L-NCG: Narrative Consensus Graphs for Robust Stance Detection in Heterogeneous Social Networks

1st Xueer Wang*
School of Software Engineering
Xiangtan University
Hunan, China 761886059@qq.com

***Abstract*—Stance detection on social media is pivotal for understanding public opinion trends and mitigating misinformation. Existing approaches predominantly rely on GNNs to model social structures. However, these methods are often limited by the homophily assumption, indiscriminately aggregating features from heterogeneous interactions and introducing substantial structural noise. Conversely, while LLMs possess superior semantic reasoning capabilities, they struggle to capture global structural patterns due to computational constraints. To address these challenges, we propose the LLM-driven Narrative Consensus Graph (L-NCG) framework. This study posits that the essence of stance consistency lies in shared narrative frames rather than mere social connectivity. Specifically, we employ an LLM with a Probabilistic Chain-of-Thought mechanism to extract fine-grained narrative role distributions from unstructured text. These narrative features function as a semantic filter to reconstruct the graph topology: edges exhibiting high narrative consensus are reinforced, while those with conflicting logic or high uncertainty are suppressed via an entropy-based penalty mechanism. Finally, a GAT performs synergistic feature aggregation on this purified topology. Experimental results on multiple benchmark datasets demonstrate that L-NCG significantly outperforms state-of-the-art baselines, effectively mitigating structural noise and enhancing robustness in complex stance detection scenarios.**

***Index Terms*—Stance Detection; Social Media Analysis; Narrative Consistency; Large Language Models; Graph Neural Networks**

I. Introduction

As social media becomes the primary arena for opinion expression, identifying user attitudes toward specific targets, known as Stance Detection, has become key to public opinion analysis [1]. Unlike traditional sentiment analysis, stance expression in modern social media is often implicit and intricate, relying on subtle narrative frames, metaphors, and complex interactions with others [2], which requires detection systems to possess deep understanding capabilities beyond surface-level text processing.

Existing research on stance detection has explored both structural modeling and semantic reasoning. Graph Neural Network (GNN)-based approaches model social structures by assuming homophily, i.e., connected users tend to share similar views [3]. While effective in relatively coherent communities, this assumption becomes less reliable in realistic social media environments characterized by heterogeneous interactions [4]. For example, replies or quote retweets may express disagreement, sarcasm, or debate rather than endorsement, resulting in complex relational patterns that challenge structure-based aggregation [5]. In parallel, Large Language Models (LLMs) have demonstrated strong zero-shot semantic reasoning ability, enabling them to capture implicit stance expressed through narrative framing, irony, and metaphor [7]. However, existing LLM-based approaches typically operate at the instance level and do not explicitly model the global interaction structure, leaving collective propagation patterns in social networks underexplored [8].

From the perspective of stance detection, existing approaches still suffer from two fundamental limitations. The first limitation is that existing approaches generally underestimate the implicit and uncertain nature of stance expression on social media. As stance is commonly implied through narrative or contextual cues rather than explicitly stated, single texts often provide weak or ambiguous stance signals. Nevertheless, most existing methods rely on the assumption of explicitly identifiable stance cues, limiting the reliability of stance modeling under implicit expressions. The second limitation is that existing methods lack an explicit mechanism to incorporate narrative-level stance consistency as a constraint on information propagation. LLM-based approaches typically focus on local semantic inference, while graph-based models propagate along observed links without modeling narrative alignment during message passing. Consequently, in short, implicit, or adversarial interactions, stance predictions are vulnerable to the amplification of noisy or contradictory signals.

To address these limitations, we propose the LLM-driven Narrative Consensus Graph (L-NCG), a stance detection framework that explicitly models narrative-level consistency to guide information propagation in social networks. The core insight is that stance alignment arises from shared narrative roles rather than raw social links. We employ a large language model to transform unstructured user texts into probabilistic narrative role representations, capturing fine-grained semantic uncertainty. These representations provide a principled basis for reconstructing the social graph, selectively reinforcing connections with consistent narratives while suppressing conflicting or uncertain relations. By performing feature aggregation along this purified graph, the model effectively integrates LLM-based semantic reasoning with graph neural network propagation, bridging the gap between local instance-level inference and global structural modeling. This design enables robust identification of stance even under implicit, polarized, or noisy social interactions, resulting in consistently strong performance across multiple benchmark datasets.

We summarize our main contributions as follows:

- We propose a narrative-consensus-based graph construction framework that leverages probabilistic narrative role distributions inferred by large language models. By jointly modeling semantic consistency and reasoning uncertainty, the framework selectively preserves structurally reliable edges while suppressing inconsistent interactions.
- We further introduce a probability-driven structural denoising mechanism that bridges LLM reasoning and GNN propagation. Through a probabilistic CoT strategy and an uncertainty-aware edge weighting scheme, semantic confidence is transformed into computable graph structures suitable for global message passing.
- Extensive experiments demonstrate that the proposed L-NCG framework consistently outperforms state-of-the-art baselines, exhibiting superior robustness in complex and noisy stance detection scenarios.

II. Related Work

A. Stance Detection

Stance detection has evolved from early neural models to reasoning-driven paradigms. Initial approaches based on CNNs, RNNs, and attention mechanisms capture target-specific dependencies but rely on static representations and struggle with semantic ambiguity [10], [11]. Pre-trained language models (PLMs) introduce contextualized representations, further enhanced by contrastive learning, prompt tuning, and knowledge infusion to address data scarcity and implicit expressions [12]–[14], yet remain fundamentally discriminative and lack interpretability. More recently, large language models (LLMs) enable reasoning-centric stance inference through techniques such as chain-of-thought prompting and multi-agent debate [2], [7], [15]. However, these methods predominantly operate at the individual post level, overlooking the social structural signals essential for modeling collective stance dynamics.

B. Graph-based Stance Detection

Integrating social structures via GNNs addresses the isolation of individual posts. Traditional methods construct user-interaction graphs based on the homophily assumption, using GCNs or target-adaptive networks to aggregate features from neighbors [3], [16]. While effective in echo chambers, these models falter in heterogeneous real-world scenarios where interactions do not imply stance alignment [4]. Indiscriminate aggregation leads to over-smoothing [17]. To mitigate this, recent studies focus on structural denoising using signed networks [18] or attention mechanisms to filter irrelevant neighbors [19]. State-of-the-art approaches attempt to synergize semantics and structure via graph transformers [5] or knowledge graphs [20]. However, a critical gap remains: existing denoising methods typically operate on surface-level semantic similarity or explicit structural features, lacking the deep narrative understanding to identify logical stance alignment. This study bridges this gap by leveraging LLM-driven reasoning to guide structural reconstruction.

III. Methodology

A. Problem Formulation and Framework Overview

We formulate stance detection in a social network as a node classification problem. Let $G = (V, E)$ denote the observed interaction graph, where $V = \{u_1, u_2, \ldots, u_N\}$ is the set of N users and $E \subseteq V \times V$ is the edge set induced by social interactions. Each user u_i is associated with a text sequence T_i expressing the user's opinion toward a given target entity. Our goal is to learn a predictor f that maps the graph and user texts to stance labels, i.e., $f : (G, \{T_i\}_{i=1}^{N}) \to \mathcal{Y}$, producing $y_i \in \mathcal{Y}$ for each node, where $\mathcal{Y}$ is the predefined set of stance categories. In contrast to topology-driven methods that depend primarily on the noisy observed edges E, we aim to derive robust node representations by inferring latent narrative structures behind user expressions.

To mitigate structural noise from heterogeneous interactions, we introduce L-NCG (Figure 1), which integrates semantic reasoning with structural modeling in three stages. First, a probabilistic chain-of-thought LLM encodes each text T_i into a continuous narrative feature capturing semantic uncertainty. Second, narrative features guide topological denoising: we compute narrative-consensus weights between users and prune inconsistent links to obtain an enhanced graph. Finally, a Graph Attention Network aggregates high-level narrative features and low-level textual embeddings over the refined structure to produce robust stance predictions.

B. Zero-Shot Narrative Role Extraction via Chain-of-Thought

To capture stance rationales toward the target entity from unstructured user text, we first construct a narrative-frame semantic space, going beyond coarse sentiment-polarity classification. We define a mutually exclusive and collectively exhaustive set of narrative roles $R = \{r_1, r_2, \ldots, r_{|R|}\}$ that covers common qualitative frames in public discourse (e.g., *victim*, *villain*, *hero*, and *perpetrator*).

LLMs may exhibit overconfidence or hallucinations in zero-shot reasoning. To mitigate this issue, we introduce Probabilistic CoT Prompting, which provides structured instructions to elicit a reasoning trace consisting of (i) focus identification, (ii) role mapping, and (iii) uncertainty assessment. The final output is parsed into a narrative distribution vector $\mathbf{p}_i \in \mathbb{R}^{|R|}$ for user u_i, satisfying $\sum_{k=1}^{|R|} \mathbf{p}_{i,k} = 1$. Each component $\mathbf{p}_{i,k}$ quantifies the confidence that u_i frames the target entity as role r_k. By converting fine-grained reasoning into computable continuous features, this module provides the quantitative basis for narrative-consensus estimation and subsequent graph reconstruction.

C. Construction of Narrative Consensus Graph

The original social network graph is denoted as $G = (V, E)$, where V represents the node set corresponding to users, and E represents the edge set established through simple

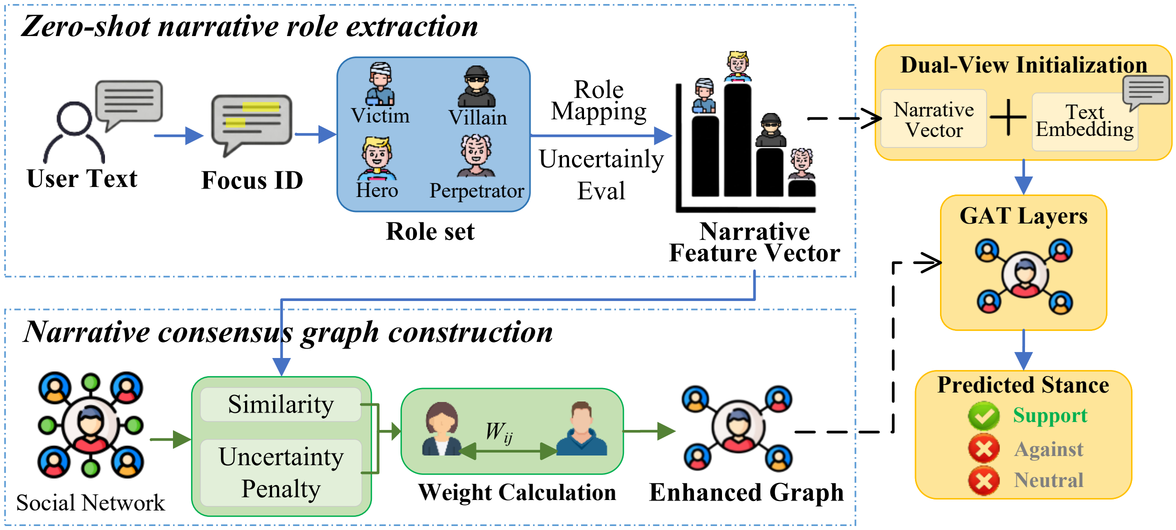

Fig. 1: The overall architecture of the proposed the L-NCG. The process evolves through three stages: extracting probabilistic narrative role vectors from unstructured text via Zero-Shot CoT prompting, reconstructing the social topology into an enhanced graph based on narrative consensus weights, and finally performing synergistic feature aggregation using a GAT for robust stance detection.

interactions. These connections frequently contain noise such as hostile interactions or irrelevant information. To address this issue, this study utilizes the narrative feature vectors $\{\mathbf{p}_i\}$ extracted in the preceding section to recompute the weights of interactions. Consequently, we construct an enhanced graph $G' = (V, E', \mathbf{W}')$. In this definition $E' \subseteq E$ denotes the set of effective edges retained after narrative filtering, while $\mathbf{W}'$ represents the weighted adjacency matrix.

1) Calculation of Narrative Consensus Weight: This study designs a composite weighting mechanism to re-evaluate social connections from the dimensions of semantic consistency and narrative certainty.

Semantic Similarity Measurement. First, we quantify the directional consistency of connected users when constructing narrative frames. Based on the narrative feature vectors, cosine similarity is adopted to calculate the basic consistency score sim_{ij} between user u_i and user u_j:

$$sim_{ij} = \frac{\mathbf{p}_i \cdot \mathbf{p}_j}{|\mathbf{p}_i||\mathbf{p}_j|} \tag{1}$$

Here, $sim_{ij} \in [-1, 1]$. When sim_{ij} approaches 1, it indicates that both parties possess highly consistent cognition regarding the narrative role of the target entity.

Uncertainty Penalty Mechanism. Considering that the probability distribution output by LLMs during zero-shot reasoning may contain uncertainty, relying solely on similarity may introduce misleading signals. Therefore, this study introduces an uncertainty penalty factor. First, Shannon entropy measures the disorder of the narrative vector $\mathbf{p}_i$ for user u_i, from which the Narrative Certainty c_i is defined as:

$$c_i = 1 - \frac{H(\mathbf{p}_i)}{\log |R|} = 1 - \frac{-\sum_{k=1}^{|R|} p_{i,k} \log p_{i,k}}{\log |R|} \tag{2}$$

Here, $\log |R|$ serves as the maximum entropy normalization term. A higher c_i indicates that the model is more certain about the narrative judgment of user u_i. Based on this, the edge-level penalty factor γ_{ij} is defined:

$$\gamma_{ij} = \min(c_i, c_j) \tag{3}$$

As long as the narrative feature of either party in the connection exhibits high entropy, the weight of that edge will be significantly suppressed.

Final Weight Generation. Integrating semantic similarity and the uncertainty penalty, the final narrative consensus weight w_{ij} is defined as follows:

$$w_{ij} = \sigma(\alpha \cdot sim_{ij}) \cdot \gamma_{ij} \tag{4}$$

In this formula, $\sigma(\cdot)$ denotes the Sigmoid function. The symbol α represents a scaling hyperparameter used to adjust the sensitivity of the model to the similarity score. This formula not only reinforces connections with consistent stances but also implements automatic down-weighting for low-confidence reasoning through the penalty factor γ_{ij}.

2) Topology Denoising and Sparsification: To remove weak and noisy edges, we sparsify the graph by thresholding the reweighted adjacency matrix $\mathbf{W}'$ at τ, defined as:

$$\mathbf{W}'ij = \begin{cases} wij, & \text{if } w_{ij} \geq \tau \\ 0, & \text{if } w_{ij} < \tau \end{cases} \tag{5}$$

Connections with weights below τ are deemed logically invalid and removed. This transformation converts the noisy graph G into a narrative-consensus-based enhanced graph G', retaining only logically self-consistent edges to provide a purified topological foundation for subsequent feature aggregation.

D. Synergistic Feature Aggregation and Stance Classification

We fuse local semantic features with global structure on the enhanced graph G' using dual-view initialization and hierarchical graph attention.

1) Dual-View Node Initialization: For any user node u_i, we construct an initial feature vector $\mathbf{h}_i^{(0)}$ by concatenating the high-level narrative vector $\mathbf{p}_i \in \mathbb{R}^{|R|}$ and the low-level textual embedding $\mathbf{t}_i \in \mathbb{R}^{d_{text}}$ (encoded by a lightweight model):

$$\mathbf{h}_i^{(0)} = [\mathbf{p}_i \parallel \mathbf{t}_i] \tag{6}$$

where $\parallel$ denotes concatenation, and $\mathbf{h}_i^{(0)} \in \mathbb{R}^{|R|+d_{text}}$. This strategy ensures complementarity between LLM reasoning and lexical semantics.

2) Hierarchical Aggregation on Denoised Graph: We adopt a GAT to execute message passing strictly on the denoised graph G'. Information propagates only along logical links with high narrative consensus ($w_{ij} \geq \tau$), blocking noise from hostile interactions. In the l-th layer, node u_i updates its representation as:

$$\mathbf{h}_i^{(l+1)} = \sigma\left(\sum_{j \in \mathcal{N}'(i)} \alpha_{ij}^{(l)} \mathbf{W}^{(l)} \mathbf{h}_j^{(l)}\right) \tag{7}$$

where $\mathcal{N}'(i)$ is the neighbor set in G', $\mathbf{W}^{(l)}$ is the transform matrix, $\sigma(\cdot)$ is the activation function, and $\alpha_{ij}^{(l)}$ is the attention coefficient calculated on effective edges. After L layers, we obtain the final representation $\mathbf{h}_i^{(L)}$.

3) Prediction and Optimization: The representation $\mathbf{h}_i^{(L)}$ is passed through an MLP and Softmax to predict stance probabilities $\hat{\mathbf{y}}_i \in \mathbb{R}^{|C|}$:

$$\hat{\mathbf{y}}_i = \text{Softmax}(\text{MLP}(\mathbf{h}_i^{(L)})) \tag{8}$$

We optimize the model using Cross-Entropy Loss with L_2 regularization:

$$\mathcal{L} = -\frac{1}{|\mathcal{V}_{train}|} \sum_{i \in \mathcal{V}_{train}} \sum_{c=1}^{|C|} y_{i,c} \log(\hat{y}_{i,c}) + \lambda \|\Theta\|_2^2 \tag{9}$$

where $\mathcal{V}_{train}$ is the training set, $y_{i,c}$ is the ground truth, and λ controls regularization.

IV. Experiment

A. Datasets and Baselines

We evaluate our model on three widely used stance detection benchmarks spanning general and political domains: SemEval-2016 Task 6 [1], a multi-class dataset with FAVOR, AGAINST, and NONE labels across five targets. P-Stance [21], a large-scale political dataset with binary labels, and COVID-19-Stance [22], which focuses on COVID-19-related discussions with three stance categories. We follow the official train/test splits for all datasets and randomly sample 10% of the training data as a validation set for hyperparameter tuning, reporting final results on the test sets. We compare our approach with six representative stance detection baselines, covering neural models, PLM-based methods, and recent LLM-based prompting frameworks. These include BiLSTM-Att [10], a target-aware BiLSTM with attention; BERT [23], fine-tuned with target–tweet concatenation; Stanceformer [24], a target-aware transformer; CoSD [25], a topic-aware method with graph-based modeling; Tree-of-Counterfactual Prompting [26], an LLM-based prompting approach with counterfactual reasoning; and MPVStance [27], an LLM-based framework with retrieval-enhanced multi-perspective verification.

B. Experimental Settings

For SemEval-2016 Task 6 and P-Stance, we follow the official splits and sample 10% of the training data as validation for hyperparameter tuning and early stopping. All inputs are constructed in a target-aware manner by concatenating the target and tweet. Our L-NCG employs a fixed LLM (GPT-4o) in a zero-shot setting to derive narrative representations via chain-of-thought prompting, which are then used to construct and refine an interaction graph with a graph attention network. All trainable parameters are optimized using Adam with early stopping on validation Macro-F1. Experiments are conducted on two NVIDIA RTX A5000 GPUs. For fair comparison, all baselines adopt the same input format when applicable; implementation details follow their original settings, including supervised training for BiLSTM-Att and BERT, and prompt-based evaluation for LLM baselines without task-specific fine-tuning. We report Accuracy and Macro-F1 on the test sets, using the validation set solely for model selection.

C. Main Results

L-NCG achieves the highest Accuracy and Macro-F1 across SemEval-2016 Task 6, COVID-19-Stance, and P-Stance, as shown in Table I, demonstrating strong generalization across topics and domains. Compared to text-only baselines like BiLSTM-Att and BERT, L-NCG significantly benefits from relational signals beyond isolated textual representations, especially for implicit stances. While advanced baselines exhibit complementary strengths, they also have clear limitations: target-aware encoders such as Stanceformer remain constrained by local textual modeling, graph-based methods like CoSD are sensitive to structural noise, and prompt-based LLM approaches, including Tree-of-Counterfactual Prompting and MPVStance, excel at instance-level reasoning but lack graph-level propagation. In contrast, L-NCG unifies LLM-derived narrative representations with graph propagation via narrative consensus, yielding consistent improvements across all datasets, notably SemEval-2016 and COVID-19-Stance.

D. Ablation Study

We conduct ablation studies to assess the contribution of key components in L-NCG by removing or modifying individual modules while keeping the rest of the architecture unchanged. All variants are evaluated under the same experimental settings as the full model. Table II reports the results on SemEval-2016 and COVID-19-Stance. Removing the narrative role extraction

Method	SemEval-2016		COVID-19-Stance		P-Stance	
	Acc.	Macro-F1	Acc.	Macro-F1	Acc.	Macro-F1
BiLSTM-Att	64.32 ± 0.36	60.09 ± 0.25	66.43 ± 0.20	63.31 ± 0.34	71.52 ± 0.48	70.18 ± 0.36
BERT	69.75 ± 0.41	66.38 ± 0.52	71.93 ± 0.72	68.82 ± 0.55	78.55 ± 0.22	77.88 ± 0.47
Stanceformer	71.22 ± 0.79	68.45 ± 0.36	73.33 ± 0.23	70.63 ± 0.34	80.08 ± 0.71	79.44 ± 0.14
CoSD	70.52 ± 0.39	67.92 ± 0.11	72.81 ± 0.19	69.99 ± 0.22	79.58 ± 0.14	78.81 ± 0.60
Tree-of-CP	68.95 ± 0.75	65.76 ± 0.12	71.08 ± 0.84	67.83 ± 0.35	77.75 ± 0.53	77.13 ± 0.80
MPVStance	70.06 ± 0.37	67.25 ± 0.52	72.43 ± 0.44	69.59 ± 0.53	79.23 ± 0.83	78.46 ± 0.33
L-NCG (Ours)	**73.61 ± 0.35**	**71.24 ± 0.44**	**85.92 ± 0.20**	**83.50 ± 0.09**	**82.41 ± 0.46**	**81.71 ± 0.10**

TABLE I: Main results on SemEval-2016 Task 6, COVID-19-Stance, and P-Stance. Accuracy (Acc.) and Macro-F1 are reported as *mean* ± *std* over five runs. Best results are highlighted in bold.

Variant	SemEval-2016		COVID-19-Stance		P-Stance	
	Acc.	Macro-F1	Acc.	Macro-F1	Acc.	Macro-F1
L-NCG (Full Model)	**73.61**	**71.24**	**85.92**	**83.50**	**82.41**	**81.71**
w/o Narrative Roles	69.85	66.92	79.86	76.94	79.32	78.41
w/o Uncertainty Penalty	71.10	68.47	83.74	81.36	81.02	80.22
w/o Topology Denoising	70.62	67.95	82.61	80.08	80.34	79.52

TABLE II: Ablation results of L-NCG on SemEval-2016 Task 6, COVID-19-Stance, and P-Stance.

Method Variant	SemEval-2016	COVID-19-Stance	P-Stance
Text + Original Graph	66.38	79.92	77.88
Narrative Semantics Only	68.74	81.36	79.26
Structure Denoising Only	69.32	82.08	79.88
L-NCG (Full Model)	**71.24**	**83.50**	**81.71**

TABLE III: Disentangling the effects of narrative semantics and structural denoising (Macro-F1).

module results in the most significant performance drop on both datasets, indicating the central role of LLM-derived narrative representations. Disabling the uncertainty penalty or topology denoising also consistently degrades performance, demonstrating the importance of uncertainty-aware weighting and noise suppression in graph construction. Overall, the results suggest that narrative role modeling, uncertainty-aware edge weighting, and topology denoising contribute in a complementary manner to robust stance detection.

E. Disentangling Narrative Semantics and Noise

Table III presents an ablation study that disentangles narrative semantic modeling and structural denoising on SemEval-2016, COVID-19-Stance, and P-Stance. Specifically, we evaluate variants that (i) incorporate narrative semantics while keeping the original graph unchanged, (ii) perform narrative-driven topology denoising without fully using narrative semantics in node representations, and (iii) the full L-NCG that integrates both components under the same settings. Each component provides consistent gains when applied alone. Narrative semantics improves over text-based variants, and topology denoising further benefits performance by refining propagation links. Combining the two yields the best results on all datasets, suggesting that narrative semantics and structural denoising are complementary. The effect is most pronounced on COVID-19-Stance, where polarized and heterogeneous interactions intensify semantic ambiguity and structural noise.

Method	Implicit Stance	High-Conflict Context	Overall
BERT	61.84	63.27	66.40
CoSD	63.95	65.88	67.90
MPVStance	64.21	66.14	67.30
L-NCG (Ours)	**68.47**	**70.02**	**71.24**

TABLE IV: Performance comparison under implicit stance expressions and high-conflict interaction contexts (Macro-F1 on SemEval-2016).

F. Implicit and High-Conflict Scenario Analysis

To evaluate robustness under challenging conditions, we test all methods on two subsets featuring implicit stance expressions and high-conflict interaction contexts, as summarized in Table IV. L-NCG achieves the best performance in both settings. Under implicit stance, text-only models such as BERT degrade markedly due to their reliance on surface lexical cues. CoSD and MPVStance partially benefit from propagation or stronger reasoning, yet remain limited by noisy social links and instance-level inference, respectively. In contrast, L-NCG uses narrative consensus to capture deeper framing signals and reconstructs the topology to suppress misleading connections, enabling robust predictions in both implicit and high-conflict scenarios. These results suggest that L-NCG is better suited to real-world stance detection where cues are subtle and interactions are noisy.

G. Sensitivity Analysis of Narrative Consensus

We evaluate L-NCG's sensitivity to the similarity scaling factor α and threshold τ on the SemEval-2016, COVID-19-Stance, and P-Stance datasets, as shown in Figure 2. Performance remains stable across a broad range of values, with moderate settings achieving the highest Macro-F1 and extreme settings causing only minor degradation. Sensitivity is slightly higher on COVID-19-Stance due to polarized interactions, yet

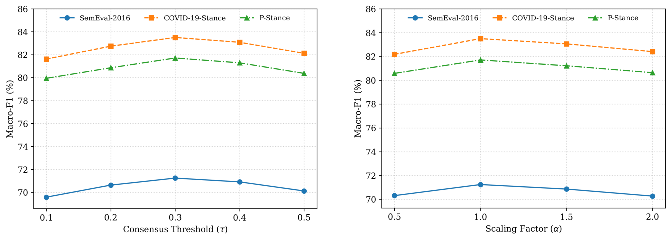

(a) Sensitivity to τ. (b) Sensitivity to α.

Fig. 2: Sensitivity analysis of L-NCG across three datasets.

variations remain limited. These results indicate that L-NCG is robust to different narrative consensus strengths.

V. Conclusion

In this paper, we present the **L**LM-driven **N**arrative **C**onsensus **G**raph (L-NCG), a novel framework for social media stance detection that explicitly addresses structural noise and semantic ambiguity. By relaxing the traditional homophily assumption, L-NCG shifts the modeling focus from surface-level social connectivity to narrative-level consensus, enabling more reliable information propagation under implicit, heterogeneous, and noisy social interactions. By integrating probabilistic narrative reasoning from large language models with graph-based message passing, the proposed framework effectively bridges local semantic inference and global structural modeling, leading to consistent performance gains across diverse benchmark datasets. Future work will investigate efficient distillation of LLM-driven narrative reasoning and temporal extensions of the narrative consensus framework.

References

[1] S. Mohammad, S. Kiritchenko, P. Sobhani, X. Zhu, and C. Cherry, "Semeval-2016 task 6: Detecting stance in tweets," in *Proceedings of the 10th international workshop on semantic evaluation (SemEval-2016)*, 2016, pp. 31–41.

[2] Y. Ji, J. Ning, Y. Zhang, Z. Liu, and H. Lin, "Llm-driven implicit target augmentation and fine-grained contextual modeling for zero-shot and few-shot stance detection," in *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 2025, pp. 5883–5895.

[3] B. Liang, Y. Fu, L. Gui, M. Yang, J. Du, Y. He, and R. Xu, "Target-adaptive graph for cross-target stance detection," in *Proceedings of the web conference 2021*, 2021, pp. 3453–3464.

[4] P. Panayotov, U. Shukla, H. T. Sencar, M. Nabeel, and P. Nakov, "Greener: Graph neural networks for news media profiling," *arXiv preprint arXiv:2211.05533*, 2022.

[5] X. Chen, B. Liu, H. Hu, Y. Cai, M. Guo, and X. Ma, "Integrating graph neural networks and large language models for stance detection via heterogeneous stance networks," *Applied Sciences*, vol. 15, no. 11, p. 5809, 2025.

[6] B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal *et al.*, "Language models are few-shot learners," *arXiv preprint arXiv:2005.14165*, vol. 1, no. 3, p. 3, 2020.

[7] X. Lan, C. Gao, D. Jin, and Y. Li, "Stance detection with collaborative role-infused llm-based agents," in *Proceedings of the international AAAI conference on web and social media*, vol. 18, 2024, pp. 891–903.

[8] R. Zhang, Y. Tian, P. Wei, D. D. Zeng, and W. Mao, "An llm-enabled knowledge elicitation and retrieval framework for zero-shot cross-lingual stance identification," in *Findings of the Association for Computational Linguistics: EMNLP 2024*, 2024, pp. 12 253–12 266.

[9] I. J. Cruickshank and L. H. X. Ng, "Prompting and fine-tuning open-sourced large language models for stance classification," *arXiv preprint arXiv:2309.13734*, 2023.

[10] I. Augenstein, T. Rocktäschel, A. Vlachos, and K. Bontcheva, "Stance detection with bidirectional conditional encoding," *arXiv preprint arXiv:1606.05464*, 2016.

[11] J. Du, R. Xu, Y. He, and L. Gui, "Stance classification with target-specific neural attention networks," in *26th International Joint Conference on Artificial Intelligence, IJCAI 2017*. International Joint Conferences on Artificial Intelligence, 2017, pp. 3988–3994.

[12] B. Liang, Z. Chen, L. Gui, Y. He, M. Yang, and R. Xu, "Zero-shot stance detection via contrastive learning," in *Proceedings of the ACM web conference 2022*, 2022, pp. 2738–2747.

[13] M. Hardalov, A. Arora, P. Nakov, and I. Augenstein, "Few-shot cross-lingual stance detection with sentiment-based pre-training," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 10, 2022, pp. 10 729–10 737.

[14] Z. He, N. Mokhberian, and K. Lerman, "Infusing knowledge from wikipedia to enhance stance detection," *arXiv preprint arXiv:2204.03839*, 2022.

[15] B. Zhang, X. Fu, D. Ding, H. Huang, G. Dai, N. Yin, Y. Li, and L. Jing, "Investigating chain-of-thought with chatgpt for stance detection on social media," *arXiv preprint arXiv:2304.03087*, 2023.

[16] Z. Yao, W. Yang, and F. Wei, "Enhancing zero-shot stance detection with contrastive and prompt learning," *Entropy*, vol. 26, no. 4, p. 325, 2024.

[17] J. Zhu, Y. Yan, L. Zhao, M. Heimann, L. Akoglu, and D. Koutra, "Beyond homophily in graph neural networks: Current limitations and effective designs," *Advances in neural information processing systems*, vol. 33, pp. 7793–7804, 2020.

[18] T. Derr, Y. Ma, and J. Tang, "Signed graph convolutional networks," in *2018 IEEE international conference on data mining (ICDM)*. IEEE, 2018, pp. 929–934.

[19] N. Kim, D. Mosallanezhad, L. Cheng, M. V. Mancenido, and H. Liu, "Robust stance detection: Understanding public perceptions in social media," in *International Conference on Advances in Social Networks Analysis and Mining*. Springer, 2024, pp. 21–37.

[20] H. Huang, B. Zhang, Y. Li, B. Zhang, Y. Sun, C. Luo, and C. Peng, "Knowledge-enhanced prompt-tuning for stance detection," *ACM Transactions on Asian and Low-Resource Language Information Processing*, vol. 22, no. 6, pp. 1–20, 2023.

[21] Y. Li, B. Jiang, K. Zhang, and C. Caragea, "P-Stance: A large dataset for stance detection in tweets," in *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, 2021, pp. 2355–2365.

[22] K. Glandt, B. Khanal, Y. Li, and C. Caragea, "Stance detection in COVID-19 tweets," in *Findings of the Association for Computational Linguistics: EMNLP 2021*. Association for Computational Linguistics, 2021, pp. 4299–4313.

[23] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, vol. 1. Association for Computational Linguistics, 2019, pp. 4171–4186.

[24] K. Garg and C. Caragea, "Stanceformer: Target-aware transformer for stance detection," *arXiv preprint arXiv:2410.07083*, 2024.

[25] Y. Cheng, Q. Zhang, C. Shi, L. Xiao, S. Hao, and L. Hu, "Cosd: Collaborative stance detection with contrastive heterogeneous topic graph learning," 2024. [Online]. Available: https://arxiv.org/abs/2404.17609

[26] M. Weinzierl and S. Harabagiu, "Tree-of-counterfactual prompting for zero-shot stance detection," in *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Aug. 2024, pp. 861–880.

[27] Z. Zhang, Z. Zhang, J. Zhang, H. Xu, and X. Cheng, "MPVStance: Mitigating hallucinations in stance detection with multi-perspective verification," in *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Jul. 2025, pp. 1053–1067.