

Group Sparse Bayesian Learning for Active Surveillance on Epidemic Dynamics

Hongbin Pei,^{1,2} Bo Yang,^{1,2*} Jiming Liu,³ Lei Dong⁴

¹College of Computer Science and Technology, Jilin University, Changchun, China

²Key Laboratory of Symbolic Computation and Knowledge Engineering of Ministry of Education, China

³Department of Computer Science, Hong Kong Baptist University, Hong Kong

⁴Institute of Remote Sensing and Geographical Information Systems, Peking University, Beijing, China

Abstract

Predicting epidemic dynamics is of great value in understanding and controlling diffusion processes, such as infectious disease spread and information propagation. This task is intractable, especially when surveillance resources are very limited. To address the challenge, we study the problem of active surveillance, i.e., how to identify a small portion of system components as sentinels to effect monitoring, such that the epidemic dynamics of an entire system can be readily predicted from the partial data collected by such sentinels. We propose a novel measure, the γ value, to identify the sentinels by modeling a sentinel network with row sparsity structure. We design a flexible group sparse Bayesian learning algorithm to mine the sentinel network suitable for handling both linear and non-linear dynamical systems by using the expectation maximization method and variational approximation. The efficacy of the proposed algorithm is theoretically analyzed and empirically validated using both synthetic and real-world data.

1 Introduction

Predicting epidemic dynamics is of great value in understanding and controlling diffusion processes. Diffusion phenomena, such as infectious disease spread and information propagation, exist widely in the real world. By the notion of a dynamical system, which is a powerful tool for characterizing dynamics (Brunton, Proctor, and Kutz 2016), the task of epidemic dynamics prediction is to estimate ensuing system states using given current states. Therefore, the foundation of this task is surveillance, which is to monitor and report the current system states in a timely manner.

However, in practice, it is often very challenging to monitor the overall components in a system because diffusion phenomena often cross over very large spatiotemporal ranges, e.g., spreading of infectious diseases in a country (Polgreen et al. 2009), air contaminant diffusion in a large city (Zheng, Liu, and Hsieh 2013), and hot topics/meme forwarding on social media (Chen et al. 2013). For large-scale diffusion phenomena, timely and all-around monitoring is hard and infeasible, especially when available surveillance resources are very limited.

Owing to a lack of systematic deployment of already limited surveillance resources, disease surveillance has long suffered from low reporting rates, biased sampling, and lengthy reporting time-lags (Gerardo-Giorda et al. 2013). For instance, Tengchong city, a malaria endemic region in China, has 18 towns (consisting of 221 villages), 167,964 households, and 658,207 residents that are distributed in a mountainous area of 5,845 square kilometers. From 2005 to 2011 in Tengchong, 7,835 confirmed malaria cases were reported, but the Tengchong Centers for Disease Control (CDC), the local disease surveillance team, could only afford to have a few staff members conducting the time-consuming case surveys.

Active surveillance is a promising strategy to address the challenge of limited surveillance resources, in which epidemic dynamics is predicted by proactively monitoring a relatively small number of sentinel components, whose state data are collected to achieve a good trade-off between prediction accuracy and surveillance cost. The key to implementing active surveillance is to determine, in a dynamical system, which components are important for dynamics prediction and how to identify them. This is a non-trivial task because the interaction structure among components, which characterizes the dynamics mechanism, is usually hidden and heterogenous, such as the social contact network in charge of disease spread (Yang et al. 2017).

Here, we address this challenge by proposing a novel importance measure, the γ value, to determine how important a component is to predicting epidemic dynamics. Based on the measure, we develop a backward-selection algorithm designated the sentinel network mining algorithm (SNMA for short) to mine a sentinel network. The sentinel network encodes the influence relationship from sentinel components to the overall system components, which is a row sparse network that contains only the influential links emitted from the sentinels. With the discovered sentinel network, one can predict the future overall system dynamics by only monitoring and feeding the sentinels' states into a system function.

Different from existing works, we model the task of sentinel network mining as a group sparse learning problem and propose an effective and flexible Bayesian learning algorithm for various dynamical systems, including the most widely used linear continuous system and logistic discrete system. The expectation-maximization (EM) and variational

*Corresponding author. ybo@jlu.edu.cn

Copyright © 2018, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

approximation methods are employed to infer the posterior distribution of the sentinel network. Particularly, we solved the scalability problem by means of proposing more efficient multiplication and inverse operations on diagonal-block metrics. Validations and comparisons were performed on both synthetic and real-world data, which show that the proposed method outperforms existing methods.

We summarize the main contributions of this paper as follows: (1) We propose a novel measure, the γ value, to identify sentinel components by modeling the sentinel network with row sparsity structure; (2) We design an effective and flexible group sparse Bayesian learning algorithm to discover the sentinel network; (3) We solve the scalability and generality problem of this algorithm to a certain degree; (4) We develop a comprehensive framework for active surveillance and validate it by performing extensive experiments on both synthetic and real-world data.

2 Related works

In the literature, most of the optimal sensor placement studies identify the sentinels' location via the framework of the budgeted maximum coverage problem (BMCP) (Khuller, Moss, and Naor 1999), which aims to maximize a specific objective by finding a set of components with the minimum budget (the budget is often defined as the set size). Examples include the timely detection of contaminated water (Krause et al. 2008) and early detection of the outbreak of Weblogs (Leskovec et al. 2007). As BMCP is NP-hard, heuristic algorithms like sub-modular maximization are often employed. Note that designing the objective functions of BMCP relies heavily on the known interaction structure among components. When the underlying interaction structure cannot be observed directly, such as the social media network in charge of information diffusion and the social contact network for disease spread, the current methods proposed within the BMCP framework cannot be used (Gomez Rodriguez, Leskovec, and Krause 2010).

In spatial statistics, Gaussian processes (GPs) can effectively represent the spatial correlation and uncertainty of the sensed field. Based on GPs, one can adopt general information criteria, typically such as mutual information (GPs-MI), to greedily select sentinels (Krause, Singh, and Guestrin 2008; Hoang et al. 2014). In this framework, both the GPs and the information criteria are model-free. They neglect the mechanism of data generation. As a result, the prior knowledge of the phenomena [e.g., the susceptible-infectious-recovered (SIR) model for infectious disease (Dimitrov and Meyers 2010)] is difficult to integrate into the method. If such available prior knowledge can be adequately incorporated, the performance of learning and prediction will be significantly improved, as shown in our experiments.

In addition to the epidemic dynamics data, some works turn to other types of available data to help identify the sentinels. For instance, socio-economic data is leveraged to estimate the malaria infection risk caused by imported cases, and further to implement an effective active surveillance plan (Yang et al. 2014). Traffic data and environmental data are used to infer real-time air quality, and further determine

the best deployment locations of new air contaminant monitoring stations (Hsieh, Lin, and Zheng 2015). It is difficult to reuse these customized methods on other types of active surveillance applications, if the domain data or domain knowledge they require are not available.

3 Active surveillance framework

For epidemic dynamics, we now propose the framework of active surveillance. It consists of three main steps.

Step 1: collect epidemic dynamics data in N components of interest.

Step 2: mine the sentinel network from the data. In the network, the number of sentinel components (i.e., sentinel nodes) k is according to a budget.

Step 3: with the sentinel network, predict future epidemic dynamics of the N components based on the data collected from the k sentinel components.

The last two steps constitute the foundation of the framework, and we will elaborate them in following sections.

3.1 Problem formulation

Consider a diffusion among N components in a dynamical system. Let matrix $\mathbf{D} \in \mathbb{R}^{T \times N} = [\mathbf{D}_1, \dots, \mathbf{D}_T]^T$ be the epidemic dynamics during a time window $[1, T]$. Specifically, $\mathbf{D}_t = [\mathbf{D}_{t,1}, \dots, \mathbf{D}_{t,N}]$, where each entry $\mathbf{D}_{t,i}$ denotes the state of component i at time t , and it may be a real number (e.g., the number of newly infected cases in the city i) or a Boolean value (e.g., whether a news is posted in the blog i). Let $\mathbf{D}^s \in \mathbb{R}^{T \times N}$ denote the surveillance data collected by k sentinel components. Specifically, $\mathbf{D}_{t,i}^s$ is equal to $\mathbf{D}_{t,i}$ when component i is a sentinel, and empty otherwise.

Let $f(\mathbf{D}_t^s; \mathbf{S})$ be the dynamical system function achieving the dynamics prediction. Let matrix $\mathbf{S} \in \mathbb{R}^{N \times N}$ denote a sentinel network, a set of key parameters in the system function. It depicts the influential relationship from the sentinels to all components, where each link $\mathbf{S}_{i,j}$ encodes the effect of sentinel i on component j by its weight. Thus, \mathbf{S} is a row sparse matrix only containing links emitted from the k sentinels. Now, the active surveillance can be formulated as to predict the future components' states based on the surveillance data \mathbf{D}^s and the sentinel network \mathbf{S} :

$$\mathbf{D}_{t+1} \approx \hat{\mathbf{D}}_{t+1} = f(\mathbf{D}_t^s; \mathbf{S}). \quad (1)$$

From Eq. 1, two computational issues need to be addressed for the goal of active surveillance:

I) Sentinel identification: How to identify the sentinels from all components and mine the sentinel network \mathbf{S} according to a given budget from the dynamics \mathbf{D} ?

II) Sentinel prediction: How to predict the future dynamics \mathbf{D}_{t+1} from the current surveillance data \mathbf{D}_t^s based on the discovered \mathbf{S} ?

3.2 Sentinel identification

Our basic idea is intuitive: *In a dynamical system, the components having little influence on others are unimportant for predicting others' states, while those exerting a heavy influence on others dominate the system dynamics and should be selected as sentinels.* In terms of the sentinel network \mathbf{S} , one

can determine whether a component is important or not by inferring row sparsity. That is, unimportant components are associated with sparse rows in \mathbf{S} , in which zeros are much more than non-zeros; on the other hand, important ones are associated with non-sparse rows. Figure 1 shows an illustration by taking a linear dynamical system as an example.

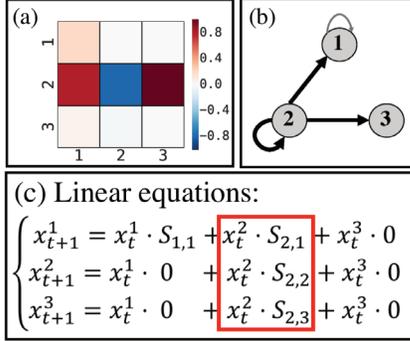


Figure 1: Unimportant = row sparse. (a) Sentinel network \mathbf{S} ; (b) the graph of \mathbf{S} ; (c) the equations of a linear dynamical system, where the component 2 dominates the system. Unimportant component 3 is associated with a sparse row.

Based on this idea, we propose a novel index, the γ value, to measure components' importance in predicting the epidemic dynamics: a component is important if it is important in both prior and posterior structures of a sentinel network. Specifically, the γ value is defined as the data-dependent hyper-parameter of the prior of the sentinel network, and also that reflecting the profiles of the posterior of the sentinel network:

$$\gamma_i = ((\boldsymbol{\mu}_s^i)^T \boldsymbol{\mu}_s^i + \text{Tr}[\boldsymbol{\Sigma}_s^i])N^{-1}, \quad (2)$$

where γ_i is the γ value of component i . In the following, we elaborate this importance measure from the perspectives of prior and posterior.

3.2.1 Prior perspective From the basic idea, a sentinel network is desired to have a row sparse structure. Thus, we adopt a zero-mean multivariate Gaussian prior for each row:

$$p(\mathbf{S}_{i,\cdot} | \gamma_i) \sim \mathcal{N}(\mathbf{0}, \gamma_i \mathbf{I}_N), \quad i = 1, \dots, N \quad (3)$$

where vector $\mathbf{S}_{i,\cdot} \in \mathbb{R}^N$ denotes the i th row of the sentinel network, and $\mathbf{I}_N \in \mathbb{R}^{N \times N}$ is an identity matrix. By doing so, γ_i controls the diversity of row i from a zero vector.

For conciseness, we vectorize matrix \mathbf{S} , i.e., let $\mathbf{s} = \text{vec}(\mathbf{S}^T)$, where operator $\text{vec}(\cdot)$ denotes the vectorization of the input matrix by stacking its columns into a column vector. By doing so, vector $\mathbf{s} \in \mathbb{R}^{N \times N}$ consists of N groups of length N , where each group is associated with a row in \mathbf{S} . Now, the row sparse structure of \mathbf{S} is equal to the group sparse structure of \mathbf{s} . In terms of the prior on \mathbf{S} (Eq. 3), the prior over \mathbf{s} is

$$p(\mathbf{s} | \boldsymbol{\gamma}) \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_0), \quad (4)$$

where vector $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_N)^T$ and the covariance matrix $\boldsymbol{\Sigma}_0 \in \mathbb{R}^{N^2 \times N^2}$ is a diagonal matrix:

$$\boldsymbol{\Sigma}_0 = \begin{bmatrix} \gamma_1 \mathbf{I}_N & & \\ & \ddots & \\ & & \gamma_N \mathbf{I}_N \end{bmatrix}. \quad (5)$$

As mentioned before, the links sent from a component reflect its effect on the system dynamics. Now, the links sent from i in \mathbf{S} (i.e., the entries of group i in \mathbf{s}) are tied together and controlled by a common data-dependent hyper-parameter γ_i . This hyper-parameter is a type of automatic relevance determination (ARD) mechanism (MacKay 1995): when γ_i is small, the group i in \mathbf{s} is sparse and vice versa; when the group i is sparse, the links sent from i are very weak in \mathbf{S} . That is, component i is unimportant and can be pruned out without losing much in prediction accuracy.

3.2.2 Posterior perspective The γ value also reflects the profile of the posterior of the sentinel network. We model the sentinel network for two kinds of dynamical systems widely used to characterize diffusion phenomena in the real world: a linear continuous system and a logistical discrete system.

Likelihood of linear system. Starting from the linear continuous system, we first give the likelihood function and illustrate the pre-processing of dynamics data. The system function of a linear continuous system is

$$\mathbf{Y} = \mathbf{X}\mathbf{S} + \mathbf{V}, \quad (6)$$

where $\mathbf{Y} = \mathbf{D}_{2:T+1} \in \mathbb{R}^{T \times N}$ and $\mathbf{X} = \mathbf{D}_{1:T} \in \mathbb{R}^{T \times N}$ are both extracted from the dynamics data \mathbf{D} . \mathbf{Y} is the epidemic dynamics later than \mathbf{X} one time-unit. Specifically, $\mathbf{Y}_{t-1} = \mathbf{X}_t = \mathbf{D}_t$ and \mathbf{V} is a Gaussian noise matrix.

For convenience, we further transform Eq. 6 into a vector form, $\mathbf{y} = \boldsymbol{\Phi}\mathbf{s} + \mathbf{v}$, where vector $\mathbf{y} = \text{vec}(\mathbf{Y}^T) \in \mathbb{R}^{TN \times 1}$, $\mathbf{s} = \text{vec}(\mathbf{S}^T) \in \mathbb{R}^{N^2 \times 1}$, and $\mathbf{v} = \text{vec}(\mathbf{V}^T) \in \mathbb{R}^{TN \times 1}$. The matrix $\boldsymbol{\Phi} = \text{Kron}(\mathbf{X}, \mathbf{I}_N) \in \mathbb{R}^{TN \times N^2}$, where the operator $\text{Kron}(\cdot, \cdot)$ represents the Kronecker product of two input matrices. Now, based on the Gaussian noise assumption, the likelihood of the linear continuous system can be given as

$$p(\mathbf{y} | \boldsymbol{\Phi}, \mathbf{s}, \lambda) \sim \mathcal{N}(\boldsymbol{\Phi}\mathbf{s}, \lambda \mathbf{I}), \quad (7)$$

where λ denotes the noise level and \mathbf{I} is an identity matrix.

Likelihood of logistical system. For the logistical discrete system, the entry in the dynamics data $\mathbf{D}_{t,i}$ is represented by a Boolean value, 0 or 1, indicating whether component i is "infected" at time t . After the same data pre-processing, we adopt a Bernoulli distribution over each entry of \mathbf{y} , i.e., \mathbf{y}_n :

$$p(\mathbf{y}_n = 1 | \boldsymbol{\Phi}_n, \mathbf{s}) = \sigma[\boldsymbol{\Phi}_n \mathbf{s}], \quad n = 1, \dots, TN, \quad (8)$$

where $\sigma[\boldsymbol{\Phi}_n \mathbf{s}] = 1/(1 + e^{-\boldsymbol{\Phi}_n \mathbf{s}})$ denotes the sigmoid function, and $\boldsymbol{\Phi}_n$ is the n th row of $\boldsymbol{\Phi}$. Then, the likelihood of the dynamics can be written as

$$p(\mathbf{y} | \boldsymbol{\Phi}, \mathbf{s}) = \prod_{n=1}^{TN} \sigma[\boldsymbol{\Phi}_n \mathbf{s}]^{y_n} (1 - \sigma[\boldsymbol{\Phi}_n \mathbf{s}])^{1-y_n}. \quad (9)$$

Now, based on the aforementioned prior and likelihood we have the following conclusion about the posterior:

Theorem 1. For both the linear system and logistical system, the posteriors of the sentinel network are a Gaussian or an approximate Gaussian: $p(\mathbf{s} | \mathbf{y}, \boldsymbol{\Phi}, \Theta) \sim \mathcal{N}(\boldsymbol{\mu}_s, \boldsymbol{\Sigma}_s)$:

(1) Linear system: hyper-parameters set $\Theta = \{\boldsymbol{\gamma}, \lambda\}$,

$$\boldsymbol{\mu}_s = \lambda^{-1} \boldsymbol{\Sigma}_s \boldsymbol{\Phi}^T \mathbf{y}, \quad \boldsymbol{\Sigma}_s^{-1} = \boldsymbol{\Sigma}_0^{-1} + \lambda^{-1} \boldsymbol{\Phi}^T \boldsymbol{\Phi};$$

(2) Logistical system: hyper-parameters set $\Theta = \{\boldsymbol{\gamma}, \boldsymbol{\xi}\}$,

$$\boldsymbol{\mu}_s = 2^{-1} \boldsymbol{\Sigma}_s \boldsymbol{\Phi}^T (2\mathbf{y} - \mathbf{1}), \quad \boldsymbol{\Sigma}_s^{-1} = \boldsymbol{\Sigma}_0^{-1} + \boldsymbol{\Phi}^T \boldsymbol{\pi}(\boldsymbol{\xi}) \boldsymbol{\Phi},$$

where $\xi = [\xi_1, \dots, \xi_{TN}]^T$ denotes variational parameters, and $\pi(\xi) \in \mathbb{R}^{TN \times TN}$ is a diagonal matrix. Specifically, $\pi(\xi)_{n,n} = -\frac{1}{2\xi_n}(\sigma[\xi_n] - \frac{1}{2})$.

Proof. See the Supporting Information¹ \square

Estimation of hyper-parameters. Theorem 1 gives the posteriors of the two kinds of systems (i.e., the Gaussian mean μ_s and covariance matrix Σ_s). In the following, we study how to iteratively update the hyper-parameters in the posteriors, because they cannot be obtained in a closed form. By treating the s as hidden variables, we can employ the EM method to estimate the hyper-parameters set Θ . When the integral in the Q function can be analytically solved (in the linear system), the EM obtains a well-formed solution to estimating the hyper-parameters; otherwise, we apply variational EM to estimate the hyper-parameters by optimizing an approximate low bound of the posterior (in the logistical system). Furthermore, the EM method can guarantee the convergence of the estimation.

Here, we only give the learning rules for the hyper-parameters since the limited space (derivation details can be found in the Supporting Information). In the linear continuous system, the learning rule for the noise parameter λ is

$$\lambda \leftarrow (TN)^{-1} (\|y - \Phi \mu_s\|_2^2 + \text{Tr}[\Sigma_s \Phi^T \Phi]). \quad (10)$$

In the logistical discrete system, we update the variational parameter vector ξ by

$$\xi_n \leftarrow \sqrt{\Phi_n (\Sigma_s + \mu_s \mu_s^T) \Phi_n^T}, \quad n = 1, \dots, TN. \quad (11)$$

It is extremely interesting that in both the linear and logistical systems the learning rule for the γ value is the same,

$$\gamma_i \leftarrow ((\mu_s^i)^T \mu_s^i + \text{Tr}[\Sigma_s^i]) N^{-1}, \quad i = 1, \dots, N, \quad (12)$$

where the vector $\mu_s^i \in \mathbb{R}^N$ and matrix $\Sigma_s^i \in \mathbb{R}^{N \times N}$ denote the i th group of μ_s and Σ_s , respectively.

Intuitively, there are two terms that contribute to the γ value according to its learning rule, Eq. 12. The first term is the inner product term $(\mu_s^i)^T \mu_s^i$, which denotes the sum of the squares of the mean weights of the overall links sent from node i in the sentinel network. In other words, it characterizes the *influence strength* of component i . The second term is the trace term $\text{Tr}[\Sigma_s^i]$, which denotes the variance of the posterior estimation on links sent from node i ; that is to say, it features the *influence uncertainty* of component i . In summary, a larger γ value corresponds to a node that has many links with large and diverse influences on other nodes.

On the whole, the γ value is an index by which the importance of a component for predicting the epidemic dynamics of an entire system can be measured. This index integrates the profiles of both the prior and posterior of a sentinel network. For a trivial component in the system, its γ value will tend to be zero due to the ARD mechanism during Bayesian learning. For an important component, whose γ value larger than zero, its γ value could indicate its monitoring priority.

¹The supporting information of this paper, *Supporting Information: Group Sparse Bayesian Learning for Active Surveillance on Epidemic Dynamics*, have posted on arXiv.org.

A component with a larger γ value exerts a great influence on other components, and its state is important to monitor for making a prediction.

Based on the γ value, we propose a backward-selection algorithm called the SNMA, as shown in Algorithm 14. It starts with all N components of interest and removes one component at a time until only k components are left (k is according the budgets). The component that is removed should be chosen as the one with the minimum γ value. The form of backward-selection algorithm is theoretically guaranteed to pick a optimal subset of components if the system perturbation is small enough (Couvreur and Bresler 2000). A trade-off between accuracy and budget is practically necessary: the more sentinels are selected, the more predictive accuracy is expected, while more cost is needed.

Algorithm 1: SNMA

Input: epidemic dynamics \mathbf{D} , quantity of components of interest N , quantity of sentinels k ;
Output: posterior structure of sentinel network, i.e., mean vector μ_s , covariance matrix Σ_s ;

- 1 Pre-processing: extract \mathbf{y} and Φ from \mathbf{D} ;
- 2 Randomly initialize γ , λ (the linear) or ξ (the logistical);
- 3 $L \leftarrow N$;
- 4 **while** $L > k$ **do**
- 5 **while** γ is not converged **do**
- 6 // The optimization step
- 7 update μ_s and Σ_s via Theorem 1;
- 8 update γ , λ (the linear) or γ , ξ (the logistical) via Eq. 12,10 and 11;
- 9 **end**
- 10 find the minimum entry i in the vector γ
- 11 // The selection step
- 12 update Φ , μ_s , Σ_s , λ , γ , ξ through pruning out the entries of i th group in them;
- 13 $L \leftarrow L - 1$;
- 14 **end**

3.3 Sentinel prediction

Once we have obtained the posterior structure of the sentinel network, the epidemic dynamics of the overall system, \mathbf{D} , can be predicted based on the surveillance data \mathbf{D}^s . Let \mathbf{D}_*^s be a new set of surveillance data, where only the values on k sentinels' locations are kept and the rest are empty. As mentioned above, we obtain Φ_*^s through the data pre-processing. Then, a predictive distribution over the following system states \mathbf{y}_* is given by

$$p(\mathbf{y}_* | \Phi_*^s, \mathbf{y}, \Phi) = \int p(\mathbf{y}_* | \Phi_*^s, s) p(s | \mathbf{y}, \Phi) ds. \quad (13)$$

Linear continuous system. In this case, the integral in Eq. 13 is a Gaussian convolution (refer to the proof of Theorem 1), whose analytical solution is a Gaussian. Then, we have

$$p(\mathbf{y}_* | \Phi_*^s, \mathbf{y}, \Phi) \sim \mathcal{N}(\mu_{y_*}, \sigma_{y_*}^2)$$

with parameters $\mu_{y_*} = \Phi_*^s \mu_s$, $\sigma_{y_*}^2 = \lambda + \Phi_*^s \Sigma_s (\Phi_*^s)^T$.

Logistical discrete system. The predictive distribution of discrete data is a Bernoulli distribution. By substituting the

term of the posterior of s in Eq. 13 for the variational approximation posterior given in Theorem 1, we have the following predictive distribution:

$$p(\mathbf{y}_{*n} = 1 | \Phi_*^s, \mathbf{y}, \Phi) \approx \int p(\mathbf{y}_{*n} = 1 | \Phi_*^s, s) q(s | \mathbf{y}, \Phi) ds,$$

where $n = 1 \dots N$.

However, the integral in this approximation is a convolution of a logistical function with a Gaussian distribution, which cannot be analytically integrated. By introducing an error function, it can be approximated as a re-parameterized logistical function (Maragakis et al. 2008):

$$\int p(\mathbf{y}_{*n}=1 | \Phi_*^s, s) q(s | \mathbf{y}, \Phi) ds \approx (1 + e^{-\tau \Phi_*^s \mu_s})^{-1}, \quad (14)$$

where $\tau = (1 + \frac{\pi}{8} \Phi_*^s \Sigma_s (\Phi_*^s)^T)^{-\frac{1}{2}}$. A very small approximation error is guaranteed in theory (Maragakis et al. 2008).

3.4 Scaling up

In the SNMA, the most expensive operations are the matrix multiplication ($\Phi^T \Phi$) and the matrix inverse (Σ_s^{-1}) for calculating the posterior parameters in Theorem 1. Considering the large size of $\Phi \in \mathbb{R}^{TN \times N^2}$ and $\Sigma_s \in \mathbb{R}^{N^2 \times N^2}$, the time complexity of one iteration will be $O(N^5 \times \max(T, N))$, making it infeasible to handle real-world problems. We now propose two fast matrix operations to solve the scalability problem. The time complexity after scaling up is reduced to $O(N^2 \times \max(T, N))$, which is faster than the competitors in the experiments.

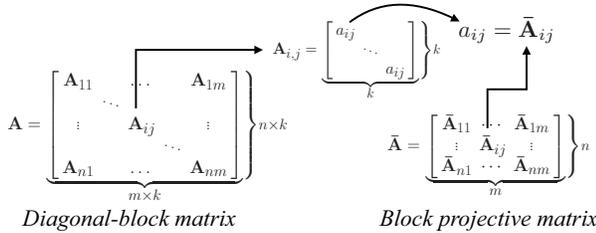


Figure 2: Illustration of the relationship between a diagonal-block matrix and block projective matrix.

Definition 1. (*Diagonal-block matrix*). Matrix $\mathbf{A} \in \mathbb{R}^{nk \times mk}$ is called a diagonal-block matrix if it consists of n by m square sub-matrices, and each sub-matrix $\mathbf{A}_{i,j} \in \mathbb{R}^{k \times k}$ is a diagonal matrix with all diagonal entries having the same value a_{ij} , $i=1 : n, j=1 : m$.

Definition 2. (*Block projective matrix*). Let $\mathbf{A} \in \mathbb{R}^{nk \times mk}$ be a diagonal-block matrix. Matrix $\bar{\mathbf{A}} \in \mathbb{R}^{n \times m}$ is called the block projective matrix of \mathbf{A} if each entry $\bar{\mathbf{A}}_{i,j} = a_{ij}$, $i=1 : n, j=1 : m$.

The structure of the two matrices and their relationship is shown in Fig. 2. The block projective matrix $\bar{\mathbf{A}}$ is much smaller than its original diagonal-block matrix \mathbf{A} . In the meantime, it contains the overall distinct elements in \mathbf{A} . Based on the definitions, we have the following theorems, which can significantly reduce the computational cost of the two kinds of operations.

Theorem 2. Let \mathbf{A}, \mathbf{B} be two diagonal-block matrices with the same size of sub-matrix. The product $\mathbf{AB} = \mathbf{C}$ is a diagonal-block matrix, and those block projective matrices satisfy $\bar{\mathbf{A}}\bar{\mathbf{B}} = \bar{\mathbf{C}}$.

Proof. See the Supporting Information. \square

Theorem 3. Let \mathbf{A} be a square diagonal-block matrix. Its inverse matrix \mathbf{A}^{-1} is also a diagonal-block matrix and satisfies $(\mathbf{A}^{-1}) = (\bar{\mathbf{A}})^{-1}$.

Proof. See the Supporting Information. \square

In the SNMA, Φ and Σ_s are two diagonal-block matrices. When we alternatively calculate the multiplication ($\Phi^T \Phi$) and the matrix inverse (Σ_s^{-1}) on the block projective matrices $\bar{\Phi} \in \mathbb{R}^{T \times N}$ and $\bar{\Sigma}_s \in \mathbb{R}^{N \times N}$ via the two theorems, the time complexity of one iteration can be reduced by 3 orders of magnitude (i.e., N^3). Note that this scaling-up technique not only solves the difficulty faced by our algorithm, but can also address other tasks involving diagonal-block matrices, such as the computational obstacle of a multiple measurement vector (MMV) model in compressed sensing (Zhang and Rao 2011). To process further large-scale data in practice, the SNMA can be readily parallelized by adopting the group testing strategy, which has been used for parallel feature selection (Zhou et al. 2014).

3.5 Embedding non-linear dynamical models

The proposed framework is flexible and can be readily extended to various dynamical systems by mean of embedding basic functions, as long as the dynamical system functions can be represented as the combinations of basic functions. Each basic function, $\psi_i(x)$, $i \in [1, \dots, m]$, can be an arbitrary function, e.g., polynomial, $\psi_i(x) = x^2 + x$; trigonometric, $\psi_i(x) = \sin(x)$; or others.

We illustrate the embedding technology based on the linear system Eq. 6 as an example, where $\mathbf{X}_{t,i}$ denotes the state of component i at time t . Let vector $\Psi(\mathbf{X}_{t,i}) = [\psi_1(\mathbf{X}_{t,i}), \dots, \psi_m(\mathbf{X}_{t,i})] \in \mathbb{R}^m$ be the mapping of $\mathbf{X}_{t,i}$ through m basic functions. Let vector $\Psi(\mathbf{X}_t) = [\Psi(\mathbf{X}_{t,1}), \dots, \Psi(\mathbf{X}_{t,N})] \in \mathbb{R}^{mN}$ denote the mapping of all components at time t , and matrix $\Psi(\mathbf{X}) = [\Psi(\mathbf{X}_1); \dots; \Psi(\mathbf{X}_T)]^T \in \mathbb{R}^{T \times mN}$ denote the mapping of all components during the time window. Then, the system function Eq.6 can be extended as $\mathbf{Y} = \Psi(\mathbf{X})\hat{\mathbf{S}} + \mathbf{V}$, where the augmented sentinel network $\hat{\mathbf{S}} \in \mathbb{R}^{mN \times N}$ characterizes the interactions from the mN mapping to the N components. This extended equation can also represent non-linear dynamical systems if non-linear basic functions are adopted, such as the SIR model, a well-studied and widely adopted disease spread model. Moreover, the basic functions of the SIR model can be constructed based on its reproduction matrix form (Wallinga, van Boven, and Lipsitch 2010).

To integrate the above extended system function into the current SNMA, just let $\Phi = \text{Kron}(\Psi(\mathbf{X}), \mathbf{I}_N) \in \mathbb{R}^{TN \times mN^2}$ and $\mathbf{s} = \text{vec}(\hat{\mathbf{S}})^T$. Meanwhile, the group size needs be changed from N to mN . The rest of the steps are the same as the SNMA shows in Algorithm 14. For the logistical system, the extended process is analogous.

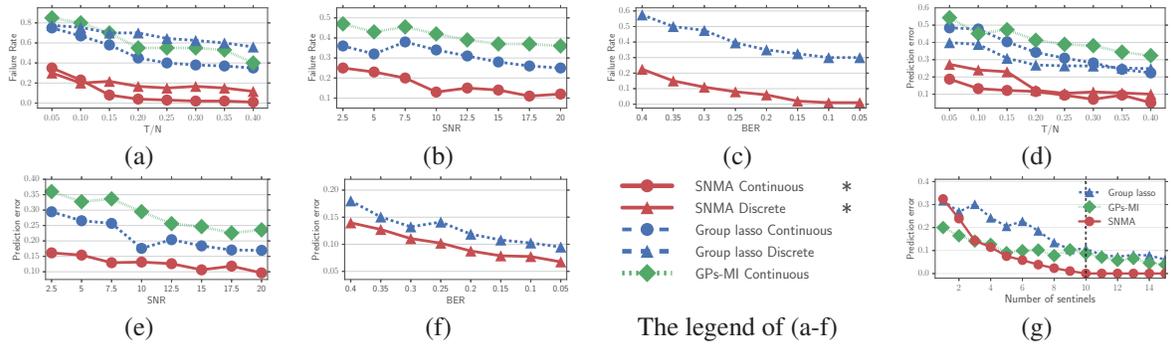


Figure 3: Results of experiments on synthetic data. (* denotes our methods).

4 Validations

Comparative Study. We validate the framework on both synthetic and three real-world data. The two most related methods are selected as competitors: group Lasso and GPs-MI. Group Lasso (Meier, Van De Geer, and Bühlmann 2008) is a typical method for group sparse learning, which is similar to that addressed by our proposed group sparse Bayesian learning algorithm. To achieve the aim of active surveillance, we use the proposed framework and replace the group sparse Bayesian learning with group lasso.

GPs-MI is a popular sensor placement method (Krause, Singh, and Guestrin 2008; Hoang et al. 2014), which is similar to the task of active surveillance addressed by our work. It outperforms the placement methods based on experiment design, such as A-, D-, and E-optimal designs. As Gaussian processes cannot directly work on discrete data, GPs-MI is only applied on experiments with continuous data.

Evaluation. Two criteria are adopted to evaluate the performance of the methods. Failure rate is used to measure whether the a method can discover the sentinels, i.e., the problem of sentinel identification. Failure rate is the percentage of wrong sentinels' locations given by a method, $1 - |\Gamma \cap \hat{\Gamma}|/|\hat{\Gamma}|$, where Γ is the set of the ground truth sentinel locations in, and $\hat{\Gamma}$ is the set discovered by, a method.

We use *root-mean-square error* (RMSE) to quantify the prediction error, i.e., the problem of sentinel prediction. Specifically, $RMSE = \frac{\|\mathbf{Y} - \hat{\mathbf{Y}}\|_2}{T/N}$, where $\hat{\mathbf{Y}}$ is the predicted epidemic dynamics based on the surveillance data.

4.1 Validations on synthetic data

Synthetic Data Generation. We generate synthetic data by imagining diffusion processes taking place in a linear continuous system or logistical discrete system. There are three steps in total: (1) Random generation of a ground truth γ -value vector with 500 entries, where 100 entries are sampled from $\mathcal{N}(0, 10)$ and the other 400 from $\mathcal{N}(0, 0.1)$, i.e., 100 sentinels and 400 trivial components; (2) Random sampling of a ground truth sentinel network \mathbf{S} via the prior Eq.3 based on the ground truth γ value; (3) Based on the \mathbf{S} , simulation of the epidemic dynamics via the linear system Eq. 6 or the logistical system Eq. 8 embedding a quadratic basic function $\phi(x) = x^2 + x$.

Environment setting. The comparisons are conducted under various data volumes and noise levels. We use the value T/N to denote the ratio of data volume to the number of parameters to estimate. *Signal-to-noise ratio* (SRN) and *bit error rate* (BER) are adopted to denote noise levels in the linear system and logistical system, respectively.

We adopt a 5-fold cross-validation strategy in experiments on synthetic data. We firstly identify 100 sentinels from the training data via the three methods. For each method, the average failure rate of sentinel identification is shown in Figs. 3 (a-c). Then, we evaluate the performance of sentinel prediction by feeding the surveillance data (collected on the 100 discovered sentinels) to the corresponding prediction model of each method, such as Eq. 13 for the proposed method. The average prediction error is shown in Figs.3(d-f). The results show that the proposed methods are superior to GPs-MI and group lasso on both failure rate and prediction error.

We evaluate the trade-off between the prediction accuracy and surveillance cost in the linear system by setting different number of sentinels as shown in Fig.3(g). In this experiment, there are only 10 ground truth sentinels in \mathbf{S} . We give the number of sentinels k (x -axis), and then evaluate the prediction error of each method (y axis). Intuitively, the more sentinels that are selected, the better the accuracy that can be obtained. However, as indicated, the improvement in prediction becomes negligible when k is over 10. Note that GPs-MI shows a better performance only when k is very small.

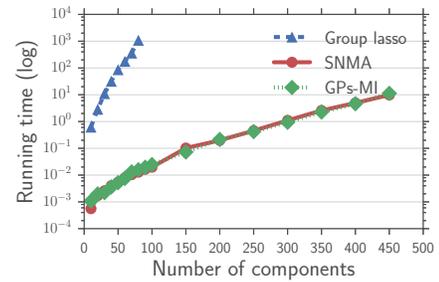


Figure 4: Comparison on running time. y -axis is the time (seconds) on a log scale. x -axis denotes the size of a system.

Fig. 4 presents the average running time of one iteration of the three methods on a PC with a 3.4GHz CPU and 8GB memory. Since GPs-MI employ forward greedy strategy (fast when only a few sentinels) but SNMA use backward selection method (fast when many sentinels), it's fair to evaluate them by comparing the running time of one iteration of each algorithm. The results show group lasso is the slowest, and GPs-MI and SNMA are almost at the same level.

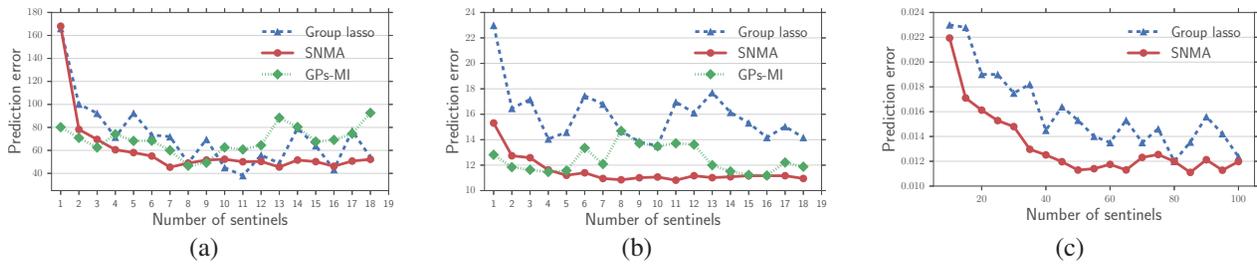


Figure 5: Comparisons on the real-world epidemic dynamics. x -axis is the number of selected sentinels and y -axis denotes the prediction error (RMSE). (a) 2009 Hong Kong H1N1 flu pandemic. (b) 2005-2009 Tengchong malaria outbreak. (c) The dynamics of hot words cascading in Baidu Tieba.

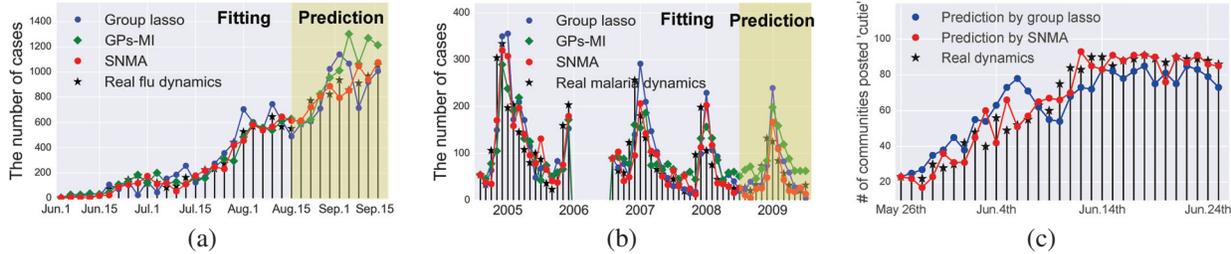


Figure 6: Model fittings and sentinel predictions for epidemic dynamics. (a) 2009 Hong Kong H1N1 flu pandemic. (b) 2005-2009 Tengchong malaria. (c) The dynamics of hot word "cutie" in the Baidu Tieba.

4.2 Validations on real diffusion data

In real cases, the failure rate cannot be evaluated because the ground truth sentinel network is unknown.

4.2.1 2009 Hong Kong H1N1 flu pandemic The cases report data of 2009 Hong Kong H1N1 influenza epidemic, was provided by Centre for Health Protection (CHP), Department of Health, Government of the Hong Kong Special Administrative Region. During this pandemic, the first imported case of human swine influenza (HSI) was confirmed on May 1, 2009. As of Sep. 2010, there were over 36,000 confirmed cases of HSI, among which about 290 were severe cases and over 80 of them died (CHP 2010). We consider the epidemic dynamics for 105 days since the disease onset in Jun. 1st 2009 based on the confirmed cases of H1N1 infection reported by CHP (www.chp.gov.hk), which gives the spatial position and infection times of each case.

Hong Kong consists of 18 administrative districts, i.e., 18 components. By merging the confirmed cases during 105 days, 3 days as a basic infectious period, we obtain the dynamics of $N=18$ and $T=35$. We set the dynamics from Jun. 1 to Aug. 15 as training data and the one from Aug. 15 to Sep. 15 as test data. According to a pre-given k , sentinel districts of Hong Kong can be identified via the methods (SNMA use the linear system setting). Then, based on the sentinels, we predict the newly cases on test data as shown in Fig 5 (a). Obviously, SNMA outperforms the competitors. Fig. 6 (a) shows a case of fittings and predictions for the dynamics (8 districts are selected as sentinels). The sentinel network of Hong Kong and the sentinels' spatial distribution can be found in Supporting Information.

4.2.2 2005-2009 Tengchong malaria outbreak Tengchong City, Yunnan Province, China, has 18 towns, and 658,207 residents that are distributed in a wide area of 5,845 km² in 2011. Because of the suitable climate for mosquito habitats, Tengchong has a quite serious malaria outbreak. Five years' (2005-2009)

monthly malaria cases data at the town level, were collected by Tengchong CDC and can be obtained from the annual reports of National Institute of Parasitic Disease, China CDC. By eliminating the missing data from 2005 Jun. to Dec., we get malaria dynamics, which contains $N = 18$ components and $T = 53$ months.

We set the malaria dynamics from 2005 to 2008 as training data and the one during 2009 as test data. Similar to the Hong Kong case, we identify the sentinel towns according to a pre-given k and predict the dynamics on the test data, as shown in Fig 5 (b). Once again, SNMA achieves the best sentinel prediction in most cases. Fig. 6 (b) shows a case of fittings and predictions for the malaria dynamics with 7 sentinels. The sentinel network of Tengchong and the sentinels' distribution is in Supporting Information.

4.2.3 Hot words diffusion in Baidu Tieba Baidu Tieba (tieba.baidu.com), one of the largest online community platforms in China, is a collection of thousands of active topic-specific communities. Tieba users can post any hot words (vocabularies that are widely used in Tieba during a short period) in any communities. Hot words often present a diffusion phenomena in Baidu Tieba: a hot word first appears in only a few forums, and then is posted gradually in many other forums by the users who are active in multiple forums. Thus, Tieba can be regarded as a logistical discrete dynamical system, where communities are components, hot words are contagions, and the infected state of a component is 0 or 1. GPs-MI cannot be applied to this case because it's base on Gaussian model and cannot directly work on discrete data.

We tracked the dynamics of 11 independent hot words cascading among the top-100 active communities in Baidu Tieba from Apr. 2014 to Oct. 2015 (18 months). Only the dynamics during the words' bursting period is preserved, and the total bursting period of the words is 738 d, i.e., $N = 100$, $T=738$ by using 1 d as a time-unit. Here, we split the training and test data in term of the hot words, i.e., we alternatively set the dynamics of one hot word as test data and the rest be the training data. Then, we identify the sen-

tinell communities in Baidu Tieba by the methods (GPs-MI cannot be applied in discrete data; SNMA use the logistic system setting). Fig 5 (c) shows the results of sentinel prediction. Fig. 6 (c) shows that the dynamics of the hot word “cutie” is successfully predicted based on the data from 66 sentinel communities.

Summarily, SNMA outperform GPs-MI and group lasso in the most experiments. Our method has two obvious advantages. (1) Our method is model-based and can readily integrate prior knowledge, which makes it more effective and easier to train. (2) Our method is more robust against noise and insufficient data owing to the Bayesian framework that can effectively handle the uncertainty from both data and model.

5 Conclusions

In this work, we addressed the challenge of epidemic dynamics prediction in cases in which surveillance resources are very limited. We proposed a novel importance measure, the γ value, by modeling a sentinel network with row sparse structure and presented an effective and flexible group sparse Bayesian learning algorithm for mining the sentinel network in two kinds of widely used dynamical systems. With the discovered sentinel network, the overall epidemic dynamics can be predicted based on partial data only collected by the few sentinels. Moreover, we significantly reduced the computational complexity of the algorithm and extended it to various nonlinear systems using basic function embedding technology. We validated the proposed framework by a set of experiments on both synthetic and real-world datasets.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under grants 61373053, 61572226, Jilin Province Natural Science Foundation under grant 20150101052JC.

References

Brunton, S. L.; Proctor, J. L.; and Kutz, J. N. 2016. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences* 113(15):3932–3937.

Chen, Y.; Amiri, H.; Li, Z.; and Chua, T.-S. 2013. Emerging topic detection for organizations from microblogs. In *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 43–52. Dublin, Ireland: ACM.

2010. Summary report on the surveillance of adverse events following HSI immunisation and expert group’s comment on the safety of hsi vaccine in Hong Kong. Technical report, Centre for Health Protection, Hong Kong.

Couvreur, C., and Bresler, Y. 2000. On the optimality of the backward greedy algorithm for the subset selection problem. *SIAM Journal on Matrix Analysis and Applications* 21(3):797–808.

Dimitrov, N. B., and Meyers, L. A. 2010. Mathematical approaches to infectious disease prediction and control. Technical report.

Gerardo-Giorda, L.; Puggioni, G.; Rudd, R. J.; Waller, L. A.; and Real, L. A. 2013. Structuring targeted surveillance for monitoring disease emergence by mapping observational data onto ecological process. *Journal of The Royal Society Interface* 10(86):20130418.

Gomez Rodriguez, M.; Leskovec, J.; and Krause, A. 2010. Inferring networks of diffusion and influence. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1019–1028. Washington, USA: ACM.

Hoang, T. N.; Low, K. H.; Jaillet, P.; and Kankanhalli, M. 2014. Nonmyopic ϵ -bayes-optimal active learning of gaussian processes.

In *Proceedings of the 24th International Conference on Machine Learning*, 739–747. Beijing, China: IMLS.

Hsieh, H.-P.; Lin, S.-D.; and Zheng, Y. 2015. Inferring air quality for station location recommendation based on urban big data. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 437–446.

Khuller, S.; Moss, A.; and Naor, J. S. 1999. The budgeted maximum coverage problem. *Information Processing Letters* 70(1):39–45.

Krause, A.; Leskovec, J.; Guestrin, C.; VanBriesen, J.; and Faloutsos, C. 2008. Efficient sensor placement optimization for securing large water distribution networks. *Journal of Water Resources Planning and Management* 134(6):516–526.

Krause, A.; Singh, A.; and Guestrin, C. 2008. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research* 9(Feb):235–284.

Leskovec, J.; Krause, A.; Guestrin, C.; Faloutsos, C.; VanBriesen, J.; and Glance, N. 2007. Cost-effective outbreak detection in networks. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 420–429.

MacKay, D. J. 1995. Probable networks and plausible predictions - a review of practical bayesian methods for supervised neural networks. *Network: Computation in Neural Systems* 6(3):469–505.

Maragakis, P.; Ritort, F.; Bustamante, C.; Karplus, M.; and Crooks, G. E. 2008. Bayesian estimates of free energies from nonequilibrium work data in the presence of instrument noise. *The Journal of Chemical Physics* 129(2):024102.

Meier, L.; Van De Geer, S.; and Bühlmann, P. 2008. The group lasso for logistic regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 70(1):53–71.

Polgreen, P. M.; Chen, Z.; Segre, A. M.; Harris, M. L.; Pentella, M. A.; and Rushton, G. 2009. Optimizing influenza sentinel surveillance at the state level. *American journal of epidemiology* 170(10):1300–1306.

Wallinga, J.; van Boven, M.; and Lipsitch, M. 2010. Optimizing infectious disease interventions during an emerging epidemic. *Proceedings of the National Academy of Sciences* 107(2):923–928.

Yang, B.; Guo, H.; Yang, Y.; Shi, B.; Zhou, X.; and Liu, J. 2014. Modeling and mining spatiotemporal patterns of infection risk from heterogeneous data for active surveillance planning. In *Proceedings of the the 28th AAAI Conference on Artificial Intelligence*, 493–499. Quebec, Canada: AAAI.

Yang, B.; Pei, H.; Chen, H.; Liu, J.; and Shang, X. 2017. Characterizing and discovering spatiotemporal social contact patterns for healthcare. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39(8):1532–1546.

Zhang, Z., and Rao, B. D. 2011. Sparse signal recovery with temporally correlated source vectors using sparse bayesian learning. *IEEE Journal of Selected Topics in Signal Processing* 5(5):912–926.

Zheng, Y.; Liu, F.; and Hsieh, H.-P. 2013. U-air: when urban air quality inference meets big data. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1436–1444. Chicago, USA: ACM.

Zhou, Y.; Porwal, U.; Zhang, C.; Ngo, H. Q.; Nguyen, X.; Ré, C.; and Govindaraju, V. 2014. Parallel feature selection inspired by group testing. In *Advances in Neural Information Processing Systems*, 3554–3562.