MLRD: Multi-Level Dynamic Graph Modeling for Rumor Detection

Anonymous ACL submission

Abstract

Rumor detection on social media has become increasingly critical due to the rapid spread of misinformation and its detrimental societal impacts. Existing approaches often model rumor propagation as a static graph. However, this formulation is suboptimal, as rumors emerge and evolve dynamically. Some studies attempt to capture the temporal evolution of the propagation process, but they typically focus on a single level or a limited dimension of information, overlooking the synergistic interactions across multiple levels. To address these limitations, we propose MLRD, a novel **M**ulti-**L**evel dynamic graph model for **R**umor **D**etection. MLRD integrates collaborative multi-path propagation and view-augmented subgraph fusion. Specifically, collaborative multi-path propagation employs agents to traverse multiple propagation paths, leveraging the interactions among these paths to capture structural-level representations. View-augmented subgraph fusion learns a continuous node-to-subgraph aggregation via end-to-end optimization and adaptively emphasizes the most discriminative features to model feature-level representations. Extensive experiments on three real-world datasets demonstrate that MLRD consistently outperforms state-of-the-art baselines.

1 Introduction

Social media platforms now play a central role in daily communication and information exchange. The pervasive use of these platforms has resulted in an unprecedented volume of user-generated content being created and disseminated online. Nevertheless, this high level of openness also enables the rapid spread of rumors, which may compromise public trust and disrupt the integrity of social discourse. Accordingly, developing effective rumor detection methods for social media has become increasingly important for alleviating their negative societal consequences.

Figure 1: Single level representations.

Early rumor detection methods predominantly rely on traditional machine learning techniques (Castillo et al., 2011; Kwon et al., 2013). With the rapid development of deep learning, neural models (Ma et al., 2016; Liu and Wu, 2018) have been introduced to automatically learn representations from textual information. More recently, Graph Neural Networks (GNNs) have gained increasing attention in rumor detection (Bian et al., 2020; Tao et al., 2024), as they naturally model propagation graphs and effectively capture the structural characteristics of rumor diffusion. Building upon this paradigm, some studies address propagation graph modeling (Cui and Jia, 2024; Liu et al., 2023; Cui et al., 2025; Ma et al., 2022; Zhu et al., 2024), which have further advanced research on rumor detection. Meanwhile, researches on subgraph structures have also been extensively explored (Zhang et al., 2023, 2025b; Tao et al., 2024). Despite their effectiveness, these propagation modeling methods are predominantly based on static graph assumptions, implicitly treating rumor diffusion as time-invariant. Such formulations overlook the inherently dynamic nature of real-world rumor propagation, where both network structure and information semantics evolve over time.

To capture temporal dynamics, some studies model rumor propagation as sequences of graph snapshots at discrete time steps (Song et al., 2021; Sun et al., 2022a), focusing on evolving network structures and emphasize either diverse topological patterns or fine-grained temporal dynamics (Xu et al., 2024; Guo et al., 2025a). Alternatively, continuous-time dynamic graph approaches have been proposed to jointly model topological struc-

ture, content semantics, and temporal information (Song et al., 2022). (Chang et al., 2024) further introduce a news-driven dynamic network to capture interactions between news articles and associated entities. More recently, (Guo et al., 2025b) propose an adaptive sliding-window mechanism with memory-augmented attention to more effectively model dynamic diffusion processes. Nevertheless, existing dynamic graph–based rumor detection methods—whether operating on discrete snapshots or continuous-time graphs—primarily focus on a single granularity or single level representations, and thus fail to jointly capture both the multi-level representations from structural- and feature-perspective, limiting detection performance. As shown in Figure 1, existing studies predominantly focus on single-level features.

To address these limitations, we propose MLRD, a novel Multi-Level dynamic graph model for Rumor Detection. MLRD leverages collaborative multi-path propagation to capture structural-level representations and view-augmented subgraph fusion to obtain feature-level representations. Specifically, collaborative multi-path propagation uses agents to traverse multiple propagation paths. In generating these paths, we enforce temporal and legality constraints to ensure causal consistency and structural plausibility. During the coordination of multiple paths, interactions between node and agent representations are exploited to derive enhanced graph representations. Meanwhile, view-augmented subgraph fusion learns a continuous node-to-subgraph aggregation and adaptively highlights the most discriminative features. This module automatically decomposes the input graph into multiple latent subgraphs, enabling the model to capture localized structural patterns. The main contributions of this paper can be summarized as follows:

- We address dynamic rumor detection by modeling both structural and feature-level information, capturing the evolving patterns of social content.
- We propose collaborative multi-path propagation for structural representations and view-augmented subgraph fusion for feature representations.
- We conduct extensive experiments on three real-world datasets, demonstrating the effectiveness of our model for rumor detection.

2 Related Work

2.1 Rumor Detection

In recent years, rumor detection methods have received increasing attention. Among these, propagation-structure-based rumor detection methods have demonstrated remarkable effectiveness. Some works investigate the global graph structure, modeling it as a user-post graph (Nguyen et al., 2020; Dou et al., 2021; Min et al., 2022; Tian et al., 2022) or a propagation graph (Wei et al., 2021; Tao et al., 2024). Meanwhile, researchers have identified several critical challenges in this line of work, including structure imbalance and class imbalance issues (Cui and Jia, 2024), the design of pre-trained strategies (Liu et al., 2023; Cui et al., 2025), as well as robustness issues (Ma et al., 2022; Zhu et al., 2024). Some studies focus on subgraph structures, aiming to distill the most salient information from the full graph while eliminating redundant or irrelevant components. (Zhang et al., 2023) propose a subgraph generation strategy to construct counterfactual evidence for rumor detection. More recently, (Zhang et al., 2025b) introduce a reinforcement learning-based method that generates contextually coherent and informative propagation patterns.

However, real-world rumor diffusion networks are inherently dynamic. Early studies jointly model spatial and temporal structures to capture message propagation patterns for rumor classification (Huang et al., 2020). Some approaches represent each news propagation process as a sequence of graph snapshots recorded at discrete time steps (Song et al., 2021; Sun et al., 2022a). These methods further emphasize either diverse topological patterns or fine-grained temporal dynamics (Xu et al., 2024; Guo et al., 2025a). In contrast, (Song et al., 2022) model news propagation using continuous-time dynamic graphs, integrating topological structure, text semantics, and temporal information. (Chang et al., 2024) propose a news-driven dynamic heterogeneous network that captures the evolving interactions between news articles and associated entities. More recently, (Guo et al., 2025b) introduce an adaptive sliding window and memory-augmented attention mechanism to model the dynamic diffusion process more effectively. While prior works have investigated various dynamic graph learning strategies, they largely neglect to capture multi-level representations, which constrains detection performance.

2.2 Dynamic Graph Neural Networks

Modeling the temporal evolution of graph-structured data has garnered increasing attention in recent years. Early studies primarily aim to capture graph dynamics by explicitly evolving model parameters over time. (Pareja et al., 2020) leverage RNN to update GCN parameters across graph snapshots over time. Beyond parameter evolution, (You et al., 2022) treat node embeddings across different GNN layers as hierarchical node states and recurrently update them to capture temporal dependencies at multiple semantic levels. More recent works explore richer paradigms for temporal modeleing in dynamic graphs. (Qi et al., 2025) propose an input snapshot fusion framework for discrete-time dynamic graphs, enabling a more comprehensive characterization of temporal evolution patterns. (Zhang et al., 2025a) introduce the paradigm of graph rationalization for dynamic GNNs, which progressively disentangles causal subgraphs from environmental subgraphs across snapshots. (He et al., 2023) investigate continual graph learning and design a sequence of graph convolution operations to mitigate the challenge of over-stability. Complementing these general modeling efforts, (Han et al., 2025) focus on identifying temporal anomalous subgraphs in dynamic financial networks, highlighting the importance of temporal structural irregularities in practical applications. Different from the above works, our study focuses on modeling rumor dynamics in real-world temporal information by leveraging agent-based traversal and subgraph learning to capture both global and local patterns.

3 Methodology

3.1 Problem Definition

Rumor detection is formulated as a classification task. Formally, we represent a rumor detection dataset as $\mathcal{R} = \{R_1, R_2, \ldots, R_{|\mathcal{R}|}\}$, where R_i denotes the i-th event and $|\mathcal{R}|$ is the total number of events. Each event is represented as $R = \{r_1, r_2, \ldots, r_n, T, G\}$, where r_1 is the source post, $r_j (j > 1)$ denotes the j-th response post, and n is the number of posts associated with event R. The set $T = \{t_1, t_2, \ldots, t_n\}$ records the publishing timestamps, where post r_i is published at time t_i. The propagation structure of the event R is modeled as a graph $G = (V, E)$ rooted at the source post r_1. V is the set of nodes corresponding to posts, and E is the set of edges capturing response relationships between posts. The adjacency matrix is denoted by $A \in \{0, 1\}^{n \times n}$, where $A_{ij} = A_{ji} = 1$ if a response relationship exists between r_i and r_j, and $A_{ij} = A_{ji} = 0$ otherwise. We further define the node feature matrix as $X \in \mathbb{R}^{n \times d_x}$, where d_x is the dimensionality of the node embeddings. The objective of rumor detection is to learn a classifier $f : \mathcal{R} \rightarrow \mathcal{Y}$ that maps each event to a label in $\mathcal{R} \in \{F, T\}$, where F and T denote Rumor and Non-Rumor, respectively.

3.2 Overview

In this section, we propose a novel Multi-Level dynamic graph model for Rumor Detection, named MLRD. As illustrated in Figure 2, MLRD contains Collaborative Multi-Path Propagation and View-Augmented Subgraph Fusion.

3.3 Collaborative Multi-Path Propagation

The collaborative multi-path propagation module employs multiple agents to traverse temporally feasible paths and engage in collaborative interactions, thereby acquiring rich node representations. This design allows the model to selectively explore information-rich structural-level representations.

3.3.1 Initialization

To model dynamic evolution, we first obtain node and edge representations that capture propagation characteristics. Given a node v with feature vector $x_v \in X$, we first project the input feature into a shared latent space to obtain the initial node representation h_v.

$$h_v = W_x x_v \tag{1}$$

where W_x is learned projection matrix.

For each edge (u, v), we encode the temporal interval between the target node v and root node as $\delta_{uv} = t_v - t_1$. This interval is then transformed into an edge representation e_{uv}.

$$e_{uv} = \cos(W_e(-\log(1 + \delta_{uv}))) \tag{2}$$

where t_v and t_1 denote the timestamps of node v and root node. W_e is a learnable matrix. The logarithmic scaling combined with the cosine transformation further captures the temporal decay and periodicity of posts over time.

3.3.2 Multi-Step Agent Traversal

We design multi-step agent traversal to acquire multiple propagation paths, ensuring uniform coverage of the graph and promoting diversity across multi-path propagation. This design enables the model to

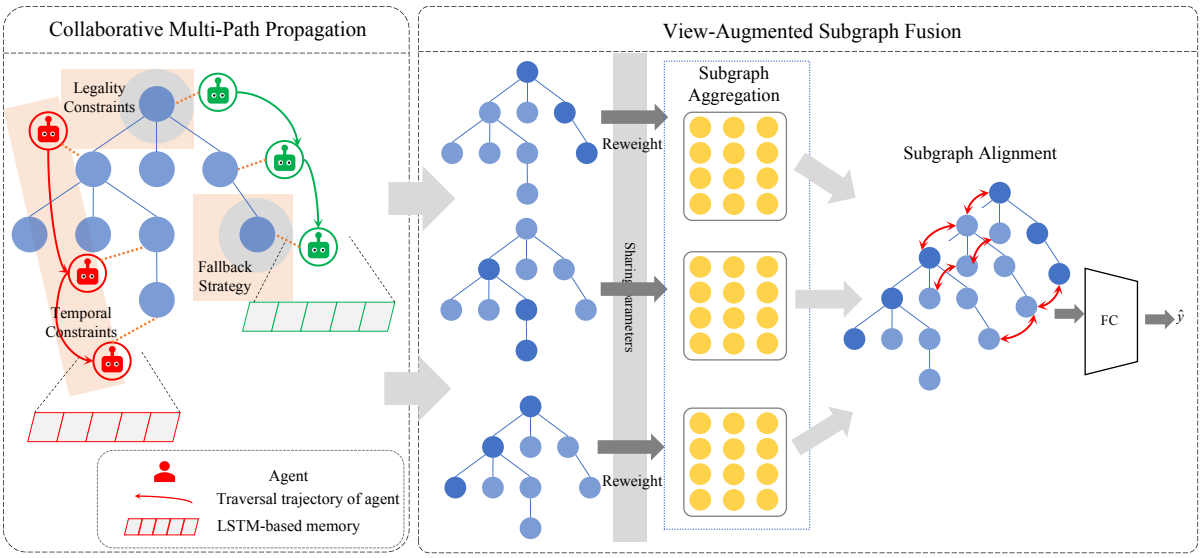

Figure 2: Overview of the proposed framework MLRD.

capture step-wise variations in the importance of nodes and edges.

Specifically, for each graph, we deploy $\mathbf{A}$ agents, each with a learnable embedding initialized as $h_{v_a}^{(0)} \in \mathbb{R}^d$. Each agent is randomly placed at node $v_a^{(0)}$ and traverses the graph for S steps, making movement decisions at each step via an attention mechanism under multiple constraints.

Attentive Selection. We compute structural compatibility scores between the agent a and candidate neighbors $\mathcal{N}(v_a^{(s)})$, which characterize the feasible choices for the agent's next position.

$$\mathbf{Q}_a = W_{\mathbf{Q}} h_{v_a}^{(s)} \tag{3}$$

$$\mathbf{K}_a = W_{\mathbf{K}} \sum_{j \in \mathcal{N}(v_a^{(s)})} [h_j^{(s)} \,||\, e_{aj} \,||\, h_{v_a}^{(s)}] \tag{4}$$

$$\alpha_{aj}^{(s)} = \frac{\langle \mathbf{Q}_a, \mathbf{K}_{aj} \rangle}{\sqrt{d}} \tag{5}$$

where $\mathbf{Q}_a$ and $\mathbf{K}_a$ are the queries and keys. $W_{\mathbf{Q}}$ and $W_{\mathbf{K}}$ are trainable parameter matrices. d is the dimension of queries and keys. $h_j^{(s)}$ denotes the neighbor embedding. e_{aj} denotes the edge representation of edge (v_a, v_j). $||$ is concatenation.

Temporal Constraints. To enforce temporal continuity and forward progression, we calculate the propagation time differences $\Delta t_{aj}^{(s)}$ between agent a and its neighboring nodes and apply temporal constraints accordingly.

$$\Delta t_{aj}^{(s)} = t_j^{(s)} - t_{v_a}^{(s)}, j \in \mathcal{N}(v_a^{(s)}) \tag{6}$$

Specifically, selections with $\Delta t_{aj}^{(s)} < 0$ are prohibited, while those satisfying $\Delta t_{aj}^{(s)} \geq 0$ are permitted. We define a binary *invalid time mask* $\mathbf{M}_t^{(s)}$ as a hard mask to explicitly encode the former constraint. For the latter case, we introduce a soft time bias $\Phi_b^{(s)}$ to penalize large forward time jumps.

$$\Phi_b^{(s)} = -\,||\,\mathrm{MLP}_t(\log(1 + \Delta t_{aj}^{(s)}))\,|| \tag{7}$$

We employ a multilayer perceptron (MLP) to encode time differences. Taking the negative absolute value ensures that the temporal penalty decays smoothly.

We then incorporate the soft time bias into the attention scores to model temporal dynamics.

$$\tilde{\alpha}_{aj}^{(s)} = \alpha_{aj}^{(s)} + \Phi_b^{(s)} \tag{8}$$

Legality constraints. To enforce temporal causality and structural validity, we introduce a legality constraint, referred to as the binary *root mask* $\mathbf{M}_r^{(s)}$. The root mask restricts each agent to visit the root node at most once during the traversal. This constraint is necessary because repeated access to the root node may introduce redundant information and degrade the quality of the learned representations.

We obtain the final attention scores between the agent and its neighbors as follows.

$$\hat{\alpha}_{aj}^{(s)} = \begin{cases} -\infty, & \mathbf{M}_t^{(s)} = 1 \\ -\infty, & \mathbf{M}_r^{(s)} = 0 \\ \tilde{\alpha}_{aj}^{(s)}, & otherwise \end{cases} \tag{9}$$

Note that if the S step has not been completed, all values in the attention scores are $-\infty$ indicating that the agent has reached a dead end. In this case, we employ a *fallback strategy*. Specifically, the agent preferentially selects short-range backward transitions based on timestamp quantiles. If no such transitions are available, the agent is forced to execute a stay action. This mechanism ensures robust traversal while preserving temporal constraints.

Agent Movement. Finally, the agent selects its next position from the neighboring nodes using Gumbel-Softmax.

$$v_{a^+}{}^{(s)} = \arg\max_j GumbelSoftmax(\hat{\alpha}_{aj}^{(s)}) \quad (10)$$

where $v_{a^+}{}^{(s)}$ denotes the selected neighbor, which corresponds to the agent's updated location at step s. We dynamically obtain the propagation paths of all agents, yielding the set of all neighbors of agents, denoted as $V_{a^+}{}^{(s)}$.

3.3.3 Multi-Path Collaborative Interaction

We enable collaborative interaction among the multiple obtained propagation paths by updating node states with information aggregated by the agents, while employing an LSTM to track the evolution of the agents' states.

Specifically, we first aggregate the incoming signals from all agents as injected representations for node updates. Then, for node $v \in V_{a^+}^{(s)}$, each agent constructs the injection for the node it occupies, enabling the node to update accordingly.

$$\hat{h}_v^{(s)} = [h_v^{(s)} \,||\, \text{AGG}(\{h_a^{(s)}\})] + \text{MLP}_z(z^{(s)}) \quad (11)$$

$$\tilde{h}_v^{(s)} = \text{LN}(h_v^{(s)} + \text{MLP}_h(\hat{h}_v^{(s)})) \quad (12)$$

where $\text{AGG}(\{h_a^{(s)}\})$ denotes the information aggregated by the agents. z^s denotes the current step embedding. LN denotes layer normalization used to stabilize training.

In addition, we apply an edge-conditioned message passing step over the node v to propagate neighborhood information. We obtain the updated node representations $h_v^{(s+1)}$.

$$\bar{h}_v^{(s)} = \tilde{h}_v^{(s)} \,||\, \text{AGG}(\{\tilde{h}_v^{(s)}, e_{uv}\}) \quad (13)$$

$$h_v^{(s+1)} = \text{LN}(\tilde{h}_v^{(s)} + \text{MLP}(\bar{h}_v^{(s)})) \quad (14)$$

We then employ an LSTM to model the agent's trajectory, updating agent representation by incorporating historical information. At each step, the agent updates its state conditioned on the newly visited node. We obtain the updated agent representation $h_a^{(s+1)}$.

$$\hat{h}_a^{(s)} = \text{MLP}([h_a^{(s)} \,||\, h_v^{(s+1)} \,||\, e_{av}] + z^{(s)}) \quad (15)$$

$$\tilde{h}_a^{(s+1)} = \text{LSTM}(\hat{h}_a^{(s)}) \quad (16)$$

$$h_a^{(s+1)} = h_a^{(s)} + \tilde{h}_a^{(s+1)} \quad (17)$$

After S steps, we aggregate the outcomes of multi-path collaborative interactions, establishing a robust and information-rich foundation for subsequent modules.

$$H_m = \frac{1}{S+1} \sum_{s=0}^{S} H^{(s)} \quad (18)$$

where $H^{(s)}$ denotes the node representations at step s. H_m denotes the enhanced representation of the entire graph after multi-path collaborative interactions.

3.4 View-Augmented Subgraph Fusion

Propagation graphs in rumor detection often exhibit rich and diverse local structures, where different subsets of nodes participate in distinct semantic patterns. Modeling this diversity using a single message-passing process can lead to representation entanglement and the loss of fine-grained feature information. To address this limitation, we propose a view-augmented subgraph fusion module that explicitly decomposes input graph into multiple latent subgraphs and learns feature-level representations from multiple perspectives.

3.4.1 Multi-View Subgraph Construction

To discover effective subgraph architectures, we incorporate a global graph representation on top of the enhanced graph obtained from the collaborative multi-path propagation module, providing reliable and comprehensive global support. Specifically, we first apply a GAT layer (Brody et al., 2021) to the original input graph, and then fuse H_m with H_g to obtain a global representation H. This fusion allows the downstream layers to leverage global enhanced semantics.

$$H_g = \text{GAT}(A, X) \quad (19)$$

$$H = \text{CONCAT}(H_m, H_g) \quad (20)$$

where CONCAT denotes the concatenation operation.

To obtain meaningful subgraph representations, we use GAT to learn a soft assignment over the node feature matrix H, producing K latent subgraphs. Specifically, each node is assigned to subgraphs with a learned probability.

$$P = \text{softmax}(\text{GAT}(A, H)), P \in \mathbb{R}^{n \times K} \quad (21)$$

where P_{ik} represents the probability that node i is associated with subgraph k. The soft assignment method used allows a node to participate in multiple subgraphs, which is crucial for modeling overlapping structural roles in real-world graphs.

3.4.2 Subgraph-Specific Feature Aggregation

Given the probability matrix P, we adaptively construct subgraph-specific contexts by reweighting the shared node representations. By assigning weights to nodes according to their subgraph-specific importance, the resulting features H_k highlight the nodes that are most relevant to target subgraph.

$$H_k = P_{:,k} \odot H, H_k \in \mathbb{R}^{n \times d} \quad (22)$$

where $\odot$ denotes element-wise multiplication.

For each latent subgraph, we perform normalized neighborhood aggregation independently:

$$\hat{H}_k(i) = \sum_{j \in \mathcal{N}(i)} \frac{1}{\sqrt{d_i d_j}} H_k(j) \quad (23)$$

where d_i denotes the in-degree of node i. This normalization ensures stable message propagation while enabling each subgraph to focus on distinct local connectivity patterns.

To maintain parameter efficiency and encourage cross-subgraph feature alignment, we apply a shared linear transformation across all subgraphs.

$$Z_k = W\hat{H}_k + b \quad (24)$$

where W and b are learnable parameters.

The subgraph-specific node representations are concatenated and projected into a unified latent space, integrating complementary information from multiple subgraphs.

$$Z = W[Z_1, Z_2, \ldots, Z_K] \quad (25)$$

3.4.3 Subgraph Alignment and Fusion

To obtain graph-level representations, we introduce a subgraph alignment pooling mechanism that explicitly models the contribution of each node to different latent subgraphs.

Statistics	TWITTER	DRWeibo	Weibo
# Events	1077	5998	4310
# Posts	60207	376659	816217
# Non-Rumors	564	3160	2187
# Rumors	513	2838	2123
# Avg. time length	416 Hours	1946 Hours	843 Hours

Table 1: Statistics of the datasets.

Specifically, we compute node-subgraph alignment scores using an additional GAT layer.

$$S = \text{softmax}(\text{GAT}_{align}(A, Z)) \quad (26)$$

where S_{ik} quantifies the importance of node i to subgraph k. This alignment mechanism allows the model to dynamically weight node contributions during graph-level aggregation.

For each subgraph k, we perform batch-wise weighted pooling.

$$g_k = \sum_{v \in V_b} S_{v,k} Z_v \quad (27)$$

where V_b denotes the set of nodes in the same graph within a mini-batch.

The final graph representation is obtained by concatenating all subgraph-level embeddings, which preserves subgraph semantics and yields a more expressive feature-level representation.

$$\mathbf{g} = [g_1, g_2, \ldots, g_K] \quad (28)$$

3.5 Training Objective

Classification. To calculate the labels of the rumors, we apply a fully connected layer followed by a softmax layer,

$$\hat{y} = \text{softmax}(W_f \mathbf{g} + b_f) \quad (29)$$

where $\hat{y}$ is the predicted probability distribution. W_f and b_f are weight and bias parameters respectively.

Loss. The aim of rumor detection is trained with the cross-entropy loss $\mathcal{L}$ of the predictions and ground truth labels.

$$\mathcal{L} = -\sum_{i}^{|\mathcal{Y}|} y_i log\hat{y}_i \quad (30)$$

where y_i denotes ground-truth label and $\hat{y}_i$ denotes the predicted probability distribution for the i-th event.

Method	Class	TWITTER				DRWeibo				Weibo			
		Acc.	Prec.	Rec.	F1	Acc.	Prec.	Rec.	F1	Acc.	Prec.	Rec.	F1
Bi-GCN	F	0.7591	0.7595	0.7419	0.7506	0.8118	0.8033	0.8031	0.8032	0.9103	0.9099	0.9110	0.9104
	T		0.7771	0.7732	0.7751		0.8256	0.8187	0.8221		0.9147	0.9109	0.9128
EBGCN	F	0.7539	0.7449	0.7428	0.7438	0.8212	0.8080	0.8116	0.8098	0.9166	0.9151	0.9157	0.9154
	T		0.7657	0.7640	0.7648		0.8333	0.8303	0.8318		0.9176	0.9173	0.9175
GACL	F	0.7609	0.7987	0.6781	0.7335	0.8541	0.8454	0.8463	0.8458	0.9367	0.9352	0.9366	0.9359
	T		0.7454	0.8377	0.7889		0.8628	0.8609	0.8618		0.9386	0.9369	0.9378
RDEA	F	0.7855	0.7942	0.7496	0.7713	0.8610	0.8564	0.8482	0.8523	0.9340	0.9294	0.9378	0.9336
	T		0.7857	0.8181	0.8016		0.8656	0.8723	0.8689		0.9393	0.9303	0.9348
DCE-RD	F	0.7796	0.7914	0.7579	0.7743	0.8527	0.8421	0.8485	0.8453	0.9303	0.9252	0.9339	0.9296
	T		0.7954	0.8010	0.7982		0.8623	0.8551	0.8587		0.9380	0.9270	0.9325
TrustRD	F	0.7695	0.7751	0.7359	0.7550	0.8500	0.8528	0.8278	0.8401	0.9312	0.9258	0.9355	0.9306
	T		0.7749	0.7974	0.7860		0.8494	0.8698	0.8595		0.9375	0.9266	0.9320
PSGT	F	0.8089	0.8148	0.7814	0.7977	0.8427	0.8342	0.8336	0.8339	0.9235	0.9171	0.9295	0.9233
	T		0.8141	0.8332	0.8235		0.8529	0.8510	0.8520		0.9315	0.9175	0.9244
RAGCL	F	0.8458	0.8444	0.8310	0.8376	0.8623	0.8576	0.8518	0.8547	0.9303	0.9218	0.9385	0.9301
	T		0.8486	0.8593	0.8539		0.8679	0.8720	0.8700		0.9396	0.9223	0.9308
DynGCN	F	0.7693	0.7647	0.7495	0.7570	0.8306	0.8197	0.8252	0.8225	0.9274	0.9120	0.9244	0.9182
	T		0.7759	0.7893	0.7825		0.8427	0.8354	0.8390		0.9203	0.9075	0.9138
FGDGNN	F	0.8408	0.8454	0.8181	0.8315	0.8516	0.8460	0.8402	0.8431	0.9406	0.9351	0.9451	0.9401
	T		0.8419	0.8617	0.8517		0.8581	0.8615	0.8598		0.9466	0.9361	0.9413
SWAM	F	0.8275	0.8356	0.7964	0.8156	0.8839	0.8792	0.8746	0.8769	0.9442	0.9407	0.9463	0.9435
	T		0.8258	0.8560	0.8406		0.8886	0.8922	0.8904		0.9479	0.9422	0.9450
MLRD	F	**0.8609**	**0.8623**	**0.8466**	**0.8544**	**0.8944**	**0.8937**	**0.8823**	**0.8880**	**0.9583**	**0.9520**	**0.9640**	**0.9580**
	T		**0.8641**	**0.8748**	**0.8694**		**0.8954**	**0.9063**	**0.9008**		**0.9646**	**0.9530**	**0.9588**

Table 2: Rumor detection results on TWITTER, DRWeibo and Weibo. Abbrev.: Rumor (F), Non-Rumor (T).

4 Experiments

4.1 Experimental Setup

Datasets. We conduct experiments on three publicly available real-world datasets, namely TWITTER (Lin et al., 2022), DRWeibo (Cui and Jia, 2024) and Weibo (Ma et al., 2016). TWITTER is collected from the Twitter platform and consists of English posts, whereas DRWeibo and Weibo are collected from the Weibo platform and contain Chinese posts. Each event in datasets contains source post and responses, as well as the propagation structure and timestamps. All datasets are labeled into two classes, Rumor (F) and Non-Rumor (T). Detailed dataset statistics are reported in Table 1.

Implementation Details. The proposed model is implemented in PyTorch and optimized using Adam. We employ BERT to encode post texts and train the model for 200 epochs with a learning rate 0.001 and batch size of 128. The hidden dimension of node representations is 128. The number of agents $\mathbf{A}$ is set to 3, 2, 2 for TWITTER, DRWeibo, and Weibo, and the number of steps S is set to 6, 6, and 10. Temperature parameter τ is set to 0.1. The number of sampled subgraphs K is set to 4. Consistent with (Bian et al., 2020), we report performance averaged over 10 runs of 5-fold cross-validation. Accuracy (Acc.), Precision (Prec.), Recall (Rec.), and F1-score (F1) are as standard metrics.

Comparison Models. We compare our proposed model with two categories of baselines: **(1) Static graph-based models**, including Bi-GCN (Bian et al., 2020), EBGCN (Wei et al., 2021), GACL (Sun et al., 2022b), RDEA (He et al., 2021), DCE-RD (Zhang et al., 2023), TrustRD (Liu et al., 2023), PSGT (Zhu et al., 2024) and RAGCL (Cui and Jia, 2024). These models investigate the roles of both graphs and subgraphs in rumor detection, while incorporating graph-based operations such as data augmentation. **(2) Dynamic graph-based models**. This category includes DynGCN (Choi et al., 2021), FGDGNN (Guo et al., 2025a) and SWAM (Guo et al., 2025b). These models designed for both discrete-time and continuous-time dynamic graphs, primarily exploring their effectiveness in rumor detection tasks.

4.2 Experimental Results

Table 2 reports the performance of different methods for rumor detection on the TWITTER, DRWeibo, and Weibo datasets. Compared with the baseline models, MLRD delivers the best results on all evaluation metrics. BiGCN and EBGCN model rumor diffusion by constructing both bidirectional propagation graphs and dispersion graphs. RDEA, TrustRD, and GACL leverage different data augmentation strategies in conjunction with contrastive learning on propagation graphs. RAGCL further

Method	TWITTER	DRWeibo	Weibo
	Acc.	Acc.	Acc.
MLRD	**0.8609**	**0.8944**	**0.9583**
w/o CMPP	0.8509	0.8855	0.9508
w/o VASF	0.8252	0.8567	0.9396
w/o dyn	0.8250	0.8525	0.9356

Table 3: Results of ablation study.

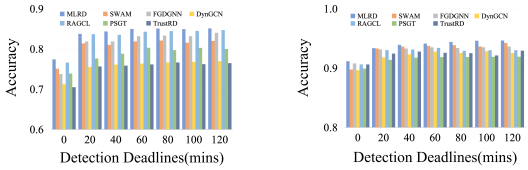

(a) TWITTER dataset (b) Weibo dataset

Figure 3: Results of early rumor detection.

investigates the wide-tree structure in rumor propagation. Beyond these approaches, DCE-RD focuses on subgraph-level method to construct causal subgraphs. PSGT aims to extract task-relevant attention subgraphs for classification.

Although these methods demonstrate strong performance, they largely overlook the inherently dynamic nature of rumor dissemination. DynGCN constructs a sequence of graph snapshots from the dynamic graph. However, it falls into a coarse-grained graph modeling paradigm. FGDGNN introduces a fine-grained framework by incorporating temporal edge information and a dynamic node update mechanism. Nevertheless, its snapshot construction still depends on predefined segments. SWAM adopts a continuous-time formulation to model sliding window. It leverages temporal characteristics of nodes to adapt to evolving rumor propagation patterns and maintains evolving states in a memorybank for iterative updates.

In contrast, MLRD employs collaborative multi-path propagation to capture structural-level representations. Multiple paths collaborative interactions between node and agent representations are exploited to derive enhanced graph representations. Simultaneously, the model leverages view-augmented subgraph fusion learns a continuous node-to-subgraph aggregation and adaptively highlights the most discriminative features, enabling the model to capture localized structural patterns.

4.3 Ablation Study

We validate the effectiveness of the MLRD framework from three perspectives, The variants are defined as follows: (1) w/o CMPP: removing Collaborative Multi-Path Propagation; (2) w/o VASF: View-Augmented Subgraph Fusion; (3) w/o dyn: instead of static graph. As shown in Table 3, removing any individual component of the MLRD framework results in a significant decline in overall performance, highlighting the importance of each component in the overall model architecture. These components jointly contribute to a more robust and generalizable framework.

4.4 Early Rumor Detection

This experiment investigates early-stage rumor detection on TWITTER and Weibo datasets, with the goal of enabling timely identification of emerging rumors. Following the protocol of (Sun et al., 2022b), we formalize this task by specifying a set of detection deadlines at $[0, 20, 40, 60, 80, 100, 120]$. Figure 3 presents the early detection performance of MLRD, in comparison with dynamic graph-based methods (SWAM, FGDGNN, and DynGCN) and static graph-based approaches (RAGCL, PSGT, and TrustRD) across varying deadlines. As the detection deadline increases, the performance of all models steadily improves. Notably, MLRD consistently surpasses all baseline methods at every deadline, demonstrating its superior effectiveness for early rumor detection.

5 Conclusion

In this paper, we propose MLRD, a novel Multi-Level dynamic graph model for Rumor Detection. MLRD leverages collaborative multi-path propagation to capture structural-level representations and view-augmented subgraph fusion to obtain feature-level representations. Collaborative multi-path propagation uses agents to traverse multiple propagation paths. We enforce temporal and legality constraints to ensure causal consistency and structural plausibility. During the coordination of multiple paths, interactions between node and agent representations are exploited to derive enhanced graph representations. Meanwhile, view-augmented subgraph fusion learns a continuous node-to-subgraph aggregation and adaptively highlights the most discriminative features. This module automatically decomposes the input graph into multiple latent subgraphs, enabling the model to capture localized structural patterns. Experiments on three public datasets demonstrate that MLRD outperforms the state-of-the-art baselines.

Limitations

The collaborative multi-path propagation mechanism introduces additional computational overhead, particularly when modeling large-scale dynamic propagation graphs with long temporal dependencies. Future work will focus on improving computational efficiency and enhancing the adaptability of the model to in the rumor detection settings.

References

Tian Bian, Xi Xiao, Tingyang Xu, Peilin Zhao, Wenbing Huang, Yu Rong, and Junzhou Huang. 2020. Rumor detection on social media with bi-directional graph convolutional networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 549–556.

Shaked Brody, Uri Alon, and Eran Yahav. 2021. How attentive are graph attention networks? *arXiv preprint arXiv:2105.14491*.

Carlos Castillo, Marcelo Mendoza, and Barbara Poblete. 2011. Information credibility on twitter. In *Proceedings of the 20th international conference on World wide web*, pages 675–684.

Ya-Ting Chang, Zhibo Hu, Xiaoyu Li, Shuiqiao Yang, Jiaojiao Jiang, and Nan Sun. 2024. Dihan: A novel dynamic hierarchical graph attention network for fake news detection. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pages 197–206.

Jiho Choi, Taewook Ko, Younhyuk Choi, Hyungho Byun, and Chong-kwon Kim. 2021. Dynamic graph convolutional networks with attention mechanism for rumor detection on social media. *Plos one*, 16(8):e0256039.

Chaoqun Cui and Caiyan Jia. 2024. Propagation tree is not deep: Adaptive graph contrastive learning approach for rumor detection. In *Proceedings of the AAAI Conference on artificial intelligence*, volume 38, pages 73–81.

Chaoqun Cui, Siyuan Li, Kunkun Ma, and Caiyan Jia. 2025. Enhancing rumor detection methods with propagation structure infused language model. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 7165–7179.

Yingtong Dou, Kai Shu, Congying Xia, Philip S Yu, and Lichao Sun. 2021. User preference-aware fake news detection. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, pages 2051–2055.

Mei Guo, Chen Chen, Chunyan Hou, Yike Wu, and Xiaojie Yuan. 2025a. Fgdgnn: Fine-grained dynamic graph neural network for rumor detection on social media. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 5676–5687.

Mei Guo, Chen Chen, Chunyan Hou, Yike Wu, and Xiaojie Yuan. 2025b. Swam: Adaptive sliding window and memory-augmented attention model for rumor detection. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 14430–14441.

Xiaolin Han, Yikun Zhang, Chenhao Ma, Lingyun Song, Reynold Cheng, and Xuequn Shang. 2025. Tempasd: Temporal anomalous subgraph discovery in large-scale dynamic financial networks. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pages 826–837.

Bowei He, Xu He, Yingxue Zhang, Ruiming Tang, and Chen Ma. 2023. Dynamically expandable graph convolution for streaming recommendation. In *Proceedings of the ACM Web Conference 2023*, pages 1457–1467.

Zhenyu He, Ce Li, Fan Zhou, and Yi Yang. 2021. Rumor detection on social media with event augmentations. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, pages 2020–2024.

Qi Huang, Chuan Zhou, Jia Wu, Luchen Liu, and Bin Wang. 2020. Deep spatial–temporal structure learning for rumor detection on twitter. *Neural Computing and Applications*, 35:12995–13005.

Sejeong Kwon, Meeyoung Cha, Kyomin Jung, Wei Chen, and Yajun Wang. 2013. Prominent features of rumor propagation in online social media. In *2013 IEEE 13th international conference on data mining*, pages 1103–1108. IEEE.

Hongzhan Lin, Jing Ma, Liangliang Chen, Zhiwei Yang, Mingfei Cheng, and Chen Guang. 2022. Detect rumors in microblog posts for low-resource domains via adversarial contrastive learning. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2543–2556.

Leyuan Liu, Junyi Chen, Zhangtao Cheng, Wenxin Tai, and Fan Zhou. 2023. Towards trustworthy rumor detection with interpretable graph structural learning. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 4089–4093.

Yang Liu and Yi-Fang Wu. 2018. Early detection of fake news on social media through propagation path classification with recurrent and convolutional networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.

Guanghui Ma, Chunming Hu, Ling Ge, Junfan Chen, Hong Zhang, and Richong Zhang. 2022. Towards robust false information detection on social networks with contrastive learning. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 1441–1450.

Jing Ma, Wei Gao, Prasenjit Mitra, Sejeong Kwon, Bernard J Jansen, Kam-Fai Wong, and Meeyoung

Cha. 2016. Detecting rumors from microblogs with recurrent neural networks.

Erxue Min, Yu Rong, Yatao Bian, Tingyang Xu, Peilin Zhao, Junzhou Huang, and Sophia Ananiadou. 2022. Divide-and-conquer: Post-user interaction network for fake news detection on social media. In *Proceedings of the ACM web conference 2022*, pages 1148–1158.

Van-Hoang Nguyen, Kazunari Sugiyama, Preslav Nakov, and Min-Yen Kan. 2020. Fang: Leveraging social context for fake news detection using graph representation. In *Proceedings of the 29th ACM international conference on information & knowledge management*, pages 1165–1174.

Aldo Pareja, Giacomo Domeniconi, Jie Chen, Tengfei Ma, Toyotaro Suzumura, Hiroki Kanezashi, Tim Kaler, Tao Schardl, and Charles Leiserson. 2020. Evolvegcn: Evolving graph convolutional networks for dynamic graphs. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 5363–5370.

QingGuo Qi, Hongyang Chen, Minhao Cheng, and Han Liu. 2025. Input snapshots fusion for scalable discrete-time dynamic graph neural networks. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1*, pages 1138–1149.

Chenguang Song, Kai Shu, and Bin Wu. 2021. Temporally evolving graph neural network for fake news detection. *Information Processing & Management*, 58(6):102712.

Chenguang Song, Yiyang Teng, Yangfu Zhu, Siqi Wei, and Bin Wu. 2022. Dynamic graph neural network for fake news detection. *Neurocomputing*, 505:362–374.

Mengzhu Sun, Xi Zhang, Jiaqi Zheng, and Guixiang Ma. 2022a. Ddgcn: Dual dynamic graph convolutional networks for rumor detection on social media. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pages 4611–4619.

Tiening Sun, Zhong Qian, Sujun Dong, Peifeng Li, and Qiaoming Zhu. 2022b. Rumor detection on social media with graph adversarial contrastive learning. In *Proceedings of the ACM Web Conference 2022*, pages 2789–2797.

Xiang Tao, Liang Wang, Qiang Liu, Shu Wu, and Liang Wang. 2024. Semantic evolvement enhanced graph autoencoder for rumor detection. In *Proceedings of the ACM Web Conference 2024*, pages 4150–4159.

Lin Tian, Xiuzhen Jenny Zhang, and Jey Han Lau. 2022. Duck: Rumour detection on social media by modelling user and comment propagation networks. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4939–4949.

Lingwei Wei, Dou Hu, Wei Zhou, Zhaojuan Yue, and Songlin Hu. 2021. Towards propagation uncertainty: Edge-enhanced Bayesian graph convolutional networks for rumor detection. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3845–3854, Online. Association for Computational Linguistics.

Shuai Xu, Jianqiu Xu, Shuo Yu, and Bohan Li. 2024. Identifying disinformation from online social media via dynamic modeling across propagation stages. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pages 2712–2721.

Jiaxuan You, Tianyu Du, and Jure Leskovec. 2022. Roland: graph learning framework for dynamic graphs. In *Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*, pages 2358–2366.

Guibin Zhang, Yiyan Qi, Ziyang Cheng, Yanwei Yue, Dawei Cheng, and Jian Guo. 2025a. Rationalizing and augmenting dynamic graph neural networks. In *International Conference on Representation Learning*, volume 2025, pages 80317–80337.

Kaiwei Zhang, Junchi Yu, Haichao Shi, Jian Liang, and Xiao-Yu Zhang. 2023. Rumor detection with diverse counterfactual evidence. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 3321–3331.

Yusong Zhang, Kun Xie, Xingyi Zhang, Xiangyu Dong, and Sibo Wang. 2025b. Rumor detection on social media with reinforcement learning-based key propagation graph generator. In *Proceedings of the ACM on Web Conference 2025*, pages 2742–2753.

Junyou Zhu, Chao Gao, Ze Yin, Xianghua Li, and Jürgen Kurths. 2024. Propagation structure-aware graph transformer for robust and interpretable fake news detection. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 4652–4663.