

# Directing Uncertainty-Aware Information Flow for Robust Diffusion Prediction

Weikang He, Yunpeng Xiao\*, Mengyang Huang, Xuemei Mou, Rong Wang, Qian Li

School of Communications and Information Engineering, Chongqing University of Posts and Telecommunications  
 heywecome@gmail.com, xiaoy@cqpt.edu.cn, huangmoonyann@gmail.com, xuemeimou00000@gmail.com,  
 wangrongl@cqpt.edu.cn, liqian@cqpt.edu.cn

## Abstract

Information diffusion prediction is crucial for understanding social network dynamics, yet existing methods often neglect user participation uncertainty. This oversight typically stems from an implicit participation homogeneity assumption, which treats all observed interactions as equally reliable propagation signals, leading to fragile inferred topologies and uncertainty contamination. To address this, we propose SIEVE, a novel framework employing two synergistic strategies. First, robust node representations are learned via controllable uncertainty injection coupled with associated contrastive learning, mitigating topological fragility. Second, an uncertainty-aware directed graph aggregation mechanism is introduced, which dynamically constructs asymmetric aggregation topologies with adaptive weighting, thereby suppressing uncertainty contamination. Experiments on four public datasets demonstrate that SIEVE significantly outperforms state-of-the-art methods, offering valuable insights for designing robust information diffusion prediction models.

**Code** — <https://github.com/HeyWeCome/BuzzBloom>

## Introduction

Predicting information diffusion in social networks is fundamental to social computing with broad applications (He et al. 2025a; Friedrich et al. 2024). Current methods primarily enhance prediction accuracy via two main avenues: (1) refining network structure representations to capture intricate social dependencies (Li et al. 2024; Jiao et al. 2024), and (2) leveraging temporal user interaction dynamics (Yang et al. 2021; Yuan et al. 2020; Sun et al. 2022; He et al. 2025b).

However, today’s media landscape of information overload and fragmented attention highlights a critical, often overlooked challenge: **inherent user participation uncertainty**. Social network theory posits widespread diffusion often relies on weak ties (Rajkumar et al. 2022), meaning many observed interactions may be low-fidelity or transient signals, not reliable indicators of consistent sharing intent.

Prevailing methods (Xu et al. 2023; Cheng et al. 2024; Jing et al. 2025), despite advancements, generally lack mechanisms to systematically quantify and mitigate this pervasive uncertainty’s impact. While attention mechanisms




Figure 1: Training paradigm and limitations of mainstream information diffusion prediction approaches.

like GAT (Zhang et al. 2024) learn task-specific neighbor relevance, they do not model the intrinsic reliability of the information source itself. This oversight is a symptom of a deeper issue: *participation homogeneity assumption* (PHA). This assumption denotes the prevalent practice of treating diverse user participation signals uniformly, failing to capture their inherent heterogeneity and varying real-world reliability. Consequently, conflating transient or low-fidelity interactions with stable propagation tendencies can incorporate spurious connections into inferred graph structures, compromising prediction reliability and robustness from these flawed representations.

Overlooking user participation uncertainty, driven by the PHA, critically impacts foundational model development (Fig.1). During diffusion graph construction, observed interactions typically become deterministic edges, often with uniform weights disregarding interaction reliability variability (Yuan et al. 2020; Sun et al. 2022; Zhong et al. 2024; Li et al. 2024; Cheng et al. 2023), implicitly treating all connections as equally dependable. Subsequently, GNN-based representation learning, using aggregation functions insensitive to connection certainty, readily incorporates signals from these unreliable edges. Consequently, node representations risk corruption by noise from uncertain connections. This neglect introduces two fundamental limitations in existing models:

**Fragile Diffusion Topology.** Reliance on the PHA inherently yields fragile diffusion topologies. Treating all observed interactions as equally valid and converting them to deterministic, often undifferentiated, connections forces models to learn from a network potentially contaminated by low-fidelity signals. This structural fragility builds predictions on an unreliable foundation, impairing capture of stable, generalizable diffusion patterns and increasing suscep-

\*Yunpeng Xiao is the corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

tibility to overfitting observed interaction noise.

**Uncertainty Contamination.** Standard diffusion graph aggregation, typically agnostic to node or edge certainty, indiscriminately pools neighbor information. In information diffusion, where aggregation pathways can be bidirectional, noise from high-uncertainty nodes or unreliable edges readily permeates the network, contaminating otherwise reliable node representations. This quality-agnostic contamination degrades learned node embedding integrity, hindering discernment of genuine propagation signals from low-fidelity interactions and compromising prediction precision.

To address these limitations, we propose *SIEVE*, an innovative information diffusion prediction framework. Departing from the PHA, *SIEVE* systematically tackles user participation uncertainty via two synergistic strategies: (1) mitigating topological fragility by learning robust node representations through controllable uncertainty injection and associated contrastive learning; and (2) suppressing uncertainty contamination via an uncertainty-aware directed aggregation with dynamic asymmetric topologies and adaptive weighting. By explicitly modeling and actively mitigating uncertainty at both representation learning and graph aggregation stages, *SIEVE* aims for more reliable and robust predictions. Overall, our contributions are as follows:

- We identify and articulate the *participation homogeneity assumption* as a critical, unaddressed limitation in existing information diffusion models, elucidating its role in fostering fragile topologies and uncertainty contamination.
- We propose *SIEVE*, a novel framework to model and suppress user participation uncertainty via two synergistic mechanisms: robust representation learning (using uncertainty injection and associated contrastive learning) to combat topological fragility, and uncertainty-aware directed aggregation to prevent uncertainty contamination.
- Extensive experiments conducted on four public datasets demonstrate that *SIEVE* outperforms state-of-the-art methods. Comprehensive ablation studies further validate the efficacy of our proposed components and offer valuable insights for designing robust diffusion modeling approaches.

## Preliminary

In this section, we introduce fundamental notations and definitions that will be used throughout the paper. We consider a set of users  $\mathcal{U}$  and a set of topics  $\mathcal{M}$ .

**Definition 1** (Topic Diffusion Cascade  $\mathcal{C}_m$ ). The propagation of a topic  $m \in \mathcal{M}$  is represented as a time-ordered sequence of user engagements, termed a *chain cascade* (Yuan et al. 2020; Sun et al. 2022). Formally, this cascade is denoted by  $\mathcal{C}_m = \langle (u_1, t_1), \dots, (u_{n_m}, t_{n_m}) \rangle$ , which is a sequence of user-timestamp pairs  $(u_i, t_i)$ . Here,  $u_i \in \mathcal{U}$  is the  $i$ -th participating user,  $t_i$  is the absolute timestamp of their engagement, and these timestamps are strictly increasing, i.e.,  $t_1 < t_2 < \dots < t_{n_m}$ .

**Definition 2** (Topic Diffusion Graph). Following prior work (Yuan et al. 2020; Li et al. 2024; Sun et al. 2022), we construct a user-topic interaction graph  $\mathcal{G}_I = (\mathcal{V}_I, \mathcal{E}_I, \mathbf{A}_I, \mathbf{X}_I)$ . Nodes  $\mathcal{V}_I = \mathcal{U} \cup \mathcal{M}$  include users and topics. Edges  $\mathcal{E}_I \subseteq \mathcal{U} \times \mathcal{M}$  represent observed user-topic

engagements. The adjacency matrix  $\mathbf{A}_I \in \{0, 1\}^{|\mathcal{V}_I| \times |\mathcal{V}_I|}$  encodes these bipartite interactions:  $\mathbf{A}_I^{u,m} = \mathbf{A}_I^{m,u} = 1$  if user  $u$  engaged with topic  $m$ . Initial node features are  $\mathbf{X}_I \in \mathbb{R}^{|\mathcal{V}_I| \times d_0}$ , where  $d_0$  is the initial feature dimension.

An  $L$ -layer Graph Neural Network (GNN) learns node representations from  $\mathcal{G}_I$ . With  $\mathbf{Z}^{(0)} = \mathbf{X}_I$ , layer  $l$ 's representation  $\mathbf{Z}^{(l)} \in \mathbb{R}^{|\mathcal{V}_I| \times d_l}$  is:

$$\mathbf{Z}^{(l)} = \sigma \left( \hat{\mathbf{A}} \cdot \mathbf{Z}^{(l-1)} \cdot \mathbf{W}^{(l)} \right), \quad l = 1, \dots, L \quad (1)$$

where  $\hat{\mathbf{A}} = \mathbf{D}^{-1/2}(\mathbf{A}_I + \mathbf{I})\mathbf{D}^{-1/2}$  is the symmetrically normalized adjacency matrix ( $\mathbf{D}$  is the degree matrix of  $\mathbf{A}_I + \mathbf{I}$ ),  $\mathbf{W}^{(l)} \in \mathbb{R}^{d_{l-1} \times d_l}$  is layer  $l$ 's trainable weight matrix mapping  $d_{l-1}$  to  $d_l$  dimensions, and  $\sigma(\cdot)$  is a non-linear activation. The final node representations  $\mathbf{Z} \in \mathbb{R}^{|\mathcal{V}_I| \times d_L}$  are obtained via a readout function  $f_{\text{readout}}$  aggregating layer-wise representations:

$$\mathbf{Z} = f_{\text{readout}} \left( \mathbf{Z}^{(0)}, \mathbf{Z}^{(1)}, \dots, \mathbf{Z}^{(L)} \right). \quad (2)$$

This GNN iteratively updates user and topic representations. While standard symmetric aggregation (Eq. 1) is common, its potential ineffectiveness in handling node uncertainty motivates our improved aggregation strategy.

**Definition 3** (Information Diffusion Prediction Task). In this paper, we target micro-cascade prediction: given a topic  $m$ 's observed cascade history up to step  $j$ ,  $\mathcal{H}_j^m = \langle (u_1^m, t_1^m), \dots, (u_j^m, t_j^m) \rangle$  ( $1 \leq j < n_m$ ), predict the next participating user  $u_{j+1}^m$ . Typically, a model uses the topic representation and  $\mathcal{H}_j^m$  to derive a cascade sequence representation  $\mathbf{O}_j^m \in \mathbb{R}^{d_h}$ , then computes a user probability distribution  $\mathbf{p}_j^m$  via softmax:

$$\mathbf{p}_j^m = \text{softmax}(\mathbf{O}_j^m \mathbf{W}_o + \mathbf{b}_o), \quad (3)$$

where  $\mathbf{W}_o \in \mathbb{R}^{d_h \times |\mathcal{U}|}$  and  $\mathbf{b}_o \in \mathbb{R}^{|\mathcal{U}|}$  are output layer parameters. Model parameters  $\Theta$  are learned by minimizing the negative log-likelihood on the training set  $\mathcal{M}_{\text{train}}$ :

$$\mathcal{L}_{\mathcal{B}} = \min_{\Theta} \left( - \sum_{m \in \mathcal{M}_{\text{train}}} \sum_{j=1}^{n_m-1} \log p(u_{j+1}^m | \mathcal{H}_j^m; \Theta) \right). \quad (4)$$

We use a full ranking approach, calculating probabilities over all users at each step, avoiding negative sampling.

## Model

Addressing fragile diffusion topology and uncertainty contamination, *SIEVE* (Fig. 2) tackles user participation uncertainty via two co-designed modules. The specifics of these modules are detailed below.

### Robust Representation Learning via Uncertainty Injection and Associated Contrastive Learning

To address fragile diffusion topology, this section details robust dynamic representation learning. Instead of directly modifying the observed graph topology (e.g.,  $\mathbf{A}_I$  or  $\hat{\mathbf{A}}$ ), we enhance node representation robustness through controllable uncertainty injection combined with contrastive learning.

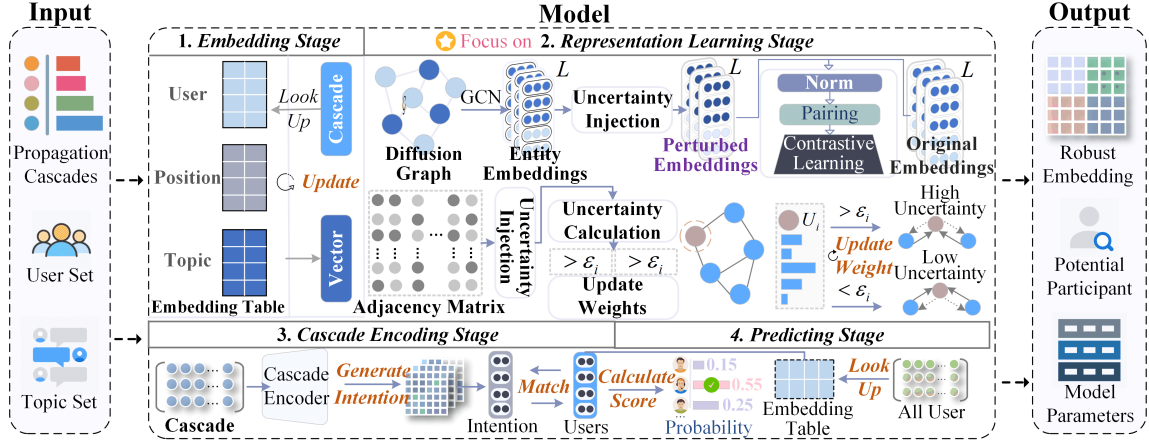


Figure 2: Overview of the *SIEVE* model architecture. It illustrates how *SIEVE* tackles user participation uncertainty through two co-designed and synergistic modules: (a) Robust Representation Learning, which utilizes uncertainty injection and contrastive learning to mitigate fragile diffusion topologies, and (b) Uncertainty-Aware Directed Aggregation, which constructs dynamic asymmetric aggregation pathways to suppress uncertainty contamination.

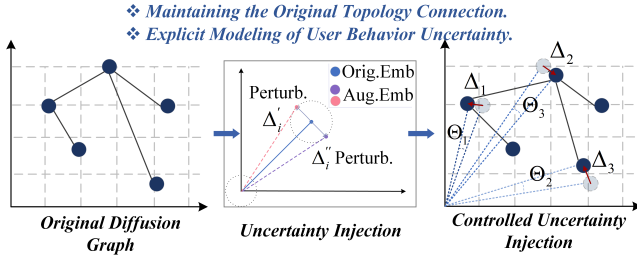


Figure 3: Visualizing the controlled uncertainty injection.

We introduce feature-level controlled uncertainty injection (Fig. 3). For node  $i$ 's representation  $\mathbf{z}_i^{(l-1)}$  at layer  $l-1$ , we add a perturbation  $\Delta_i^{(l)}$  generated under dual constraints. First, its  $L_2$  norm is governed by a layer-specific learnable parameter  $\epsilon_l$ , ensuring  $\|\Delta_i^{(l)}\|_2 = \epsilon_l$ . Second, its sign pattern is constrained to match that of the original representation  $\mathbf{z}_i^{(l-1)}$ . This aligns with the assumption that uncertainty primarily affects the strength of a user's latent state rather than its core semantic direction, which is encoded by  $\text{sign}(\mathbf{z}_i^{(l-1)})$ . The perturbation is thus computed using a base random vector  $\bar{\Delta}_i^{(l)} \sim \mathcal{N}(0, 1)$  as follows:

$$\Delta_i^{(l)} = \epsilon_l \cdot \frac{|\bar{\Delta}_i^{(l)}| \odot \text{sign}(\mathbf{z}_i^{(l-1)})}{\|\bar{\Delta}_i^{(l)} \odot \text{sign}(\mathbf{z}_i^{(l-1)})\|_2 + \delta}, \quad (5)$$

where  $\delta$  is a small constant for numerical stability. This mechanism models inherent randomness in user participation. The learnable  $\epsilon_l$  offers standardized, layer-adaptive control over perturbation intensity, stabilizing subsequent contrastive learning. The sign constraint simulates marginal fluctuations around the core state, compelling the model to learn robustness against such common randomness. This tar-

geted data augmentation enhances robustness specifically against feature-level fluctuations tied to user participation uncertainty.

This uncertainty injection is incorporated within the GNN's layer-wise propagation. For layer  $l$ , aggregation uses an asymmetric, uncertainty-aware adjacency matrix  $\hat{\mathbf{S}}^{(l-1)}$  (dynamically computed from  $\mathbf{Z}^{(l-1)}$ , detailed in the following Section), not the standard  $\hat{\mathbf{A}}$ . *SIEVE*'s GNN update for layer  $l$  is:

$$\mathbf{Z}^{(l)} = \sigma \left( \hat{\mathbf{S}}^{(l-1)} \cdot \left( \mathbf{Z}^{(l-1)} + \Delta^{(l)} \right) \cdot \mathbf{W}^{(l)} \right), \quad (6)$$

where  $\Delta^{(l)}$  contains node perturbations (Eq. 5) with layer-specific  $\epsilon_l$ ,  $\mathbf{W}^{(l)}$  are learnable weights, and  $\sigma(\cdot)$  is non-linear activation. Eq. 6 concurrently applies feature perturbations via  $\Delta^{(l)}$  for robustness and employs adaptive topology  $\hat{\mathbf{S}}^{(l-1)}$  to suppress flow from high-uncertainty neighbors.

To effectively use injected uncertainty and learn robust representations, we employ associated contrastive learning, aiming for representations invariant to this augmentation, enhancing robustness and discriminability. We design one complementary tasks: node-level ( $\mathcal{L}_{\text{nd}}$ ) contrastive learning.

Node-level contrastive learning ( $\mathcal{L}_{\text{nd}}$ ) uses two independent forward passes with uncertainty injection on the same batch, yielding final layer representations  $\mathbf{Z}^{(L)'} and \mathbf{Z}^{(L)''}$ . It maximizes similarity for the same node's views and minimizes it for others:

$$\mathcal{L}_{\text{nd}} = \sum_{i \in \mathcal{B}} -\log \frac{\exp \left( \text{sim}(\mathbf{z}_i^{(L)'}, \mathbf{z}_i^{(L)''}) / \tau_{\text{nd}} \right)}{\sum_{j \in \mathcal{B}} \exp \left( \text{sim}(\mathbf{z}_i^{(L)'}, \mathbf{z}_j^{(L)''}) / \tau_{\text{nd}} \right)}, \quad (7)$$

where  $\text{sim}(\cdot, \cdot)$  is cosine similarity,  $\tau_{\text{nd}}$  is temperature, and  $\mathcal{B}$  is the batch node set. Minimizing  $\mathcal{L}_{\text{nd}}$  incentivizes learning node representations consistent and discriminative despite

feature perturbations ( $\Delta^{(l)}$ ) and dynamic topology ( $\hat{\mathbf{S}}^{(l-1)}$ ), enhancing uncertainty handling.

### Uncertainty-Aware Directed Graph Aggregation

To mitigate uncertainty contamination from symmetric aggregation, we propose uncertainty-aware directed graph aggregation. This quantifies node uncertainty and dynamically constructs an asymmetric aggregation topology  $\hat{\mathbf{S}}$  to suppress uncertainty propagation.

While Shannon entropy  $H(p_i) = -\sum_k p_{ik} \log p_{ik}$  is a theoretical uncertainty measure, inferring  $p_i$  from high-dimensional embeddings  $\mathbf{z}_i$  is challenging. We thus use a proxy based on node representation stability, assuming low uncertainty for nodes whose representations are stable against minor feature perturbations (Eq. 5).

For layer  $l$  aggregation, we compute node  $v_i$ 's uncertainty proxy  $U_i^{(l-1)}$  using layer  $l-1$  representations  $\mathbf{Z}^{(l-1)}$ .  $U_i^{(l-1)}$  measures  $\mathbf{z}_i^{(l-1)}$ 's stability against its perturbed version  $\tilde{\mathbf{z}}_i^{(l-1)}$  in its local neighborhood:

$$U_i^{(l-1)} = \frac{1}{2} \left( l(\mathbf{z}_i^{(l-1)}, \tilde{\mathbf{z}}_i^{(l-1)}) + l(\tilde{\mathbf{z}}_i^{(l-1)}, \mathbf{z}_i^{(l-1)}) \right), \quad (8)$$

where  $l(z, \tilde{z})$  is an InfoNCE-like contrastive loss. Negative samples are restricted to perturbed representations of node  $i$  and its first-order neighbors  $\mathcal{N}(i)$ .

$$l(z, \tilde{z}) = -\log \frac{\exp(\text{sim}(z, \tilde{z})/\tau_u)}{\sum_{k \in \mathcal{N}(i) \cup \{i\}} \exp(\text{sim}(z, \tilde{\mathbf{z}}_k^{(l-1)})/\tau_u)}. \quad (9)$$

Here,  $\text{sim}$  is cosine similarity,  $\tau_u$  is temperature, and  $\tilde{\mathbf{z}}_k^{(l-1)}$  is perturbed  $\mathbf{z}_k^{(l-1)}$ . This local design captures stability within the immediate environment; larger  $U_i^{(l-1)}$  implies higher uncertainty.

Based on node uncertainty proxies  $U_j^{(l-1)}$ , we dynamically compute the aggregation weight matrix  $\hat{\mathbf{S}}^{(l-1)}$  for layer  $l$ . For a path from node  $j$  to  $i$ , the asymmetric weight  $\hat{\mathbf{S}}_{ij}^{(l-1)}$  depends on the original connection  $A_I^{i,j}$ , source uncertainty  $U_j^{(l-1)}$ , and target's learnable tolerance  $\varepsilon_i$ . As shown in Fig. 4, the weight is:

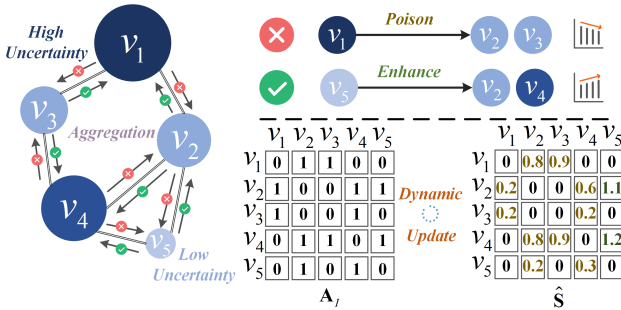


Figure 4: Illustration of asymmetric propagation topology reconstruction based on node uncertainty proxy metrics.

$$\hat{\mathbf{S}}_{ij}^{(l-1)} = A_I^{i,j} \cdot \psi \left( e^{-U_j^{(l-1)}} - \varepsilon_i \right), \quad (10)$$

where  $e^{-U_j^{(l-1)}} \in (0, 1]$  is source  $j$ 's confidence, and  $\varepsilon_i$  is a learnable sensitivity threshold for the target node  $i$ . The function  $\psi(x)$  non-linearly adjusts weights:

$$\psi(x) = \begin{cases} \gamma^{(l-1)} \cdot s(x), & \text{if } x \geq 0, \\ \beta, & \text{if } x < 0, \end{cases} \quad (11)$$

where  $s(x)$  is Sigmoid, and  $\beta$  is a small non-negative hyperparameter suppressing unreliable connections. The amplification  $\gamma^{(l-1)}$  is dynamically adjusted based on layer  $l-1$ 's median uncertainty  $U_{median}^{(l-1)} = \text{median}(\{U_k^{(l-1)}\}_{k \in \mathcal{V}_I})$  (robust to outliers):

$$\gamma^{(l-1)} = \gamma_0 \cdot e^{-U_{median}^{(l-1)}}, \quad (12)$$

where  $\gamma_0$  is a base scaling factor. This yields conservative amplification (smaller  $\gamma^{(l-1)}$ ) for high overall uncertainty, and vice versa. This dynamic, asymmetric  $\hat{\mathbf{S}}^{(l-1)}$  replaces  $\hat{\mathbf{A}}$  in the  $l$ -th layer GNN update (Eq. 6), suppressing uncertainty contamination by guiding propagation along more reliable directed paths.

### Basic Sequence Encoder

Given our focus on user participation uncertainty, we use a basic self-attention mechanism for sequence feature extraction. This choice facilitates clearer variable isolation, enabling more direct evaluation of our core mechanisms' impact on diffusion prediction, without confounding from an overly complex sequence model.

For a topic  $m$ 's history  $\mathcal{H}_j^m = \langle (u_1^m, t_1^m), \dots, (u_j^m, t_j^m) \rangle$ , input to self-attention is the sequence of user representations  $\{\mathbf{z}_{u_1^m}, \dots, \mathbf{z}_{u_j^m}\}$ , each augmented with positional embeddings. Self-attention processes this sequence yielding  $\mathbf{O}_j^m \in \mathbb{R}^{d_h}$ , representing sequential information up to  $t_j^m$ . This  $\mathbf{O}_j^m$  is used to predict  $u_{j+1}^m$ , and the encoder is trained via  $\mathcal{L}_B$  (Eq. 4).

### Joint Learning

We employ joint learning: the two co-designed modules serve as auxiliary objectives to enhance user and topic embeddings, ultimately benefiting diffusion prediction. The overall loss is:

$$\mathcal{L} = \mathcal{L}_B + \lambda_{cl} \mathcal{L}_{nd}, \quad (13)$$

where  $\lambda_{cl}$  is a hyperparameter balancing the loss terms.

## Experiments

To validate the effectiveness of *SIEVE*, we conduct thorough experiments and present the detailed results below.

| Dataset             | Users  | Cascades | Interaction | $\bar{L}$ | Density |
|---------------------|--------|----------|-------------|-----------|---------|
| <b>Douban</b>       | 12,232 | 3,475    | 75,619      | 21.76     | 0.18%   |
| <b>Christianity</b> | 1,651  | 589      | 15,327      | 26.02     | 1.58%   |
| <b>Android</b>      | 2,927  | 678      | 28,512      | 42.05     | 1.44%   |
| <b>Memetracker</b>  | 4,709  | 10,130   | 162,464     | 16.04     | 0.34%   |

Table 1: Statistics of the datasets.  $\bar{L}$  denotes the average cascade length.

| Dataset      | Metrics | DyHGCN | MSHGAT | CEGCN  | DisenIDP | MINDS  | GRASS         | GODEN         | PMRCA         | SIEVE         |
|--------------|---------|--------|--------|--------|----------|--------|---------------|---------------|---------------|---------------|
| Android      | Hit@10  | 0.0996 | 0.1005 | 0.1015 | 0.0954   | 0.1096 | 0.1101        | 0.1134        | <u>0.1249</u> | <b>0.1532</b> |
|              | MAP@10  | 0.0619 | 0.0584 | 0.0594 | 0.0555   | 0.0677 | 0.0680        | 0.0641        | <u>0.0699</u> | <b>0.0814</b> |
|              | Hit@50  | 0.1962 | 0.1908 | 0.1928 | 0.1988   | 0.1989 | 0.1999        | 0.2076        | <u>0.2187</u> | <b>0.2450</b> |
|              | MAP@50  | 0.0663 | 0.0625 | 0.0635 | 0.0599   | 0.0716 | 0.0712        | <u>0.0795</u> | 0.0737        | <b>0.0917</b> |
| Christianity | Hit@10  | 0.2406 | 0.2722 | 0.2730 | 0.2584   | 0.3214 | 0.3321        | 0.3415        | <u>0.3866</u> | <b>0.4316</b> |
|              | MAP@10  | 0.1505 | 0.1718 | 0.1728 | 0.1575   | 0.1955 | 0.2013        | 0.2120        | <u>0.2576</u> | <b>0.2831</b> |
|              | Hit@50  | 0.4142 | 0.4477 | 0.4490 | 0.4359   | 0.4778 | 0.4887        | 0.4935        | <u>0.5385</u> | <b>0.5799</b> |
|              | MAP@50  | 0.1589 | 0.1794 | 0.1805 | 0.1655   | 0.2037 | 0.2144        | 0.2230        | <u>0.2648</u> | <b>0.2906</b> |
| Douban       | Hit@10  | 0.1982 | 0.1957 | 0.1967 | 0.1504   | 0.1956 | 0.2295        | 0.2423        | <u>0.2766</u> | <b>0.3077</b> |
|              | MAP@10  | 0.1016 | 0.1074 | 0.1085 | 0.0641   | 0.1142 | 0.1392        | 0.1488        | <u>0.1524</u> | <b>0.1832</b> |
|              | Hit@50  | 0.3288 | 0.3349 | 0.3360 | 0.3119   | 0.3087 | 0.3663        | 0.3901        | <u>0.4191</u> | <b>0.4439</b> |
|              | MAP@50  | 0.1077 | 0.1140 | 0.1152 | 0.0712   | 0.1199 | 0.1455        | 0.1556        | <u>0.1592</u> | <b>0.1984</b> |
| Memetracker  | Hit@10  | 0.2707 | 0.3452 | 0.3689 | 0.2955   | 0.3915 | 0.4202        | 0.4117        | <u>0.4904</u> | <b>0.5172</b> |
|              | MAP@10  | 0.1462 | 0.2008 | 0.2153 | 0.1683   | 0.2388 | 0.2618        | 0.2438        | <u>0.2761</u> | <b>0.3007</b> |
|              | Hit@50  | 0.3487 | 0.5120 | 0.5530 | 0.4130   | 0.5755 | <u>0.6024</u> | 0.5939        | 0.5654        | <b>0.6213</b> |
|              | MAP@50  | 0.1516 | 0.2095 | 0.2240 | 0.1751   | 0.2467 | <u>0.2705</u> | 0.2527        | <u>0.2812</u> | <b>0.3046</b> |

Table 2: Performance Comparison of Models on Different Datasets. The highest score in each row is bolded, and the second-highest score is underlined. We use five random seeds for the experiment and take the average as the final result.

## Experimental Setup

**Datasets** We evaluate *SIEVE* on four widely adopted public datasets: Douban (Zhong et al. 2012), Christianity (Sankar et al. 2020a), Android (Sankar et al. 2020a), and Memetracker (Leskovec, Backstrom, and Kleinberg 2009) (statistics in Table 1). Following prior work (Yang et al. 2021; Yuan et al. 2020; Sun et al. 2022), cascades are randomly split into training (80%), validation (10%), and test (10%) sets. Cascades with fewer than 5 users are excluded.

**Baseline Methods** We compare *SIEVE* against eight representative open-source baselines: DyHGCN (Yuan et al. 2020), MSHGAT (Sun et al. 2022), CEGCN (Wang et al. 2022), GRASS (Li et al. 2024), MINDS (Jiao et al. 2024), DisenIDP (Cheng et al. 2023), GODEN (Wang, Zhou, and Hu 2024) and PMRCA (He et al. 2025a). These baselines have shown superior performance over classic methods like FOREST (Yang et al. 2021), Inf-VAE (Sankar et al. 2020a), and TAN (Wang, Yang, and Shi 2021), which are thus not directly compared.

**Experimental Details** Following prior work (He et al. 2025a; Yuan et al. 2020; Sun et al. 2022; Li et al. 2024), we evaluate prediction performance using Hit@K and MAP@K. All methods are optimized using Adam with carefully tuned hyperparameters for fair comparison.

## Overall Performance

Table 2 presents a comparison of the performance of *SIEVE* and other baseline methods across the four datasets. From the table, we make several observations:

(1) *SIEVE* consistently achieves the leading performance across all evaluated datasets and metrics. This robust and comprehensive superiority underscores the fundamental efficacy of its design in tackling the inherent uncertainties within user participation, a critical aspect often overlooked in conventional diffusion modeling.

(2) Baseline methods exhibit varied efficacy. PMRCA consistently ranks as a strong competitor, while other GNN-based approaches show performance differences attributable to their distinct architectural designs, such as dynamic graph learning or specialized cascade encoding. GRASS also demonstrates notable strength in certain scenarios.

(3) *SIEVE*’s significant and consistent margin over even the most competitive baselines, including PMRCA, underscores the critical limitation of the “participation homogeneity assumption.” By explicitly tackling this assumption, *SIEVE* achieves a new level of prediction accuracy, revealing a key avenue for advancement in diffusion modeling.

On Memetracker, *SIEVE*’s performance margin is smaller than on other datasets. We attribute this to two factors. First, its short average cascades ( $\bar{L} = 16.04$ ) may mitigate the error accumulation that *SIEVE* is designed to address. Second, its propagation may be dominated by content virality rather than nuanced user uncertainty. In such scenarios, strong baselines like PMRCA can effectively capture the primary diffusion signal, reducing the marginal gain from *SIEVE*’s more sophisticated uncertainty modeling. Nonetheless, *SIEVE*’s leading performance validates its robustness and suggests that the dominant drivers of diffusion vary across networks.

## Further Analysis of *SIEVE*

**Ablation Study** To assess the individual and collective contributions of *SIEVE*’s core components, we evaluate variants excluding: (1) Robust Representation Learning (RRL), which omits uncertainty injection and contrastive objectives; (2) Uncertainty-Aware Aggregation (UAA), reverting to standard GNN aggregation; and (3) *w/o Both* components. Results are presented in Table 3.

The ablation results confirm the critical and synergistic contributions of our proposed modules. The substantial performance degradation of the *w/o Both* variant across

all datasets underscores the imperative of addressing user participation uncertainty, which is SIEVE’s foundational premise. Furthermore, individually ablating either RRL or UAA markedly impairs performance, confirming their distinct, indispensable roles: RRL forges resilient node embeddings against feature perturbations, while UAA constructs reliable, dynamic aggregation pathways. While the precise leverage of each module can subtly adapt to dataset-specific characteristics, their combined efficacy is consistently paramount. This highlights their synergistic interplay in robustly navigating the multifaceted challenges of participation uncertainty in information diffusion.

| Dataset      | Model        | Hits@10       | MAP@10        | Hits@20       | MAP@20        |
|--------------|--------------|---------------|---------------|---------------|---------------|
| Android      | <i>SIEVE</i> | <b>0.1532</b> | <b>0.0814</b> | <b>0.2023</b> | <b>0.0847</b> |
|              | w/o RRL      | 0.1355        | 0.0713        | 0.1786        | 0.0743        |
|              | w/o UAA      | 0.1418        | 0.0691        | 0.1832        | 0.0719        |
|              | w/o Both     | 0.0642        | 0.0342        | 0.0833        | 0.0366        |
| Christianity | <i>SIEVE</i> | <b>0.4316</b> | <b>0.2831</b> | <b>0.4855</b> | <b>0.2868</b> |
|              | w/o RRL      | 0.4031        | 0.2612        | 0.4616        | 0.2657        |
|              | w/o UAA      | 0.3967        | 0.2287        | 0.4678        | 0.2344        |
|              | w/o Both     | 0.2024        | 0.1335        | 0.2424        | 0.1377        |
| Douban       | <i>SIEVE</i> | <b>0.3077</b> | <b>0.1832</b> | <b>0.3631</b> | <b>0.1874</b> |
|              | w/o RRL      | 0.2759        | 0.1537        | 0.3256        | 0.1582        |
|              | w/o UAA      | 0.2722        | 0.1496        | 0.3210        | 0.1542        |
|              | w/o Both     | 0.1717        | 0.0924        | 0.2174        | 0.0971        |
| Memetracker  | <i>SIEVE</i> | <b>0.5172</b> | <b>0.3007</b> | <b>0.5942</b> | <b>0.3059</b> |
|              | w/o RRL      | 0.4829        | 0.2697        | 0.5727        | 0.2745        |
|              | w/o UAA      | 0.4795        | 0.2599        | 0.5678        | 0.2661        |
|              | w/o Both     | 0.2799        | 0.1499        | 0.3392        | 0.1547        |

Table 3: Ablation Study Results.

**Analysis of Uncertainty Injection Strategy** To scrutinize our uncertainty injection mechanism (Eq. 5), we benchmark it against diverse perturbation alternatives. Evaluated strategies include: (1) **Ours**; (2) **NoPert** (baseline,  $\Delta = 0$ ); (3) **UUI**, employing a normalized uniform random vector  $\epsilon \cdot \frac{\mathbf{r}_{\text{base}}}{\|\mathbf{r}_{\text{base}}\|_2 + \delta_0}$  ( $r_{\text{base},k} \sim U(0, 1)$ ) without sign-alignment; (4) **AGN**,  $\Delta = \epsilon \cdot \mathbf{g}$  ( $g_k \sim \mathcal{N}(0, 1)$ ); (5) **AUN**,  $\Delta = \epsilon \cdot \mathbf{u}$  ( $u_k \sim U(-1, 1)$ ); (6) **MBN**, multiplicative dropout  $\tilde{z}_k = \frac{m_k}{1-p} z_k$  ( $m_k \sim \text{Bernoulli}(1-p)$ ,  $p = \epsilon$ ); (7) **FUSI**, where USI’s base uniform random vector  $\mathbf{r}_{\text{base}}$  is pre-generated and reused; and (8) **FAGN**, where AGN’s base Gaussian noise vector  $\mathbf{g}$  is pre-generated, reused, and scaled by  $\epsilon$ .

Fig. 5 shows **Ours** outperforming all alternatives, validating its design for robust representation learning. While most augmentations improve over **NoPert**, **Ours** sophisticated approach yields superior gains. Performance degradation upon removing sign-alignment (**UUI**) or simplifying to basic additive noises (**AGN**, **AUN**) confirms these components’ criticality. Multiplicative noise (**MBN**) offers only moderate utility. Moreover, dynamic resampling of perturbations in **Ours** is demonstrably more effective than fixed noise patterns (**FUSI**, **FAGN**) for learning generalizable, resilient features. Thus, **Ours** is a highly robust uncertainty injection strategy.

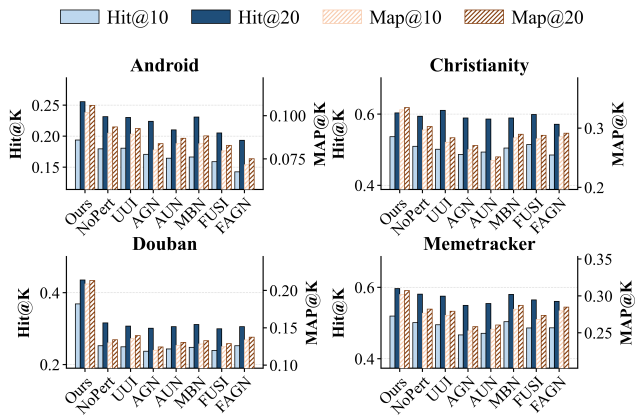


Figure 5: Analysis of uncertainty injection strategies.

**Robustness to Participation Uncertainty** To evaluate SIEVE’s robustness to user participation uncertainty, we inject synthetic noisy edges into  $\mathcal{G}_I$ . These edges represent spurious participations and are randomly added between user-topic pairs without prior interaction. The degree of this uncertainty is controlled by  $\rho \in [0.05, 0.95]$ , the fraction of such non-interacting pairs receiving an injected edge, modelling escalating levels of uncertainty. SIEVE and key baselines are then benchmarked on these graphs.

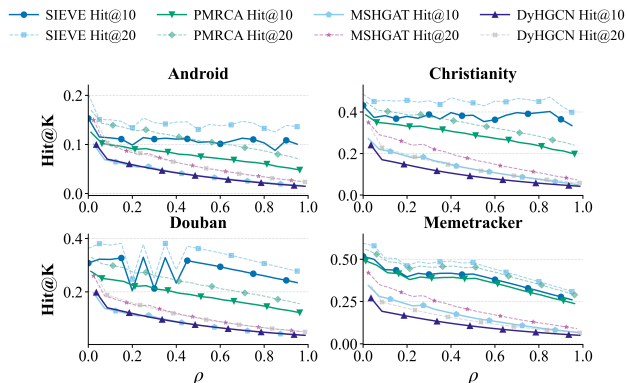


Figure 6: Robustness of SIEVE and baselines to user participation uncertainty.

The results (Fig. 6) underscore SIEVE’s pronounced resilience to this induced participation uncertainty. While all models’ accuracy declines with increasing  $\rho$ , SIEVE consistently degrades more gracefully. Crucially, even at high levels of participation uncertainty where baseline performance sharply deteriorates, SIEVE often maintains a clear advantage, or performs comparably to baselines under conditions of significantly lower uncertainty. This sustained superiority highlights the efficacy of SIEVE’s dedicated mechanisms, robust representation learning and uncertainty-aware directed aggregation, in mitigating spurious participation signals, vital for reliable prediction.

**Analysis of Perturbation Magnitude  $\epsilon$**  We perform a sensitivity analysis on the perturbation magnitude  $\epsilon$ , finding its optimal value is intrinsically linked to dataset-specific structural properties (Fig. 7). On dense graphs, the model tolerates a larger  $\epsilon$  due to robust intrinsic signals, while on sparse ones, performance peaks at a small value.

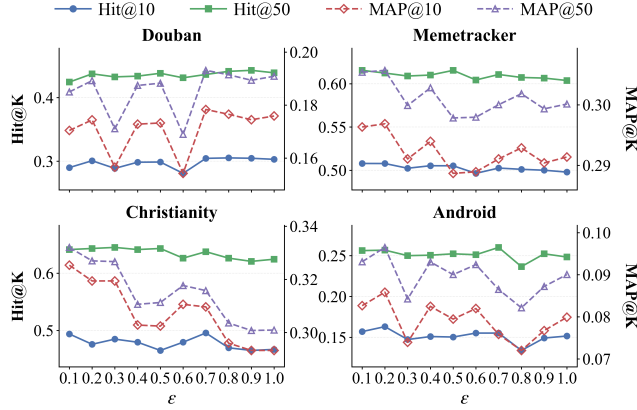


Figure 7: Model performance vs. perturbation magnitude  $\epsilon$ .

However, density is not the sole determinant: the dense Android dataset’s noise sensitivity stems from its deep propagation paths ( $\bar{L} = 42.05$ ), while on the sparse Douban, a large  $\epsilon$  acts as a powerful regularizer against overfitting. These complex interactions prove a fixed  $\epsilon$  is suboptimal, validating our learnable, layer-adaptive  $\epsilon_l$  designed to autonomously find the optimal balance for each environment.

**Analysis of Uncertainty Suppression Factor  $\beta$**  We then analyze the uncertainty suppression factor  $\beta$ , which penalizes unreliable connections. The results in Fig. 8 reveal that the optimal strategy is profoundly dataset-dependent and defies simple heuristics.

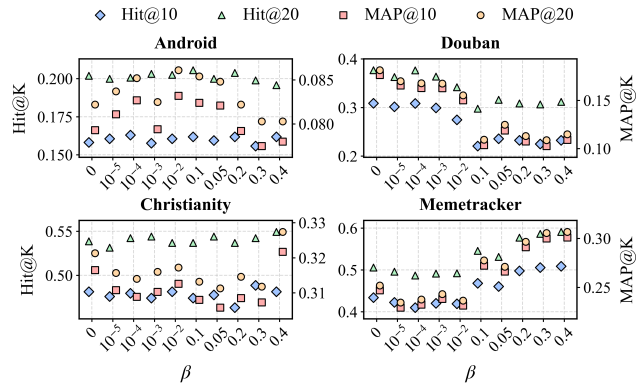


Figure 8: Model performance vs. suppression factor  $\beta$ .

Counter-intuitively, the extremely sparse Douban dataset favors the most aggressive strategy: hard-pruning of edges ( $\beta = 0$ ). This suggests that for this graph, the harm from misleading shortcuts outweighs the risk of severing uncertain paths. Conversely, on Memetracker, performance im-

proves with larger  $\beta$ , implying a misalignment between our stability-based proxy and the dataset’s chaotic dynamics; a larger  $\beta$  effectively mitigates the damage of misclassifying valuable paths as uncertain. The denser Android dataset conforms to expectations, where soft-pruning optimally balances noise suppression and connectivity preservation.

These findings confirm that a one-size-fits-all suppression strategy is suboptimal. The ideal  $\beta$  is a complex function of graph-specific noise characteristics and the proxy’s efficacy, validating its role as a crucial model parameter.

## Related Work

This section briefly reviews related work and highlights the innovations of *SIEVE*.

**Information Diffusion Prediction Models** Information diffusion prediction research has largely focused on refining network structure representations (Jing et al. 2025; Cheng et al. 2024; Li et al. 2024; Jiao et al. 2024; He et al. 2025a) or leveraging interaction dynamics (Yang et al. 2021; Yuan et al. 2020; Sun et al. 2022; Sankar et al. 2020b; Cheng et al. 2023, 2025). However, these methods often operate under an implicit PHA, overlooking the heterogeneity and varying reliability of user participation signals, which lead to fragile diffusion topologies. *SIEVE* distinguishes itself by directly addressing the PHA, aiming to systematically model and mitigate user participation uncertainty.

**Uncertainty in Graph Neural Networks** Recent advances in GNNs have increasingly focused on quantifying and leveraging uncertainty to enhance model robustness and reliability (Han et al. 2025; Chen et al. 2025; Fuchsgruber, Wollschläger, and Günnemann 2024). While these methods are powerful, they are often not specifically address the unique challenge of *user participation uncertainty* inherent in information diffusion. *SIEVE* addresses this by tackling the underlying PHA with an innovative dual strategy. It proactively immunizes representations against low-fidelity signals via uncertainty injection and contrastive learning, while its uncertainty-aware aggregation suppresses contamination during message passing. This integrated approach, unlike reactive adjustments in models like UnGSL, simultaneously targets both fragile topologies and uncertainty contamination, the key consequences of PHA.

## Conclusion

This paper addresses user participation uncertainty in diffusion prediction, a challenge overlooked by the common Participation Homogeneity Assumption. We propose *SIEVE*, a framework that learns robust representations via controllable uncertainty injection and contrastive learning. *SIEVE* then uses an uncertainty-aware aggregation mechanism to mitigate noise propagation. Extensive experiments validate that *SIEVE* consistently outperforms state-of-the-art methods. Future work could extend *SIEVE* to model the temporal dynamics of user uncertainty, potentially leading to more robust and temporally-aware diffusion models.

## Acknowledgements

This paper is partially supported by the National Natural Science Foundation of China (Grant No.62576057), Natural Science Foundation of Chongqing (Grant No. CSTB2025NSCQ-LZX0148, CSTB2025NSCQ-GPX1254), Major Scientific and Technological Research Program of Chongqing Municipal Education Commission (Grant No.KJZD-M202500603), Doctor Student Innovative Talent Program of CQUPT (Grant No.BYJS202414) and Chongqing Postgraduate Research and Innovation Project (Grant No.CYB25253).

## References

- Chen, Q.; Li, S.; Liu, Y.; Pan, S.; Webb, G. I.; and Zhang, S. 2025. Uncertainty-Aware Graph Neural Networks: A Multi-hop Evidence Fusion Approach. *IEEE Transactions on Neural Networks and Learning Systems*, 1–15.
- Cheng, Z.; Liu, Y.; Zhong, T.; Zhang, K.; Zhou, F.; and Yu, P. S. 2025. Disentangling Inter- and Intra-Cascades Dynamics for Information Diffusion Prediction. *IEEE Transactions on Knowledge and Data Engineering*, 37(8): 4548–4563.
- Cheng, Z.; Ye, W.; Liu, L.; Tai, W.; and Zhou, F. 2023. Enhancing Information Diffusion Prediction with Self-Supervised Disentangled User and Cascade Representations. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 3808–3812.
- Cheng, Z.; Zhou, F.; Xu, X.; Zhang, K.; Trajcevski, G.; Zhong, T.; and Yu, P. S. 2024. Information cascade popularity prediction via probabilistic diffusion. *IEEE Transactions on Knowledge and Data Engineering*.
- Friedrich, T.; Göbel, A.; Klodt, N.; Krejca, M. S.; and Pappik, M. 2024. The irrelevance of influencers: Information diffusion with re-activation and immunity lasts exponentially long on social network models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 17389–17397.
- Fuchsgruber, D.; Wollschläger, T.; and Günnemann, S. 2024. Energy-based Epistemic Uncertainty for Graph Neural Networks. In Globerson, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J.; and Zhang, C., eds., *Advances in Neural Information Processing Systems*, volume 37, 34378–34428. Curran Associates, Inc.
- Han, S.; Zhou, Z.; Chen, J.; Hao, Z.; Zhou, S.; Wang, G.; Feng, Y.; Chen, C.; and Wang, C. 2025. Uncertainty-Aware Graph Structure Learning. In *Proceedings of the ACM on Web Conference 2025*, WWW '25, 4863–4874. New York, NY, USA: Association for Computing Machinery. ISBN 9798400712746.
- He, W.; Xiao, Y.; Huang, M.; Mou, X.; Wang, R.; and Li, Q. 2025a. A Pattern-Driven Information Diffusion Prediction Model Based on Multisource Resonance and Cognitive Adaptation. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '25, 592–601. New York, NY, USA: Association for Computing Machinery. ISBN 9798400715921.
- He, W.; Xiao, Y.; Mou, X.; Li, T.; Wang, R.; and Li, Q. 2025b. An Information Diffusion Prediction Model Aligning Multiple Propagation Intentions With Dynamic User Cognition. *IEEE Transactions on Computational Social Systems*, 1–16.
- Jiao, P.; Chen, H.; Bao, Q.; Zhang, W.; and Wu, H. 2024. Enhancing Multi-Scale Diffusion Prediction via Sequential Hypergraphs and Adversarial Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 8571–8581.
- Jing, X.; Jing, Y.; Lu, Y.; Deng, B.; Chen, X.; and Yang, D. 2025. CasFT: Future Trend Modeling for Information Popularity Prediction with Dynamic Cues-Driven Diffusion Models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 11906–11914.
- Leskovec, J.; Backstrom, L.; and Kleinberg, J. 2009. Meme-tracking and the dynamics of the news cycle. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, 497–506.
- Li, H.; Xia, C.; Wang, T.; Wang, Z.; Cui, P.; and Li, X. 2024. GRASS: Learning Spatial–Temporal Properties From Chainlike Cascade Data for Microscopic Diffusion Prediction. *IEEE Transactions on Neural Networks and Learning Systems*, 35(11): 16313–16327.
- Rajkumar, K.; Saint-Jacques, G.; Bojinov, I.; Brynjolfsson, E.; and Aral, S. 2022. A causal test of the strength of weak ties. *Science*, 377(6612): 1304–1310.
- Sankar, A.; Zhang, X.; Krishnan, A.; and Han, J. 2020a. Inf-VAE: A variational autoencoder framework to integrate homophily and influence in diffusion prediction. In *Proceedings of the 13th international conference on web search and data mining*, 510–518.
- Sankar, A.; Zhang, X.; Krishnan, A.; and Han, J. 2020b. Inf-VAE: A Variational Autoencoder Framework to Integrate Homophily and Influence in Diffusion Prediction. In *Proceedings of the 13th International Conference on Web Search and Data Mining*, WSDM '20, 510–518. New York, NY, USA.
- Sun, L.; Rao, Y.; Zhang, X.; Lan, Y.; and Yu, S. 2022. MS-HGAT: memory-enhanced sequential hypergraph attention network for information diffusion prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 4156–4164.
- Wang, D.; Wei, L.; Yuan, C.; Bao, Y.; Zhou, W.; Zhu, X.; and Hu, S. 2022. Cascade-enhanced graph convolutional network for information diffusion prediction. In *International Conference on Database Systems for Advanced Applications*, 615–631. Springer.
- Wang, D.; Zhou, W.; and Hu, S. 2024. Information Diffusion Prediction with Graph Neural Ordinary Differential Equation Network. In *Proceedings of the 32nd ACM International Conference on Multimedia*, MM '24, 9699–9708. New York, NY, USA: Association for Computing Machinery. ISBN 9798400706868.
- Wang, H.; Yang, C.; and Shi, C. 2021. Neural information diffusion prediction with topic-aware attention network. In

*Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 1899–1908.

Xu, X.; Zhou, F.; Zhang, K.; Liu, S.; and Trajcevski, G. 2023. CasFlow: Exploring Hierarchical Structures and Propagation Uncertainty for Cascade Prediction. *IEEE Transactions on Knowledge and Data Engineering*, 35(4): 3484–3499.

Yang, C.; Wang, H.; Tang, J.; Shi, C.; Sun, M.; Cui, G.; and Liu, Z. 2021. Full-scale information diffusion prediction with reinforced recurrent networks. *IEEE Transactions on Neural Networks and Learning Systems*, 34(5): 2271–2283.

Yuan, C.; Li, J.; Zhou, W.; Lu, Y.; Zhang, X.; and Hu, S. 2020. DyHGNCN: A Dynamic Heterogeneous Graph Convolutional Network to Learn Users' Dynamic Preferences for Information Diffusion Prediction. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2020, Ghent, Belgium, September 14–18, 2020, Proceedings, Part III*, 347–363. Berlin, Heidelberg.

Zhang, H.-K.; Zhang, Y.-G.; Zhou, Z.; and Li, Y.-F. 2024. HONGAT: graph attention networks in the presence of high-order neighbors. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 16750–16758.

Zhong, E.; Fan, W.; Wang, J.; Xiao, L.; and Li, Y. 2012. Comsoc: adaptive transfer of user behaviors over composite social network. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, 696–704.

Zhong, T.; Zhang, J.; Cheng, Z.; Zhou, F.; and Chen, X. 2024. Information Diffusion Prediction via Cascade-Retrieved In-context Learning. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '24*, 2472–2476. New York, NY, USA: Association for Computing Machinery. ISBN 9798400704314.