Source Localization in Continuous-Time Propagation via Spectral ODE Modeling

Dongpeng Hou
hdp0918@mail.nwpu.edu.cn
Northwestern Polytechnical University
Xi'an, Shaanxi, China

Yuchen Wang
wany810@mail.nwpu.edu.cn
Northwestern Polytechnical University
Xi'an, Shaanxi, China

Giulio Cimini
giulio.cimini@roma2.infn.it
University of Rome Tor Vergata
Roma, Italy

Roberto Benzi
Roberto.benzi@uniroma2.eu
University of Rome Tor Vergata
Roma, Italy

Huixiang Zhang*
zhanghuixiang@nwpu.edu.cn
Northwestern Polytechnical University
Xi'an, Shaanxi, China

Zhen Wang
w-zhen@nwpu.edu.cn
Northwestern Polytechnical University
Xi'an, Shaanxi, China

Chao Gao
cgao@nwpu.edu.cn
Northwestern Polytechnical University
Xi'an, Shaanxi, China

Abstract

Source localization has attracted increasing attention in recent years due to its vital role in governing the harmful propagation. However, existing localization methods do not fully consider the temporal characteristics in propagation and struggle to leverage the continuous-time information of real-world propagation scenarios. Moreover, the aggregation ability of GNN based localization models is limited by the structural noise commonly present in complicated real-world topologies. To address these challenges, a Spectral Neural Ordinary Differential Equation (SNODE) is proposed to infer the source in real-world continuous-time scenarios. First, the forward propagation is formulated as a flow based ODE system, and the source localization problem is transformed into an inverse ODE modeling task. Second, a neural process based on a graph variational autoencoder is introduced to encode global latent propagation patterns as a conditioning variable for the ODE system. Third, a spectral graph optimization is performed to suppress topological noise by filtering out high-frequency components that degrade the quality of graph aggregation in the neural process. Comprehensive experiments demonstrate that SNODE not only outperforms the optimal baseline in real-world continuous-time propagation scenarios with an average performance improvement of 43.8%, but also achieves consistently superior performance in synthetic discrete-time datasets with an improvement of 4.5%, highlighting its strong generalization ability in different propagation settings. Our code is available at https://github.com/cgao-comp/SNODE.

*Corresponding author.

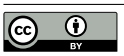

WWW '26, Dubai, United Arab Emirates.

ACM ISBN 979-8-4007-2307-0/2026/04
https://doi.org/10.1145/3774904.3792402

CCS Concepts

• **Human-centered computing** → **Social network analysis**; **Social networks**.

Keywords

Source Localization, Social Networks, Graph Mining, Continuous-time Propagation, Spectral Optimization

ACM Reference Format:
Dongpeng Hou, Yuchen Wang, Giulio Cimini, Roberto Benzi, Huixiang Zhang, Zhen Wang, and Chao Gao. 2026. Source Localization in Continuous-Time Propagation via Spectral ODE Modeling. In *Proceedings of the ACM Web Conference 2026 (WWW '26), April 13–17, 2026, Dubai, United Arab Emirates.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3774904.3792402

1 Introduction

Social networks significantly facilitate information spreading, influencing decision-making processes, user behavior, and public opinions [27]. The ease and rapidity of information propagation in social networks also enable harmful information such as rumors or misinformation to propagate swiftly and widely, often leading to significant negative impacts [6]. Therefore, the task of identifying the initial source user of harmful information has become critically important for timely interventions and effective social management [18, 36].

In recent years, more and more methods have been proposed to locate the sources in different scenarios with various settings [15, 17, 38]. Traditional methods typically rely on source prominence or heuristic rules [1, 34, 37, 40]. Although such methods have shown theoretical advances under fixed propagation parameters or specific network structures [45], these strong assumptions limit their generalization in different scenarios.

Further, graph deep learning techniques have provided new ideas for modeling propagation dynamics [44]. By learning implicit correlations between network topology and propagation features end-to-end [42], graph neural networks (GNNs) partially reduce the reliance on prior knowledge [8]. However, due to the inherent reliance on neighborhood aggregation, GNNs are fundamentally sensitive to the structural characteristics of the underlying graph [12, 19]. Therefore, when applied to real-world propagation scenarios with complicated interactions and noisy topologies, vanilla GNN modules struggle to learn robust representations. Challenges emerge mainly from two aspects:

(1) **Propagation Dynamics:** Many existing source localization methods are developed and evaluated under discrete-time simulation propagation environments, such as the Susceptible-Infected (SI) or Independent Cascade (IC) models, which simplify real-world propagation dynamics into discrete intervals. As a result, source localization methods built on such environments often do not consider and fail to capture the fine-grained temporal characteristics of continuous-time propagation dynamics, limiting their applicability and generalization in real-world scenarios.
(2) **Network Topology:** In contrast to the simplified yet robust assumptions of classical synthetic networks, such as Random Tree (RT) models with a single parameter or the Barabási-Albert (BA) network with stable scale-free characteristics, real-world propagation exhibits inherently irregular structures and high-frequency noise.[1] This stems from the complex decision-making behaviors of users and the uncertainty in interpersonal connections, which introduce significant randomness into the observed cascades. Therefore, it is difficult for GNNs to effectively learn robust representations compromised by such structural noise, limiting their performance in real-world propagation.

Motivated by these critical limitations, we propose a novel Spectral Neural ODE (SNODE) framework for continuous-time source localization. SNODE integrates continuous-time modeling via ordinary differential equations (ODEs), neural variational inference for capturing global propagation dynamics, and spectral optimization to robustly filter out high-frequency noise inherent in real-world networks. Specifically, by leveraging the invertibility of ODE flows [3], SNODE allows the inference task of source localization to be conducted in reverse, starting from observed states and tracing back to the initial source. To encode the latent graph and propagation dynamics context within the ODE system, we further introduce a neural process based on a graph variational autoencoder (GVAE), serving as a global latent variable. Moreover, to ensure the latent distributional posterior aligns closely with the observed prior distributions, we employ a Kullback-Leibler (KL) divergence constraint. Finally, recognizing that structural noise in real-world data may impair the graph aggregation quality within the GVAE module, we apply spectral graph optimization that selectively preserves meaningful low-frequency patterns while suppressing detrimental high-frequency signals. Our main contributions are summarized as follows:

- A closed-form ODE flow is derived for source localization under continuous-time propagation scenarios, enabling reverse inference from observed snapshots to the initial states.
- A latent neural process modeled via graph variational inference is introduced to embed the global propagation dynamics and graph topology into the ODE system, enabling robust prior conditioning during the inference phase.
- A Laplacian spectral optimization strategy is developed to filter high-frequency structural noise in real-world social propagation data, enhancing GVAE's aggregation and robustness.

2 Related Work

In practice, snapshot data capturing the network state at convenient time points is usually available, making snapshot-based localization a mainstream research direction. Early works such as LPSI [40] estimate source prominence, while GCNSI [8] applies graph convolutional networks to identify multiple rumor sources without predefined propagation models. Subsequent methods, including IVGD [35], MCGNN [32], and SL_VAE [22], extract dynamic propagation representations before source inference. Encoder–decoder frameworks [13] have been introduced to learn user influence matrices, which are then utilized for source prediction. DSLF [39] explores cross-platform source localization by leveraging shared topic-level propagation patterns between heterogeneous social networks. Recent research has incorporated diffusion and contrastive learning techniques. Huang et al.[14] formulated a denoising diffusion-based localization framework to address the ill-posed nature of the inverse problem. Yan et al.[43] proposed a discrete diffusion model using reversible residual blocks with GCNs. Cheng et al.[4] developed a heuristic GNN framework addressing user heterogeneity, multisource detection, and class imbalance. Bao et al.[2] introduced a self-supervised approach combining graph contrastive learning and feature augmentation.

Despite these advances, most existing methods still rely on static underlying topologies or synthetic propagation using simple propagation models as their experimental testbed. Such simplified environments neglect the noise and uncertainty in real-world social propagation. Consequently, conventional GNN-based modules often suffer from degraded aggregation quality and limited robustness when handling irregular, noisy network structures in real-world propagation scenarios characterized by random and complicated user interactions. In addition, these methods are still designed for discrete-time settings and cannot capture the fine-grained characteristics of real social propagation that evolves in continuous time, which reduces their performance in continuous-time source localization tasks.

3 Preliminary and Problem Formulation

A set of snapshots with continuous timestamps $\{X_{>0}, G\}$ is observed based on an undirected network $G = (V, E, \mathcal{F})$, where V, E, and $\mathcal{F}$ are the node (user) set, edge (relationship) set, and node attributes, respectively. Defining that X_0 is the initial snapshot in which only the source node is activated, our objective is to solve X_0 given the

[1] The indicator specific heat quantifies entropy fluctuations in the network topology and provides a spectral perspective for evaluating the impact of topological high-frequency noise on structural stability and robustness. And our spectral analysis reveals that the specific heat curve of real data lacks the stable plateau typically observed in classical networks (e.g., BA and RT networks). Detailed interpretation is provided in Sec. 4.3.

observed snapshots $X_{>0}$ and network structure G. Formally, this inverse source inference task can be represented as:

$$X_0^* = \arg\max_{X_0} \; p(X_0 \mid X_{>0}, G). \tag{1}$$

The variational inference is employed to derive the evidence lower bound (ELBO) of Eq. (1):

$$\begin{aligned}
&\log p(X_0|X_{>0}, G) = \log \int_z p(X_0, z \mid X_{>0}, G)\, dz \\
&= \log \int_z q(z|X, G) \frac{p(X_0, z \mid X_{>0}, G)}{q(z \mid X, G)}\, dz \\
&\geq E_{q(z|X,G)} \left[\log \frac{p(X_0, z|X_{>0}, G)}{q(z|X, G)} \right] \quad \text{(Jensen's inequality)} \\
&= E_{q(z|X,G)} \log p(X_0|X_{>0}, z, G) - D_{\mathrm{KL}}(q(z|X, G)||p(z|X_{>0}, G))
\end{aligned}, \tag{2}$$

where the learnable latent variable z is introduced to encode the global context of the propagation dynamics. The first term evaluates how well the model can reconstruct X_0 given the observations and latent context z, while the KL divergence regularizes $q(z \mid X, G)$ to approximate the unknown true posterior. Since the backward likelihood $p(X_0 \mid X_{>0}, z, G)$ is still hard to compute directly, we train a neural ODE-based forward model $p(X_{>0} \mid X_0, z, G)$ in supervised settings, leveraging the reversibility of neural ODEs [3] to recover X_0 during inference:

$$X(t) = X(t + \Delta t) - \int_t^{t+\Delta t} f(X(t'), G, \mathcal{F}, z)\, dt', \tag{3}$$

where f defines the ODE dynamics system.

4 Method

SNODE leverages the flow nature of ODEs to enable reverse inference for source localization [3]. Section 4.1 derives the forward and inverse conditional formulation of the ODE for the localization task. Section 4.2 introduces a neural process that provides global context to guide the ODE system. Section 4.3 further optimizes the graph aggregation process within the neural process via spectral filtering. The framework of SNODE is demonstrated in Fig. 1.

4.1 ODE based Source Inference

Sec. 3 transforms the inverse source localization problem into a forward propagation modeling task leveraging the flow nature of the ODE. To further model the propagation dynamics in the ODE rigorously, we adopt a continuous-time exponential propagation process [11, 16, 30]. Specifically, under the introduced latent variable z, the propagation process along edge $\langle v_i, v_j \rangle \in E$ from v_i to v_j is determined by a propagation probability density function (PDF) $p(v_i, v_j, z, t) = \begin{cases} f(v_i, v_j, z) e^{-f(v_i, v_j, z)t}, & \text{if } t \geq 0 \\ 0, & \text{if } t < 0 \end{cases}$, where $f(v_i, v_j, z) \in [0, \infty)$ represents the propagation rate determining the strength of node v_i influencing v_j controlled by the global latent variable z. Accordingly, the cumulative distribution function (CDF) $F(v_i, v_j, z, t) = \int_0^t p(v_i, v_j, z, t')\, \mathrm{d}t' = 1 - \mathrm{e}^{-f(v_i, v_j, z)t}$ is the probability that node v_i succeeds in influencing v_j along edge $\langle v_i, v_j \rangle$ by time t.

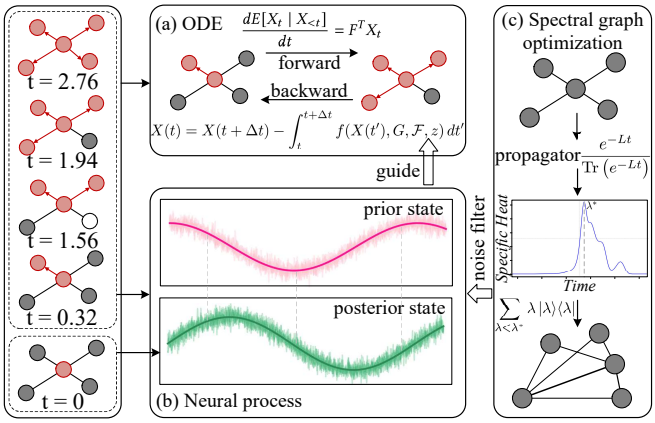

Figure 1: The framework of SNODE includes three critical components. (a) The derived ODE is the framework for continuous-time source localization. (b) As a global-context encoder, the neural process guides the ODE framework in sensing the overall propagation dynamics at different timestamps. (c) Considering that the topological noise can influence the graph aggregation in the neural process, a spectral graph optimization is introduced to integrate out the high frequencies.

Lemma 1. *Given historical observation up to time t, the probability that node v_i becomes activated or influenced during a time interval τ can be expressed as:*

$$p(X_{t+\tau}(v_i) = 1 \mid X_{\leq t}, z, G) = 1 - e^{-\tau \sum_{j \in \mathcal{N}(v_i)} f(v_i, v_j, z)}, \tag{4}$$

where $\mathcal{N}(v_i)$ represents the neighbors of node v_i.

Consequently, the instantaneous activation rate for all nodes in the network is established by the following ordinary differential equation in the forward timeline:

$$\frac{dE[X_t \mid X_{<t}]}{dt} = F^T X_t, \tag{5}$$

where $F \in \mathbb{R}+^{N \times N}$ is the activation rate matrix with elements $F_{ij} = f(v_i, v_j, z)$. Given that the state vector X_t must reside within the range [0,1], it is clipped based on $\hat{X}_t = \max(0, \min(1, X_t))$.

Integrating over the continuous time interval from 0 to an arbitrary time T in the forward timeline, given the initial state X_0, we obtain the expected state at time T:

$$E[X_T \mid X_0] = X_0 + \int_0^T \frac{dE[X_t \mid X_{<t}]}{dt} dt \quad = X_0 + \int_0^T F^T \hat{X}_t dt. \tag{6}$$

In the theory of ODEs, the solution defines a flow on the state space that can be uniquely integrated both forward and backward in time under appropriate conditions [3]. This reversibility enables us to trace the dynamics in either time direction, making ODEs particularly suited for the source inference task, where we reverse the propagation process to recover the initial activation state.

Based on **Lemma 1**, we can get instantaneous activation rate $\gamma_t(v_i) = \sum_{v_j \in \mathcal{N}(v_i)} f(v_i, v_j, z) X_t(v_j)$. Then, when defining the reverse infinitesimal time interval $|\overleftarrow{\tau}| \to 0^+$, we further obtain the following integral form for source inference:

$$E[X_0 \mid X_T] = X_T + \int_0^T -F^T \hat{X}_t dt. \tag{7}$$

Lemma 2. *For a given node $v_i \in V$, considering the reverse-time propagation process at time t, the normalized expected deactivation rate is given by:*

$$\Gamma_t(v_i) = -\mathbb{E}\left[X_t(v_i) \sum_{v_j \in \mathcal{N}(v_i)} f(v_j, v_i, z) X_t(v_j) \mid X_{>t}\right]. \tag{8}$$

Lemma 2 gives the expected deactivation rate.We then prove that it conforms to the inverse ODE.

Theorem 1. *Given the random observation X_t at time t, consider the infinitesimal interval $\overleftarrow{\tau} \to 0^+$, we establish the following ODE for inference.*

$$\frac{dE[X_t|X_{>t}]}{dt} = -X_t \odot (F^T X_t), \tag{9}$$

where $\odot$ is the Hadamard product.

Then, combined with the clip strategy, the normalized integral form for source inference is obtained.

$$E[X_0 \mid X_T] = X_T + \int_0^T -\hat{X}_t \odot F^T \hat{X}_t dt. \tag{10}$$

4.2 Neural Process for Propagation Rate Evaluation

The inverse ODE system for source inference is introduced in Sec. 4.1, where the propagation rate matrix F models the intensity of information spread from node v_i to v_j. The learnable edge function $f(\cdot)$ can be instantiated with deep modules, e.g., $f(v_i, v_j) = \text{MLP}([\mathcal{F}(v_i) \| \mathcal{F}(v_j)])$, where $\|$ denotes vector concatenation. However, modeling propagation solely based on local node features may overlook the spatio-temporal dynamics. To address this and capture heterogeneous influence patterns, we incorporate a shared latent variable z, inspired by neural latent variable models [9, 28]. The edge function is thus redefined as $f(v_i, v_j, z) = \text{MLP}([\mathcal{F}(v_i) \| \mathcal{F}(v_j) \| z])$, enabling the model to learn edge-wise propagation rates influenced by both local attributes and the global context z.

As z is unobserved, we learn it via graph variational autoencoder. An encoder is introduced to approximate the posterior of z from observations $\{X, G, \mathcal{F}\}$, where X encodes dynamic features and $\mathcal{F}$ encodes inherent features. The encoder employs a GNN over G to aggregate $\{X, \mathcal{F}\}$, with temporal encodings integrated into the input. Node-level representations h_i are computed and aggregated into a graph-level embedding $h_{\text{graph}} = \text{Pool}(\{h_i\}_{i=1}^{|V|})$, e.g., using attention pooling [20]. The variational posterior of z is then modeled as a multivariate Gaussian with diagonal covariance.

$$z \sim q_\phi(z|\Psi(X, G, \mathcal{F})) = \mathcal{N}(z; \mu, \sigma^2 I), \tag{11}$$

where Ψ is the graph representation operator, μ represents the mean vector of the posterior distribution of the global latent variable z, and σ^2 represents the variance vector of the posterior distribution.

A critical strategy for the neural process in SNODE is to align the prior distribution of observations in the test scenario (i.e., X_0 is unavailable) with the posterior distribution of observations in the training scenario (i.e., X_0 is available). Since Eq. (2) has derived the KL divergence item, this similarity is achieved through the KL divergence loss:

$$\begin{aligned}
&\mathcal{L}_{\text{KL}} = D_{\text{KL}}\big(\mathcal{N}(\mu_{\text{post}}, \sigma^2_{\text{post}} I) \,\|\, \mathcal{N}(\mu_{\text{prior}}, \sigma^2_{\text{prior}} I)\big) \\
&= D_{\text{KL}}\big(q_\phi(z \mid \Psi(\{X_0, X_{>0}\}, G, \mathcal{F})) \,\|\, q_\phi(z \mid \Psi(X_{>0}, G, \mathcal{F}))\big). \\
&\text{s.t. } \big[\mu_{\text{post}}, \sigma_{\text{post}}\big] \leftarrow \big[\mu_{\text{post}}, \sigma_{\text{post}}\big] - \eta \nabla_{[\mu_{\text{post}} \sigma_{\text{post}}]} \mathcal{L}_{\text{pred}}, \\
&\big[\mu_{\text{prior}}, \sigma_{\text{prior}}\big] \leftarrow \big[\mu_{\text{prior}}, \sigma_{\text{prior}}\big] - \eta \nabla_{[\mu_{\text{prior}} \sigma_{\text{prior}}]} \mathcal{L}_{\text{KL}}.
\end{aligned} \tag{12}$$

Here, the posterior parameters $(\mu_{\text{post}}, \sigma_{\text{post}})$ are optimized through global gradient descent with respect to the main classification loss $\mathcal{L}_{\text{pred}}$ for source localization, while the prior parameters $(\mu_{\text{prior}}, \sigma_{\text{prior}})$ are optimized through local gradient descent for KL alignment from VAE module. During the training phase, $z \sim \mathcal{N}_{\text{post}}$ is deployed in the edge function f, and the KL loss is explicitly used to push the prior distribution toward the posterior distribution. During the test phase, when the initial observation X_0 is unavailable, the GVAE has been optimized such that the prior latent distribution $p(z \mid X_{>0})$ closely approximates the posterior latent distribution $p(z \mid \{X_0, X_{>0}\})$.

4.3 Spectral Optimization for Social Graph Aggregation

As demonstrated in Sec. 4.2, the performance of the global latent distribution is critically conditional on the quality of the GNN aggregation. However, the real-world social cascades are formed based on complicated interactions. Spectral analysis indicates that the topology of real-world cascades differs markedly from that of classical benchmark networks. As seen in Fig. 2, unlike BA and RT networks, which exhibit a clear plateau in their specific heat curves due to their regular structure with robust assumptions or stable scale-free topology, the specific heat of real-world propagation graphs displays dramatic fluctuations and instability. This phenomenon indicates significant high-frequency noise and redundant connections inherent in real data, which disrupts the representational capacity and efficiency of the GNN encoder.

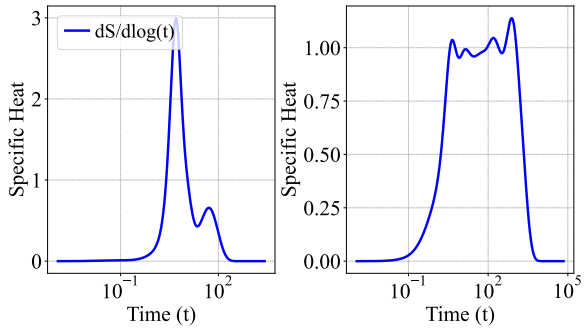

(a) Real-world network (b) Barabasi Albert network

Figure 2: The specific heat of the real-world network and the Barabási-Albert network.

To address this issue, we use the spectral decomposition that optimizes the network Laplacian to reduce its rank and filter out redundant connections with high-frequency noise [33]. Considering the information propagation process in the network from the

perspective of entropy, its dynamic equation converging to stability can be described as the network propagator $\hat{K} = e^{-Lt}$.

Lemma 3. *The network propagator $\hat{K} = e^{-Lt}$ is a spectral filter that suppresses high-frequency components and preserves low-frequency structural information over time.*

In parallel, the canonical system entropy of a network is defined as:

$$\begin{aligned} S\big[\hat{\rho}(t)\big] &= \mathrm{Norm}\left(-\mathrm{Tr}\left[\hat{\rho}(t)\log\hat{\rho}(t)\right]\right) \\ &= -\frac{1}{\log(N)}\sum_{i=1}^{N}\mu_i(t)\log\mu_i(t), \\ \hat{\rho}(t) &= \frac{\hat{K}}{\mathrm{Tr}(\hat{K})} = \frac{e^{-Lt}}{\mathrm{Tr}\left(e^{-Lt}\right)}, \end{aligned} \tag{13}$$

where $\hat{\rho}(t)$ is tantamount to the canonical (i.e., regularized) density operator in statistical physics (or to the functional over fields configurations) [10, 29], and μ_i represents the specific $\hat{\rho}(t)$ set of eigenvalues.

Lemma 4. *The entropy $S \in [0,1]$ quantifies entropic transitions (or information propagation transitions) over the network. As the propagation time t increases from 0 to ∞, $S[\hat{\rho}(t)]$ decreases monotonically from 1 (highly localized and heterogeneous phase) to 0 (fully mixed and homogeneous phase), reflecting the transition from localized to global information propagation.*

The temporal derivative of the entropy, $\frac{\mathrm{d}S}{\mathrm{d}(\log t)}$, represents the specific heat of the system. In particular, a constant specific heat reflects the scale-invariant nature of the network (e.g., BA network and RT network in Fig. 2). While in our scenarios, S indicates the system variant from the disordered state to the stable state. Hence, $\frac{\mathrm{d}S}{\mathrm{d}(\log t)}$ measures the rate at which the system approaches stability per unit time interval. Since S is monotonically decreasing over time, implying a negative slope. We define the specific heat as the absolute form $C(t) = -\frac{\mathrm{d}S}{\mathrm{d}(\log t)}$ to quantify the magnitude of change. This definition ensures that $C(t) \in \mathbb{R}^+$, i.e., the specific heat is always a non-negative real quantity that reflects the intensity of entropy dissipation in the system over time. And a higher $C(t)$ indicates a rapid transition to stability, while a lower $C(t)$ implies a slower convergence toward stability.

If t is very large, then only the pure low frequency is held, resulting in excessive loss of information on the graph. Without loss of generality, let us consider the case in which we can preserve sufficient topology information but also filter out more noise on the graph, at a meaningful time point t^* (for example, where C shows a maximum at the first time). And only modes on topology information larger than t^* are kept. Then, we use the bra-ket formalism in which the Laplacian operator is $\sum_\lambda \lambda\,|\lambda\rangle\langle\lambda|$, where $|\lambda\rangle$ denotes the Laplacian eigenvector to its eigenvalue λ. We then identify and integrate out the $n < N$ eigenvalues $\lambda \geq \lambda^* = \frac{1}{t^*}$ and the corresponding eigenvectors $|\lambda\rangle$.

Intuitively, the first peak achieves the best balance between "removing as much noise as possible" and "preserving as much original structure as possible". This peak corresponds to the moment when the highest-frequency structural noise is removed for the first time on a large scale. At this point, $\hat{K}$ suppresses the unstable high-frequency Laplacian modes while fully retaining the low-frequency components that encode the essential graph topology. Actually, any value of t^* between the first and the last peak is theoretically admissible, as it simply determines the resolution scale of the spectral filtering process [24]. However, choosing a later peak would remove a large fraction of the low-frequency eigenmodes, thereby discarding much of the topological information. In downstream GNN tasks, resulting severe oversmoothing and performance degradation.

By reducing the Laplacian operator to the contribution of the N-n slow eigenvectors with $\lambda < \lambda^*$, the spectral filtering strategy of the Laplacian can be performed:

$$L' = \sum_{\lambda<\lambda^*} \lambda\,|\lambda\rangle\langle\lambda|\,. \tag{14}$$

Finally, we reconstruct an optimized graph input matrix G' from the filtered Laplacian L'. Specifically, for $i \neq j$ we set $G'_{ij} = -L'_{ij}$, and define the diagonal elements as $L'_{ii} = \sum_{j\neq i} G'_{ij}$, $G'_{ii} = 0$.

Lemma 5. *When eigenvectors corresponding to high-frequency eigenvalues are removed during spectral filtering, the resulting truncated Laplacian L' may yield a disconnected adjacency matrix $A' = -L' + \mathrm{diag}(L')$. The isolated components in A' break GNN's message-passing capability.*

To avoid this structural degeneration, we impose a conservative constraint strategy to preserve connectivity when selecting the spectral cutoff. Specifically, we aim to maximize the spectral threshold $\lambda^* = 1/t^*$, while ensuring that the resulting graph G' remains connected:

$$\begin{aligned} L' = \sum_{\lambda<\lambda^*} \lambda\,|\lambda\rangle\langle\lambda| \max_{\lambda^*} \quad \text{s.t.} \\ \mathrm{Graph}\big(-L' + \mathrm{diag}(L')\big) \text{ remains connected.} \end{aligned} \tag{15}$$

This spectral optimization process effectively suppresses high-frequency noise and redundant connections in the original graph. By feeding the optimized matrix G' into the GNN encoder, we ensure that the learned latent representations are based on the stable, low-frequency structure of the network. In turn, this improves the robustness and performance of our neural process for source inference, especially when applied to real-world propagation graphs with unstable structures.

5 Experiments

5.1 Datasets and Setting

We use three real-world datasets collected from Weibo and Twitter platforms for source localization, namely Weibo [26], Twitter15, and Twitter16 [23, 25]. The node attributes in the dataset include user description, blue verification status, location, registration date, number of posts, fans list, and followings list. The relevant statistical information of the three datasets is shown in Tab. 3.

To validate the generalizability of SNODE in discrete-time scenarios, we follow common experimental settings in prior work and use synthetic datasets simulated by SI and IC models [8, 22, 35]. The Facebook network with 4,039 nodes and 88,234 edges (#Avg(degree): 43.69) [21] and the Power Grid network with 4,941 nodes and 6,594 edges (#Avg(degree): 2.67) [41] are used as the underlying topology. Based on the network, 10% nodes are selected as ground-truth

Table 1: Source localization performance on the real-world cascades in Twitter15, Twitter16, Weibo. The bold values represent the best results, while underlined values denote the second-best.

Dataset	Twitter15		Twitter16		Weibo	
Metrics	F1	AED	F1	AED	F1	AED
DDMSL [43]	<u>0.663</u>	<u>0.442</u>	<u>0.651</u>	<u>0.450</u>	<u>0.622</u>	<u>0.672</u>
GIN-SD [5]	0.578	0.541	0.564	0.556	0.523	0.912
HFSD [4]	0.468	0.645	0.490	0.630	0.454	0.987
Diff [14]	0.612	0.532	0.605	0.539	0.578	0.774
TGASI [13]	0.515	0.637	0.511	0.642	0.497	0.921
IVGD [35]	0.368	0.851	0.342	0.884	0.330	1.211
SL_VAE [22]	0.355	0.855	0.342	0.871	0.346	1.271
SNODE	**0.943**	**0.074**	**0.936**	**0.077**	**0.905**	**0.185**

Table 2: Source localization performance on the synthetic cascades in Facebook and Power Grid network. The bold values represent the best results, while underlined values denote the second-best.

Dataset	Facebook				Power Grid			
Propagation	IC		SI		IC		SI	
Metrics	F1	AED	F1	AED	F1	AED	F1	AED
DDMSL [43]	<u>0.671</u>	<u>0.666</u>	<u>0.667</u>	<u>0.674</u>	<u>0.728</u>	<u>0.654</u>	<u>0.705</u>	<u>0.699</u>
GIN-SD [5]	0.592	0.811	0.589	0.818	0.571	1.036	0.554	1.080
HFSD [4]	0.350	1.298	0.325	1.340	0.447	1.281	0.411	1.335
Diff [14]	0.605	0.779	0.552	0.882	0.675	0.767	0.623	0.845
TGASI [13]	0.671	0.668	0.647	0.695	0.728	0.675	0.703	0.784
IVGD [35]	0.403	1.200	0.376	1.211	0.460	1.255	0.441	1.288
SL_VAE [22]	0.364	1.227	0.318	1.323	0.627	0.903	0.603	0.910
SNODE	**0.713**	**0.589**	**0.702**	**0.595**	**0.745**	**0.615**	**0.733**	**0.595**

Table 3: Statistics of the real datasets. The user scale of each cascade ranges from 10 to 57,186 (avg. 883). The depth spans 2 to 22 (avg. 6.15).

Statistic	Twitter15	Twitter16	Weibo
#cascades	1490	818	4664
#rumor	370	205	2244
#non-rumor	746	412	2082
#nodes	480,405	289,504	2,856,519
#edges	565,948	334,603	3,508,596

sources.And 1,000 sets of propagation data are independently generated for each network, and each set includes several snapshots with discrete timestamps [13]. The localization performance is evaluated using the standard F1 score [31] (F1) metric and average error distance (AED) [7] metric.

$$F1\text{-}score = \frac{2 * Precision * Recall}{Precision + Recall}, \tag{16}$$

$$\Delta_{\text{AED}} = \frac{1}{K} \sum_{k=1}^{K} d\left(r^*, r\right), \tag{17}$$

where *Recall* is the proportion of ground-truth sources that are successfully predicted. And *Precision* is the proportion of ground-truth sources in the predicted nodes. Here, *Recall* and *Precision* have the same weight in the standard F1. And $d(r^*, r)$ is the shortest distance between the predicted source and the ground-truth source.

To demonstrate the validity and novelty of the proposed localization methods, we consider DDMSL [43], GIN-SD [5], HFSD [4], Diff [14], TGASI [13], IVGD [35], SL_VAE [22] for comparison. These SOTA methods are listed as follows.

- DDMSL [43] proposes a discrete diffusion model for the SIR infection process and designs a reversible residual block with graph convolutional networks to infer the source.
- GIN-SD [5] considers incomplete user data scenarios and utilizes a positional embedding module to distinguish incomplete nodes in the source inference process.
- HFSD [4] is a GNN based heuristic framework for source detection in social networks, tackling user diversity, multi-source detection, and class imbalance.
- Diff [14] proposes a two-stage denoising diffusion model to quantify the uncertainty of the propagation process and improve the detection accuracy of source detection.
- TGASI [13] learns the influence matrix between any two users and considers the interaction dynamics in the source inference process.

- IVGD [35] is a generic framework of invertible graph diffusion models for source localization in the networks.
- SL_VAE [22] is a probabilistic model that uses deep generative models in the forward diffusion estimation model to approximate the diffusion source distribution in an uncertainty quantification way.

In our experiments, we employ a 10-fold cross-validation strategy to divide the training and test datasets. Further, SNODE reports the average prediction performance in the test datasets.

5.2 Performance in Real-world Cascades

The source localization performance based on the real-world datasets is illustrated in Tab. 1. Among the baselines, DDMSL consistently achieves the strongest localization performance and thus serves as the optimal baseline for comparison. This is mainly because it adopts an adaptive diffusion model with a reversible residual network, enabling end-to-end source inference and diffusion path reconstruction[43]. Compared with DDMSL, SNODE demonstrates superior improvements across all three datasets and evaluation metrics. Specifically, in terms of the F1 score, SNODE achieves improvements of 42.2%, 43.8%, and 45.5% (avg: 43.8%) in Twitter15, Twitter16, and Weibo datasets, respectively. Meanwhile, the corresponding reductions in average error distance reach 83.3%, 82.9%, and 72.5% (avg: 79.5%).

Such improvements and stability mainly include three innovations: (1) An ODE-based modeling framework that naturally supports continuous-time propagation and reverse inference, making it more consistent with the temporal characteristics of real-world scenarios. (2) A variational neural process that encodes global latent dynamics to enhance the performance of the interaction function within the ODE framework. (3) A spectral graph optimization technique that filters high-frequency structural noise, thereby improving the aggregation and representation ability of GNNs under complicated and uncertain real-world topologies. These results and analysis clearly demonstrate the superior localization performance of SNODE in continuous-time real-world propagation scenarios.

5.3 Performance in Synthetic Cascades

To verify that SNODE remains competitive in widely used discrete-time scenarios, we evaluate the localization performance in the synthetic cascades summarized in Tab. 2. Because simulated graphs provide no user-profile attributes, we construct basic features for each node using centrality and infection statistics of the dynamics. Nevertheless, SNODE continues to exhibit superior performance under such settings. However, compared with the optimal baseline DDMSL, the average F1 score improvements decrease from 43.8% in real-world cascades to 4.5% in synthetic cascades, while the reductions in AED drop from 79.5% to about 11.0%. Such a decrease can be attributed to two primary reasons. (1) Synthetic datasets are generated under discrete-time propagation models, which diminish the advantage of the continuous-time ODE flow model. (2) Synthetic datasets lack user profile diversity compared with the real-world datasets, reducing the effectiveness of the variational module in modeling propagation dynamics from a user feature perspective. This highlights the importance of profiles for evaluating source localization methods under real-world conditions.

Interestingly, the degree of performance improvement with SNODE still varies depending on the underlying graph structure. In the Facebook network, SNODE achieves an average F1 score improvement of 5.75%, whereas on the Power Grid network, the improvement is only 3.15%. Spectral analysis reveals that the specific heat curve of the Facebook network exhibits greater fluctuation and lacks the clear plateau observed in the Power Grid topology. This suggests that the Facebook graph contains more high-frequency structural noise, allowing the spectral optimization module in SNODE to play a more significant role in enhancing feature aggregation under less stable topological conditions.

5.4 Ablation Study

To further prove the contributions of the designed components in SNODE, we further study the influence of each component on the source detection performance. The critical modules of the SNODE model include the ODE solution $-X_t \odot (F^T X_t)$, the neural process z, the KL divergence of prior and posterior for z, and the spectral optimization. Based on these components, five variant models of SNODE are developed as follows.

(1) *w/o Continuous Integration (CI)* eliminates the continuous-time integration process and directly applies the discrete one-step update $-X_t \odot (F^T X_t)$ to predict the previous state. In contrast to the full model, which accumulates multiple infinitesimal updates $\left(-X_t \odot (F^T X_t)\right) \cdot dt$ to approximate a continuous trajectory, this variant ignores the temporal continuity of propagation and thus simplifies the ODE flow into a static discrete approximation.

(2) *w/o Expected Solution (ES)* replaces the expected deactivation rate of the backward flow $-X_t \odot (F^T X_t)$ in Eq.(8) with an intuitive form $-F^T X_t$ in Eq.(7). By removing the expectation operator, this variant no longer follows the closed-form formulation consistent with the localization setting, where uninfected or weakly infected nodes cannot act as spreaders in preceding states. Consequently, it produces a less faithful reconstruction of the propagation inverse scenarios.

(3) *w/o z* removes the global guidance from the neural process module and defines the propagation rate between nodes only as local node pairs $f(v_i, v_j) = \text{MLP}\left([\mathcal{F}(v_i) \| \mathcal{F}(v_j)]\right)$. This simplification removes the global latent variable that encodes overall propagation context, forcing the model to rely only on local node interactions.

(4) *w/o KL* removes the Kullback–Leibler divergence regularization term that aligns the prior and posterior distributions of the latent variable z. Without this constraint, the consistency between training and inference phases is weakened, leading to instability in the learned latent representations and degraded generalization and robustness performance.

(5) *w/o Spectral Optimization (SO)* performs message passing directly on the unfiltered topology. The absence of spectral filtering preserves high-frequency noise in the Laplacian spectrum, which interferes with low-frequency structural signals and degrades the effectiveness of GNN aggregation. This variant effectively demonstrates the contribution of spectral graph optimization to robust structural learning.

Table 4: Ablation study on real-world datasets. The bold values represent the best results.

Variant	Twitter15		Weibo	
	F1	AED	F1	AED
w/o CI	0.891 (↓ 5.51%)	0.122 (↑ 64.86%)	0.827 (↓ 8.62%)	0.304 (↑ 64.32%)
w/o ES	0.865 (↓ 8.27%)	0.153 (↑ 106.76%)	0.832 (↓ 8.07%)	0.291 (↑ 57.30%)
w/o z	0.912 (↓ 3.29%)	0.108 (↑ 45.95%)	0.860 (↓ 4.97%)	0.245 (↑ 32.43%)
w/o KL	0.930 (↓ 1.38%)	0.081 (↑ 9.46%)	0.863 (↓ 4.64%)	0.240 (↑ 29.73%)
w/o SO	0.924 (↓ 2.01%)	0.080 (↑ 8.11%)	0.876 (↓ 3.20%)	0.217 (↑ 17.30%)
SNODE	**0.943**	**0.074**	**0.905**	**0.185**

As shown in Tab. 4, all variant models exhibit varying degrees of performance degradation in both F1 score and AED compared to the complete SNODE. This highlights the effectiveness of each component in enhancing the model's localization performance.

5.5 Effect of the Spectral Threshold λ^* on Performance

The graph spectral optimization described in Sec.4.3 directly influences the quality of graph representation and consequently affects localization accuracy. To analyze the effect of spectral filtering, we vary the spectral cutoff threshold $\lambda^* = \frac{1}{t^*}$, which determines the number of low-frequency eigenmodes preserved during optimization. The comparison of localization performance under different thresholds is shown in Fig.3.

In SNODE, the first peak of the specific heat curve is selected as the optimal propagation time t^*, yielding the best source localization results. When $\hat{t} > t^*$ with a smaller $\hat{\lambda}^*$, more Laplacian eigenmodes are filtered out, which progressively reduces the Laplacian rank. This may keep fewer structural characteristics or even fragment the graph into weakly connected components, thereby hindering message propagation in the GNN. As $t \to \infty$, all non-zero modes vanish ($\lambda|\lambda\rangle\langle\lambda| = 0$), causing complete loss of structural characteristics, a full-connected graph remaining, and minimal localization performance. In contrast, when $t < t^*$, the filter becomes looser, retaining additional high-frequency components. At t = 0, no spectral filtering is applied, so the graph fully retains structural noise, and performance degrades to that of the *w/o SO* variant. The monotonic trends on both sides of t^* demonstrate that the chosen spectral threshold achieves a favorable trade-off between suppressing high-frequency noise and preserving informative structural patterns.

Furthermore, one insight is whether other spectral peaks beyond the first one could also provide meaningful thresholds for optimization. To investigate this, we examined subsequent local maxima of the specific heat curve. However, the results indicate that later peaks contribute little additional benefit. For instance, on the Twitter dataset, localization accuracy decreases from 0.943 at the first peak to 0.913 at the second. This decline suggests that the first spectral peak offers the most effective balance between noise suppression and information retention, while later peaks correspond to oversmoothed spectral states that gradually lose much structural fidelity.

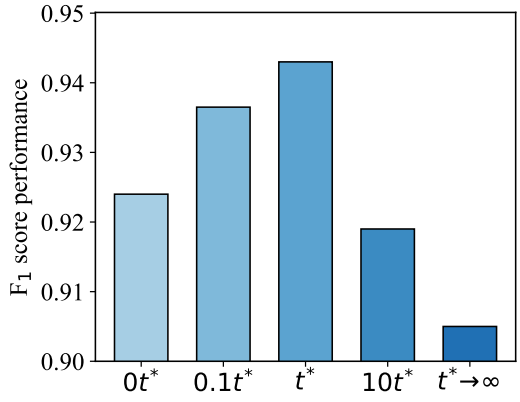

Figure 3: The impact of propagation time t^* on localization performance.

6 Conclusion

This paper proposes SNODE, a novel source localization framework that integrates continuous-time modeling via ODEs, global latent inference through neural processes, and structural noise suppression via graph spectral optimization. Extensive experiments in both real-world and synthetic datasets demonstrate that SNODE significantly outperforms SOTA methods. What's more, in real-world scenarios with continuous-time propagation, SNODE exhibits significantly greater improvements compared to synthetic datasets. Because the continuous-time modeling capability of the ODE framework enables accurate reconstruction of fine-grained temporal propagation trajectories. And the variational neural process effectively leverages user profile features that are inherently available in real-world cascades but absent in synthetic simulations, enhancing the informativeness of global latent representations. Moreover, for networks with distinct structural noise, where the specific heat curve lacks a stable plateau, the spectral optimization can offer greater gains in graph aggregation by filtering out the high-frequency modes, further boosting localization performance.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (Nos. 62471403, U22B2036), the Technological Innovation Team of Shaanxi Province (No. 2025RS-CXTD-009), the

International Cooperation Project of Shaanxi Province (No. 2025GH-YBXM-017), the Key Research and Development Program of Shaanxi (Program No. 2024GX-ZDCYL-01-20), and the Tencent Foundation and XPLORER PRIZE.

References

[1] Syed Shafat Ali, Tarique Anwar, Ajay Rastogi, and Syed Afzal Murtaza Rizvi. 2019. EPA: Exoneration and Prominence based Age for Infection Source Identification. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 891–900.

[2] Qing Bao, Ying Jiang, Wang Zhang, Pengfei Jiao, and Jing Su. 2024. Graph contrastive learning for source localization in social networks. *Information Sciences* 679 (2024), 121090.

[3] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. 2018. Neural ordinary differential equations. *Advances in neural information processing systems* 31 (2018).

[4] Le Cheng, Peican Zhu, Chao Gao, Zhen Wang, and Xuelong Li. 2024. A Heuristic Framework for Sources Detection in Social Networks via Graph Convolutional Networks. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 54, 11 (2024), 7002–7014.

[5] Le Cheng, Peican Zhu, Keke Tang, Chao Gao, and Zhen Wang. 2024. GIN-SD: source detection in graphs with incomplete nodes via positional encoding and attentive fusion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 55–63.

[6] Anneliese Depoux, Sam Martin, Emilie Karafillakis, Raman Preet, Annelies Wilder-Smith, and Heidi Larson. 2020. The pandemic of social media panic travels faster than the COVID-19 outbreak. *Journal of Travel Medicine* 27, 3 (2020), taaa031.

[7] Ming Dong, Bolong Zheng, Guohui Li, Chenliang Li, Kai Zheng, and Xiaofang Zhou. 2022. Wavefront-Based Multiple Rumor Sources Identification by Multi-Task Learning. *IEEE Transactions on Emerging Topics in Computational Intelligence* 6, 5 (2022), 1068–1078.

[8] Ming Dong, Bolong Zheng, Nguyen Quoc Viet Hung, Han Su, and Guohui Li. 2019. Multiple Rumor Source Detection with Graph Convolutional Networks. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 569–578.

[9] Marta Garnelo, Dan Rosenbaum, Christopher Maddison, Tiago Ramalho, David Saxton, Murray Shanahan, Yee Whye Teh, Danilo Rezende, and SM Ali Eslami. 2018. Conditional neural processes. In *International Conference on Machine Learning*. PMLR, 1704–1713.

[10] Walter Greiner, Ludwig Neise, and Horst Stöcker. 2012. *Thermodynamics and statistical mechanics*. Springer Science & Business Media.

[11] Shushan He, Hongyuan Zha, and Xiaojing Ye. 2020. Network diffusions via neural mean-field dynamics. *Advances in Neural Information Processing Systems* 33 (2020), 2171–2183.

[12] Dongpeng Hou, Yuchen Wang, Chao Gao, Xianghua Li, and Zhen Wang. 2024. New Localization Frameworks: User-centric Approaches to Source Localization in Real-world Propagation Scenarios. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*. 839–848.

[13] Dongpeng Hou, Zhen Wang, Chao Gao, and Xuelong Li. 2023. Sequential attention source identification based on feature representation. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*. 4794–4802.

[14] Bosong Huang, Weihao Yu, Ruzhong Xie, Jing Xiao, and Jin Huang. 2023. Two-Stage Denoising Diffusion Model for Source Localization in Graph Inverse Problems. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 325–340.

[15] Chunlin Huang, Xingwu Liu, Minghua Deng, Yang Zhou, and Dongbo Bu. 2018. A survey on algorithms for epidemic source identification on complex networks. *Chinese Journal of Computers* 41, 06 (2018), 1156–1179.

[16] Keke Huang, Ruize Gao, Bogdan Cautis, and Xiaokui Xiao. 2024. Scalable continuous-time diffusion framework for network inference and influence estimation. In *Proceedings of the ACM Web Conference 2024*. 2660–2671.

[17] Jiaojiao Jiang, Sheng Wen, Bo Liu, Shui Yu, Yang Xiang, and Wanlei Zhou. 2019. *Malicious attack propagation and source identification*. Springer, Gewerbestrasse 11, 6330 Cham, Switzerland.

[18] Jiaojiao Jiang, Sheng Wen, Shui Yu, Yang Xiang, and Wanlei Zhou. 2016. Identifying propagation sources in networks: State-of-the-art and comparative studies. *IEEE Communications Surveys & Tutorials* 19, 1 (2016), 465–481.

[19] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).

[20] Junhyun Lee, Inyeop Lee, and Jaewoo Kang. 2019. Self-attention graph pooling. In *International Conference on Machine Learning*. PMLR, 3734–3743.

[21] Jure Leskovec and Julian J Mcauley. 2012. Learning to discover social circles in ego networks. In *Advances in Neural Information Processing Systems*. 539–547.

[22] Chen Ling, Junji Jiang, Junxiang Wang, and Zhao Liang. 2022. Source Localization of Graph Diffusion via Variational Autoencoders for Graph Inverse Problems. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1010–1020.

[23] Xiaomo Liu, Armineh Nourbakhsh, Quanzhi Li, Rui Fang, and Sameena Shah. 2015. Real-time Rumor Debunking on Twitter. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*. 1867–1870.

[24] Lorenzo Lucarini, Giulio Cimini, and Pablo Villegas. 2025. Geometric Criticality in Scale-Invariant Networks. *arXiv preprint arXiv:2507.11348* (2025).

[25] Jing Ma, Wei Gao, Prasenjit Mitra, Sejeong Kwon, Bernard J. Jansen, Kam-Fai Wong, and Cha Meeyoung. 2016. Detecting Rumors from Microblogs with Recurrent Neural Networks. In *the 25th International Joint Conference on Artificial Intelligence*. 3818–3824.

[26] Jing Ma, Wei Gao, and Kam-Fai Wong. 2017. Detect Rumors in Microblog Posts Using Propagation Structure via Kernel Learning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, Vol. 1. 708–717.

[27] Priyanka Meel and Dinesh Kumar Vishwakarma. 2020. Fake news, rumor, information pollution in social media and web: A contemporary survey of state-of-the-arts, challenges and opportunities. *Expert Systems with Applications* 153 (2020), 112986.

[28] Alexander Norcliffe, Cristian Bodnar, Ben Day, Jacob Moss, and Pietro Liò. 2021. Neural ODE Processes. In *International Conference on Learning Representations*. 1–19.

[29] RK Pathria and PAUL D Beale. 2011. Statistical Mechanics. 3rd.

[30] Jean Pouget-Abadie and Thibaut Horel. 2015. Inferring graphs from cascades: A sparse recovery framework. In *International Conference on Machine Learning*. PMLR, 977–986.

[31] Marina Sokolova, Nathalie Japkowicz, and Stan Szpakowicz. 2006. Beyond accuracy, F-score and ROC: A family of discriminant measures for performance evaluation. In *Australasian Joint Conference on Artificial Intelligence*. 1015–1021.

[32] Kuangchi Sun, Zhenfeng Huang, Hanling Mao, Aisong Qin, Xinxin Li, Weili Tang, and Jianbin Xiong. 2021. Multi-Scale Cluster-Graph Convolution Network With Multi-Channel Residual Network for Intelligent Fault Diagnosis. *IEEE Transactions on Instrumentation and Measurement* 71 (2021), 1–12.

[33] Pablo Villegas, Tommaso Gili, Guido Caldarelli, and Andrea Gabrielli. 2023. Laplacian renormalization group for heterogeneous networks. *Nature Physics* 19, 3 (2023), 445–450.

[34] Hongjue Wang and Kaijia Sun. 2020. Locating source of heterogeneous propagation model by universal algorithm. *Europhysics Letters* 131, 4 (2020), 48001.

[35] Junxiang Wang, Junji Jiang, and Liang Zhao. 2022. An Invertible Graph Diffusion Neural Network for Source Localization. In *Proceedings of the ACM Web Conference*. 1058–1069.

[36] Ranran Wang, Yin Zhang, Wenchao Wan, Min Chen, and Mohsen Guizani. 2024. Distributed Rumor Source Detection via Boosted Federated Learning. *IEEE Transactions on Knowledge and Data Engineering* 36, 11 (2024), 5986–6001.

[37] Zhen Wang, Dongpeng Hou, Chao Gao, Xiaoyu Li, and Xuelong Li. 2023. Lightweight source localization for large-scale social networks. In *Proceedings of the ACM Web Conference 2023*. 286–294.

[38] Zhen Wang, Dongpeng Hou, Shu Yin, Chao Gao, and Xianghua Li. 2024. Joint source localization in different platforms via implicit propagation characteristics of similar topics. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*. 2424–2432.

[39] Zhen Wang, Dongpeng Hou, Shu Yin, Chao Gao, and Xianghua Li. 2024. Joint source localization in different platforms via implicit propagation characteristics of similar topics. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24, Kate Larson (Ed.). International Joint Conferences on Artificial Intelligence Organization*. 2424–2432.

[40] Zheng Wang, Chaokun Wang, Jisheng Pei, and Xiaojun Ye. 2017. Multiple source detection without knowing the underlying propagation model. In *Proceedings of the AAAI Conference on Artificial Intelligence*. PALO ALTO, CA 94303 USA, San Francisco, CA, 217–223.

[41] Duncan J Watts and Steven H Strogatz. 1998. Collective dynamics of 'small-world'networks. *nature* 393, 6684 (1998), 440–442.

[42] Max Welling and Thomas N Kipf. 2016. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*. 1–14.

[43] Xin Yan, Hui Fang, and Qiang He. 2024. Diffusion model for graph inverse problems: Towards effective source localization on complex networks. *Advances in Neural Information Processing Systems* 36 (2024).

[44] Ye Yan, Yi Qian, Hamid Sharif, and David Tipper. 2012. A survey on smart grid communication infrastructures: Motivations, requirements and challenges. *IEEE Communications Surveys & Tutorials* 15, 1 (2012), 5–20.

[45] Dengyong Zhou, Olivier Bousquet, Thomas Lal, Jason Weston, and Bernhard Schölkopf. 2004. Learning with local and global consistency. *Advances in neural information processing systems* 16 (2004), 321–328.