

Inferring Patient Zero on Temporal Networks via Graph Neural Networks

Xiaolei Ru¹, Jack Murdoch Moore¹, Xin-Ya Zhang¹, Yeting Zeng², Gang Yan^{1*}

¹ School of Physical Science and Engineering, and National Key Laboratory of Autonomous Intelligent Unmanned Systems, Tongji University, Shanghai, China

² Zhongshan Hospital, Fudan University, Shanghai, China
{ruxl,jackmoore,xinyazhang,gyan}@tongji.edu.cn, ytceng21@m.fudan.edu.cn

Abstract

The world is currently seeing frequent local outbreaks of epidemics, such as COVID-19 and Monkeypox. Preventing further propagation of the outbreak requires prompt implementation of control measures, and a critical step is to quickly infer patient zero. This backtracking task is challenging for two reasons. First, due to the sudden emergence of local epidemics, information recording the spreading process is limited. Second, the spreading process has strong randomness. To address these challenges, we tailor a gnn-based model to establish the inverse statistical association between the current and initial state implicitly. This model uses contact topology and the current state of the local population to determine the possibility that each individual could be patient zero. We benchmark our model on data from important epidemiological models on five real temporal networks, showing performance significantly superior to previous methods. We also demonstrate that our method is robust to missing information about contact structure or current state. Further, we find the individuals assigned higher inferred possibility by model are closer to patient zero in terms of core number and the activity sequence recording the times at which the individual had contact with other nodes.

Introduction

Artificial intelligence is providing increasingly valuable support to a diverse range of applications in medicine (Buch, Ahmed, and Maruthappu 2018) by assisting humans to sift through massive amounts of medical data and uncover the relevant underlying patterns (Jittprasong 2022). Here, we propose an AI-based solution for source detection in contagion processes, i.e., inferring patient zero, the individual who initially introduces the virus into the local population. At present, countries around the world are plagued by highly contagious diseases, e.g., COVID-19 and Monkeypox. During an incubation period, the initially infected population can be contagious without exhibiting detectable symptoms, which can lead to the unforeseeable outbreak of an epidemic. After an outbreak, source detection is a necessary measure to investigate pathogenesis, cut off the route of transmission and block further propagation. An efficient detection algorithm can use limited data to determine where

a virus comes from and guide governments to develop targeted procedures for epidemic prevention.

The backtracking problem's practical significance means it has received sustained attention across several communities, including epidemiology, network science and computer science. However, many algorithms require detailed information about the spreading process, such as the exact moment when each individual was infected or the person by whom each individual was infected, and the sudden and unexpected onset of epidemics make this data difficult to obtain in practice. Our algorithm is designed for the more common scenario in which the available information is limited to three aspects (see Fig. 1). First, a snapshot of the distribution of the disease throughout the population at the final moment. This final state is usually obtained by carrying out a full screening, in COVID-19, for example, by nucleic acid testing. Second, the contact relation of the local population during the contagion processes, which can be modelled as a temporal network in which nodes represent individuals and links represent physical proximity within a distance which risks transmission. These contact networks can be reconstructed from survey or modern location technologies, such as mobile signaling data. Third, an estimated time interval – not necessarily a precise time – when the outbreak starts.

In many epidemic models, individuals are infected with a certain probability when they come into contact with a patient, with the infection rate depending on the type of disease. Thus, the spreading process is strongly random. Even with the same patient zero, restarting the spreading process many times can result in very different snapshots. For example, when patient zero encounters an individual who is a hub node in the contact network, whether this encounter becomes an effective transmission has a great influence on the final size of infected population. On the other hand, similar final snapshots can emerge from distinct patient zeros, especially when their topological distance on the temporal network is small. Soft Margin Monte Carlo estimator (Antulov-Fantulin et al. 2015) runs extensive simulations of spreading taking each infected individual as a potential patient zero in turn, and quantifies the similarity between the simulated and real snapshots. The individual with maximum similarity is identified as patient zero. However, this method provides limited insights into how the behavior of the spreading dynamics system depends on the initial state (Strogatz 2018),

*Corresponding author

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Modeling of contagion process in local population. Contacts between individuals change over time and can be represented by a temporal network G . The SIR model is applied to simulate the spreading process on this temporal contact network. The individual carrying the virus at the initial moment is called patient zero. Our problem is to infer who is the patient zero from the current (final) state of the population.

and fails when contact networks are dense or the estimated parameters used in simulation deviate slightly from those of the real contagion process.

To provide implicit understanding and improve concrete inferences, we tailor a graph neural network (GNN) (San-toro et al. 2017; Battaglia et al. 2018; Zhou et al. 2020; Wu et al. 2020) to establish the statistical association between initial sources and final snapshots under the stochastic spreading environment. The source detection problem can be mapped to a multi-classification problem in which each individual in the population is a separate category. The gnn-based model processes the encoded spreading information, and estimates the possibility of each individual to be patient zero (see Fig. 2). Experiments on simulation data generated by classic epidemiological models on five real temporal networks indicate that the performance of our method is significantly superior to established methods. We also design two scenarios with missing information to show the robustness of the proposed algorithm, including partially missing contact structure and partial non-observation of population state. Further, we give some insight into the interior reasoning of gnn-based models based on topological properties of patient zero. Specifically, we find the closeness of the activity sequence (which records the times at which a node had an active edge) and core number (Batagelj and Zaversnik 2003; Carmi et al. 2007) to those of the true patient zero tend to increase with inferred possibility to be the infection source, but the same is not true of degree.

Related Work

Several methods of source detection (Huang et al. 2018) are based on calculating measures of node centrality on the sampled network, such as distance centrality, Jordan centrality (Zhu and Ying 2014), closeness centrality, modified betweenness centrality (Comin and da Fontoura Costa 2011), dynamical age (Fioriti, Chinnici, and Palomo 2014), and rumor centrality (Shah and Zaman 2010, 2011), which estimates the probability of a node as the source as proportional to the count of all possible sequences of infection starting from that node.

The strategy of Pinto, Thiran, and Vetterli (2012) requires a set of monitoring nodes which record from which neighbor and at which times information is received, and adopt

a maximum likelihood estimator for potential sources that averages over two different sources of randomness: the uncertainty in the paths that the information takes to reach the monitors and the uncertainty in the time that the information takes to cross the edges. For every possible source of the epidemic, Lokhov et al. (2014) use a fast dynamic message-passing method (DMP) to estimate the probability that all nodes in the network are the state in the observed snapshot, and a mean-field like approximation to compute the probability of the observed snapshot as a product of the marginal probabilities. Finally, the potential sources are ranked according to that probability. To overcome some drawbacks of DMP, the belief propagation equations are derived (Altarelli et al. 2014) for the posterior distribution of the time evolution of the state of the system conditioned on some observations. GNN has been used to study contagion processes for learning the reverse dynamics and predicting patient zero, yet on static networks (Shah et al. 2020). Time-reversal backward spreading (TRBS) (Shen et al. 2016) requires the arrival time of certain signals at monitoring nodes and network structure weighted to represent time delay of propagation. TRBS starts from a monitor o_k and spreads to all nodes in the network along the reversed direction of links to yield a reversed arrival time $t_{o_k} - \hat{t}(i, o_k)$ at a node i , where $\hat{t}(i, o_k)$ is the shortest time delay from o_k to i . Then, the node i gets a vector $\mathbf{T}_i = [t_{o_1} - \hat{t}(i, o_1), t_{o_2} - \hat{t}(i, o_2), \dots, t_{o_m} - \hat{t}(i, o_m)]^T$ from the set of monitors. The node with the minimum variance of $\mathbf{T}_i$ is the inferred source. Backward temporal diffusion process (Huang et al. 2017) extends TRBS to temporal networks.

Most algorithms are either developed for static networks or require a considerable number of monitoring nodes to precisely record information during the spreading process. This is not applicable to our problem due to the time-varying temporal contact network and the unexpected nature of the outbreak making infeasible the timely installation of monitors.

Method

Problem Formulation

The local population represents a typical complex system. To represent the architecture of the complex system we use a temporal network in which nodes are the system's components and temporal links capture changing interactions




Figure 2: Schematic demonstration of inference by backtracking network. A three-dimensional vector to characterize the state of an individual: susceptible—[1 0 0], infectious—[0 1 0] and recovered—[0 0 1]. The slice sequence of posited temporal network $\tilde{G}$ is integrated into an aggregation. A binary vector, with length equal to the time span of $\tilde{G}$, characterizes the activation pattern of each edge. The BN gives the probability each individual is patient zero.

between these components (Holme and Saramäki 2012; Li et al. 2017). The temporal network $G = \{G_0, G_1, \dots, G_T\}$ is an ordered sequence of slice networks G_t on the same node set V , as shown in Fig. 1. The evolution of states on slice networks are represented as $S = \{S_0, S_1, \dots, S_T\}$, where S_t is the state of the node set on slice G_t . Since an interval of possible initial times is known rather than an exact initial time, the posited temporal network $\tilde{G}$ may start either earlier or later than reality G . Then, our problem is expressed concisely as: inferring S_0 given $\{S_T, \tilde{G}\}$. Assuming only one node is infected in S_0 , the problem is further refined into that of inferring patient zero.

Modeling of Contagion Process

Epidemic modeling describes the dynamical evolution of the contagion process within the local population as a function of time (Pastor-Satorras et al. 2015). Epidemic models generally assume that the population can be divided into different classes depending on the stage of the disease (Diekmann and Heesterbeek 2000), such as susceptible S who are healthy currently and can be exposed to infection, infectious I who are sick and can transmit the disease to contacts, and recovered R who have previously had the disease and are now non-infectious and immune. Here we adopt the classic three-state susceptible-infected-recovered (SIR) model which defines basic individual-level processes governing the two possible transitions of individuals from one class to another as:

$$S(i) + I(j) \xrightarrow{p} I(i) + I(j), \quad (1)$$

$$I(i) \xrightarrow{q} R(i). \quad (2)$$

The first transition Eq. 1, denoted $S \rightarrow I$, occurs when a susceptible individual i has contact with an infectious individual j and becomes infected with infection rate p . The second transition Eq. 2, denoted $I \rightarrow R$, occurs when the infectious individual i recovers from the disease with recovery

rate q and obtains immunity which lasts longer than the time scale of the epidemic outbreak.

Encoding Available Information

The limited information available for source detection includes the final state and structure of temporal networks. These are designed as a representation that can be exploited in machine learning models.

a. Final state. When the SIR model is applied to simulate the spreading process, the state of an individual C_i can transform among three classes: susceptible, infectious and recovered. The final state of the complex system is a macro aggregate comprising the microstates of the node set, $S_T = \{C_0, C_1, \dots, C_N\}$. The one-hot encoding technique is used to encode the individual state as a node feature (see Fig. 2).

b. Structure of temporal networks. The aggregated network $\tilde{G}_a(V, E)$ is a superposition of all slices in the temporal network $\tilde{G}$, containing any edge which appears in at least one slice. The adjacent matrix A_a of $\tilde{G}_a$ represents the static topology properties of interactions between nodes, but cannot provide temporal properties depending on the times when interactions are active. An additional feature vector X_e , of dimension equal to the length of the sequence $\{\tilde{G}_0, \tilde{G}_1, \dots, \tilde{G}_T\}$, is attached to each edge e to record its activation pattern (see Fig. 2). The set of all edge features is $X_a = \{X_e \mid \forall e \in E\}$.

Backtracking Network

We tailor a simple graph neural network focusing on source detection, called backtracking network (BN), as shown in Fig. 2. Compared with the spreading dynamic which drives the system from an initial state forward to its final state, the role of BN is to approximate an inverse function mapping the final state back to its initial state: $S_0 = \mathbf{BN}(S_T, \tilde{G})$. Given conditions S_T , A_a and X_a , the BN computes with

the kernel-based convolutional operator (Gilmer et al. 2017; Simonovsky and Komodakis 2017):

$$v \rightarrow e: \quad g_{(i,j)}^{l+1} = f^e(g_{(i,j)}^l, h_j^l), \quad (3)$$

$$e \rightarrow v: \quad h_i^{l+1} = f^v(h_i^l) + \sum_{j \in \mathcal{N}(i)} g_{(i,j)}^{l+1}, \quad (4)$$

where $g_{i,j}^l$ and h_i^l are the hidden features of edge (i, j) and node i at layer l respectively. Initially, $g_{(i,j)}^0 = p^e(X_{(i,j)})$ and $h_i^0 = p^v(C_i)$, where p^e/p^v is the projection layer. The message function f^e delivers the information on node to the edge originating from it. The update function f^v integrates the information about disease states of neighbors and activation patterns of their contacts. f^e and f^v are represented by two independent fully-connected layers with activation function ReLU. The objective function is formulated as:

$$\mathcal{L} = - \sum_{i=1}^N y_i \log \frac{\exp(h_i)}{\sum_{j=1}^N \exp(h_j)}, \quad (5)$$

where $h_i \in \mathbb{R}^1$ is the return of another projection layer following the gnn-based layers, and label $y_i = 1$ if node i is the real source else $y_i = 0$. The softmax function inside the log function maps h_i to the detection probability of node i being the source.

Expert Knowledge

Expert knowledge can complement machine, improving performance by injecting human prior knowledge into the machine reasoning process, which is particularly valuable in limited sample scenarios (Raissi, Perdikaris, and Karniadakis 2019; Hao et al. 2022). In SIR model, it is obvious that individuals whose states are susceptible at final moment are certainly not the source. If an individual is the source being infectious at initial moment, his state will not transition to be susceptible (see Eq. 1-2). For humans, this knowledge is easy to obtain because we know the mechanism of the contagion process underlying the data generation whereas the machine does not. What BN learns is that those susceptible individuals have never been the sources in training samples (statistical association – thus gives them low detection probabilities), not that they cannot be the sources (spreading mechanism – these probabilities should be zero). As the sample size increases, the given probabilities get close to zero gradually. However, it may cause the BN to make wrong inferences when sample size is not sufficient enough. Therefore, we add a huge negative value on the susceptible individuals to force their detection probabilities to become zero after applying softmax function according to:

$$h_s \leftarrow h_s - \infty. \quad (6)$$

Then, h_s is the final output on susceptible node s combined with expert knowledge. The overall framework of our method for source detection is summarized as Algorithm 1.

Experiment

We evaluate the performance of the proposed method on five real temporal networks, check its robustness in two scenar-

Algorithm 1: Inference of patient zero

Input: Final snapshot of spreading in population S_T , epidemic model e.g. SIR, estimated start interval of propagation, backtracking network BN

Output: Patient zero

- 1: Generate training set with size M .
 - 2: **while** M has not meet **do**
 - 3: Select start time t_0 in estimated interval and a individual activated at t_0 as patient zero randomly.
 - 4: Simulate spreading process using SIR on posited contact network $\tilde{G}$
 - 5: Collect the final snapshot and patient zero as a pair of sample and label.
 - 6: **end while**
 - 7: Encode the available information in samples
 - 8: Train BN combined with expert knowledge
 - 9: **return** Patient zero in S_T inferred by BN
-

ios where available spreading information is partially missing, and analyse the relationship between topological attributes in contact network and the detectability of patient zero.

Experiment Steps

Temporal networks. Statistics of the following five networks are provided in Tab. 3: (a) Sexual contact (Rocha, Liljeros, and Holme 2011), is an empirical temporal network of sexual contacts in Brazil; (b) Hospital (Isella et al. 2011), contains contacts among patients and health-care workers in a hospital ward in Lyon, France; (c) European email (Leskovec, Kleinberg, and Faloutsos 2007), is generated using email data from a large European research institution; (d) Bitcoin (Kumar et al. 2016), is a who-trusts-whom network of people who trade using Bitcoin on a platform called Bitcoin OTC; (e) Message (Panzarasa, Opsahl, and Carley 2009), comprises private messages sent on an online social network at the University of California, Irvine. The two face-to-face networks, Sexual contact and Hospital, are undirected, while others are directed and transmission only occurs along the direction of an edge.

Datasets. We generate 20,000 samples (19,600/200/200 for training/validation/test) for each temporal network by using the SIR model with two parameter combinations of mean infection and recovery rate ($\bar{p} = 0.3, \bar{q} = 0.01$ and

Name	Nodes	Edges	tmin/tmax/tend	<k>	<ka>	Directed
Sex	5762	9319	1/10/101	1.2	3.23	N
Hospital	75	1139	1/5/71	4.4	30.4	N
Eu email	625	2113	1/3/11	1.2	3.38	Y
Bitcoin	1302	4053	1/10/81	0.8	3.11	Y
Message	1522	8764	1/3/11	1.9	5.76	Y

Table 1: Statistics of real temporal networks. $\langle k \rangle$ and $\langle k_a \rangle$ is the mean degree in slice and aggregate networks respectively. Times are given in days.




Figure 3: Distribution of final number of affected individuals in datasets. The value on abscissa represents the final number of Infectious + Recovered individuals, among whom we infer the single patient zero. There are 20,000 independent samples in each dataset.

$\bar{p} = 0.8, \bar{q} = 0.05$) to simulate the spreading process. Thus, we obtain ten experimental datasets in total, and the distribution of the final number of individuals affected by the epidemic is shown in Fig. 3. Each simulation individually chooses the start time $t_0 \in \mathbb{N}_+$ from a fixed interval $[t_{\min}, t_{\max}]$ randomly and selects patient zero randomly from among all individuals who are active (i.e., participate in an active edge) at time t_0 . Simulations end at time t_{end} . When using these samples, the start time t_0 is masked and only the estimated interval $[t_{\min}, t_{\max}]$ is known. Thus, we unify the time span of the posited temporal networks $\tilde{G}$ input into BN as $[t_{\min}, t_{\text{end}}]$ for all samples. For each simulation the infection rate p and recovery rate q are independently sampled from the uniform distribution $U_p \sim [\bar{p} - 0.1, \bar{p} + 0.1]$ and $U_q \sim [\bar{q} - 0.01, \bar{q} + 0.01]$. This choice accommodates possible differences between reality and the calibrated epidemic model.

Baselines. We compared BN with five baselines: (a) Jordan centrality (Zhu and Ying 2014); (b) Modified betweenness centrality (Comin and da Fontoura Costa 2011); (c) Dynamical age (Fioriti, Chinnici, and Palomo 2014); (d) Backtracking network without expert knowledge (BNW); (e) Soft Margin Monte Carlo estimator (SM). For a test sample, SM runs a large number M (200 in our experiment) of epidemic simulations taking each infectious and recovered individual in turn as the potential patient zero, and the inferred probability is proportional to the similarity between the real and simulated snapshots calculated as (Antulov-Fantulin et al. 2015):

$$P(i) \propto \frac{1}{M} \sum_{m=1}^M \exp\left(-[\varphi(\vec{r}_*, \vec{r}_{i,m}) - 1]^2\right), \quad (7)$$

where the vectors $\vec{r}_*$ and $\vec{r}_{i,m} \in \mathbb{R}^N$ record the state of the

population in the real snapshot and m -th simulated one with individual i as source. φ is the Jaccard similarity function $\varphi : \mathbb{R}^N \times \mathbb{R}^N \rightarrow [0, 1]$ given by the ratio of the size of the intersection of the sets of infected individuals in $\vec{r}_*$ and $\vec{r}_i$ to the size of their union; (f) Soft Margin Monte Carlo estimator with exact start time (SMT). All methods share test sets.

Evaluation metrics. Under each evaluation metric, a success is as follows: (a) Top-1. The real patient zero is the individual with maximal inferred probability to be patient zero; (b) Top-5. The real patient zero is among the five individuals with the highest inferred probability to be patient zero; (c) Hop-1. The topological distance between real and inferred patient zero on the aggregate network G_a is less than or equal to one. In other words, the inferred patient zero is within one hop of the true patient zero (either the true patient zero or the true patient zero’s direct neighbor).

Training details. We train BN by the ADAM (Kingma and Ba 2015) optimizer with learning rate of 10^{-3} (we report the variation of performance with the number of layers in Results). The batch size is 128 and training epoch is 200. We compose our code on PyTorch framework and run all experiments in a local machine with two NVIDIA Tesla V100 32GB GPUs.

Results

On all datasets, the performance of our gnn-based method is significantly superior (see Tab. 2) to SM and SMT, which for Top-5 only achieves accuracy above 0.6 on the small-scale Hospital network, with 75 nodes. The failure of SM and SMT indicates the similarity between the snapshots generated by multiple simulations on the same patient zero is not distinctly higher than the similarity for distinct patient zeros with small topological distance on the temporal network. The superiority of SMT (i.e. SM with exact start time) to SM shows that even if the spreading origins from the same patient zero, starting a few days earlier or later results in different epidemic processes and final snapshots. We find that the accuracy of BN is usually better than that of BNW in metric Top-1. When BNW recognizes as suspicious a set of individuals including both the true patient zero and susceptible individuals, BNW (without expert knowledge) needs more features than BN to exclude the susceptible individuals.

We check the effect of receptive fields of nodes on BN’s performance. As the number of layers (one layer comprises both Eq. 2 and Eq. 3) increases, the node can gather information from its neighbors with farther topological distance, i.e., the receptive field becomes larger. Taking epidemic spreading on sexual contact network as an example, the variation of performance with a number of layers from 1 to 7 is shown in Tab.3. We find the performance approaches its maximum when the number of layers is 5, and further expanding the receptive field provides no obvious improvement to the performance. In other datasets, the performance even declines since the information gathered on adjacent nodes become homogeneous as their receptive field becomes too large.

$\bar{p}=0.3 \bar{q}=0.01$	Sex	Hospital	Bitcoin	Eu email	Message
Jordan	0.02 / 0.08 / 0.19	0.12 / 0.12 / 0.82	0.04 / 0.06 / 0.25	0.07 / 0.24 / 0.41	0.02 / 0.05 / 0.16
Betweenness	0.01 / 0.03 / 0.01	0.00 / 0.00 / 0.17	0.01 / 0.01 / 0.02	0.00 / 0.00 / 0.02	0.00 / 0.02 / 0.01
Dynamic Age	0.03 / 0.08 / 0.19	0.10 / 0.26 / 0.97	0.00 / 0.04 / 0.30	0.09 / 0.25 / 0.58	0.00 / 0.01 / 0.35
SM	0.10 / 0.22 / 0.28	0.44 / 0.79 / 0.98	0.08 / 0.23 / 0.40	0.16 / 0.44 / 0.65	0.08 / 0.13 / 0.30
SMT	0.12 / 0.28 / 0.29	0.64 / 0.91 / 0.98	0.07 / 0.26 / 0.43	0.17 / 0.49 / 0.67	0.09 / 0.16 / 0.27
BNW	0.50 / 0.91 / 0.87	1.00 / 1.00 / 1.00	0.41 / 0.88 / 0.85	0.39 / 0.84 / 0.76	0.25 / 0.52 / 0.54
BN	0.50 / 0.92 / 0.88	1.00 / 1.00 / 1.00	0.45 / 0.87 / 0.89	0.44 / 0.85 / 0.83	0.28 / 0.59 / 0.53
$\bar{p}=0.8 \bar{q}=0.05$	Sex	Hospital	Bitcoin	Eu email	Message
Jordan	0.01 / 0.02 / 0.04	0.09 / 0.09 / 0.86	0.00 / 0.01 / 0.21	0.05 / 0.10 / 0.23	0.01 / 0.01 / 0.14
Betweenness	0.02 / 0.02 / 0.02	0.00 / 0.00 / 0.29	0.00 / 0.01 / 0.01	0.00 / 0.01 / 0.01	0.01 / 0.02 / 0.02
Dynamic Age	0.00 / 0.02 / 0.04	0.19 / 0.38 / 0.95	0.00 / 0.07 / 0.28	0.02 / 0.06 / 0.29	0.02 / 0.03 / 0.37
SM	0.04 / 0.13 / 0.14	0.33 / 0.65 / 0.98	0.03 / 0.12 / 0.36	0.25 / 0.52 / 0.61	0.12 / 0.21 / 0.34
SMT	0.04 / 0.10 / 0.14	0.31 / 0.65 / 0.98	0.04 / 0.13 / 0.37	0.25 / 0.52 / 0.62	0.13 / 0.24 / 0.37
BNW	0.40 / 0.86 / 0.82	1.00 / 1.00 / 1.00	0.47 / 0.86 / 0.85	0.49 / 0.86 / 0.78	0.35 / 0.62 / 0.56
BN	0.45 / 0.92 / 0.84	1.00 / 1.00 / 1.00	0.46 / 0.85 / 0.79	0.55 / 0.91 / 0.82	0.36 / 0.66 / 0.58

Table 2: Performance comparison using dataset from SIR model with mean infection and recovery rates (top: $\bar{p} = 0.3, \bar{q} = 0.01$; bottom: $\bar{p} = 0.8, \bar{q} = 0.05$) on five real temporal networks. Each element contains three values representing the accuracy in metric Top-1/Top-5/Hop-1 respectively.

Metric	L1	L2	L3	L4	L5	L6	L7
Top-1	0.39	0.45	0.45	0.46	0.48	0.45	0.49
Top-5	0.78	0.82	0.88	0.88	0.92	0.89	0.91
Hop-1	0.77	0.81	0.82	0.86	0.87	0.88	0.86

Table 3: Performance versus receptive field of node on dataset of sexual contact. In column L_n , BN uses n layers.

Robustness

We design two more practical scenarios in which spreading information is missing, as shown in Fig. 4. (a) Links missing. The incomplete network structure generates fictitious temporal patterns of contacts and activity rhythms of individuals in $\tilde{G}$. The original propagation paths are obscured and clues to infer by whom patients are infected became blurred. To investigate this, we adjust the proportion of missing links from mild to severe, and remove the links randomly from each slice of the temporal network. (b) Partial observation. Only a fraction of node states in final state S_T is observed. This leads to four classes of individual states in $\tilde{S}_T$: susceptible, infectious, recovered and unknown. Confused by the unobserved individuals, we even fail to judge the final number of affected individuals. In this case, the problem is expressed as: inferring S_0 given $\{\tilde{S}_T, \tilde{G}\}$. Concretely, we adjust the proportion of observed nodes from dense to sparse, and select observed nodes randomly from the population. Taking epidemic spreading on sexual contact network as an example, the robustness of our algorithm is shown in Fig. 5. The faster decline in the first scenario supports the assertion that the underlying structure of the temporal network has a




Figure 4: Graphic illustration of two types of information missing. (a) Contact missing. The light gray dashed lines represent the missing links between the nodes. (b) Partial observation.

drastic impact on spreading dynamics (Holme and Saramäki 2012; Bansal et al. 2010). When the information about network structure is insufficient, BN is challenged to establish correctly the association between the final state and initial state which represents the inversion of spreading dynamics. In contrast, partial observation has a substantially weaker influence on inference via BN. Even if eighty percent of final states are unknown, BN still achieves around 0.7 accuracy according to the metric Top-5.

Topological Insight into the Interior of BN

To give some insight into the implicit inverse association $S_0 = \text{BN}(S_T, \tilde{G})$, we analyze from this perspective: what are the common characteristics among the individuals which BN assigns top probability of being the source? We visualize two cases of the final snapshot of the epidemic on temporal network of sex contacts, and observe that the individ-




Figure 5: Robustness to missing information. The values on the abscissa represent the proportion of either missing links (ml) or unobserved individuals (ui). The filled areas indicate the gap of robustness between two scenarios.

uals given higher probabilities by BN always play a role in the spreading process matched with that of the real patient zero, as shown in Fig. 6. (a) Passersby. The patient zero, a fringe node with 1 degree, activates only in one day, infects a hub neighbor by chance ($\bar{p}=0.3$) and never activates again. The outbreak is caused mainly by the subsequent propagation from this neighbor of patient zero, which is a hub having 47 neighbors. (b) Participant. The patient zero is a hub with frequent activation, and contacts 91 other nodes. The patient zero makes a substantial contribution to the outbreak on its own. Comparing (a) with (b), we notice that although these two patient zeros play completely different roles in the spreading process, in either case BN can find alternative nodes which could fulfil a role similar to the true patient zero. (a) top1 (source), top2 and top5 are all fringe nodes and connect to the same hub node (top3) directly. BN recognizes that top3 is the main contributor of this outbreak, but recognises that instead of being the source, it may instead be infected on a certain slice of the temporal network. (b) top1~5 are all hubs, but are not simply the nodes in the network with the highest degrees.

For further understanding, we quantify the match between topological properties of the nodes which BN estimates have the highest probability to be patient zero and topological




Figure 6: The role of patient zero plays in spreading process (a) Passersby and (b) Participant. The five nodes given top probabilities are enlarged in size and filled with different shades of red. The two local aggregated networks shown here only contains the second-order inner neighbors of patient zero whose states are infectious or recovered in the final snapshot.




Figure 7: Quantification of match between patient zero and nodes given top inferred probability (a) Mean similarity of activity sequence and (b) MAD of k-shell and degree on aggregated network $\tilde{G}_a$.

properties of the true patient zero. The three metrics we consider are mean similarity of activity sequences, and absolute mean absolute deviation (MAD) between degree and k-shell on the aggregated network, as shown in Fig. 7. The similarity of activation sequences is defined as $(v_i \cdot v_j) / (|v_i| |v_j|)$, where v_i is the activity sequence of node i on temporal network. Sorting nodes in descending order by inferred probability, we observe the progressive decrease of similarity, and also an increasing trend for MAD of k-shell but not degree, which initially drops rapidly but then fluctuates. This indicates that nodes with degree similar to patient zero are not necessarily given high inferred probability. The BN prefers to extract k-shell as topological inference identifier rather than degree, corresponding to the conclusion (Kitsak et al. 2010) that core in the network determines efficiency of spreaders.

Conclusion

Our research involves backtracking the spreading of epidemics on temporal networks and inferring patient zero with the help of artificial intelligence. The gnn-based method we propose opens a novel path to deal with this kind of inverse problem in random environments. We show that this problem is solved efficiently by a gnn-based method even in the case of missing information about network structure and node states.

Source detection is not only restricted to epidemiology, but extends to other fields such as rumor in sociology and pollution in ecology. The emergency of any of these phenomena will lead to confusion in the local area, and locating the source is crucial for subsequent interventions to eradicate negative events. The proposed method can be easily transferred to solve these problems. Moreover, the inference of patient zero is just a special case of initial state inference for complex systems the entities of which can have a diverse mixture of states described by a mixture of discrete and continuous values, and many important scientific questions can be formulated in this way. Hence, extending our research to inference of more general types of initial states is an interesting direction.

Ethics Statement

Our problem entails privacy and ethical concerns: the identity of spreading sources, and the data which would allow identification of these sources, should be kept safe.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (grant no. T2225022, 12161141016, 62088101, 11875043 and 12150410309), Shanghai Municipal Science and Technology Major Project (grant no. 2021SHZDZX0100), Shanghai Municipal Commission of Science and Technology Project (grant no. 19511132101), and Fundamental Research Funds for the Central Universities.

References

- Altarelli, F.; Braunstein, A.; Dall'Asta, L.; Lage-Castellanos, A.; and Zecchina, R. 2014. Bayesian inference of epidemics on networks via belief propagation. *Physical Review Letters*, 112(11): 118701.
- Antulov-Fantulin, N.; Lančić, A.; Šmuc, T.; Štefančić, H.; and Šikić, M. 2015. Identification of patient zero in static and temporal networks: Robustness and limitations. *Physical Review Letters*, 114(24): 248701.
- Bansal, S.; Read, J.; Pourbohloul, B.; and Meyers, L. A. 2010. The dynamic nature of contact networks in infectious disease epidemiology. *Journal of Biological Dynamics*, 4(5): 478–489.
- Batagelj, V.; and Zaversnik, M. 2003. An $O(m)$ algorithm for cores decomposition of networks. *arXiv preprint cs/0310049*.
- Battaglia, P. W.; Hamrick, J. B.; Bapst, V.; Sanchez-Gonzalez, A.; Zambaldi, V.; Malinowski, M.; Tacchetti, A.; Raposo, D.; Santoro, A.; Faulkner, R.; et al. 2018. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*.
- Buch, V. H.; Ahmed, I.; and Maruthappu, M. 2018. Artificial intelligence in medicine: current trends and future possibilities. *British Journal of General Practice*, 68(668): 143–144.
- Carmi, S.; Havlin, S.; Kirkpatrick, S.; Shavitt, Y.; and Shir, E. 2007. A model of Internet topology using k-shell decomposition. *Proceedings of the National Academy of Sciences*, 104(27): 11150–11154.
- Comin, C. H.; and da Fontoura Costa, L. 2011. Identifying the starting point of a spreading process in complex networks. *Physical Review E*, 84(5): 056105.
- Diekmann, O.; and Heesterbeek, J. A. P. 2000. *Mathematical epidemiology of infectious diseases: model building, analysis and interpretation*, volume 5. John Wiley & Sons.
- Fioriti, V.; Chinnici, M.; and Palomo, J. 2014. Predicting the Sources of an Outbreak with a Spectral Technique. *Applied Mathematical Sciences*, 8(135): 6775–6782.
- Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; and Dahl, G. E. 2017. Neural message passing for quantum chemistry. In *International Conference on Machine Learning*, 1263–1272. PMLR.
- Hao, Z.; Liu, S.; Zhang, Y.; Ying, C.; Feng, Y.; Su, H.; and Zhu, J. 2022. Physics-Informed Machine Learning: A Survey on Problems, Methods and Applications. *arXiv preprint arXiv:2211.08064*.
- Holme, P.; and Saramäki, J. 2012. Temporal networks. *Physics Reports*, 519(3): 97–125.
- Huang, C.; Liu, X.; Deng, M.; Zhou, Y.; and Bu, D. 2018. A Survey on Algorithms for Epidemic Source Identification on Complex Networks [J]. *Chinese Journal of Computers*, 41(06): 1156–1179.
- Huang, Q.; Zhao, C.; Zhang, X.; and Yi, D. 2017. Locating the source of spreading in temporal networks. *Physica A: Statistical Mechanics and its Applications*, 468: 434–444.
- Isella, L.; Romano, M.; Barrat, A.; Cattuto, C.; Colizza, V.; Van den Broeck, W.; Gesualdo, F.; Pandolfi, E.; Ravà, L.; Rizzo, C.; et al. 2011. Close encounters in a pediatric ward: measuring face-to-face proximity and mixing patterns with wearable sensors. *PloS One*, 6(2): e17144.
- Jittprasong, C. 2022. Artificial Intelligence and Medicine: A literature review. *arXiv preprint arXiv:2205.00322*.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In Bengio, Y.; and LeCun, Y., eds., *3rd International Conference on Learning Representations*.
- Kitsak, M.; Gallos, L. K.; Havlin, S.; Liljeros, F.; Muchnik, L.; Stanley, H. E.; and Makse, H. A. 2010. Identification of influential spreaders in complex networks. *Nature Physics*, 6(11): 888–893.
- Kumar, S.; Spezzano, F.; Subrahmanian, V.; and Faloutsos, C. 2016. Edge weight prediction in weighted signed networks. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*, 221–230. IEEE.
- Leskovec, J.; Kleinberg, J.; and Faloutsos, C. 2007. Graph evolution: Densification and shrinking diameters. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 1(1): 2–es.
- Li, A.; Cornelius, S. P.; Liu, Y.-Y.; Wang, L.; and Barabási, A.-L. 2017. The fundamental advantages of temporal networks. *Science*, 358(6366): 1042–1046.
- Lokhov, A. Y.; Mézard, M.; Ohta, H.; and Zdeborová, L. 2014. Inferring the origin of an epidemic with a dynamic message-passing algorithm. *Physical Review E*, 90(1): 012801.
- Panzarasa, P.; Opsahl, T.; and Carley, K. M. 2009. Patterns and dynamics of users' behavior and interaction: Network analysis of an online community. *Journal of the American Society for Information Science and Technology*, 60(5): 911–932.
- Pastor-Satorras, R.; Castellano, C.; Van Mieghem, P.; and Vespignani, A. 2015. Epidemic processes in complex networks. *Reviews of Modern Physics*, 87(3): 925.
- Pinto, P. C.; Thiran, P.; and Vetterli, M. 2012. Locating the source of diffusion in large-scale networks. *Physical Review Letters*, 109(6): 068702.
- Raissi, M.; Perdikaris, P.; and Karniadakis, G. E. 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving

nonlinear partial differential equations. *Journal of Computational physics*, 378: 686–707.

Rocha, L. E.; Liljeros, F.; and Holme, P. 2011. Simulated epidemics in an empirical spatiotemporal network of 50,185 sexual contacts. *PLoS Computational Biology*, 7(3): e1001109.

Santoro, A.; Raposo, D.; Barrett, D. G.; Malinowski, M.; Pascanu, R.; Battaglia, P.; and Lillicrap, T. 2017. A simple neural network module for relational reasoning. *Advances in Neural Information Processing Systems*, 30.

Shah, C.; Dehmamy, N.; Perra, N.; Chinazzi, M.; Barabási, A.-L.; Vespignani, A.; and Yu, R. 2020. Finding patient zero: Learning contagion source with graph neural networks. *arXiv preprint arXiv:2006.11913*.

Shah, D.; and Zaman, T. 2010. Detecting sources of computer viruses in networks: theory and experiment. In *Proceedings of the ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, 203–214.

Shah, D.; and Zaman, T. 2011. Rumors in a network: Who’s the culprit? *IEEE Transactions on Information Theory*, 57(8): 5163–5181.

Shen, Z.; Cao, S.; Wang, W.-X.; Di, Z.; and Stanley, H. E. 2016. Locating the source of diffusion in complex networks by time-reversal backward spreading. *Physical Review E*, 93(3): 032301.

Simonovsky, M.; and Komodakis, N. 2017. Dynamic edge-conditioned filters in convolutional neural networks on graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3693–3702.

Strogatz, S. H. 2018. *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering*. CRC press.

Wu, Z.; Pan, S.; Chen, F.; Long, G.; Zhang, C.; and Philip, S. Y. 2020. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1): 4–24.

Zhou, J.; Cui, G.; Hu, S.; Zhang, Z.; Yang, C.; Liu, Z.; Wang, L.; Li, C.; and Sun, M. 2020. Graph neural networks: A review of methods and applications. *AI Open*, 1: 57–81.

Zhu, K.; and Ying, L. 2014. Information source detection in the SIR model: A sample-path-based approach. *IEEE/ACM Transactions on Networking*, 24(1): 408–421.