≈ 0.9–1.0), and

• controversial or divisive content with lower agreement levels

These findings highlight a critical engagement credibility mismatch, where content visibility is driven by factors such as emotional intensity, controversy, or polarization rather than factual accuracy.

5.4 Temporal Dynamics of Misinformation Activity

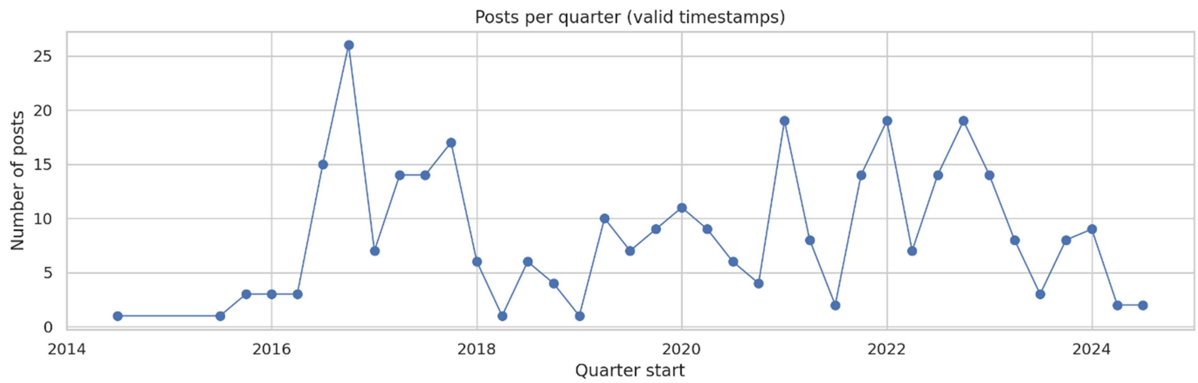

Figure 6. Posts per quarter (valid timestamps)

Temporal analysis reveals a non-stationary process characterized by intermittent bursts of activity. The dataset shows minimal activity during early years (2014–2016), followed by rapid growth (2016–2017), decline (2018–2019), renewed bursts (2020–2022), and eventual stabilization (2023–2024)

The Mann–Kendall test is given by

$$S = \sum_{i=1}^{n-1} \sum_{j=i+1}^{n} \text{sign}(x_j - x_i) \qquad (5)$$

The Mann Kendall test indicates no significant monotonic trend ($p > 0.05$), confirming that the observed dynamics are not linear but episodic.

These results suggest a hybrid temporal process composed of:

• event-driven surges triggered by external socio-political factors, and

• persistent baseline activity

Such burst-like behavior indicates that misinformation propagation is highly reactive and influenced by exogenous events.

5.5 Network Structure and Growth Dynamics

The subreddit interaction network exhibits clear scale-free characteristics, with a small number of highly connected hubs coexisting with many sparsely connected nodes. The degree distribution follows a power law

$$P(k) \sim k^{-\gamma} \quad (6)$$

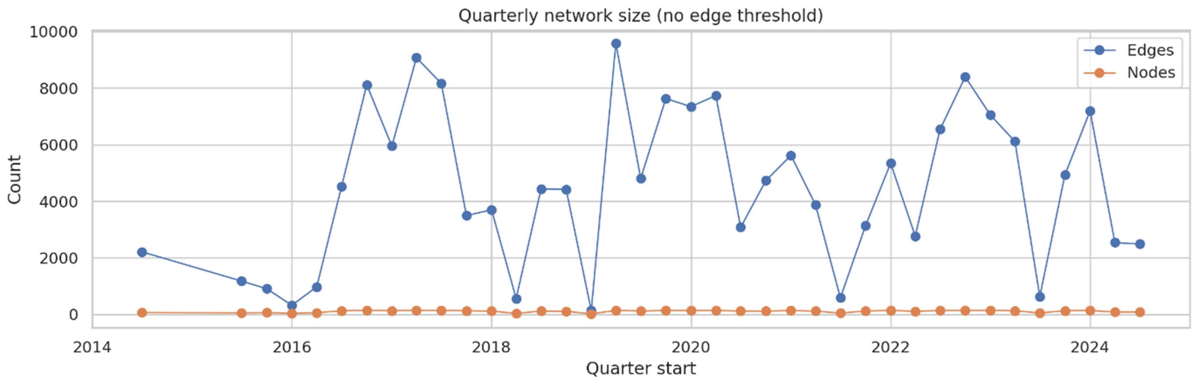

Figure 7. Quarterly network size (nodes and edges)

k: node degree

γ: scaling exponent

($\gamma \approx 2.3$, $R^2 \approx 0.87$), while clustering (~0.62) and modularity (~0.48) indicate strong community structure.

A key observation is that network evolution is driven by densification rather than expansion:

- Node growth remains relatively stable over time
- Edge growth shows significant variability

This indicates that misinformation spreads through increasingly dense interactions within existing communities, rather than through the addition of new participants. Such behavior reinforces existing structures and intensifies local propagation.

5.6 Structural Stability and Evolution

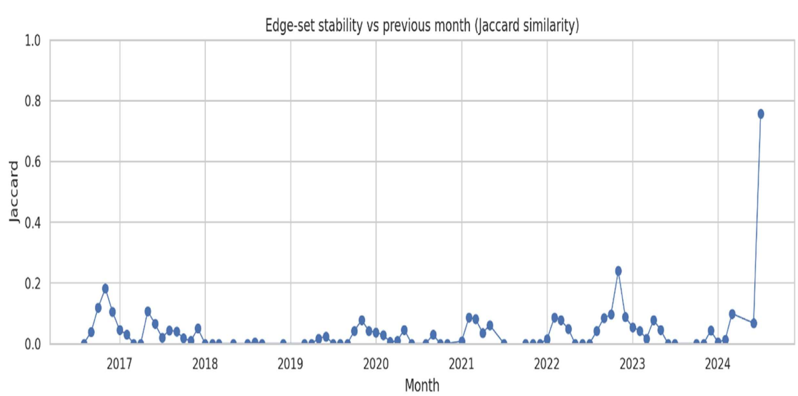

Figure 8. Edge-set stability over time (Jaccard similarity)

Temporal stability analysis reveals a divergence between node persistence and edge volatility.

Jaccard similarity is computed as:

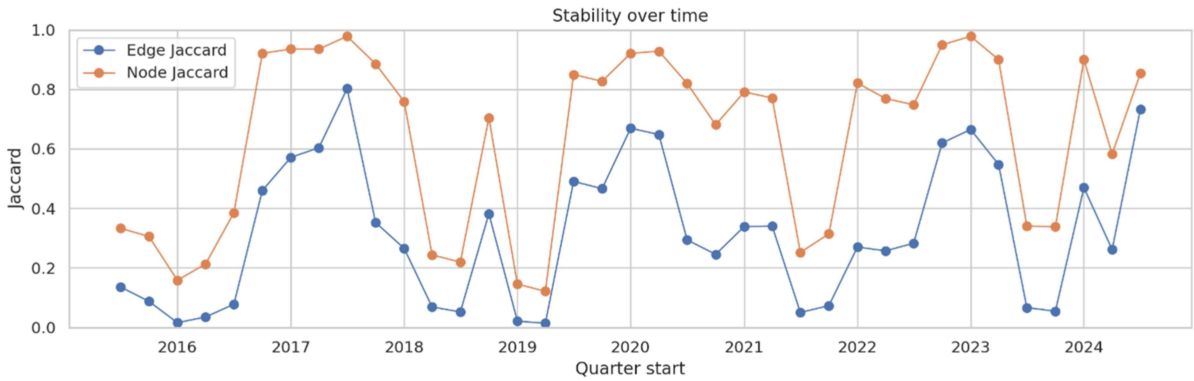

Figure 9. Node vs edge stability comparison

$$J(A,B) = \frac{|A \cap B|}{|A \cup B|} \qquad (7)$$

Node-level Jaccard similarity remains high (≈0.7–0.95), indicating that the set of participating communities and vocabulary remains relatively stable over time.

In contrast, edge-level similarity varies widely (≈ 0.0–0.8), indicating frequent reconfiguration of relationships between nodes.

This suggests that while the actors remain persistent, their interactions evolve dynamically. Periodic increases in stability may correspond to coordinated campaigns or sustained narrative phases.

5.7 Interaction Intensity and Content Richness

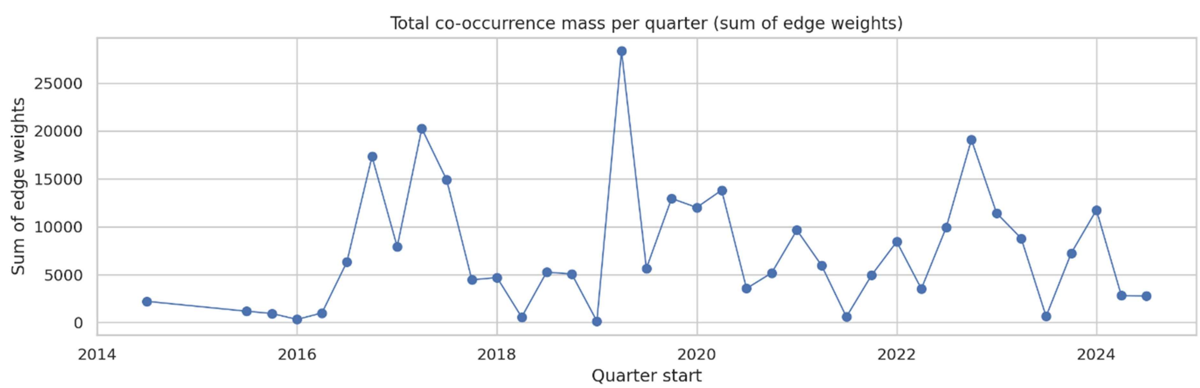

Figure 10. Co-occurrence mass per post

Co-occurrence intensity is computed as:

$$\text{Weight per Post} = \frac{\sum w_{ij}}{N_{\text{posts}}} \qquad (8)$$

Where:

- w_{ij}: edge weight between terms i, j

- N_p: number of posts

The co-occurrence mass per post, representing normalized interaction intensity, shows significant temporal variability. Peaks correspond to periods of high semantic density, while troughs indicate fragmented discourse.

Importantly, these fluctuations are not directly aligned with post volume, indicating that interaction strength not content quantity drives amplification. Event-driven peaks reflect periods of concentrated attention and repeated interactions, whereas later trends suggest gradual stabilization and increased semantic richness.

5.8 Cross-Community Propagation Patterns

Extending the structural and temporal observations discussed in previous sections, cross-community interactions provide deeper insight into how misinformation transitions between localized clusters and broader network structures.

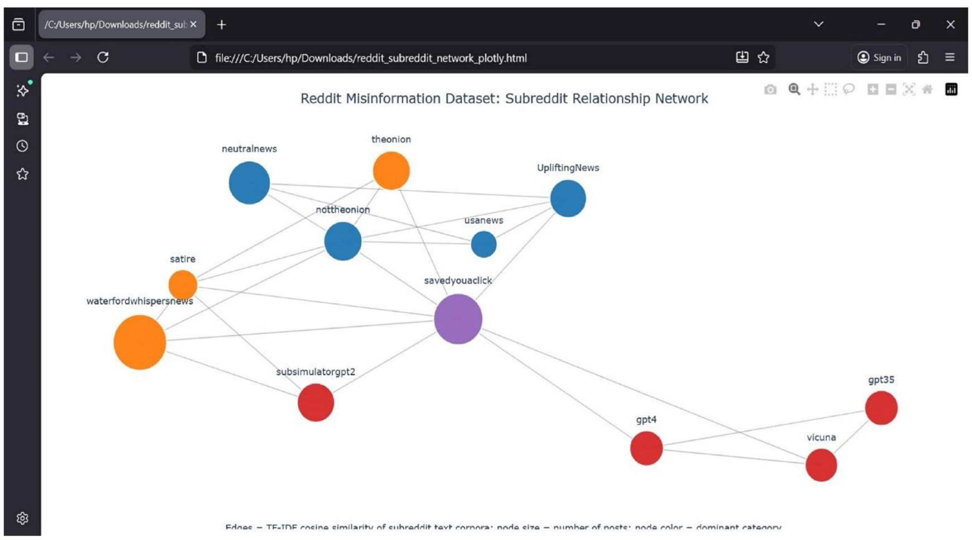

Figure 11. Reddit Misinformation Dataset: Subreddit Relationship Network

To visually validate the hybrid propagation mechanism, Figure 11 presents the subreddit relationship network constructed using similarity-based connections. The network reveals clearly defined clusters corresponding to topical or ideological communities, indicating strong intra-community cohesion. At the same time, a limited number of bridge nodes (e.g., aggregator or cross-posting communities) facilitate inter-community information flow. This structural pattern supports the coexistence of localized echo chambers and selective cross-community diffusion pathways.

This structure supports a two-stage propagation mechanism:

1. local amplification within tightly connected communities

2. controlled diffusion across communities via bridge nodes

edge density comparison is

$$D = \frac{2E}{N(N-1)} \quad (9)$$

Where:

- E: number of edges
- N: number of nodes

Statistical analysis confirms that intra-community edge density significantly exceeds inter-community density ($p<0.001$), reinforcing the coexistence of echo chamber effects and selective cross-community spread.

5.9 Semantic Network Structure

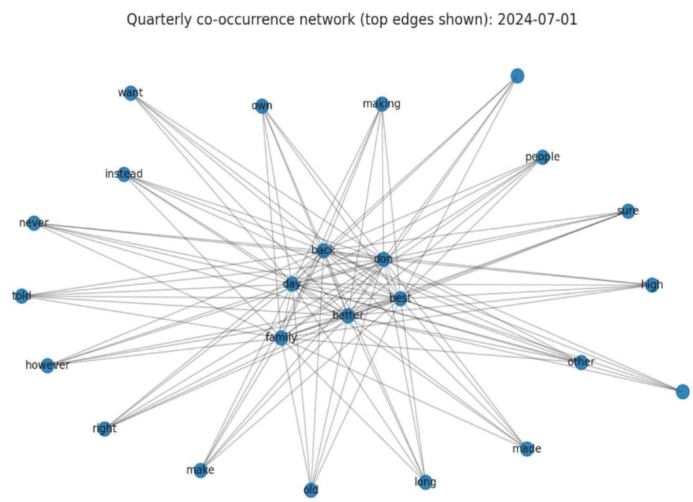

Figure 12. Quarterly co-occurrence network
(top edges)

The semantic co-occurrence network exhibits a hub and spoke topology, where a small number of high frequency terms dominate the structure. Central nodes reflect evaluative and conversational language, indicating that misinformation discourse is often opinion driven.

Peripheral nodes contribute contextual diversity without significantly altering the core structure. The presence of socially grounded terms suggests that misinformation is embedded within everyday narratives rather than isolated informational content.

Overall, the semantic network reflects a cohesive yet flexible structure, balancing dominant themes with contextual variation.

5.10 Correlation Structure of Network Metrics

Correlation analysis is computed by

$$r = \frac{\sum(x_i-\bar{x})(y_i-\bar{y})}{\sqrt{\sum(x_i-\bar{x})^2\sum(y_i-\bar{y})^2}} \quad (10)$$

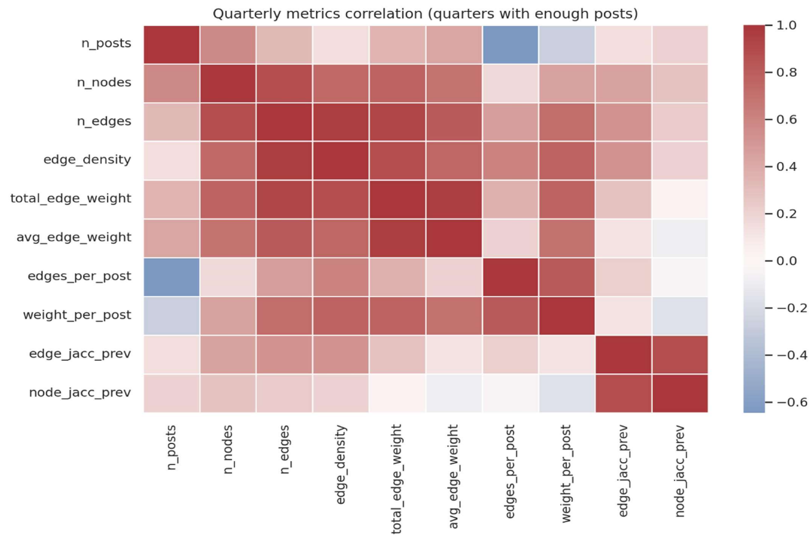

Figure 13. Correlation heatmap of quarterly network metrics

Correlation analysis reveals strong positive relationships among structural metrics, including nodes, edges, and density, indicating that larger networks tend to be more interconnected.

Normalized interaction measures (e.g., edges per post) exhibit internal consistency, while stability metrics show weak correlations with structural properties. A notable negative relationship between post volume and interaction intensity suggests that increased activity may dilute the strength of connectivity.

These findings highlight a distinction between network scale and interaction quality, emphasizing that growth does not necessarily imply stronger engagement.

5.11 Integrated Interpretation

The results of the multi-dimensional analysis collectively reveal that misinformation propagation is not governed by a single dominant mechanism, but rather emerges from the interaction of structural, temporal, and behavioral dynamics within online communities.

First, the concentration of misinformation in a limited number of subreddits highlights structural inequality in information production. This aligns with preferential attachment dynamics, where high-activity communities accumulate disproportionate visibility and influence. These hubs act as primary amplifiers, shaping the overall information landscape.

Second, the observed engagement credibility mismatch indicates a critical decoupling between popularity and reliability. The weak to moderate correlation between post score and upvote ratio suggests that visibility is influenced more by emotional resonance, controversy, or novelty than factual correctness. This reinforces the notion that misinformation thrives in attention-driven ecosystems, where engagement signals are not reliable indicators of truth.

Third, the network exhibits a dual nature of stability and dynamism. While node-level persistence confirms that core communities remain active over time, edge level volatility demonstrates that relationships between them are continuously reconfigured. This dynamic rewiring enables misinformation narratives to adapt, evolve, and re-emerge across contexts without requiring new participants.

Fourth, the densification-driven growth pattern suggests that misinformation spreads primarily through intensified intra network interactions rather than network expansion. This indicates that existing communities become increasingly interconnected, reinforcing internal narratives and strengthening echo chamber effects.

Fifth, the temporal analysis reveals a hybrid process combining baseline activity with event-driven bursts. These bursts are likely triggered by external socio-political events, leading to rapid surges in activity and interaction intensity. Such episodic amplification underscores the reactive nature of misinformation ecosystems.

Finally, the hybrid propagation mechanism characterized by strong intra-community cohesion and selective inter-community diffusion provides a unifying explanation for the observed patterns. Echo chambers enable sustained local reinforcement, while bridge nodes facilitate controlled cross-community transmission. This dual mechanism allows misinformation to maintain both depth (within communities) and breadth (across communities).

Taken together, these findings suggest that misinformation propagation is best understood as a self-reinforcing, temporally adaptive, and structurally embedded process, rather than a purely content-driven phenomenon.

5.12 Implications for Misinformation Mitigation

The insights derived from this study have important implications for the design of effective misinformation mitigation strategies. Traditional approaches that focus solely on content verification or post-level moderation are insufficient, as they overlook the networked and dynamic nature of misinformation spread.

A key implication is the need to target high centrality hub communities, which act as primary drivers of visibility. Interventions at these nodes such as content ranking adjustments, moderation prioritisation, or credibility labelling can significantly reduce large scale amplification.

Equally important is the identification and monitoring of bridge nodes, which enable cross-community diffusion. These nodes represent critical control points in the network, where misinformation transitions from localized clusters to broader audiences. Limiting the reach or influence of such nodes can effectively contain propagation across community boundaries.

The presence of an engagement credibility mismatch suggests that platforms should move beyond engagement-based ranking systems. Incorporating credibility-aware ranking mechanisms, such as weighting content by consensus or source reliability, can help reduce the visibility of misleading content that gains traction through polarization or sensationalism.

Temporal dynamics further imply that mitigation strategies must be adaptive and time-sensitive. Since misinformation spreads in bursts, real-time monitoring systems that detect sudden spikes in activity can enable early intervention, thereby reducing downstream amplification.

Additionally, the densification of interactions within communities highlights the importance of disrupting echo chambers. This can be achieved through exposure diversification strategies, such as recommending cross-cutting content or promoting interactions between heterogeneous communities.

Finally, the findings emphasize the need for multi-level intervention frameworks that integrate content, network structure, and temporal signals. Effective mitigation requires a shift from static moderation policies to dynamic, system-level approaches that account for how misinformation evolves and propagates over time.

6. Limitations and Future Work

Despite providing comprehensive insights into the dynamics of misinformation, this study has several limitations that should be acknowledged.

First, the dataset lacks explicit ground-truth labels for toxicity, intent, or factual accuracy, relying instead on category-based annotations (e.g., satire, misleading content). This may introduce label ambiguity and limit the precision of behavioral interpretations.

Second, the dataset is derived exclusively from Reddit, which introduces platform-specific bias. The structural and interaction patterns observed may not generalize to other platforms such as Twitter, Facebook, or emerging decentralized networks, where user behavior and algorithmic curation differ significantly.

Third, the analysis is constrained by temporal granularity, as aggregation at quarterly intervals may obscure fine-grained dynamics such as rapid cascades or short-lived viral events. Higher-resolution temporal data could provide deeper insights into micro-level propagation mechanisms.

Fourth, the study primarily relies on descriptive and statistical analysis, without incorporating predictive or causal modeling. While the findings reveal strong associations, they do not establish causal relationships between network properties and the spread of misinformation.

Fifth, the absence of user-level behavioral features (e.g., user credibility, interaction history) limits the ability to distinguish between organic and coordinated misinformation campaigns.

Future work can address these limitations through several extensions. The integration of Graph Neural Networks (GNNs) and Temporal Graph Networks (TGNs) can enable predictive modeling of misinformation propagation. Incorporating explainable AI techniques would further enhance interpretability, allowing researchers to identify key drivers of misinformation spread.

Additionally, expanding the analysis to cross-platform datasets can provide a more holistic understanding of misinformation ecosystems. The inclusion of fine-grained temporal data and user-level features would enable more detailed modeling of propagation dynamics and coordination patterns.

Finally, future research may explore intervention simulation frameworks in which mitigation strategies are tested in controlled environments to evaluate their effectiveness before real-world deployment.

6.1 Concluding Remarks

Building upon the empirical findings and their theoretical implications discussed above, the following concluding remarks synthesize the key contributions of this study. This study demonstrates that misinformation propagation is governed by a complex interplay of temporal bursts, structural persistence, and interaction dynamics. Rather than being driven solely by content volume or individual behaviour, misinformation emerges as a network-mediated phenomenon, shaped by the organisation and evolution of online communities.

The findings highlight that misinformation systems are inherently adaptive, combining stable participation structures with dynamically evolving interactions. This enables misinformation to persist over time while continuously adjusting to new contexts and events.

Moreover, the coexistence of echo chamber reinforcement and selective cross-community diffusion underscores the dual nature of misinformation spread localized yet scalable. This duality explains how misinformation can remain deeply embedded within communities while still achieving widespread reach.

Ultimately, the study emphasizes the need for a paradigm shift in misinformation research and mitigation from static, content-centric approaches to dynamic, network-aware frameworks. Understanding misinformation as a temporal and structural process is essential for developing effective, scalable, and sustainable solutions in the evolving digital information landscape.

Acknowledgement: This work is not funded by any grant.

Declaration: Conflict of Interest is none. The original complete dataset is deposited to the journal.

References

[1] Vosoughi, S., Roy, D., Aral, S. (2018). The spread of true and false news online. *Science* 359, 1146–1151.

[2] Ferrara, E., Cresci, S., Luceri, L. (2020). Misinformation, manipulation, and abuse on social media in the era of covid-19. *J. Comput. Soc. Sci*. 3, 271–277.

[3] Ma, Y., Qu, D., Wang, Y. (2026). Tracking evolving communities in fake news cascades using temporal graphs. *Sci Rep 16*, 4952.

[4] Ramya, S. P., Eswari, R. (2024). A regularisation-based simple shallow perceptron network for the detection of fake news in social networks. *Multimedia Tools and Applications* 83:77617-77637.

[5] Jiang, Y., Yu, X., Wang, Y., Xu, X., Song, X., Maynard, D. (2023). Similarity-aware multimodal prompt lear- ning for fake news detection. *Information Sciences* 647:119446.

[6] Su, X., Yang, J., Wu, J., Qiu, Z. (2024). Debunking fake news in online social networks without text analysis. *In: 2024 IEEE international conference on data mining (ICDM). Piscataway. IEEE*. 450-459.

[7] Wei, Wang; Yuanhui, Ma; Tao, Wu; Yang, Dai; Xingshu, Chen Lidia A., Braunstei. (2019). Containing

misinformation spreading in temporal social networks, *Chaos* 29, 123131.

[8] Dou, Y., Shu, K., Xia, C., Yu, P. S., Sun, L. (2021). User preference-aware fake news detection. In: Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval. Virtual event Canada. New York. ACM. 2051-2055.

[9] Shu, K., Wang, S., Liu, H. (2019). Beyond news contents: the role of social context for fake news detection. In: Proceedings of the twelfth ACM international conference on web search and data mining. WSDM '19. New York, NY, USA. *Association for Computing Machinery*. 312-320.

[10] Park, M., Chai, S. (2023). Constructing a user-centered fake news detection model by using classification algorithms in machine learning techniques. *IEEE Access* 11:71517-71527.

[11] Zhang, L., Zhang, X., Zhou, Z., Zhang, X., Wang, S., Yu, P. S., Li, C. (2025). Early detection of multimodal fake news via reinforced propagation path generation. *IEEE Transactions on Knowledge and Data Engineering* 37:613-625.

[12] Zhang, L., Zhang, X., Zhou, Z., Huang, F., L. i., C. (2024). Reinforced adaptive knowledge learning for multimodal fake news detection. *Proceedings of the AAAI Conference on Artificial Intelligence* 38:16777-16785.

[13] Zhu, K., Fan, C., Tao, J., Xue, J., Xie, H., Liu, X., Li, Y., Wen, Z, L. v., Z. (2024). Dual-view multimodal interaction in multimodal sentiment analysis. *In: 2024 IEEE international conference on multimedia and expo (ICME). Piscataway. IEEE*. 1-6.

[14] Wang, L., Zhang, C., Xu, H., Xu, Y., Xu, X., Wang, S. (2023). Cross-modal contrastive learning for multimodal fake news detection. *In: Proceedings of the 31st ACM international conference on multimedia. New York. ACM*. 5696-5704

[15] Luvembe, A. M., L. i., W, L. i., S, Liu, F, Wu, X. (2024). CAF-ODNN: complementary attention fusion with optimized deep neural network for multimodal fake news detection. *Information Processing & Management* 61:103653.

[16] Peng L, Jian S, Kan Z, Qiao L, Li D. (2024). a. Not all fake news is semantically similar: contextual semantic representation learning for multimodal fake news detection. *Information Processing Management* 61:103564.

[17] Hu, J., Zhang, J., L i Z. (2025). Tracing truth: dynamic temporal networks for multi-modal fake news detection. *PeerJ Computer Science* 11:e2998.

[18] Xue, J., Wang Y., Tian, Y., L. i., Y, Shi, L, Wei, L. (2021). Detecting fake news by exploring the consistency of multimodal data. *Information Processing Management* 58:102610.

[19] Yadav, A., Gupta, A. (2024). An emotion-driven, transformer-based network for multimodal fake news detection. *International Journal of Multimedia Information Retrieval* 13:7.

[20] Zhang, L. et al. (2025). Early Detection of Multimodal Fake News via Reinforced Propagation Path Generation, *IEEE Transactions on Knowledge and Data Engineering*, vol. 37 (2), p. 613-625, Feb.

[21] Sivasankari, S., Vadivu, G. (2022). Tracing the fake news propagation path using social network analysis. *Soft Comput* 26, 12883–12891.

[22] Jing, J., Li, F., Song, B., Zhang, K. K. R. Choo. (2023). Disinformation Propagation Trend Analysis and Identification Based on Social Situation Analytics and Multilevel Attention Network," *In IEEE Transactions on Computational Social Systems*, vol. 10 (2), p. 507-522, April.

[23] Zhang, Q., Cook, J., Yilmaz, E. (2021). Detecting and Forecasting Misinformation via Temporal and Geometric Propagation Patterns. In: Hiemstra, D., Moens, MF., Mothe, J., Perego, R., Potthast, M., Sebastiani, F. (eds) Advances in Information Retrieval. ECIR 2021. Lecture Notes in Computer Science(), vol 12657. *Springer, Cham*.

[24] Joan, Plepi, Flora, Sakketou, Henri-Jacques, Geiss, Lucie, Flek. (2022). Temporal Graph Analysis of Misinformation Spreaders in Social Media, *In: Proceedings of TextGraphs-16: Graph-based Methods for Natural Language Processing, Edited by Joan Plepi et al.* p.89-104.

[25] Stam, Caleb., Saldanha, Emily., Halappanavar, Mahantesh., Acharya, Anurag. (2025). Leveraging Language Modeling and Dynamic Social Network Analysis to recognize Patterns in the Spread of COVID-19 Misinformation, *In: PETRA '25: Proceedings of the 18th ACM International Conference on PErvasive Technologies Related to Assistive Environments*, Pages 479 - 485.

[26] Ghosh, S., Mitra, P., Nakov, P. (2024). Clock against Chaos: Dynamic Assessment and Temporal Intervention in Reducing Misinformation Propagation. *Proceedings of the International AAAI Conference on Web and Social Media,* 18(1), 462-473.

[27] Victor Ethan, Tabitha Dennis .(2024). Temporal Graph Neural Networks For Early Detection Of Coordinated Fake News Campaigns Across Social Media Platforms, https://www.researchgate.net/profile/Jerry-Samuel-3/publication/397822373.

[28] Sankar, A., Wu, Y., Gou, L., Zhang, W. Yang, H. (2020). Dysat: deep neural representation learning on dynamic graphs via self-attention networks. *In Proceedings of the 13th ACM International Conference on Web Search and Data Mining (WSDM)* 519–527.

[29] Xu, D., Ruan, C., Körpeoglu, E., Kumar, S. Achan, K. (2020). Inductive representation learning on temporal graphs. *In Proceedings of the 8th International Conference on Learning Representations* (ICLR 2020).

[30] Rossi, E. et al. (2020). Temporal graph networks for deep learning on dynamic graphs. arXiv preprint arXiv:2006.10637.

[31] Pareja, A. et al. (2020). Evolvegcn: Evolving graph convolutional networks for dynamic graphs. *In: Proceedings of the AAAI Conference on Artificial Intelligence,* vol. 34 5363–5370.

[32] Gupta, S., Kundu, S. (2025). Communities in streaming graphs: small space data structure, benchmark data generation, and linear algorithm. *ACM Trans. Knowl. Discov. Data* 19, 1–24 (2025).