

---

# Scaling Epidemic Inference on Contact Networks: Theory and Algorithms

---

Guanghui Min   Yinhan He   Chen Chen  
University of Virginia  
{jjm8vr, nee7ne, zrh6du}@virginia.edu

## Abstract

Computational epidemiology is crucial in understanding and controlling infectious diseases, as highlighted by large-scale outbreaks such as COVID-19. Given the inherent uncertainty and variability of disease spread, Monte Carlo (MC) simulations are widely used to predict infection peaks, estimate reproduction numbers, and evaluate the impact of non-pharmaceutical interventions (NPIs). While effective, MC-based methods require numerous runs to achieve statistically reliable estimates and variance, which suffer from high computational costs. In this work, we present a unified theoretical framework for analyzing disease spread dynamics on both directed and undirected contact networks, and propose an algorithm, **RAPID**, that significantly improves computational efficiency. Our contributions are threefold. First, we derive an asymptotic variance lower bound for MC estimates and identify the key factors influencing estimation variance. Second, we provide a theoretical analysis of the probabilistic disease spread process using linear approximations and derive the convergence conditions under non-reinfection epidemic models. Finally, we conduct extensive experiments on six real-world datasets, demonstrating our method’s effectiveness and robustness in estimating the nodes’ final state distribution. Specifically, our proposed method consistently produces accurate estimates aligned with results from a large number of MC simulations, while maintaining a runtime comparable to a single MC simulation. Our code and datasets are available at <https://github.com/GuanghuiMin/RAPID>.

## 1 Introduction

In recent years, outbreaks of infectious diseases such as COVID-19 have highlighted the critical need for accurate modeling and prediction of disease dynamics [15, 8, 5, 49]. Existing epidemic studies tend to model the disease spread process using differential equations at the population level, and often assume homogeneous populations with random mixing of initial infections within contact networks [27]. Consequently, such work can only provide a rough estimate of the overall dynamics of epidemics, while missing important details about individual- or local-level infection dynamics in the networks. In practice, individual-level infection dynamic inference has important practical implications. For example, individuals with a high probability of transmitting the disease can be prioritized for vaccination, thus reducing the overall infected population [22]. In addition, understanding the results of local infections can provide a more accurate projection of the global trajectory of the epidemic. Although some recent studies have enabled advances in understanding the interplay of network structure and epidemic spread process, they can only provide an insight about the epidemic at a high level by determining whether an outbreak will escalate into a complete pandemic [50, 16, 43], while neglecting the impact of distribution of initial cases and offering limited information on the infection states and expected size of the infected population.

Currently, inferring the infectious state of individual nodes in the epidemic network heavily depends on the Monte Carlo (MC) simulations, as they do not rely on assumptions of population homogeneity or specific distributions of initial infections [60, 14, 10, 11, 21, 19]. For example, during the COVID-19 pandemic, MC simulations have been widely employed to predict infection peaks, estimate reproduction numbers, and assess the impact of non-pharmaceutical interventions (NPIs) [60, 14]. Due to the simplicity and robustness of MC simulations, they are widely adopted in various related application domains such as information dissemination and cybersecurity, where network-based contagion models play a crucial role [26, 47, 57]. However, a non-negligible drawback for MC simulations is their high computational cost, as it requires hundreds or even thousands of runs of simulations to reach a statistically reliable result due to the inherent uncertainty and variability of the disease spread process [43].

In this paper, we address the high computational cost of MC simulations for modeling the disease spread process by presenting an efficient framework to estimate the infection dynamics on large networks while ensuring high estimation accuracy on predicting the infection states of individual nodes at convergence. To our knowledge, we are the first to theoretically analyze how the network structure impacts the variance of epidemic simulations and provide a linear approximation scheme for nonlinear epidemic dynamics with provable performance guarantees. Our key contributions are as follows:

- **Theoretical Analysis:** We present three key results: (1) we derive the variance lower bound of the Monte Carlo estimator on large contact networks and reveal how the reproduction number, average out-degree, initial infection ratio, and the number of simulations affect the variance; (2) we provide a linear approximation of the nonlinear epidemic dynamics, clarifying the contribution of neighborhoods to disease spread; and (3) we establish global convergence conditions for probability-based propagation under non-reinfection epidemic models.
- **Algorithm and Results:** Building on these theoretical insights, we propose **RAPID** (**R**esidual-**A**ware **P**ropagation for **I**nfection **D**ynamics), a local approximation algorithm designed to efficiently estimate the node-level state distribution in large networks while preserving interpretability and reliability throughout the computational process. In particular, RAPID consistently produces accurate estimates aligned with results from a large number of MC simulations, while maintaining a runtime comparable to a single MC simulation.

## 2 Preliminary

**Notations.** We denote a general graph<sup>1</sup> as  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{A})$ , where  $\mathcal{V}$  is the set of nodes and  $N = |\mathcal{V}|$  is the number of nodes;  $\mathcal{E}$  denotes the set of edges, and the edge between nodes  $v_i$  and  $v_j$  is denoted as  $e_{ij} = (v_i, v_j)$ ;  $\mathbf{A} \in \{0, 1\}^{N \times N}$  is the adjacency matrix, where  $\mathbf{A}_{ij} = 1$  indicates that an edge exists between  $v_i$  and  $v_j$ , otherwise  $\mathbf{A}_{ij} = 0$ . In this work, we assume a standard SIR (Susceptible-Infected-Recovered) model with a constant transmission probability  $\beta$  and recovery probability  $\gamma$  governs the spread of a virus on this network. All notations and symbols are summarized in Table 1.

**Definition 1** (Individual-Level Epidemic Inference). *Given a directed contact network  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{A})$ , epidemic parameters  $\beta$  and  $\gamma$ , and an initial infected node set  $\mathcal{I}_0$ , the goal of individual-level epidemic inference is to estimate the steady-state (converged) state probability distribution  $P^i(\infty) = [P_S^i(\infty), P_I^i(\infty), P_R^i(\infty)]^T$  for each node  $i \in \mathcal{V}$ .*

This problem fundamentally differs from typical graph-based spatio-temporal prediction tasks, which aim to forecast future node-level observations (such as traffic or sensor signals) based on historical data. In contrast, node-level epidemic inference seeks to estimate the final distribution of susceptible probabilities at convergence with given initial conditions and epidemic parameters, and the contagion process follows the stochastic infection dynamics without relying on past observation sequences.

---

<sup>1</sup>In this work, the graph or network can be either directed or undirected. Without loss of generality, we derive our analysis on directed graphs as undirected graphs can be considered as a special case of directed graphs.

Table 1: Summary of key notations

|                                                           |                                                                                                                     |
|-----------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Graph</b>                                              |                                                                                                                     |
| $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{A})$    | Directed contact network with node set $\mathcal{V}$ , edge set $\mathcal{E}$ , and adjacency matrix $\mathbf{A}$ . |
| $N$                                                       | Number of nodes, $ \mathcal{V} $ .                                                                                  |
| $\bar{k}, D$                                              | Average out-degree, graph diameter.                                                                                 |
| $\mathcal{N}^{\text{in}}(i), \mathcal{N}^{\text{out}}(i)$ | In-/out-neighbors of node $i$ .                                                                                     |
| <b>Epidemic parameters</b>                                |                                                                                                                     |
| $\beta, \gamma$                                           | Transmission probability, recovery probability in unit time.                                                        |
| $\mathcal{I}_0, \alpha$                                   | Initial infected node set, with fraction $\alpha =  \mathcal{I}_0 /N$ .                                             |
| $p_e$                                                     | Edge-level transmission probability, $p_e = \beta/(\beta + \gamma)$ .                                               |
| <b>State variables</b>                                    |                                                                                                                     |
| $P_S^i(t), P_I^i(t), P_R^i(t)$                            | Probabilities that node $i$ is susceptible, infected, or recovered at time $t$ .                                    |
| $P^i(t)$                                                  | Full state vector of node $i$ : $[P_S^i(t), P_I^i(t), P_R^i(t)]^T$ .                                                |
| $P_S(t), P_I(t), P_R(t)$                                  | Vectors over all nodes, e.g., $P_S(t) = [P_S^1(t), \dots, P_S^N(t)]^T$ .                                            |
| $P(t)$                                                    | Full network state: $[P^1(t), \dots, P^N(t)]^T$ .                                                                   |
| $R_{\text{res}}(i)$                                       | Propagation residual at node $i$ .                                                                                  |

## 2.1 SIR Dynamics

The classical SIR model describes disease transmission as a continuous-time Markov process, governed by the following system of ordinary differential equations [27]:

$$\frac{dS}{dt} = -\beta \frac{SI}{N}, \quad \frac{dI}{dt} = \beta \frac{SI}{N} - \gamma I, \quad \frac{dR}{dt} = \gamma I. \quad (1)$$

These nonlinear equations are generally difficult to solve analytically, and the implicit assumption of homogeneous random mixing may not hold in complex contact networks. Notably, the SIR framework serves as a representative example of general non-reinfection epidemic models, where once nodes transition into the recovered (or removed) state, they no longer re-enter the infection dynamics [46, 28].

## 2.2 Probabilistic Disease Propagation in Neighborhood

Given a directed contact network  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{A})$ , epidemic parameters  $\beta$  and  $\gamma$ , and the full network state  $P(t)$  at time  $t$ , we model the SIR process as a non-reinfection system. Specifically, for each node  $i \in \mathcal{V}$ , the probability of being susceptible at time  $t+1$ , denoted  $P_S^i(t+1)$ , is given by its value at time  $t$  multiplied by the probability of not being infected during the interval  $(t, t+1]$ . The probability of being infected at time  $t+1$ ,  $P_I^i(t+1)$ , consists of two parts: the probability of becoming infected during  $(t, t+1]$  while being susceptible at  $t$ , and the probability

of remaining infected by not transitioning to the recovered state. Similarly, the probability of being recovered at  $t+1$ ,  $P_R^i(t+1)$ , accounts for nodes that either recovered during  $(t, t+1]$  or remained in the recovered state from the previous time step. Assuming that infection events are independent across neighbors within each unit time interval [43, 25], the resulting update equations correspond to the derivation of the message passing formulation in [24]:

$$\begin{aligned} P_S^i(t+1) &= P_S^i(t) \prod_{j \in \mathcal{V}} \mathbf{A}_{ji} (1 - \beta P_I^j(t)) \\ P_I^i(t+1) &= P_S^i(t) [1 - \prod_{j \in \mathcal{V}} \mathbf{A}_{ji} (1 - \beta P_I^j(t))] + (1 - \gamma) P_I^i(t) \\ P_R^i(t+1) &= P_R^i(t) + \gamma P_I^i(t) \end{aligned} \quad (2)$$

Based on these update equations, we develop a baseline inference algorithm, **PID** (**P**ropagation for **I**nfection **D**ynamics), as shown in Algorithm 1.

### Algorithm 1 Sketch of PID

**Require:** Graph  $\mathcal{G}$ , rates  $\beta, \gamma$ , initial infected set  $\mathcal{I}_0$ , threshold  $\varepsilon$

- 1: Initialize  $P_S, P_I, P_R, t$
- 2: **while**  $\|P(t) - P(t-1)\|_2 > \varepsilon$  **do**
- 3:   **for all**  $i \in \mathcal{V}$  **do**
- 4:     Update  $P_S^i(t), P_I^i(t), P_R^i(t)$  with Eq. (2)
- 5:   **end for**
- 6:    $t \leftarrow t + 1$
- 7: **end while**
- 8: **return**  $P_S, P_I, P_R$

### 3 Theoretical Analysis

In this section, we provide a thorough analysis on the impact of network structure and epidemic parameters on the variance of Monte Carlo estimators; then we derive a linear approximation of nonlinear epidemic dynamics that highlights neighborhood propagation, and establish convergence guarantees for probability-based propagation in non-reinfection models.

#### 3.1 Monte Carlo Estimator Variance Analysis

In many existing epidemic dynamic studies, Monte Carlo simulations are often repeated  $10^4$ – $10^5$  times without examining their necessity or computational trade-offs. In this subsection, we aim to theoretically investigate this important question: *under what conditions does the variance of the steady-state infection probability across the network become large when performing Monte Carlo simulations of epidemic dynamics, such that more simulation runs are needed to achieve sufficient accuracy?* To the best of our knowledge, this question has never been explored in prior work. Addressing this overlooked gap not only offers practical guidance for real-world simulation tasks but also motivates our development of scalable methods for epidemic inference.

**Theorem 3.1** (Monte Carlo Estimation Variance Lower Bound). *Given a directed contact network  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  with  $N = |\mathcal{V}|$ , average out-degree  $\bar{k}$ , and diameter  $D$ . Let  $\mathcal{I}_0 \subseteq \mathcal{V}$  be the initially infected node set with fraction  $\alpha := |\mathcal{I}_0|/N$ . Assume an SIR model parameterized by infection probability  $\beta$  and recovery probability  $\gamma$ . Using  $M$  independent Monte Carlo simulations to estimate each node's infection probability  $p_i$ , the average variance of the estimator  $\hat{p}_i$  satisfies:*

$$\frac{1}{N} \sum_{i=1}^N \text{Var}(\hat{p}_i - p_i) \gtrsim \frac{1}{2M} \min\{1 - (1 - p_0)^{c\bar{k}\alpha}, (1 - p_0)^{c\bar{k}\alpha}\}, \quad (3)$$

where

$$p_0 := \left(\frac{\beta}{\beta + \gamma}\right)^\ell, \quad \ell := \min\{D, \frac{\log N}{\log \bar{k}}\}, \quad (4)$$

and  $c > 0$  is a constant depending on the network structure.

*Proof.* We briefly outline the key argument here; the full derivation is in the Appendix. Each node's infection probability  $p_i$  is estimated via Monte Carlo simulation as

$$\hat{p}_i = \frac{1}{M} \sum_{m=1}^M \mathbb{I}_i^{(m)}, \quad (5)$$

where  $\mathbb{I}_i^{(m)}$  indicates infection in simulation  $m$ . Its variance is straightforwardly given by  $\text{Var}[\hat{p}_i] = p_i(1 - p_i)/M$ . Under the continuous-time SIR model, the transmission probability along any edge is characterized by competing exponential infection and recovery processes, resulting in  $p_e = \beta/(\beta + \gamma)$  [28, 46]. Assuming independence and sparse network structure, node-level infection probabilities are approximated as  $p_i \approx 1 - \prod_{j=1}^{m_i} (1 - p_e^{\ell_{ij}})$ , where  $\ell_{ij}$  represents the path length between nodes  $i$  and  $j$  and we neglect higher-order interactions and overlapping paths [25, 43].

The number of disjoint infection paths  $m_i$  scales approximately as  $c\bar{k}\alpha$ , reflecting the influence of average degree  $\bar{k}$  and initial infection fraction  $\alpha$  [43, 25, 28]. Typical path lengths  $\ell$  can be approximated as  $\ell = \min\{D, \log N / \log \bar{k}\}$  [9, 42], leading to the simplified expression  $p_i \approx 1 - (1 - p_0)^{c\bar{k}\alpha}$ , with  $p_0 := (p_e)^\ell$ . Given the concavity of  $p_i(1 - p_i)$ , we apply a standard lower bound, yielding  $p_i(1 - p_i) \geq \frac{1}{2} \min\{p_i, 1 - p_i\}$ . Averaging this lower bound across all nodes, we obtain the claimed variance lower bound for the Monte Carlo estimator:

$$\frac{1}{N} \sum_{i=1}^N \text{Var}(\hat{p}_i - p_i) \gtrsim \frac{1}{2M} \min\{1 - (1 - p_0)^{c\bar{k}\alpha}, (1 - p_0)^{c\bar{k}\alpha}\}. \quad (6)$$

□

**Remark 1.** *Through Theorem 3.1, we observe that the variance of the Monte Carlo estimator is primarily influenced by four factors:  $\beta/\gamma$ ,  $\bar{k}$ ,  $\alpha$ , and  $M$ . The base probability  $p_0 := (\beta/(\beta + \gamma))^\ell$  increases with  $\beta/\gamma$ , since the ratio  $\beta/(\beta + \gamma)$  grows monotonically. Meanwhile,  $\ell := \min\{D, \log N / \log \bar{k}\}$  decreases as  $\bar{k}$  increases, leading to an indirect increase of  $p_0$  with  $\bar{k}$ .*




Figure 1: Influence of key factors on MC estimator variance.

As  $p_0$ ,  $\bar{k}$ , or  $\alpha$  increase, the term  $(1 - p_0)^{c\bar{k}\alpha}$  in the variance bound decreases. Since the bound depends on  $\min\{1 - (1 - p_0)^{c\bar{k}\alpha}, (1 - p_0)^{c\bar{k}\alpha}\}$ , it achieves its maximum when this term approaches 0.5. Therefore, estimator variance increases initially with  $\beta/\gamma$ ,  $\bar{k}$ , or  $\alpha$ , reaching a peak when the infection probability nears 0.5, and then decreases as the infection probability saturates. In contrast,  $M$  appears in the denominator and monotonically reduces  $\bar{v}$  as  $M$  increases. These factor effects are empirically verified in Figure 1.

### 3.2 Neighborhood Propagation Analysis

In Section 2.2, we show how diseases propagate probabilistically through local neighborhoods. However, using PID directly for inference can lead to great computational inefficiencies, especially on large-scale networks. A key bottleneck is the *path-level inefficiency*: in each update round, all nodes with non-zero infected probabilities must be updated, resulting in a propagation complexity of approximately  $\mathcal{O}(|\mathcal{E}|)$  as more nodes become affected. In this section, we provide a theoretical analysis of neighborhood-level approximations and global convergence conditions, laying the groundwork for scalable inference methods.

**Remark 2.** When using simulation-based methods for inference, it is common to approximate the continuous-time dynamics with a discrete-time process [18, 17]. Here,  $\beta$  and  $\gamma$  denote the per-unit-time transmission probability and healing probability, respectively. However, when  $\beta$  and  $\gamma$  take relatively large values, it becomes necessary to scale them down to smaller time units, i.e., setting  $\beta := \beta\Delta_t$  and  $\gamma := \gamma\Delta_t$ . This adjustment ensures consistency with the continuous infection rate:

$$\lim_{\Delta_t \rightarrow 0} 1 - (1 - \beta\Delta_t)^{1/\Delta_t} = \beta + \mathcal{O}(\beta^2), \quad (7)$$

which approximates linear growth for small  $\beta$ . Otherwise, if  $\beta$  and  $\gamma$  are large, the non-linear higher-order terms become non-negligible and can introduce significant errors in inference. Therefore, in the following discussions, we assume scaled-down  $\beta$  and  $\gamma$  satisfying  $\beta \ll 1$  and  $\gamma \ll 1$ .

A major challenge in SIR inference is that the underlying dynamic system is highly nonlinear, as shown in Eq. (2). Remark 2 guarantees that the higher-order errors remain sufficiently small. Following the linearization technique of [47], we locally linearize the nonlinear dynamics and derive the following theorem.

**Theorem 3.2** (Linear Approximation of Infection Dynamics). *Given a directed contact network  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{A})$ . Assume an SIR model parameterized by infection probability  $\beta$  and recovery probability  $\gamma$ . The infection probability update equation in Eq. (2)*

$$P_I^i(t+1) = P_S^i(t) - P_S^i(t+1) + (1 - \gamma)P_I^i(t) \quad (8)$$

can be approximated by the following linear expression:

$$P_I^i(t+1) \approx P_S^i(t)(\beta \sum_{j \in \mathcal{V}} \mathbf{A}_{ji} P_I^j(t)) + (1 - \gamma)P_I^i(t). \quad (9)$$

*Proof.* We sketch the main argument; detailed derivation is provided in the Appendix. Under the SIR model on a directed contact network  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{A})$ , let  $P_I^i(t)$  and  $P_S^i(t)$  denote the marginal probabilities that node  $i$  is infected or susceptible at time  $t$ . The infection update equation follows from accounting for new infections and recoveries in Eq. (2):

$$P_I^i(t+1) = P_S^i(t)[1 - \prod_{j \in \mathcal{V}} \mathbf{A}_{ji}(1 - \beta P_I^j(t))] + (1 - \gamma)P_I^i(t). \quad (10)$$

Here,  $P_S^i(t) - P_S^i(t+1)$  represents the probability that node  $i$  transitions from susceptible to infected at time  $t+1$ , while  $(1-\gamma)P_I^i(t)$  reflects staying infected without recovery. We then apply the first-order approximation  $\log(1-x) \approx -x$  for small  $x$  to linearize the product term  $\prod_j (1 - \beta \mathbf{A}_{ji} P_I^j(t))$ . This yields the following approximation for the probability of remaining susceptible:

$$P_S^i(t+1) \approx P_S^i(t)(1 - \beta \sum_{j \in \mathcal{V}} \mathbf{A}_{ji} P_I^j(t)). \quad (11)$$

Substituting into the update equation, we obtain:

$$P_I^i(t+1) \approx P_S^i(t)(\beta \sum_{j \in \mathcal{V}} \mathbf{A}_{ji} P_I^j(t)) + (1-\gamma)P_I^i(t). \quad (12)$$

This linear form captures the cumulative infection pressure from neighbors and the retention of infection due to non-recovery with probability  $1-\gamma$ .  $\square$

**Theorem 3.3** (Convergence Condition for Non-Reinfection Epidemic Models). *Consider a generalized non-reinfection epidemic model (e.g.,  $S^*I^2V^*$  [50]) over a directed contact network  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{A})$ , where each susceptible node becomes infected with probability  $\beta$  upon contact with an infected neighbor, and each infected node transitions to an absorbing state with probability  $\gamma$  at each time step. Then, the system's state distribution  $P(t)$  converges if and only if*

$$\max_{i \in \mathcal{V}} P_I^i(t) \rightarrow 0. \quad (13)$$

*Proof.* We sketch our proof as follows; detailed derivations are provided in the Appendix. We divide the argument into necessity and sufficiency. For necessity, we observe that if the state difference  $\|P(t) - P(t-1)\|_2$  tends to zero, then by the update equation of the absorbing state,  $\|P_{V^*}(t) - P_{V^*}(t-1)\|_2 = \gamma\|P_I(t-1)\|_2$  must also vanish. Since this term is bounded by the total state difference, we immediately obtain  $\|P_I(t-1)\|_2 \rightarrow 0$ , and thus  $\max_{i \in \mathcal{V}} P_I^i(t-1) \rightarrow 0$ .

For sufficiency, assume  $\max_{i \in \mathcal{V}} P_I^i(t-1) \rightarrow 0$ , which implies that  $\|P_I(t-1)\|_2 \rightarrow 0$ . Since the absorbing state update satisfies  $\|P_{V^*}(t) - P_{V^*}(t-1)\|_2 = \gamma\|P_I(t-1)\|_2$ , it follows that  $\|P_{V^*}(t) - P_{V^*}(t-1)\|_2 \rightarrow 0$ . Moreover, we show in the Appendix that  $\|P_I(t) - P_I(t-1)\|_2 \rightarrow 0$ . By conservation of total probability, we then have:

$$\|P_{S^*}(t) - P_{S^*}(t-1)\|_2 \leq \|P_I(t) - P_I(t-1)\|_2 + \|P_{V^*}(t) - P_{V^*}(t-1)\|_2. \quad (14)$$

Combining these bounds shows that all components in  $\|P(t) - P(t-1)\|_2$  vanish as  $t$  increases, proving the claim.  $\square$

**Remark 3.** *Theorem 3.3 implies that under non-reinfection epidemics models such as SIR, SEIR, or SIRD, the process inevitably converges as the infected probability at every node vanishes asymptotically, i.e.,  $\forall i \in \mathcal{V}, P_I^i(t) \rightarrow 0$ . Consequently, the entire disease dynamics can be equivalently analyzed as a transition from susceptible to non-susceptible states ( $S^* \rightarrow V^*$ ), without needing to explicitly model the transient infection state in the long term.*

## 4 Methodology

In this section, we first introduce the proposed algorithm **RAPID**, and then analyze its computational complexity. Our design ensures high efficiency, while the theoretical results in [24, 37] guarantee the accuracy of inference on sparse or locally tree-like epidemic contact networks, which are commonly encountered in practice.

### 4.1 The Proposed Algorithms

As our theoretical analysis in section 3.2 suggested that PID may suffer from computational inefficiency in the neighborhood propagation process, we propose an accelerated counterpart, **RAPID** (**R**esidual-**A**ware **P**ropagation for **I**nfection **D**ynamics), to efficiently scale inference on large contact networks. Our key idea is to define a *propagation residual* that quantifies the potential change in a node's infection probability due to local propagation. Nodes with large residuals are marked as active, and infection updates are computed only for these nodes, significantly reducing the overall computational cost.

First, we analyze the source of a one-step update on a node's infection probability. According to Theorem 3.2, we have

$$P_I^i(t+1) \approx P_S^i(t)(\beta \sum_{j \in \mathcal{V}} \mathbf{A}_{ji} P_I^j(t)) + (1-\gamma)P_I^i(t). \quad (15)$$

Taking the difference between consecutive time steps, we get:

$$\begin{aligned} P_I^i(t+1) - P_I^i(t) &\approx P_S^i(t)(\beta \sum_{j \in \mathcal{V}} \mathbf{A}_{ji} P_I^j(t)) - P_S^i(t-1)(\beta \sum_{j \in \mathcal{V}} \mathbf{A}_{ji} P_I^j(t-1)) \\ &\quad + (1-\gamma)(P_I^i(t) - P_I^i(t-1)) \\ &\leq P_S^i(t-1)(\beta \sum_{j \in \mathcal{V}} \mathbf{A}_{ji} (P_I^j(t) - P_I^j(t-1))) + (1-\gamma)(P_I^i(t) - P_I^i(t-1)) \\ &\leq \underbrace{\beta \sum_{j \in \mathcal{V}} \mathbf{A}_{ji} (P_I^j(t) - P_I^j(t-1))}_{\text{Change due to propagation}} + \underbrace{(1-\gamma)(P_I^i(t) - P_I^i(t-1))}_{\text{Change due to retained infection}}. \end{aligned} \quad (16)$$

We aim to update a node's state through local propagation only when the external influence is sufficiently strong. Based on the derivation in Eq. (16), we define

$$R_{\text{res}}(i) = \beta \sum_{j \in \mathcal{V}} \mathbf{A}_{ji} (P_I^j - \tilde{P}_I^j), \quad (17)$$

as the *propagation residual* at node  $i$ , where  $P_I^j$  denotes the current infected probability of node  $j$ , and  $\tilde{P}_I^j$  denotes its cached value before the most recent update.

An overview of **RAPID** is shown in Algorithm 2. We first perform  $p$  rounds of probabilistic disease propagation using Eq.(2), which serves as a preheating phase to initialize node states and compute propagation residuals. After preheating, we maintain a max-heap priority queue (referred to as the active queue) sorted by node residuals. At each iteration, the node  $i$  with the highest residual is popped and its full state vector  $\tilde{P}^i$  is updated using Eq.(2) in line 13.

The resulting change in infection probability,  $P_I^i - \tilde{P}_I^i$ , is linearly propagated to its out-neighbors by incrementing their propagation residuals. If a neighbor's updated residual exceeds the threshold, it is pushed into the active queue for future propagation as is shown from line 15 to line 19. After propagation, the residual of the current node is set to zero, indicating it is no longer active. The algorithm terminates when the active queue becomes empty. Since  $R_{\text{res}}(i) \leq \beta \sum_{j \in \mathcal{V}} \mathbf{A}_{ji} P_I^j$ , Theorem 3.3 guarantees that once the system converges ( $\max_{i \in \mathcal{V}} P_I^i \rightarrow 0$ ), the residuals vanish ( $R_{\text{res}}(i) \rightarrow 0$ ), ensuring convergence of **RAPID**.

## 4.2 Complexity Analysis

**Theorem 4.1** (Worst-Case Time Complexity of RAPID). *Consider RAPID on a directed contact network  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  with  $N = |\mathcal{V}|$  nodes and average out-degree  $\bar{k}$ . Let the infection probability be  $\beta$ , the recovery probability be  $\gamma$ , and the residual threshold be  $\varepsilon > 0$ . Then, under worst-case conditions, the total time complexity of RAPID is*

$$\mathcal{O}(N \cdot \bar{k} \cdot \min(\frac{1}{\varepsilon}, \frac{1}{\beta})). \quad (18)$$

---

### Algorithm 2 Sketch of RAPID

---

**Require:** Graph  $\mathcal{G}$ , threshold  $\varepsilon$ , rates  $\beta, \gamma$ , initial infected set  $\mathcal{I}_0$ , preheat steps  $p$

- 1: Initialize  $P_S, P_I, P_R$
  - 2: Preheat system for  $p$  rounds and record  $\tilde{P}_I$
  - 3: Initialize max heap  $\mathcal{H} \leftarrow \emptyset$
  - 4: **for**  $v \in \mathcal{V}$  **do**
  - 5:   Compute  $(R_{\text{res}}(v))$  with Eq. (17)
  - 6:   Push  $(R_{\text{res}}(v), v)$  into  $\mathcal{H}$
  - 7: **end for**
  - 8: **while**  $\mathcal{H}$  not empty **do**
  - 9:    $(R_{\text{res}}(i), i) \leftarrow \text{popTop}(\mathcal{H})$
  - 10:   **if**  $R_{\text{res}}(i) \leq \varepsilon$  **then break**
  - 11:   **end if**
  - 12:   **if**  $P_I^i > \varepsilon$  **then**
  - 13:     Update  $P_S^i, P_I^i, P_R^i, \tilde{P}_I^i$
  - 14:     Set  $\delta \leftarrow P_I^i - \tilde{P}_I^i$
  - 15:     **for**  $j \in \mathcal{N}^{\text{out}}(i)$  **do**
  - 16:        $R_{\text{res}}(j) \leftarrow R_{\text{res}}(j) + \beta \cdot \delta$
  - 17:       **if**  $R_{\text{res}}(j) > \varepsilon$  **then push**  $j$  **into**  $\mathcal{H}$
  - 18:       **end if**
  - 19:     **end for**
  - 20:   **end if**
  - 21:    $R_{\text{res}}(i) \leftarrow 0$
  - 22: **end while**
  - 23: **return**  $P_S, P_I, P_R$
-




Figure 2: Trade-off between Kendall-Tau and Runtime across six datasets.

*Proof.* The key insight is that a node becomes active only if its residual exceeds  $\varepsilon$ , and the residual update it can receive in a single iteration approximates  $\bar{k} \cdot \beta$ , since infection can be passed through  $\bar{k}$  neighbors on average at rate  $\beta$ . In the worst case, all  $N$  nodes become active, and each may trigger residual pushes to up to  $\bar{k}$  neighbors. When  $\varepsilon$  is large (i.e.,  $\varepsilon \gg \bar{k}\beta$ ), only large residuals are propagated, and the number of total pushes is upper bounded by  $\mathcal{O}(N \cdot \bar{k} \cdot \beta / \varepsilon)$ . However, when  $\varepsilon$  becomes very small (i.e.,  $\varepsilon \ll \bar{k}\beta$ ), further reducing  $\varepsilon$  does not lead to more node activations, since the residuals themselves are bounded. Therefore, the total number of pushes saturates at  $\mathcal{O}(N \cdot \bar{k})$ . Combining both regimes gives the overall time complexity. This scaling behavior mirrors residual-based approximations in PPR push methods [3], but here it arises from the probabilistic infection dynamics of epidemic modeling.  $\square$

## 5 Empirical Study

In this section, we conduct extensive experiments to investigate the effectiveness and efficiency of the proposed method.

### 5.1 Experiment Settings

**Datasets.** We use graphs from diverse domains, including a real-world hospital contact network (carilion-Hospital [1]), a real-world HIV transmission network (hiv-Trans [40]), communication networks (email-Enron [35, 29], email-EuAll [34]), and social networks (soc-Epinions [52], soc-Pokec [54]). Among them, carilion-Hospital and email-Enron are undirected, while the others are directed. Nodes with no incoming or outgoing edges are removed and each undirected edge is represented as two reciprocal directed edges. Detailed statistics are summarized in Table 2. We also include a Bitcoin trade network (bitcoin-Alpha [33, 32]) in the appendix.

**Baselines.** To serve as baselines, we implemented a broad range of inference models. A detailed description is provided in the Appendix.

- *Network centrality measures:* Degree, Eigenvector [44], PageRank [45], Betweenness [6] and Closeness [53]. These methods assess structural importance and are evaluated by ranking.
- *MC-based inference:* 5-run and 10-run Monte Carlo simulations (MC-5 and MC-10).
- *Probability-based inference:* Propagation for Infection Dynamics (PID), which retains all propagation steps during disease spread process, and is an unaccelerated invariant of **RAPID**.

Table 2: Datasets Statistics

| Dataset           | #Nodes    | #Edges     | Avg Deg |
|-------------------|-----------|------------|---------|
| carilion-Hospital | 11,810    | 30,994     | 2.62    |
| hiv-Trans         | 35,230    | 34,243     | 0.97    |
| email-Enron       | 36,692    | 367,662    | 10.02   |
| email-EuAll       | 265,214   | 420,045    | 1.58    |
| soc-Epinions      | 75,879    | 508,837    | 6.71    |
| soc-Pokec         | 1,632,803 | 30,622,564 | 18.75   |




Figure 3: Classification metrics (precision, recall, F1) on hiv-Trans.




Figure 4: **RAPID** runtime grows sublinearly with graph size and density.

**Parameter Settings.** To ensure the practical relevance of our study, we consulted the literature to obtain the infectious period and basic reproduction number ( $R_0$ ) for various diseases, as summarized in the Appendix. In our experiments, the ground-truth infection probability for each node is estimated by averaging its infection frequency over 50-run Monte Carlo simulations<sup>2</sup>. For reporting and baseline comparisons, we adopt the parameter setting of Nipah Virus ( $\beta = 1/18$ ,  $\gamma = 1/9$ ) with an initial infection fraction of  $\alpha = 0.01$ . Results for other disease parameters are similar and are provided in the Appendix.

## 5.2 Effectiveness and Efficiency Analysis

**Effectiveness of RAPID.** We evaluate the effectiveness of **RAPID** and baseline methods from three perspectives: the ranking of node-level infection probabilities, the classification performance of nodes whether infected, and the accuracy of infection probability estimation.

Figure 2 shows the trade-off between Kendall- $\tau$  and runtime across six datasets. **RAPID** achieves the best or near-best rank correlation on all of them while keeping the runtime consistently fast. We treat nodes with inferred infection probability above 0.5 as infected. Figure 3 presents the classification results on hiv-Trans, where **RAPID** significantly improves precision over baselines, achieving the best overall performance.

Table 3 reports the mean absolute error (MAE) for infection probability estimation. **RAPID** consistently outperforms MC-5, MC-10, and PID across all datasets. As noted in [37], when contact networks contain many short loops, nodes within these loops cyclically influence each other’s states over time, which reduces the accuracy of the original PID algorithm. **RAPID** addresses this issue by em-

Table 3: MAE comparison (lower is better). All values are scaled by  $10^{-2}$ . Best results are in **bold**.

| Dataset           | MC-5             | MC-10            | PID             | <b>RAPID</b>           |
|-------------------|------------------|------------------|-----------------|------------------------|
| carilion-Hospital | 13.01 $\pm$ 0.35 | 10.78 $\pm$ 0.44 | 5.67 $\pm$ 0.51 | <b>2.64</b> $\pm$ 0.49 |
| hiv-Trans         | 5.12 $\pm$ 0.12  | 3.24 $\pm$ 0.11  | 6.43 $\pm$ 0.52 | <b>1.27</b> $\pm$ 0.47 |
| email-Enron       | 7.60 $\pm$ 0.02  | 5.98 $\pm$ 0.03  | 8.70 $\pm$ 0.04 | <b>4.66</b> $\pm$ 0.00 |
| email-EuAll       | 2.09 $\pm$ 0.01  | 1.63 $\pm$ 0.01  | 1.36 $\pm$ 0.02 | <b>1.03</b> $\pm$ 0.01 |
| soc-Epinions      | 5.48 $\pm$ 0.03  | 4.31 $\pm$ 0.02  | 4.92 $\pm$ 0.03 | <b>2.77</b> $\pm$ 0.00 |
| soc-Pokec         | 4.50 $\pm$ 0.00  | 3.54 $\pm$ 0.00  | 3.32 $\pm$ 0.00 | <b>2.32</b> $\pm$ 0.00 |

employing asynchronous max-heap iterations to accelerate convergence and selective updates to mitigate the effect of short loops, thereby achieving higher inference accuracy than PID. Notably, although the Kendall- $\tau$  on hiv-Trans is lower than on other datasets, largely due to the large number of uninfected nodes at convergence, **RAPID** still achieves superior classification accuracy and lower MAE compared to all baselines.

**Runtime Comparison.** Table 4 shows the runtime of **RAPID** and baselines on each dataset. On the six datasets, **RAPID** achieves an average speedup of  $5.11\times$ ,  $10.67\times$ , and  $8.52\times$  over MC-5, MC-10, and PID, respectively. Compared to PID, the speedup is more pronounced on denser graphs, reaching  $12.91\times$  on soc-Epinions. Against MC-10, **RAPID** exhibits stable acceleration ranging from  $7.64\times$  to  $12.58\times$ , consistently maintaining a runtime comparable to that of single-run MC across all datasets.

<sup>2</sup>On carilion-Hospital and hiv-Trans, we adopt 1000-run MC simulations as the ground truth for acceptable estimator variance.

Table 4: Runtime comparison across datasets (seconds, lower is better).  $\Delta$  indicates the speedup factor relative to RAPID, computed as  $\Delta = \frac{\text{Baseline time}}{\text{RAPID time}}$ .

|                  |          | carilion-Hospital <sup>2</sup>  | hiv-Trans <sup>2</sup>          | email-Enron                     | email-EuAll                      | soc-Epinions                     | soc-Pokec                         |
|------------------|----------|---------------------------------|---------------------------------|---------------------------------|----------------------------------|----------------------------------|-----------------------------------|
| MC-5             | $t$      | 5.81 $\pm$ 0.51                 | 21.87 $\pm$ 3.15                | 29.84 $\pm$ 1.43                | 169.46 $\pm$ 8.23                | 59.69 $\pm$ 1.26                 | 1241.00 $\pm$ 18.79               |
|                  | $\Delta$ | 5.43 $\times$                   | 5.16 $\times$                   | 5.41 $\times$                   | 5.06 $\times$                    | 5.82 $\times$                    | 3.78 $\times$                     |
| MC-10            | $t$      | 13.46 $\pm$ 1.14                | 49.94 $\pm$ 1.74                | 56.45 $\pm$ 0.79                | 330.97 $\pm$ 6.81                | 122.31 $\pm$ 2.44                | 2506.60 $\pm$ 45.46               |
|                  | $\Delta$ | 12.58 $\times$                  | 11.78 $\times$                  | 10.24 $\times$                  | 9.88 $\times$                    | 11.91 $\times$                   | 7.64 $\times$                     |
| MC-50            | $t$      | 1234.73 $\pm$ 13.13             | 4678.18 $\pm$ 8.96              | 279.26 $\pm$ 3.09               | 1659.57 $\pm$ 23.55              | 614.58 $\pm$ 2.36                | 12782.37 $\pm$ 237.30             |
|                  | $\Delta$ | 1153.95 $\times$                | 1103.34 $\times$                | 50.66 $\times$                  | 49.52 $\times$                   | 59.86 $\times$                   | 38.93 $\times$                    |
| PID              | $t$      | 3.56 $\pm$ 0.01                 | 17.91 $\pm$ 0.14                | 66.95 $\pm$ 0.29                | 206.18 $\pm$ 0.65                | 132.60 $\pm$ 0.62                | 4056.89 $\pm$ 6.40                |
|                  | $\Delta$ | 3.33 $\times$                   | 4.22 $\times$                   | 12.14 $\times$                  | 6.15 $\times$                    | 12.91 $\times$                   | 12.36 $\times$                    |
| <b>RAPID</b> $t$ |          | <b>1.07<math>\pm</math>0.00</b> | <b>4.24<math>\pm</math>0.03</b> | <b>5.51<math>\pm</math>0.04</b> | <b>33.50<math>\pm</math>0.05</b> | <b>10.27<math>\pm</math>0.09</b> | <b>328.28<math>\pm</math>0.66</b> |

**Complexity Validation.** Figure 4 shows the empirical log-scale runtime of **RAPID** versus  $\log(N\bar{k})$  on six real-world networks. The fitted complexity runtime  $\propto (N\bar{k})^{0.74}$  is better than the worst-case bound established in Theorem 4.1, confirming the validity of our analysis.

## 6 Conclusion

In this paper, we systematically studied the variance of Monte Carlo simulations for modeling the disease spread process in contact networks, and introduce a linear approximation for infection propagation under non-reinfection models with a provable convergence guarantee. Based on the theoretical findings, we propose **RAPID**, a residual-driven framework to infer node-level infection probability distribution with high estimation accuracy and low computational cost. Experiments on six real-world networks show that **RAPID** matches the accuracy of multi-run Monte Carlo within the runtime of a single simulation. Future work includes extensions to models with reinfection, time-varying parameters, and dynamic networks.

## Acknowledgments

The authors would like to thank the anonymous reviewers for their constructive comments. This work was supported in part by the NVIDIA Academic Grant Program, the Commonwealth Cyber Initiative (CCI) under Award No. HV-2Q25-032, and the National Science Foundation under Grant No. 2331315. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

## References

- [1] B. Adhikari, B. Lewis, A. Vullikanti, J. M. Jiménez, and B. A. Prakash. Fast and near-optimal monitoring for healthcare acquired infection outbreaks. *PLoS computational biology*, 15(9):e1007284, 2019.
- [2] M. Al-Raei. The study of human monkeypox disease in 2022 using the epidemic models: Herd immunity and the basic reproduction number case. *Annals of Medicine and Surgery*, 85(2):316–321, 2023.
- [3] R. Andersen, F. Chung, and K. Lang. Local graph partitioning using pagerank vectors. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, pages 475–486. IEEE, 2006.
- [4] R. Andersen, F. R. K. Chung, and K. J. Lang. Local graph partitioning using pagerank vectors. In *Proceedings of the 47th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 475–486, 2006.
- [5] P. Block, M. Hoffman, I. J. Raabe, J. B. Dowd, C. Rahal, R. Kashyap, and M. C. Mills. Social network-based distancing strategies to flatten the covid-19 curve in a post-lockdown world. *Nature human behaviour*, 4(6):588–596, 2020.
- [6] U. Brandes. A faster algorithm for betweenness centrality. *Journal of mathematical sociology*, 25(2):163–177, 2001.
- [7] Centers for Disease Control and Prevention (CDC). Infectious diseases overview, 2023. Accessed: 2024-12-06.

- [8] M. Chinazzi, J. T. Davis, M. Ajelli, C. Gioannini, M. Litvinova, S. Merler, A. Pastore y Piontti, K. Mu, L. Rossi, K. Sun, et al. The effect of travel restrictions on the spread of the 2019 novel coronavirus (covid-19) outbreak. *Science*, 368(6489):395–400, 2020.
- [9] F. Chung and L. Lu. The average distances in random graphs with given expected degrees. *Proceedings of the National Academy of Sciences*, 99(25):15879–15882, 2002.
- [10] D. R. Cox and V. Isham. *Point processes*, volume 12. CRC Press, 1980.
- [11] M. J. Crawley. *The R book*. John Wiley & Sons, 2012.
- [12] M. Eichner, S. F. Dowell, and N. Firese. Incubation period of ebola hemorrhagic virus subtype zaire. *Osong public health and research perspectives*, 2(1):3–7, 2011.
- [13] G. Elidan, I. McGraw, and D. Koller. Residual belief propagation: Informed scheduling for asynchronous message passing. In *Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 165–173, 2006.
- [14] N. M. Ferguson, D. Laydon, G. Nedjati-Gilani, N. Imai, K. Ainslie, M. Baguelin, S. Bhatia, A. Boonyasiri, Z. Cucunubá, G. Cuomo-Dannenburg, et al. Impact of non-pharmaceutical interventions (npis) to reduce covid-19 mortality and healthcare demand. imperial college covid-19 response team. *Imperial College COVID-19 Response Team*, 20(10.25561):77482, 2020.
- [15] S. Flaxman, S. Mishra, A. Gandy, H. J. T. Unwin, T. A. Mellan, H. Coupland, C. Whittaker, H. Zhu, T. Berah, J. W. Eaton, et al. Estimating the effects of non-pharmaceutical interventions on covid-19 in europe. *Nature*, 584(7820):257–261, 2020.
- [16] A. Ganesh, L. Massoulié, and D. Towsley. The effect of network topology on the spread of epidemics. In *Proceedings IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies.*, volume 2, pages 1455–1466. IEEE, 2005.
- [17] F. Gao, J. Liu, and Y. Zhao. Recurrent dynamic message passing with loops for epidemics on networks. *Chaos, Solitons & Fractals*, 184:115045, 2024.
- [18] F. Gao, J. Zhang, and Y. Zhang. Neural enhanced dynamic message passing. In *International Conference on Artificial Intelligence and Statistics*, pages 10471–10482. PMLR, 2022.
- [19] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin. *Bayesian data analysis*. Chapman and Hall/CRC, 1995.
- [20] R. Grant, L.-B. L. Nguyen, and R. Breban. Modelling human-to-human transmission of monkeypox. *Bulletin of the World Health Organization*, 98(9):638, 2020.
- [21] X. He, E. H. Lau, P. Wu, X. Deng, J. Wang, X. Hao, Y. C. Lau, J. Y. Wong, Y. Guan, X. Tan, et al. Temporal dynamics in viral shedding and transmissibility of covid-19. *Nature medicine*, 26(5):672–675, 2020.
- [22] Y. He, C. Chen, S. Wang, G. Min, and J. Li. Demystify epidemic containment in directed networks: Theory and algorithms. In *Proceedings of the Eighteenth ACM International Conference on Web Search and Data Mining*, pages 905–913, 2025.
- [23] G. Jeh and J. Widom. Scaling personalized web search. In *Proceedings of the 12th International Conference on World Wide Web (WWW)*, pages 271–279, 2003.
- [24] B. Karrer and M. E. Newman. Message passing approach for general epidemic models. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 82(1):016101, 2010.
- [25] D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146, 2003.
- [26] J. O. Kephart and S. R. White. Directed-graph epidemiological models of computer viruses. In *Computation: the micro and the macro view*, pages 71–102. World Scientific, 1992.
- [27] W. O. Kermack and A. G. McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character*, 115(772):700–721, 1927.
- [28] I. Z. Kiss, J. C. Miller, P. L. Simon, et al. Mathematics of epidemics on networks. *Cham: Springer*, 598(2017):31, 2017.

- [29] B. Klimt and Y. Yang. Introducing the enron corpus. In *CEAS*, volume 4, page 1, 2004.
- [30] F. R. Kschischang, B. J. Frey, and H. Loeliger. Factor graphs and the sum-product algorithm. *IEEE Trans. Inf. Theory*, 47(2):498–519, 2001.
- [31] A. Kucharski and C. L. Althaus. The role of superspreading in middle east respiratory syndrome coronavirus (mers-cov) transmission. *Eurosurveillance*, 20(25):21167, 2015.
- [32] S. Kumar, B. Hooi, D. Makhija, M. Kumar, C. Faloutsos, and V. Subrahmanian. Rev2: Fraudulent user prediction in rating platforms. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pages 333–341. ACM, 2018.
- [33] S. Kumar, F. Spezzano, V. Subrahmanian, and C. Faloutsos. Edge weight prediction in weighted signed networks. In *Data Mining (ICDM), 2016 IEEE 16th International Conference on*, pages 221–230. IEEE, 2016.
- [34] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graph evolution: Densification and shrinking diameters. *ACM transactions on Knowledge Discovery from Data (TKDD)*, 1(1):2–es, 2007.
- [35] J. Leskovec, K. J. Lang, A. Dasgupta, and M. W. Mahoney. Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters. *Internet Mathematics*, 6(1):29–123, 2009.
- [36] P. Lofgren, S. Banerjee, A. Goel, and C. Seshadhri. Fast-ppr: Scaling personalized pagerank estimation for large graphs. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2014.
- [37] A. Y. Lokhov, M. Mézard, and L. Zdeborová. Dynamic message-passing equations for models with unidirectional dynamics. *Physical Review E*, 91(1):012811, 2015.
- [38] S. P. Luby. The pandemic potential of nipah virus. *Antiviral research*, 100(1):38–43, 2013.
- [39] V. P. Martínez, N. Di Paola, D. O. Alonso, U. Pérez-Sautu, C. M. Bellomo, A. A. Iglesias, R. M. Coelho, B. López, N. Periolo, P. A. Larson, et al. “super-spreaders” and person-to-person transmission of andes virus in argentina. *New England Journal of Medicine*, 383(23):2230–2241, 2020.
- [40] M. Morris and R. Rothenberg. Hiv transmission network metastudy project: An archive of data from eight network studies, 1988–2001. 2011.
- [41] K. P. Murphy, Y. Weiss, and M. I. Jordan. Loopy belief propagation for approximate inference: An empirical study. In *Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 467–475. Morgan Kaufmann, 1999.
- [42] M. Newman. *Networks: An Introduction*. Oxford University Press, 03 2010.
- [43] M. E. Newman. Spread of epidemic disease on networks. *Physical review E*, 66(1):016128, 2002.
- [44] M. E. Newman. The mathematics of networks. *The new palgrave encyclopedia of economics*, 2(2008):1–12, 2008.
- [45] L. Page. The pagerank citation ranking: Bringing order to the web. Technical report, Technical Report, 1999.
- [46] R. Pastor-Satorras, C. Castellano, P. Van Mieghem, and A. Vespignani. Epidemic processes in complex networks. *Reviews of modern physics*, 87(3):925–979, 2015.
- [47] R. Pastor-Satorras and A. Vespignani. Epidemic spreading in scale-free networks. *Physical review letters*, 86(14):3200, 2001.
- [48] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.
- [49] S. Pei, S. Kandula, and J. Shaman. Differential effects of intervention timing on covid-19 spread in the united states. *Science advances*, 6(49):eabd6370, 2020.
- [50] B. A. Prakash, D. Chakrabarti, N. C. Valler, M. Faloutsos, and C. Faloutsos. Threshold conditions for arbitrary cascade models on arbitrary networks. *Knowledge and information systems*, 33:549–575, 2012.
- [51] M. Reuter, S. Biasotti, D. Giorgi, G. Patanè, and M. Spagnuolo. Discrete laplace–beltrami operators for shape analysis and segmentation. *Computers & Graphics*, 33(3):381–390, 2009.

- [52] M. Richardson, R. Agrawal, and P. Domingos. Trust management for the semantic web. In *International semantic Web conference*, pages 351–368. Springer, 2003.
- [53] G. Sabidussi. The centrality index of a graph. *Psychometrika*, 31(4):581–603, 1966.
- [54] L. Takac and M. Zabovsky. Data analysis in public social networks. In *International scientific conference and international workshop present day trends of innovations*, volume 1, 2012.
- [55] P. A. Vial, F. Valdivieso, G. Mertz, C. Castillo, E. Belmar, I. Delgado, M. Tapia, and M. Ferrés. Incubation period of hantavirus cardiopulmonary syndrome. *Emerging Infectious Diseases*, 12(8):1271, 2006.
- [56] H. Wang, Z. Wei, J. Gan, Y. Yuan, X. Du, and J. Wen. Edge-based local push for personalized pagerank. *Proc. VLDB Endow.*, 15(7):1376–1389, 2022.
- [57] Y. Wang, D. Chakrabarti, C. Wang, and C. Faloutsos. Epidemic spreading in real networks: An eigenvalue viewpoint. In *22nd International Symposium on Reliable Distributed Systems, 2003. Proceedings.*, pages 25–34. IEEE, 2003.
- [58] Z. Wong, C. Bui, A. Chughtai, and C. Macintyre. A systematic review of early modelling studies of ebola virus disease in west africa. *Epidemiology & Infection*, 145(6):1069–1094, 2017.
- [59] World Health Organization (WHO). Fact sheets on infectious diseases, 2023. Accessed: 2024-12-06.
- [60] G. Xie. A novel monte carlo simulation procedure for modelling covid-19 spread over time. *Scientific reports*, 10(1):13120, 2020.
- [61] J. S. Yedidia, W. T. Freeman, and Y. Weiss. Constructing free-energy approximations and generalized belief propagation algorithms. *IEEE Trans. Information Theory*, 51(7):2282–2312, 2005.
- [62] H. Zhang, P. Lofgren, and A. Goel. Approximate personalized pagerank on dynamic graphs. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2016.

## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: The abstract and introduction in this paper accurately reflect the paper's contributions and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: We discuss the limitations of the work in the supplementary material.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: In this paper, all proofs of theorems are provided and all assumptions are clearly stated or referenced in the statement of any theorems.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: All the results in this paper can be reproduced.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

#### 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The paper specify all the experimental details necessary to understand the results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments in this paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

#### 8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: For each experiment, the paper provide sufficient information on the computer resources needed to reproduce the experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

#### 9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

#### 10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: The paper discuss both potential positive societal impacts and negative societal impacts of the work performed in the supplementary material.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

#### 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

#### 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: The creators or original owners of assets used in the paper are properly credited and the license and terms of use are explicitly mentioned and properly respected.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.

- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

### 13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

### 15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [No]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

#### 16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

# Supplemental Material

**Organization.** This supplementary material is organized as follows. Appendix A reviews related work and relevant theoretical background. Appendix B presents detailed proofs of the main theorems. Appendix C provides complete algorithm descriptions and highlights our design rationale from a **Jacobian perturbation** perspective, from which we derive the **classical epidemic threshold** and offer a physical interpretation of the propagation residual. Appendix D outlines implementation details and presents additional experiments, including: (i) how **Monte Carlo variance affects the reliability of MC-based inference**, (ii) the **extension of our method to the SEIR model**, (iii) the **impact of initial infection distributions**, (iv) experiments under alternative disease settings, and (v) a **sensitivity analysis of hyperparameters**. Appendix E discusses current limitations and potential directions for future research.

## A Related Work

**Belief Propagation on Graphs.** Belief propagation (BP), also known as the sum–product algorithm, is a foundational message-passing technique for probabilistic inference on graphs. Originally introduced by Pearl for exact inference in tree-structured Bayesian networks [48], BP was later applied to graphs with cycles as an approximate method, i.e. loopy BP [41]. The algorithm iteratively transmits local “messages” along edges of a factor graph to compute marginal probability estimates [30]. Despite the lack of general convergence guarantees, loopy BP often produces accurate results and has achieved remarkable success across domains (e.g. decoding of error-correcting codes and vision networks). Numerous extensions have enhanced BP’s accuracy and robustness. For instance, Yedidia et al. reinterpreted BP in terms of variational free-energy minimization and proposed generalized BP for higher-order region graphs [61]. Other work has focused on improving BP’s convergence, such as methods that dampen oscillations or prioritize message updates based on residual magnitude [13]. These advances have solidified BP as a versatile approximate inference framework in graphical models.

**Local Push Algorithms and Variants.** LocalPush algorithms are a class of residual-based methods for approximate graph propagation, widely used in computing node influence scores or similarities on large graphs. A prime example is the estimation of Personalized PageRank (PPR) vectors, which measure the influence of a source node across a graph. Early approaches [23] found that by maintaining a “residual” at each node and pushing this probability mass outwards along edges, one can obtain accurate PPR estimates efficiently without full graph traversal. Andersen et al. formalized this idea by introducing a local push procedure with rigorous error bounds, enabling fast local graph partitioning via PageRank [4]. The LocalPush strategy restricts computation to the source’s local neighborhood, yielding near-linear time updates for sparse networks. Subsequent research has built on this foundation to handle more complex settings. Notably, bidirectional and multi-target push algorithms improve efficiency for queries between specific node pairs [36]. Dynamic extensions have been developed to maintain PPR results under streaming graph updates, by locally adjusting residuals after each edge insertion or deletion [62]. Further optimizations address weighted graphs and precision guarantees, as in the edge-based LocalPush variant that decomposes pushes per edge to reduce redundant work [56]. Collectively, these residual-based propagation techniques have become essential tools for scalable graph mining, enabling approximate inference of importance or similarity scores with provable accuracy on massive networks.

## B Theoretical Proof

### B.1 Proof of Theorem 3.1

**Theorem 3.1** (Monte Carlo Estimation Variance Lower Bound) Given a directed contact network  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  with  $N = |\mathcal{V}|$ , average out-degree  $\bar{k}$ , and diameter  $D$ . Let  $\mathcal{I}_0 \subseteq \mathcal{V}$  be the initially infected node set with fraction  $\alpha := |\mathcal{I}_0|/N$ . Assume an SIR model parameterized by infection probability  $\beta$  and recovery probability  $\gamma$ . Using  $M$  independent Monte Carlo simulations to estimate each node’s

infection probability  $p_i$ , the average variance of the estimator  $\hat{p}_i$  satisfies:

$$\frac{1}{N} \sum_{i=1}^N \text{Var}(\hat{p}_i - p_i) \gtrsim \frac{1}{2M} \min\{1 - (1 - p_0)^{c\bar{k}\alpha}, (1 - p_0)^{c\bar{k}\alpha}\}, \quad (19)$$

where

$$p_0 := \left(\frac{\beta}{\beta + \gamma}\right)^\ell, \quad \ell := \min\{D, \frac{\log N}{\log \bar{k}}\}, \quad (20)$$

and  $c > 0$  is a constant depending on the network structure.

*Proof.* Each node's infection probability  $p_i$  is estimated by Monte Carlo simulation:

$$\hat{p}_i := \frac{1}{M} \sum_{m=1}^M \mathbb{I}_i^{(m)}, \quad (21)$$

where  $\mathbb{I}_i^{(m)} \in \{0, 1\}$  indicates whether node  $i$  is infected in simulation  $m$ . The variance is:

$$\text{Var}[\hat{p}_i] = \frac{1}{M} p_i (1 - p_i). \quad (22)$$

Under the continuous-time SIR model, infection along an edge competes with recovery via independent exponential clocks. The transmission probability per edge is [28, 46]:

$$p_e := \Pr[T_{\text{infect}} < T_{\text{recover}}] = \frac{\beta}{\beta + \gamma}. \quad (23)$$

Assuming that infection is mainly driven by multiple independent disjoint paths [25, 43], the infection probability of node  $i$  can be approximated as:

$$p_i \approx 1 - \prod_{j=1}^{m_i} (1 - p_e^{\ell_{ij}}), \quad (24)$$

where we neglect overlapping paths and higher-order interactions in sparse networks.

In sparse graphs, the number of disjoint infection paths to a node is heuristically assumed to scale with the average degree  $\bar{k}$  and the initial infection fraction  $\alpha$ , yielding  $m_i \approx c\bar{k}\alpha$ . This scaling reflects the intuition that a higher initial infection fraction increases the number of potential sources and thus the likelihood of multiple disjoint infection paths [43, 25, 28]. The typical path length approximates  $\ell := \min\{D, \log N / \log \bar{k}\}$  [9, 42]. Substituting these yields:

$$p_i \approx 1 - (1 - p_0)^{c\bar{k}\alpha}, \quad p_0 := (p_e)^\ell. \quad (25)$$

Since  $p_i(1 - p_i)$  is concave with maximum at  $p_i = \frac{1}{2}$ , we lower bound it by:

$$p_i(1 - p_i) \geq \frac{1}{2} \min\{p_i, 1 - p_i\}. \quad (26)$$

Averaging this bound over all nodes (the bound is node-independent and thus unchanged by averaging), we obtain the claimed variance lower bound for the Monte Carlo estimator:

$$\frac{1}{N} \sum_{i=1}^N \text{Var}(\hat{p}_i - p_i) \gtrsim \frac{1}{2M} \min\left\{1 - (1 - p_0)^{c\bar{k}\alpha}, (1 - p_0)^{c\bar{k}\alpha}\right\}. \quad (27)$$

□

## B.2 Proof of Theorem 3.2

**Theorem 3.2** (Linear Approximation of Infection Dynamics) Given a directed contact network  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{A})$ . Assume an SIR model parameterized by infection probability  $\beta$  and recovery probability  $\gamma$ . The infection probability update equation

$$P_I^i(t+1) = P_S^i(t) - P_S^i(t+1) + (1 - \gamma)P_I^i(t) \quad (28)$$

can be approximated by the following linear expression:

$$P_I^i(t+1) \approx P_S^i(t) \left( \beta \sum_{j \in \mathcal{V}} \mathbf{A}_{ji} P_I^j(t) \right) + (1 - \gamma)P_I^i(t). \quad (29)$$

*Proof.* Substitute the susceptible update:

$$P_S^i(t+1) := P_S^i(t) \prod_j (1 - \beta \mathbf{A}_{ji} P_I^j(t)) \quad (30)$$

into

$$P_I^i(t+1) = P_S^i(t) - P_S^i(t+1) + (1 - \gamma)P_I^i(t). \quad (31)$$

Take logarithm:

$$\log \left( \prod_j (1 - \beta \mathbf{A}_{ji} P_I^j(t)) \right) = \sum_j \log(1 - \beta \mathbf{A}_{ji} P_I^j(t)). \quad (32)$$

Approximate  $\log(1 - x) \approx -x$  and exponentiate both sides:

$$\prod_j (1 - \beta \mathbf{A}_{ji} P_I^j(t)) \approx e^{-\beta \sum_j \mathbf{A}_{ji} P_I^j(t)}. \quad (33)$$

According to Remark 2, we have  $\beta \ll 1$ . Use  $1 - e^{-x} \approx x$ :

$$1 - e^{-\beta \sum_j \mathbf{A}_{ji} P_I^j(t)} \approx \beta \sum_j \mathbf{A}_{ji} P_I^j(t). \quad (34)$$

Substitute into  $P_I^i(t+1)$  to obtain the linear form.  $\square$

### B.3 Proof of Theorem 3.3

**Lemma B.1.**  $\|P_I(t-1)\|_2 \rightarrow 0$  is sufficient for  $\|P_I(t) - P_I(t-1)\|_2 \rightarrow 0$ .

*Proof.* This problem is equivalent to proving that for all  $\varepsilon > 0$ , there exists  $\delta > 0$  such that if  $\|P_I(t-1)\|_2 < \varepsilon$ , then  $\|P_I(t) - P_I(t-1)\|_2 < \delta$ .

We have:

$$\max_{i \in \mathcal{V}} P_I^i(t-1) \leq \|P_I(t-1)\|_2 \leq \varepsilon. \quad (35)$$

First, for the absorbing state  $P_{V^*}(t)$ :

$$\|P_{V^*}(t) - P_{V^*}(t-1)\|_2 = \|\gamma P_I(t-1)\|_2 = \gamma \|P_I(t-1)\|_2 \leq \gamma \varepsilon. \quad (36)$$

Next, for  $P_S(t)$ :

$$\begin{aligned} P_S^i(t-1) - P_S^i(t) &= P_S^i(t-1) - P_S^i(t-1) \prod_{j \in \mathcal{V}} (1 - \beta \mathbf{A}_{ji} P_I^j(t-1)) \\ &\leq 1 - \prod_{j \in \mathcal{V}} (1 - \beta \mathbf{A}_{ji} P_I^j(t-1)) \\ &\leq 1 - (1 - \beta \varepsilon)^{d_{in}(i)}. \end{aligned}$$

By Bernoulli's inequality, if  $0 < x < 1$ , then  $(1 - x)^{d_{in}(i)} \geq e^{-d_{in}(i)x}$ . Letting  $x = \beta \varepsilon$ , we get:

$$(1 - \beta \varepsilon)^{d_{in}(i)} \geq e^{-d_{in}(i)\beta \varepsilon}. \quad (37)$$

Thus:

$$P_S^i(t-1) - P_S^i(t) \leq 1 - e^{-d_{in}(i)\beta \varepsilon}. \quad (38)$$

Since  $P_S^i(t-1) - P_S^i(t) \in (0, 1)$ , we can further bound:

$$1 - e^{-d_{in}(i)\beta \varepsilon} \leq (1 - e^{-N\beta})\varepsilon. \quad (39)$$

Therefore:

$$P_S^i(t-1) - P_S^i(t) \leq (1 - e^{-N\beta})\varepsilon. \quad (40)$$

Aggregating:

$$\|P_S(t-1) - P_S(t)\|_2 \leq \sqrt{N}(1 - e^{-N\beta})\varepsilon. \quad (41)$$

Finally, for  $P_I(t) - P_I(t-1)$ :

$$\begin{aligned}\|P_I(t) - P_I(t-1)\|_2 &= \|1 - P_S(t) - P_{V^*}(t) - (1 - P_S(t-1) - P_{V^*}(t-1))\|_2 \\ &= \|P_S(t-1) - P_S(t) + P_{V^*}(t-1) - P_{V^*}(t)\|_2 \\ &\leq \|P_S(t-1) - P_S(t)\|_2 + \|P_{V^*}(t) - P_{V^*}(t-1)\|_2 \\ &\leq \sqrt{N}(1 - e^{-N\beta})\varepsilon + \gamma\varepsilon.\end{aligned}$$

Let:

$$\delta := \left(\sqrt{N}(1 - e^{-N\beta}) + \gamma\right)\varepsilon. \quad (42)$$

Then from  $\max_{i \in \mathcal{V}} P_I^i(t-1) \leq \varepsilon$ , we conclude:

$$\|P_I(t) - P_I(t-1)\|_2 \leq \delta. \quad (43)$$

□

**Theorem 3.3** (Convergence Condition for Non-Reinfection Epidemic Models) Consider a generalized non-reinfection epidemic model (e.g.,  $S^*I^2V^*$  [50]) over a directed contact network  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{A})$ , where each susceptible node becomes infected with probability  $\beta$  upon contact with an infected neighbor, and each infected node transitions to an absorbing state with probability  $\gamma$  at each time step. Then, the system's state distribution  $P(t)$  converges if and only if

$$\max_{i \in \mathcal{V}} P_I^i(t) \rightarrow 0. \quad (44)$$

*Proof.* We first prove necessity. We aim to show that for any  $\varepsilon > 0$ , if  $\|P(t) - P(t-1)\|_2 \leq \varepsilon$ , then  $\exists \delta > 0$  such that  $\|P_I(t-1)\|_2 \leq \delta$ . Since

$$\|P_{V^*}(t) - P_{V^*}(t-1)\|_2 = \gamma\|P_I(t-1)\|_2 \leq \|P(t) - P(t-1)\|_2 \leq \varepsilon, \quad (45)$$

it follows that

$$\|P_I(t-1)\|_2 \leq \frac{\varepsilon}{\gamma}. \quad (46)$$

Therefore,

$$\max_{i \in \mathcal{V}} P_I^i(t-1) \leq \|P_I(t-1)\|_2 \leq \delta := \frac{\varepsilon}{\gamma}. \quad (47)$$

Next we prove sufficiency. For any  $\varepsilon > 0$ , assume  $\max_{i \in \mathcal{V}} P_I^i(t-1) \leq \varepsilon$ . Then

$$\|P_I(t-1)\|_2 \leq \sqrt{N} \max_{i \in \mathcal{V}} P_I^i(t-1) \leq \sqrt{N}\varepsilon. \quad (48)$$

From the update of  $P_{V^*}$ ,

$$\|P_{V^*}(t) - P_{V^*}(t-1)\|_2 = \|\gamma P_I(t-1)\|_2 = \gamma\|P_I(t-1)\|_2 \leq \gamma\sqrt{N}\varepsilon. \quad (49)$$

By Lemma B.1, since  $\max_i P_I^i(t-1) \leq \varepsilon$ , there exists  $\delta' > 0$  such that  $\|P_I(t) - P_I(t-1)\|_2 \leq \delta'$ .

Since  $P_S(t) + P_I(t) + P_{V^*}(t) = \mathbf{1}$ ,

$$\begin{aligned}\|P_S(t) - P_S(t-1)\|_2 &= \|1 - P_I(t) - P_{V^*}(t) - (1 - P_I(t-1) - P_{V^*}(t-1))\|_2 \\ &= \|P_I(t-1) - P_I(t) + P_{V^*}(t-1) - P_{V^*}(t)\|_2 \\ &\leq \|P_I(t) - P_I(t-1)\|_2 + \|P_{V^*}(t) - P_{V^*}(t-1)\|_2 \\ &\leq \delta' + \gamma\sqrt{N}\varepsilon.\end{aligned}$$

Therefore,

$$\begin{aligned}\|P(t) - P(t-1)\|_2^2 &= \|P_S(t) - P_S(t-1)\|_2^2 + \|P_I(t) - P_I(t-1)\|_2^2 + \|P_{V^*}(t) - P_{V^*}(t-1)\|_2^2 \\ &\leq (\delta' + \gamma\sqrt{N}\varepsilon)^2 + (\delta')^2 + (\gamma\sqrt{N}\varepsilon)^2.\end{aligned}$$

Taking square roots:

$$\|P(t) - P(t-1)\|_2 \leq \sqrt{(\delta' + \gamma\sqrt{N}\varepsilon)^2 + (\delta')^2 + (\gamma\sqrt{N}\varepsilon)^2} := \delta. \quad (50)$$

Thus, for any  $\varepsilon > 0$ ,  $\|P_I(t-1)\|_2 \leq \varepsilon$  implies  $\|P(t) - P(t-1)\|_2 \leq \delta$ . □

## C Details of Algorithms

### C.1 Baselines

All centrality-based baselines (including Degree, Eigenvector [44], PageRank [45], Betweenness [6] and Closeness [53]) are implemented using Python's `networkx`<sup>3</sup> library.

### C.2 MC for SIR

#### Algorithm 1: Monte Carlo Simulation for SIR

**Require:** Graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , infection rate  $\beta$ , recovery rate  $\gamma$ , initial infected set  $\mathcal{I}_0$ , number of trials  $M$

**Ensure:** Estimated final-state probabilities  $P_S^i, P_I^i, P_R^i$  for all  $i \in \mathcal{V}$

```

1:  $\text{count}_S^i \leftarrow 0, \text{count}_I^i \leftarrow 0, \text{count}_R^i \leftarrow 0$  for all  $i \in \mathcal{V}$  ▷ Initialize counters
2: for  $m = 1$  to  $M$  do ▷ Run  $M$  independent trials
3:    $\text{state}^i \leftarrow S$  for all  $i \in \mathcal{V}$ 
4:   for  $i \in \mathcal{I}_0$  do  $\text{state}^i \leftarrow I$ 
5:   end for
6:   while some node  $i$  has  $\text{state}^i = I$  do
7:     Initialize  $\text{next\_state}^i \leftarrow \text{state}^i$  for all  $i$ 
8:     for all  $i \in \mathcal{V}$  do
9:       if  $\text{state}^i = I$  then
10:        for all  $j \in \mathcal{N}^{\text{out}}(i)$  do
11:          if  $\text{state}^j = S$  and  $\text{rand}() < \beta$  then
12:             $\text{next\_state}^j \leftarrow I$ 
13:          end if
14:        end for
15:        if  $\text{rand}() < \gamma$  then
16:           $\text{next\_state}^i \leftarrow R$ 
17:        end if
18:      end if
19:    end for
20:     $\text{state}^i \leftarrow \text{next\_state}^i$  for all  $i$ 
21:  end while ▷ Update final counts
22:  for all  $i \in \mathcal{V}$  do
23:    if  $\text{state}^i = S$  then
24:       $\text{count}_S^i \leftarrow \text{count}_S^i + 1$ 
25:    end if
26:    if  $\text{state}^i = I$  then
27:       $\text{count}_I^i \leftarrow \text{count}_I^i + 1$ 
28:    end if
29:    if  $\text{state}^i = R$  then
30:       $\text{count}_R^i \leftarrow \text{count}_R^i + 1$ 
31:    end if
32:  end for
33: end for
34:  $P_S^i \leftarrow \text{count}_S^i / M, P_I^i \leftarrow \text{count}_I^i / M, P_R^i \leftarrow \text{count}_R^i / M$  ▷ Normalize to obtain final-state probabilities
35: return  $P_S, P_I, P_R$ 

```

<sup>3</sup><https://networkx.org/>

### C.3 Details of PID Algorithm

#### Algorithm 2: PID: Probabilistic Infection Dynamics

**Require:** Graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , infection rate  $\beta$ , recovery rate  $\gamma$ , initial infected set  $\mathcal{I}_0$ , threshold  $\varepsilon$

**Ensure:** Marginal probabilities  $P_S(t), P_I(t), P_R(t)$

```

1:  $P_S^i(0) \leftarrow 1, P_I^i(0) \leftarrow 0, P_R^i(0) \leftarrow 0, \forall i \in \mathcal{V}$ 
2: for  $i \in \mathcal{I}_0$  do
3:    $P_S^i(0) \leftarrow 0, P_I^i(0) \leftarrow 1$ 
4: end for
5:  $t \leftarrow 0$ 
6: repeat
7:   for all  $i \in \mathcal{V}$  do
8:      $P_S^i(t+1) \leftarrow P_S^i(t) \cdot \prod_{j \in \mathcal{N}^{\text{in}}(i)} (1 - \beta \cdot P_I^j(t))$ 
9:      $P_I^i(t+1) \leftarrow P_S^i(t) - P_S^i(t+1) + (1 - \gamma) \cdot P_I^i(t)$ 
10:     $P_R^i(t+1) \leftarrow 1 - P_S^i(t+1) - P_I^i(t+1)$ 
11:   end for
12:    $t \leftarrow t + 1$ 
13: until  $\|P_S(t) - P_S(t-1)\|_2 + \|P_I(t) - P_I(t-1)\|_2 + \|P_R(t) - P_R(t-1)\|_2 \leq \varepsilon$ 
14: return  $P_S(t), P_I(t), P_R(t)$ 

```

### C.4 Details of RAPID Algorithm

#### Algorithm 3: RAPID: Residual-Accelerated Propagation for Infection Dynamics

**Require:** Graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , rates  $\beta, \gamma$ , initial infected set  $\mathcal{I}_0$ , threshold  $\varepsilon$ , preheat steps  $p$

**Ensure:** State probabilities  $P_S, P_I, P_R$

```

1:  $P_S^v \leftarrow 1, P_I^v \leftarrow 0, P_R^v \leftarrow 0$  for all  $v \in \mathcal{V}$  ▷ Initialization
2: for  $v \in \mathcal{I}_0$  do
3:    $P_S^v \leftarrow 0, P_I^v \leftarrow 1$ 
4: end for
5:  $\tilde{P}_I^v \leftarrow P_I^v$  for all  $v$  ▷ Snapshot for residuals
6: ▷ Preheat phase to reduce approximation error
7: for  $t = 1$  to  $p$  do
8:   for all  $v \in \mathcal{V}$  do
9:      $\tilde{P}_S^v \leftarrow P_S^v$ 
10:     $P_S^v \leftarrow P_S^v \cdot \prod_{u \in \mathcal{N}^{\text{in}}(v)} (1 - \beta \cdot P_I^u)$ 
11:     $P_I^v \leftarrow (1 - \gamma) \cdot P_I^v + (\tilde{P}_S^v - P_S^v)$ 
12:     $P_R^v \leftarrow 1 - P_S^v - P_I^v$ 
13:   end for
14: end for
15: ▷ Residual initialization
16: for all  $v \in \mathcal{V}$  do
17:    $R_{\text{res}}(v) \leftarrow \beta \cdot \sum_{u \in \mathcal{N}^{\text{in}}(v)} (P_I^u - \tilde{P}_I^u)$ 
18: end for
19: Initialize max-heap  $\mathcal{H}$  with all  $v$  such that  $R_{\text{res}}(v) > \varepsilon$ 
20: while  $\mathcal{H}$  not empty do ▷ Residual-driven propagation
21:    $(R_{\text{res}}(i), i) \leftarrow \text{popTop}(\mathcal{H})$ 
22:   if  $R_{\text{res}}(i) \leq \varepsilon$  then break
23:   end if
24:   if  $P_I^i > \varepsilon$  then
25:      $\tilde{P}_S^i \leftarrow P_S^i$ 
26:   end if

```

```

27:    $P_S^i \leftarrow P_S^i \cdot \prod_{u \in \mathcal{N}^{\text{in}}(i)} (1 - \beta \cdot P_I^u)$ 
28:    $P_I^i \leftarrow (1 - \gamma) \cdot P_I^i + (\tilde{P}_S^i - P_S^i)$ 
29:    $P_R^i \leftarrow 1 - P_S^i - P_I^i$ 
30:    $\delta \leftarrow P_I^i - \tilde{P}_I^i$ 
31:    $\tilde{P}_I^i \leftarrow P_I^i$ 
32:                                      $\triangleright$  Propagate residual to out-neighbors
33:   for  $j \in \mathcal{N}^{\text{out}}(i)$  do
34:      $R_{\text{res}}(j) \leftarrow R_{\text{res}}(j) + \beta \cdot \delta$ 
35:     if  $R_{\text{res}}(j) > \varepsilon$  then push  $(R_{\text{res}}(j), j)$  into  $\mathcal{H}$ 
36:   end if
37: end for
38: end if
39:    $R_{\text{res}}(i) \leftarrow 0$ 
40: end while
41: return  $P_S, P_I, P_R$ 

```

### C.5 Design of Propagation Residual

The residual term is derived from the one-step temporal difference of the infected probability, as formulated in Eq. (15):

$$P_I^i(t+1) - P_I^i(t) \lesssim \underbrace{\beta \sum_{j \in \mathcal{V}} \mathbf{A}_{ji} (P_I^j(t) - P_I^j(t-1))}_{\text{Change due to propagation}} + \underbrace{(1 - \gamma)(P_I^i(t) - P_I^i(t-1))}_{\text{Change due to retained infection}}. \quad (51)$$

Although the propagation residual was originally motivated by the temporal dynamics of infection spread, its formulation is not heuristic. Instead, it admits a principled interpretation as a first-order Jacobian approximation of neighbor-induced perturbations under a linearized infection model.

In the following, we formally derive this connection, show how it naturally leads to the classical epidemic threshold, and provide a physical interpretation of the residual-driven max-heap update mechanism.

**Theorem C.1** (Jacobian Interpretation of Propagation Residual). *In Eq. (16), we define the propagation residual at node  $i$  as:*

$$R_{\text{res}}(i) = \beta \sum_{j \in \mathcal{V}} \mathbf{A}_{ji} (P_I^j - \tilde{P}_I^j), \quad (52)$$

where  $P_I^j$  and  $\tilde{P}_I^j$  denote the current and previously cached infection probabilities of node  $j$ . This residual can be interpreted as a local Jacobian-vector product:

$$R_{\text{res}}(i) = \sum_{j \in \mathcal{V}} \frac{\partial P_I^i}{\partial P_I^j} \cdot \delta_j, \quad (53)$$

where  $\delta_j = P_I^j - \tilde{P}_I^j$ , and the partial derivative  $\frac{\partial P_I^i}{\partial P_I^j} = \beta \mathbf{A}_{ji}$  arises from the linear approximation of the infection dynamics. Thus,  $R_{\text{res}}(i)$  quantifies the first-order propagation response of node  $i$  to perturbations in its neighbors.

**Remark 4** (Recovery of Epidemic Threshold). Denote  $\Delta P_I(t+1) = P_I^i(t+1) - P_I^i(t)$ . Starting from Eq. (C.5), we can express the system in matrix form as:

$$\Delta P_I(t+1) \lesssim (\beta \cdot \mathbf{A}^T + (1 - \gamma) \cdot \mathbf{I}) \cdot \Delta P_I(t), \quad (54)$$

which defines a linear dynamical system. For this system to remain stable, the spectral radius of the Jacobian matrix must satisfy:

$$\rho(\beta \cdot \mathbf{A}^T + (1 - \gamma) \cdot \mathbf{I}) < 1, \quad (55)$$

which leads to the classical epidemic threshold condition:

$$\frac{\beta}{\gamma} < \frac{1}{\lambda_{\max}(\mathbf{A})}. \quad (56)$$

This derivation illustrates that our residual-based formulation naturally recovers fundamental conditions in epidemic dynamics.

**Remark 5** (Connection to Local Lipschitz Bound). The Jacobian matrix  $\mathbf{J} = \beta \cdot \mathbf{A}^T$  governs the propagation sensitivity with respect to neighbor infection states. It induces a local Lipschitz upper bound:

$$\|\Delta P_I(t+1)\| \lesssim \|\mathbf{J}\| \cdot \|\Delta P_I(t)\|. \quad (57)$$

The residual term  $R_{\text{res}}$  captures the Jacobian-induced local propagation energy, and thus provides a structure-aware signal for selective inference.

**Remark 6** (Connection to Discrete Diffusion Operators). Although our model targets stochastic epidemic diffusion, the propagation residual exhibits a structural resemblance to classical diffusion operators. Specifically, the residual at node  $i$ ,

$$R_{\text{res}}(i) = \beta \sum_{j \in \mathcal{V}} \mathbf{A}_{ji} (P_I^j - \tilde{P}_I^j), \quad (58)$$

aggregates state changes from incoming neighbors, analogous to the discrete Laplace-Beltrami operator [51]:

$$\Delta f_i = \frac{1}{d_i} \sum_{j \in \mathcal{N}_i} w_{ij} (f_i - f_j), \quad (59)$$

which measures local imbalance across a mesh. While not a Laplacian in the strict sense, our formulation shares a key geometric intuition: it emphasizes update prioritization where local dynamics remain active—analogous to high-gradient regions in physical diffusion.

## D Experiments

### D.1 Experimental Settings

**Infrastructure and Implementation.** All experiments were conducted on a machine equipped with a Intel(R) Xeon(R) Gold 6248 CPU @ 2.50GHz processor with 376 GB Memory. All algorithms are implemented in Python with the Networkx library.

**Setup of Figure 1.** The setup for each subfigure of Figure 1 in the main paper is as follows: **(a)** We run 20 Monte Carlo simulations on a 1000-node asymmetric Erdős-Rényi graph ( $\bar{k} = 10$ ), varying  $\beta$  and  $\gamma$  across a  $[0.01, 1.0]$  grid. Each point reports the average node-level infection variance, and we plot its relationship with  $\log(\beta/\gamma)$  under a fixed initial infection ratio of 10%. **(b)** We simulate disease spread on 1000-node asymmetric Erdős-Rényi graphs with varying target average degree  $\bar{k} \in \{0.5, 1, \dots, 20\}$ . For each graph, we randomly infect 10% of nodes and compute the average node-level variance over a  $6 \times 6$  grid of  $(\beta, \gamma) \in [0.01, 0.2]$ . We then analyze the relationship between the actual average degree and the mean variance of the estimator. **(c)** We run 10 Monte Carlo simulations on a 1000-node asymmetric Erdős-Rényi graph ( $\bar{k} = 10$ ), across a  $6 \times 6$  grid of  $(\beta, \gamma) \in [0.01, 0.2]$ . For each initial infection ratio  $\alpha \in \{0.01, 0.02, \dots, 0.30\}$ , we compute the mean variance of node-level infection probabilities. **(d)** We run Monte Carlo simulations on a 1000-node asymmetric Erdős-Rényi graph ( $\bar{k} = 10$ ), using an initial infection ratio of 1%. For each sample size  $M \in \{5, 10, \dots, 50\}$ , we evaluate the mean estimator variance across a  $6 \times 6$  grid of  $(\beta, \gamma) \in [0.01, 0.2]$ .

**Hyperparameter Setting in Experiments.** For **RAPID**, the number of preheat steps  $p$  is set to 20, and the propagation residual threshold is set to  $10^{-3}$ .

**Full results for Figure 3.** We treat nodes with inferred infection probability above 0.5 as infected. Figure 3 in the main paper shows the classification results on hiv-Trans using MC-5, MC-10, PID, and **RAPID**. Here, we present the full results across all six datasets in Figure 5, illustrating each method’s ability to distinguish highly infected individuals.

**SIR Parameter Setting of More Diseases.** To ensure the practical relevance of our research, we reviewed the literature and compiled the infectious period and basic reproduction number ( $R_0$ ) data for various infectious diseases, as summarized in Table 5. It is important to note that we assume a well-equipped healthcare environment where infected individuals are isolated and treated promptly upon showing symptoms, effectively preventing further transmission. As a result, for diseases listed




Figure 5: Classification metrics (precision, recall, F1) on six datasets.

in the table other than the common cold, the infectious period is taken to be the incubation period of the disease.

In traditional epidemiological studies,  $\gamma$  can also be defined as the rate at which a disease transitions from being infectious to non-infectious. Therefore,  $\gamma$  can be calculated as the reciprocal of the infectious period. The SIR model parameters are derived as follows:

$$\gamma = \frac{1}{\text{Infectious Period}}$$

$$\beta = R_0 \cdot \gamma$$

For simplicity, when the infectious period is reported as a range, we use the average of the minimum and maximum values for calculations.

Table 5: Summary of some infectious diseases with their parameters in the SIR model.

| Disease          | Infectious Period               | Transmission                                                      | $R_0$                 | $\gamma$ [1/day] | $\beta$ [1/day] |
|------------------|---------------------------------|-------------------------------------------------------------------|-----------------------|------------------|-----------------|
| Nipah Virus      | 4–14 days [59]                  | Body fluids                                                       | 0.5 [38]              | 0.1111           | 0.0556          |
| Andes Hantavirus | 7–39 days (median 18 days) [55] | Respiratory droplets and body fluids                              | 1.2 (0.8–1.6) [39]    | 0.0556           | 0.0667          |
| MERS             | 5 days (2–14 days) [7]          | Respiratory droplets                                              | 0.5 (0.3–0.8) [31]    | 0.2              | 0.1             |
| Ebola            | Average 12.7 days [12]          | Body fluids                                                       | 1.8 (1.4–1.8) [58]    | 0.0787           | 0.1417          |
| Mpox             | 4–14 days [59]                  | Physical contact, body fluids, respiratory droplets, sexual (MSM) | 2.1 (1.1–2.7) [20, 2] | 0.1111           | 0.2333          |

## D.2 Influence of Variance on the Reliability of MC Inference

In this subsection, we aim to bridge the conclusion of Theorem 3.1 with empirical evidence, by demonstrating how a high variance in Monte Carlo (MC) estimators can lead to degraded inference accuracy.

We investigate the relationship between Kendall- $\tau$  and ground-truth per-node infection standard deviation on the `hiv-Trans` dataset by varying the initial infection ratio  $\alpha$ , as shown in Figure 6. We observe that while MC-based methods exhibit performance degradation as uncertainty increases, RAPID and PID maintain relatively stable ranking performance across high-variance regimes.

## D.3 Extension to Two-stage Non-reinfection Epidemic Model

Although Theorem 3.3 in the main paper shows that our algorithm, while developed under the SIR model, naturally extends to generalized non-reinfection models, in this subsection we provide empirical results on the SEIR model to demonstrate its effectiveness in a more complex setting.




Figure 6: Scatter plot and LOESS-smoothed trend lines of Kendall- $\tau$  versus ground-truth per-node infection standard deviation.

**SEIR Model Dynamics.** The SEIR model extends the classic SIR framework by introducing an intermediate *exposed* (E) state to capture the incubation period of infection. It models a two-stage infection process: individuals first become *exposed* after contact with infectious nodes, and then transition to the *infected* state after a latency period. The continuous-time dynamics are governed by the following ODEs:

$$\begin{aligned}\frac{dS}{dt} &= -\beta SI, \\ \frac{dE}{dt} &= \beta SI - \sigma E, \\ \frac{dI}{dt} &= \sigma E - \gamma I, \\ \frac{dR}{dt} &= \gamma I,\end{aligned}\tag{60}$$

where  $\beta$  is the infection rate,  $\sigma$  is the incubation (exposure-to-infection) rate, and  $\gamma$  is the recovery rate. The probability-based update equations are given by:

$$\begin{aligned}P_S^i(t+1) &= P_S^i(t) \prod_j (1 - \beta A_{ji} P_I^j(t)), \\ P_E^i(t+1) &= (1 - \sigma) P_E^i(t) + P_S^i(t) \left[ 1 - \prod_j (1 - \beta A_{ji} P_I^j(t)) \right], \\ P_I^i(t+1) &= (1 - \gamma) P_I^i(t) + \sigma P_E^i(t), \\ P_R^i(t+1) &= P_R^i(t) + \gamma P_I^i(t).\end{aligned}\tag{61}$$

**Extend RAPID to SEIR.** Similar to the idea adopted in the SIR model, for general non-reinfection models, global convergence requires that both  $P_E \rightarrow 0$  and  $P_I \rightarrow 0$ . In particular, for the two main transition stages  $S \rightarrow E$  and  $E \rightarrow I$ , we define the following residuals:

- Exposure-stage residual (from  $S \rightarrow E$ ):

$$R_{\text{res}}^E(i) = \beta \sum_{j \in V} A_{ji} (P_I^j(t) - P_I^j(t-1)).\tag{62}$$

- Infection-stage residual (from  $E \rightarrow I$ ):

$$R_{\text{res}}^I(i) = \sigma (P_E^i(t) - P_E^i(t-1)).\tag{63}$$

Note that once  $P_E^i \rightarrow 0$ , the infection-stage residual  $R_{\text{res}}^I(i) \rightarrow 0$  naturally follows. Moreover,  $R_{\text{res}}^I(i)$  depends solely on the change in the node's own state, whereas our main paper defines the propagation residual to be structure-aware. Hence, we define the global propagation residual for non-reinfection models as:

$$R_{\text{res}}(i) := R_{\text{res}}^E(i) = \beta \sum_{j \in V} A_{ji} (P_I^j - \tilde{P}_I^j),\tag{64}$$

which is exactly the same as our derivation in the main paper for SIR! To verify the effectiveness of our method, we conduct similar experiment in the main paper on hiv-Trans.

**Experiment.** We adopt the same parameter setting as in the Nipah Virus scenario, with  $\beta = 1/18$  and  $\gamma = 1/9$ , and further set  $\sigma = 0.8$  to model the transition probability from  $E \rightarrow I$ . The initially infected ratio  $\alpha = 0.01$ . We use 1000-run Monte Carlo simulation as the ground truth. All other hyperparameter settings are kept consistent with those in the main paper.




Figure 7: Kendall- $\tau$  vs. Runtime on hiv-Trans under SEIR model.




Figure 8: Classification metrics on hiv-Trans under SEIR model.

Table 6: SEIR Results on hiv-Trans with  $\alpha = 0.01$ . All values are scaled by  $10^{-2}$ . Best results are in bold.

| Method       | hiv-Trans                         |                                   |
|--------------|-----------------------------------|-----------------------------------|
|              | MAE $\downarrow$                  | $t$ (s) $\downarrow$              |
| MC-5         | $6.13 \pm 0.25$                   | $3.55 \pm 0.21$                   |
| MC-10        | $4.01 \pm 0.35$                   | $6.79 \pm 0.47$                   |
| PID          | $6.23 \pm 0.51$                   | $10.50 \pm 0.14$                  |
| <b>RAPID</b> | <b><math>0.82 \pm 0.23</math></b> | <b><math>2.43 \pm 0.12</math></b> |

Figure 7 and Figure 8 present the Kendall- $\tau$  versus runtime trade-off and the radar chart of precision, F1 and recall rate under the SEIR model on the hiv-Trans dataset, respectively. Table 6 further reports the detailed MAE and runtime values. These results closely mirror those under the SIR setting in the main paper, confirming that our RAPID method remains effective across general non-reinfection models beyond SIR.

#### D.4 Influence of Distribution of Initially Infected Nodes

As highlighted in the Introduction of the main paper, traditional epidemic modeling often relies on assumptions of random mixing and population homogeneity, which limit its applicability in real-world scenarios. In contrast, **RAPID** takes the initial infected set  $\mathcal{I}_0$  as an explicit input, allowing flexible adaptation to different initial conditions. In this subsection, we demonstrate that **RAPID** remains robust under varying distributions of  $\mathcal{I}_0$ .

**Impact of Initial Infection Phase  $\alpha$ .** In this part, we investigate the impact of different values of  $\alpha$ , which reflect different stages of an epidemic at the initial time. While the main paper presents results for  $\alpha = 0.01$  to simulate the early stage which is the most common setting, we additionally evaluate  $\alpha = 0.5$  and  $\alpha = 0.8$  to represent mid-stage and late-stage outbreaks, respectively, and report the corresponding results.

- $\alpha = 0.5$ . The trade-off between Kendall- $\tau$  and runtime is illustrated in Figure 9, while the MAE scores are summarized in Table 7.
- $\alpha = 0.8$ . The trade-off between Kendall- $\tau$  and runtime is illustrated in Figure 10, while the MAE scores are summarized in Table 8.

Based on the comparison across different epidemic phases, we make the following observations: (1) Consistent with the results in the main paper, our **RAPID** algorithm maintains high accuracy across




Figure 9: Trade-off between Kendall-Tau and Runtime across six datasets. ( $\alpha = 0.5$ )

Table 7: MAE comparison ( $\alpha = 0.5$ , lower is better). All values are scaled by  $10^{-2}$ . Best results are in **bold**.

| Dataset       | MC-5            | MC-10           | PID             | <b>RAPID</b>                    |
|---------------|-----------------|-----------------|-----------------|---------------------------------|
| bitcoin-Alpha | 4.47 $\pm$ 0.02 | 3.39 $\pm$ 0.04 | 2.65 $\pm$ 0.10 | <b>2.27<math>\pm</math>0.04</b> |
| hiv-Trans     | 4.15 $\pm$ 0.11 | 2.14 $\pm$ 0.03 | 4.24 $\pm$ 0.20 | <b>1.02<math>\pm</math>0.20</b> |
| email-Enron   | 3.99 $\pm$ 0.03 | 3.14 $\pm$ 0.00 | 2.79 $\pm$ 0.01 | <b>2.26<math>\pm</math>0.01</b> |
| email-EuAll   | 1.33 $\pm$ 0.01 | 1.02 $\pm$ 0.00 | 0.73 $\pm$ 0.00 | <b>0.68<math>\pm</math>0.00</b> |
| soc-Epinions  | 2.99 $\pm$ 0.02 | 2.36 $\pm$ 0.01 | 1.91 $\pm$ 0.01 | <b>1.57<math>\pm</math>0.01</b> |
| soc-Pokec     | 2.24 $\pm$ 0.00 | 1.76 $\pm$ 0.00 | 1.35 $\pm$ 0.00 | <b>1.16<math>\pm</math>0.00</b> |




Figure 10: Trade-off between Kendall-Tau and Runtime across six datasets. ( $\alpha = 0.8$ )

all  $\alpha$  settings, consistently accelerating both MC-based methods and PID, while preserving a runtime comparable to a single-run MC simulation. (2) As the initial infection ratio increases, centrality-based baselines become increasingly unreliable, and by  $\alpha = 0.8$ , most of them exhibit almost no correlation with ground truth.

**Impact of Initial Infection Clustering Patterns.** In this part, we divide the initial infection set into several disjoint and disconnected clusters to examine the impact on different methods. We conduct experiments on the `bitcoin-Alpha` dataset with an initial infection ratio of  $\alpha = 0.01$ . Figure 11 demonstrates the robustness of **RAPID** under varying distributions of initial infection clusters.

Table 8: MAE comparison ( $\alpha = 0.8$ , lower is better). All values are scaled by  $10^{-2}$ . Best results are in **bold**.

| Dataset       | MC-5            | MC-10           | PID                    | <b>RAPID</b>           |
|---------------|-----------------|-----------------|------------------------|------------------------|
| bitcoin-Alpha | 1.84 $\pm$ 0.03 | 1.34 $\pm$ 0.01 | 0.96 $\pm$ 0.02        | <b>0.94</b> $\pm$ 0.01 |
| hiv-Trans     | 1.68 $\pm$ 0.03 | 1.26 $\pm$ 0.00 | <b>0.31</b> $\pm$ 0.05 | 0.45 $\pm$ 0.07        |
| email-Enron   | 1.62 $\pm$ 0.02 | 1.29 $\pm$ 0.01 | 0.92 $\pm$ 0.01        | <b>0.87</b> $\pm$ 0.01 |
| email-EuAll   | 0.57 $\pm$ 0.00 | 0.45 $\pm$ 0.00 | 0.30 $\pm$ 0.00        | <b>0.29</b> $\pm$ 0.00 |
| soc-Epinions  | 1.21 $\pm$ 0.00 | 0.95 $\pm$ 0.00 | 0.68 $\pm$ 0.01        | <b>0.63</b> $\pm$ 0.01 |
| soc-Pokec     | 0.89 $\pm$ 0.00 | 0.70 $\pm$ 0.00 | <b>0.49</b> $\pm$ 0.00 | <b>0.47</b> $\pm$ 0.00 |




Figure 11: Expected MAE vs. Number of Disjoint Infection Clusters.

## D.5 Results for more Diseases

According to Remark 2, since the discrete SIR simulation aims to approximate a continuous-time process, we require  $\beta \ll 1$  and  $\gamma \ll 1$ . Under this regime, the simulation behavior depends primarily on the ratio  $R_0 = \beta/\gamma$ , rather than the specific values of  $\beta$  and  $\gamma$ , which aligns with classical epidemiological theory. Therefore, in our additional experiments, we evaluate the algorithm's performance using representative values of  $R_0$  listed in Table 5.

**Andes Hantavirus** ( $R_0 = 1.2$ ). The Kendall- $\tau$  versus runtime trade-off is shown in Figure 12, and the MAE comparison across different datasets is presented in Figure 9.




Figure 12: Trade-off between Kendall-Tau and Runtime across six datasets. (**Andes Hantavirus**)

**MERS** ( $R_0 = 0.5$ ). As MERS share the same  $R_0$  with Nipah Virus according to Table 5, the results are the same as those in main paper.

**Ebola** ( $R_0 = 1.8$ ). The Kendall- $\tau$  versus runtime trade-off is shown in Figure 13, and the MAE comparison across different datasets is presented in Figure 10.

Table 9: MAE comparison (**Andes Hantavirus**, lower is better). All values are scaled by  $10^{-2}$ . Best results are **in bold**.

| Dataset       | MC-5             | MC-10                  | PID              | <b>RAPID</b>           |
|---------------|------------------|------------------------|------------------|------------------------|
| bitcoin-Alpha | 8.21 $\pm$ 0.10  | 6.53 $\pm$ 0.03        | 8.27 $\pm$ 0.06  | <b>6.38</b> $\pm$ 0.05 |
| hiv-Trans     | 14.82 $\pm$ 0.30 | <b>4.99</b> $\pm$ 0.12 | 16.79 $\pm$ 0.35 | 5.40 $\pm$ 0.36        |
| email-Enron   | 6.82 $\pm$ 0.08  | 5.26 $\pm$ 0.01        | 9.31 $\pm$ 0.02  | <b>4.52</b> $\pm$ 0.08 |
| email-EuAll   | 2.13 $\pm$ 0.01  | 1.69 $\pm$ 0.02        | 2.05 $\pm$ 0.03  | <b>1.17</b> $\pm$ 0.02 |
| soc-Epinions  | 5.47 $\pm$ 0.03  | 4.33 $\pm$ 0.05        | 6.96 $\pm$ 0.01  | <b>4.31</b> $\pm$ 0.03 |
| soc-Pokec     | 3.64 $\pm$ 0.00  | <b>2.86</b> $\pm$ 0.00 | 3.68 $\pm$ 0.01  | 3.51 $\pm$ 0.01        |




Figure 13: Trade-off between Kendall-Tau and Runtime across six datasets. (**Ebola**)

Table 10: MAE comparison (**Ebola**, lower is better). All values are scaled by  $10^{-2}$ . Best results are **in bold**.

| Dataset       | MC-5             | MC-10           | PID             | <b>RAPID</b>           |
|---------------|------------------|-----------------|-----------------|------------------------|
| bitcoin-Alpha | 8.79 $\pm$ 0.09  | 6.82 $\pm$ 0.08 | 7.10 $\pm$ 0.11 | <b>4.77</b> $\pm$ 0.08 |
| hiv-Trans     | 10.50 $\pm$ 0.44 | 6.29 $\pm$ 0.23 | 3.31 $\pm$ 0.72 | <b>1.87</b> $\pm$ 0.62 |
| email-Enron   | 7.49 $\pm$ 0.04  | 5.98 $\pm$ 0.01 | 9.27 $\pm$ 0.04 | <b>4.39</b> $\pm$ 0.01 |
| email-EuAll   | 2.10 $\pm$ 0.00  | 1.67 $\pm$ 0.01 | 1.46 $\pm$ 0.01 | <b>1.02</b> $\pm$ 0.01 |
| soc-Epinions  | 5.51 $\pm$ 0.02  | 4.30 $\pm$ 0.01 | 5.41 $\pm$ 0.02 | <b>2.99</b> $\pm$ 0.01 |
| soc-Pokec     | 4.33 $\pm$ 0.00  | 3.40 $\pm$ 0.00 | 3.42 $\pm$ 0.00 | <b>2.68</b> $\pm$ 0.00 |

**Mpox** ( $R_0 = 2.1$ ). The Kendall- $\tau$  versus runtime trade-off is shown in Figure 14, and the MAE comparison across different datasets is presented in Figure 11.

Table 11: MAE comparison (**Mpox**, lower is better). All values are scaled by  $10^{-2}$ . Best results are **in bold**.

| Dataset       | MC-5            | MC-10           | PID             | <b>RAPID</b>           |
|---------------|-----------------|-----------------|-----------------|------------------------|
| bitcoin-Alpha | 5.44 $\pm$ 0.03 | 4.20 $\pm$ 0.07 | 2.71 $\pm$ 0.02 | <b>2.68</b> $\pm$ 0.01 |
| hiv-Trans     | 7.44 $\pm$ 0.03 | 6.03 $\pm$ 0.04 | 1.84 $\pm$ 0.61 | <b>1.78</b> $\pm$ 0.66 |
| email-Enron   | 5.57 $\pm$ 0.01 | 4.37 $\pm$ 0.02 | 2.83 $\pm$ 0.01 | <b>2.77</b> $\pm$ 0.01 |
| email-EuAll   | 0.96 $\pm$ 0.00 | 0.76 $\pm$ 0.00 | 0.48 $\pm$ 0.00 | <b>0.47</b> $\pm$ 0.00 |
| soc-Epinions  | 2.84 $\pm$ 0.01 | 2.24 $\pm$ 0.00 | 1.47 $\pm$ 0.01 | <b>1.45</b> $\pm$ 0.01 |
| soc-Pokec     | 6.74 $\pm$ 0.00 | 5.29 $\pm$ 0.00 | 3.58 $\pm$ 0.00 | <b>3.52</b> $\pm$ 0.00 |

**Discussion.** (1) Across all epidemic settings and datasets, our RAPID method consistently maintains high Kendall- $\tau$  and low MAE, often achieving the best performance while significantly improving computational efficiency. (2) Under the Andes Hantavirus setting ( $R_0 = 1.2$ ), we observe notably higher variance on the hiv-Trans dataset, aligning with Theorem 3.1 in the main paper. However, this increased stochasticity also challenges the fidelity of our reference distribution. In our experi-




Figure 14: Trade-off between Kendall-Tau and Runtime across six datasets. (**Mpx**)

ments, we estimate the ground truth infection probability  $\bar{P}_{GT}$  using 50 Monte Carlo simulations (on *hiv-Trans*, we use 1000-run Monte Carlo simulation as ground-truth). While this yields accurate approximations on most datasets, the elevated randomness in *hiv-Trans* makes the reference distribution inherently noisy. As a result, the MAE values should be interpreted with caution, as they may partially reflect the variance of the reference rather than the bias of the estimators.

## D.6 Sensitivity of Hyperparameters

**RAPID** has two key hyperparameters: the number of preheat steps  $p$  used to initialize the inference via PID, and the propagation residual threshold  $\varepsilon$  used to control message updates. Figure 15 shows how these hyperparameters affect performance in terms of mean Kendall- $\tau$  and MAE.

We observe that increasing  $p$  initially improves both Kendall- $\tau$  and MAE, as more preheating facilitates better initialization. However, beyond a moderate value, the performance begins to degrade, indicating that excessive preheating may lead to over-propagation and reduced inference quality. For the residual threshold  $\varepsilon$ , we observe that overly loose settings (i.e., large  $\varepsilon$ ) lead to reduced performance, as many small but informative updates are skipped. However, as  $\varepsilon$  becomes very small, its influence diminishes. This is because the scale of local residuals is largely determined by  $\bar{k}\beta$ , so further reducing  $\varepsilon$  without changing  $\beta$  has limited effect. Overall, **RAPID** maintains stable performance across a wide range of settings, demonstrating strong robustness to both hyperparameters.




Figure 15: Sensitivity analysis on preheat steps  $p$  and propagation residual threshold  $\varepsilon$ .

## E More Discussion

### E.1 Limitation

Although **RAPID** achieves high inference accuracy and significant runtime reduction across various networks, several limitations remain. First, our framework assumes a discrete-time approximation of the continuous-time infection process, and the linearization results rely on sufficiently small

per-time-step transmission and recovery probabilities, i.e.,  $\beta, \gamma \ll 1$  (see Remark 2). In practical scenarios with coarse time resolution or unscaled parameters, this assumption may not strictly hold, potentially introducing approximation bias. Second, **RAPID** is designed under the assumption of non-reinfection epidemic models (e.g., SIR, SEIR), where infections are irreversible. Extending the method to reinfection settings such as SIS or SIRS models remains non-trivial, as it would require fundamentally different update dynamics and potentially new convergence guarantees.

## E.2 Future Work

**Extension to General Epidemic Settings.** While our method is developed under the setting of non-reinfection models, extending it to more general epidemic settings is a promising direction. For reinfective models (e.g., SIS or SIRS), where nodes can return to the susceptible state, the propagation process becomes recurrent and requires modifying the residual formulation to account for cyclical transitions. Moreover, real-world epidemics often involve time-varying parameters—such as dynamic infection or recovery rates  $\beta(t)$  and  $\gamma(t)$ —which can be incorporated into our framework by adapting the update rules to operate on temporal sequences of residuals. Finally, in dynamic graphs where nodes and edges evolve over time (e.g., contact networks or mobility patterns), our residual-driven approach can be extended by computing propagation residuals on each snapshot and reinitializing the active node queue accordingly. Together, these extensions would enable our method to handle a wider range of realistic and complex epidemic processes.

**Extension to Graph Neural Networks (GNNs).** Our theoretical framework of linear approximation, residual-driven selection, and asynchronous propagation can be naturally extended to enhance Graph Neural Networks. Specifically, message passing in GNNs often involves nonlinear transformations, which can similarly benefit from Jacobian-based linear approximations (explained in Theorem C.1) to quantify node-level sensitivity to neighbor perturbations. By defining propagation residuals based on node feature dynamics, one can design prioritized, residual-guided asynchronous GNNs that selectively update nodes according to their local propagation significance. Such an approach promises significant computational efficiency improvements on large-scale graphs and offers theoretical insights into the interpretability and stability of GNNs. Exploring this connection presents an exciting avenue for future research.