

Influential Node Detection in Implicit Social Networks using Multi-task Gaussian Copula Models

Qunwei Li

Syracuse University

QLI33@SYR.EDU

Bhavya Kailkhura

Lawrence Livermore National Labs

KAILKHURA1@LLNL.GOV *

Jayaraman J. Thiagarajan

Lawrence Livermore National Labs

JJAYARAM@LLNL.GOV

Zhenliang Zhang

Intel Labs

ZHENLIANG.ZHANG@INTEL.COM

Pramod K. Varshney

Syracuse University

VARSHNEY@SYR.EDU

Editor: Oren Anava, Marco Cuturi, Azadeh Khaleghi, Vitaly Kuznetsov, Alexander Rakhlin

Abstract

Influential node detection is a central research topic in social network analysis. Many existing methods rely on the assumption that the network structure is completely known *a priori*. However, in many applications, network structure is unavailable to explain the underlying information diffusion phenomenon. To address the challenge of information diffusion analysis with incomplete knowledge of network structure, we develop a multi-task low rank linear influence model. By exploiting the relationships between contagions, our approach can simultaneously predict the volume (i.e. time series prediction) for each contagion (or topic) and automatically identify the most influential nodes for each contagion. The proposed model is validated using synthetic data and an ISIS twitter dataset. In addition to improving the volume prediction performance significantly, we show that the proposed approach can reliably infer the most influential users for specific contagions.

1. Introduction

Information emerges dynamically and diffuses quickly via agent interactions in complex networks (e.g. social networks) (López-Pintado, 2008). Consequently, understanding and prediction of information diffusion mechanisms are challenging. There is a rapidly growing interest in exploiting knowledge of the information dynamics to better characterize the factors influencing spread of diseases, planned terrorist attacks, and effective social marketing campaigns, etc (Guille and Hacid, 2012). The broad applicability of this problem in social network analysis has led to focused research on the following questions: (I) Which contagions are the most popular and can diffuse the most? (II) Which members of the network

*. This work was supported in part by ARO under Grant W911NF-14-1-0339. This work was performed under the auspices of the U.S. Dept. of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

are influential and play important roles in the diffusion process? (III) What is the range over which the contagions can diffuse (Guille et al., 2013)? While attempting to answer these questions, one is confronted with two crucial challenges. First, a descriptive diffusion model, which can mimic the behavior observed in real world data, is required. Second, efficient learning algorithms are required for inferring influence structure based on the assumed diffusion model.

A variety of information diffusion prediction frameworks have been developed in the literature (Yang and Leskovec, 2010; Wang et al., 2013; Guille et al., 2013; Du et al., 2013; Zhang et al., 2016). A typical assumption in many of these approaches is that a connected network graph and knowledge of the corresponding structure are available *a priori*. However, in practice, the structure of the network can be implicit or difficult to model, e.g., modeling the structure of the spread of infectious disease is almost impossible. As a result, network structure unaware diffusion prediction models have gained interest. For example, (Yang and Leskovec, 2010), Yang *et. al.* proposed a linear influence model, which can effectively predict the information volume by assuming that each of the contagions spreads with the same influence in an implicit network. Subsequently, in (Wang et al., 2013), the authors extended LIM by exploiting the sparse structure in the influence function to identify the influential nodes. Though the relationships between multiple contagions can be used for more accurate modeling, most of the existing approaches ignore that information.

In this paper, we address the above issues by augmenting linear influence models with complex task dependency information. More specifically, we consider the dependency of different contagions in the network, and characterize their relationships using Copula Theory. Furthermore, by imposing a low-rank regularizer, we are able to characterize the clustering structure of the contagions and the nodes in the network. Through this novel formulation, we attempt to both improve the accuracy of the prediction system and better regularize the influence structure learning problem. Finally, we develop an efficient algorithm based on proximal mappings to solve this optimization problem. Experiments with synthetic data reveal that the proposed approach fairs significantly better than a state-of-the-art multi-task variant of LIM both in terms of volume prediction and influence structure estimation performance. In addition, we demonstrate the superiority of the proposed method in predicting the time-varying volume of tweets using the ISIS twitter dataset¹.

2. Background

In this section, we present the formulation of linear influence model (LIM) (Yang and Leskovec, 2010) and discuss its limitations. Consider a set of N nodes that participate in an information diffusion process of K different contagions over time. Node $u \in \{1, \dots, N\}$ can be infected by contagion $k \in \{1, \dots, K\}$ at time $t \in \{0, 1, \dots, T\}$. The volume $V_k(t)$ is defined as the total number of nodes that get infected by the contagion k at time t . Let the indicator function $M_{u,k}(t) = 1$ represent the event that node u got infected by contagion k at time t , and 0 otherwise. LIM models the volume $V_k(t)$ as a sum of influences of nodes u

1. ISIS dataset from Kaggle is available at <https://www.kaggle.com/kzaman/how-isis-uses-twitter>.

that got infected before time t :

$$V_k(t+1) = \sum_{u=1}^N \sum_{l=0}^{L-1} M_{u,k}(t-l) I_u(l+1), \quad (1)$$

where each node u has a particular non-negative influence function $I_u(l)$. One can simply think of $I_u(l)$ as the number of follow-up infections l time units after u got infected. The value of L is set to indicate that the influence of a node drops to 0 after L time units. Thus, the influence of node u is denoted by the vector $\mathbf{I}_u = (I_u(1), \dots, I_u(L))^T \in \mathbb{R}^{L \times 1}$. Next, using the notation $\mathbf{V}_k = (V(1), \dots, V(T))^T \in \mathbb{R}^{T \times 1}$ and $\mathbf{I} = (\mathbf{I}_1^T, \dots, \mathbf{I}_N^T)^T \in \mathbb{R}^{LN \times 1}$, the inference procedure of LIM can be formulated as follows

$$\text{minimize} \quad \sum_{k=1}^K \|\mathbf{V}_k - \mathbf{M}_k \cdot \mathbf{I}\|_2^2 + \mathbb{1}(\mathbf{I}), \quad (2)$$

where \mathbf{M}_k is obtained via concatenation of $M_{u,k}$, $\|\cdot\|_2$ denotes the Euclidean norm, and $\mathbb{1}(\mathbf{I})$ is an indicator function that is zero when $I_{uk}(l) \geq 0$ and $+\infty$ otherwise. Though LIM has been effective in predicting the future volume for each contagion, it assumes that each node has the same influence across all the contagions. Consequently, to achieve contagion-sensitive node selection in an implicit network, the LIM model was extended and the multitask sparse linear influential model (MSLIM) was proposed in (Wang et al., 2013).

The influence function is defined by extending \mathbf{I}_u in LIM into contagion-sensitive $\mathbf{I}_{u,k} \in \mathbb{R}^{L \times 1}$, which is a L -length vector representing the influence of the node u for the contagion k . For each contagion k , let $\mathbf{I}^k \in \mathbb{R}^{LN \times 1}$ be the vector obtained by concatenating $\mathbf{I}_{1k}, \dots, \mathbf{I}_{Nk}$. For each node u , the influence matrix for the node u is defined: $\mathbf{I}_u = (\mathbf{I}_{u1}, \dots, \mathbf{I}_{uK}) \in \mathbb{R}^{L \times K}$. Using these notations, the inference procedure to estimate $\mathbf{I}_{u,k}$ was formulated as follows

$$\text{minimize} \quad \frac{1}{2} \sum_{k=1}^K \|\mathbf{V}_k - \mathbf{M}_k \cdot \mathbf{I}^k\|_2^2 + \lambda \sum_{u=1}^N \|\mathbf{I}_u\|_F + \gamma \sum_{u=1}^N \sum_{k=1}^K \|\mathbf{I}_{uk}\|_2 + \mathbb{1}(\mathbf{I}), \quad (3)$$

where $\|\cdot\|_F$ denotes the Frobenius norm. The penalty term $\|\mathbf{I}_u\|_F$ was used to encourage the entire matrix \mathbf{I}_u to be zero altogether, which means that the node u is non-influential for all different contagions. If the estimated $\|\mathbf{I}_u\|_F > 0$ (i.e., the matrix \mathbf{I}_u is non-zero), a fine-grained selection is performed by the penalty $\sum_{u=1}^N \sum_{k=1}^K \|\mathbf{I}_{uk}\|_2$, which is essentially a group-Lasso penalty and can encourage the sparsity of vectors $\{\mathbf{I}_{uk}\}$. For a specific contagion k , one can identify the most influential nodes by finding the optimal solution $\{\hat{\mathbf{I}}_{uk}\}$ of (3). However, the penalty terms used in MSLIM encourages that certain nodes have no influence over all the contagions which may not be true in practice. Furthermore, for most of the real world applications, there exists complex dependencies among the contagions. In order to alleviate these shortcomings, we propose a novel probabilistic multi-task learning framework and develop efficient optimization strategies.

3. Proposed Approach

Probabilistic Multi-Contagion Modeling of Diffusion: We assume a linear regression model for each task: $\mathbf{V}_k = \mathbf{M}_k \mathbf{I}^k + \mathbf{n}_k$, where $\mathbf{V}_k, \mathbf{M}_k$ and \mathbf{I}^k are defined as before, and

$\mathbf{n}_k \in \mathbb{R}^{T \times 1}$ is an i.i.d. zero-mean Gaussian noise vector with the covariance matrix $\mathbf{\Sigma}_k$. The distribution for \mathbf{V}_k given \mathbf{M}_k , \mathbf{I}^k and $\mathbf{\Sigma}_k$ can be expressed as

$$\mathbf{V}_k | \mathbf{M}_k, \mathbf{I}^k, \mathbf{\Sigma}_k \sim \mathcal{N}(\mathbf{M}_k \mathbf{I}^k, \mathbf{\Sigma}_k) = \frac{\exp\left(-\frac{1}{2}(\mathbf{V}_k - \mathbf{M}_k \mathbf{I}^k)^T \mathbf{\Sigma}_k^{-1} (\mathbf{V}_k - \mathbf{M}_k \mathbf{I}^k)\right)}{(2\pi)^{\frac{T}{2}} |\mathbf{\Sigma}_k|^{\frac{1}{2}}}. \quad (4)$$

Assuming that the influence for a single contagion is also Gaussian distributed, we can express the marginal distributions as $\mathbf{I}^k | \mathbf{m}_k, \mathbf{\Theta}_k \sim \mathcal{N}(\mathbf{m}_k, \mathbf{\Theta}_k)$, where $\mathbf{m}_k \in \mathbb{R}^{LN \times 1}$ is the mean vector and can be expressed as $\mathbf{m}_k = [\mathbf{m}_{1,k}^T, \dots, \mathbf{m}_{N,k}^T]^T$, and $\mathbf{\Theta}_k \in \mathbb{R}^{LN \times LN}$ is the covariance matrix of \mathbf{I}^k . For a node u and contagion k , we assume that the variables in the influence \mathbf{I}_{uk} have the same mean, i.e., $\mathbf{m}_{u,k} = m_{u,k} \mathbf{1}_{L \times 1}$, where $m_{u,k}$ is a scalar and $\mathbf{1}_{L \times 1}$ is a vector of all ones with dimension $L \times 1$. Let $\mathbf{m}' \in \mathbb{R}^{N \times K}$ represent the mean matrix with entries $m_{u,k}$, and it is connected as $\mathbf{m} = (\mathbf{m}_1, \dots, \mathbf{m}_K) = \mathbf{Q} \mathbf{m}'$, where $\mathbf{Q} \in \mathbb{R}^{LN \times N} = \mathcal{I}_{N \times N} \otimes \mathbf{1}_{L \times 1}$ and $\mathcal{I}_{N \times N}$ is the identity matrix with dimension $N \times N$ and \otimes is the Kronecker product operator.

3.1 Dependence Structure Modeling Using Copulas

Consider a general case where the contagions are correlated. We construct a new influence matrix $\mathbf{I} = [\mathbf{I}^1, \dots, \mathbf{I}^K] \in \mathbb{R}^{LN \times K}$. In our formulation, \mathbf{I}^k 's are assumed to be correlated and the joint distribution of \mathbf{I} is not a simple product of all the marginal distributions of \mathbf{I}^k as is adopted by most multi-task learning formulations. Here, we propose to use a multi-task copula that is obtained by tailoring the copula model for the multi-task learning problem.

Theorem 1 (*Sklar's Theorem*). *Consider an N -dimensional distribution function F with marginal distribution functions F_1, \dots, F_N . Then there exists a copula C , such that for all x_1, \dots, x_N in $[-\infty, \infty]$, $F(x_1, \dots, x_N) = C(F_1(x_1), \dots, F_N(x_N))$. If F_n is continuous for $1 \leq n \leq N$, then C is unique, otherwise it is determined uniquely on $\text{Ran} F_1 \times \dots \times \text{Ran} F_N$ where $\text{Ran} F_n$ is the range of F_n . Conversely, given a copula C and univariate CDFs F_1, \dots, F_N , F is a valid multivariate CDF with marginals F_1, \dots, F_N .*

As a direct consequence of Sklar's Theorem, for continuous distributions, the joint probability density function (PDF) $f(x_1, \dots, x_N)$ is obtained by,

$$f(x_1, \dots, x_N) = \left(\prod_{n=1}^N f_n(x_n) \right) c(F_1(X_1), \dots, F_N(X_N)), \quad (5)$$

where $f_n(\cdot)$ is the marginal PDF and c is termed as the copula density given by

$$c(v) = \frac{\partial^N C(v_1, \dots, v_N)}{\partial v_1, \dots, \partial v_N} \quad (6)$$

where $v_n = F_n(x_n)$. We extend the copula theory to multi-task learning and express the joint distribution of \mathbf{I} as follows:

$$p(\mathbf{I}^1, \mathbf{I}^2, \dots, \mathbf{I}^K) = \left(\prod_{k=1}^K \mathcal{N}(\mathbf{m}_k, \mathbf{\Theta}_k) \right) c(F_1(\mathbf{I}^1), F_2(\mathbf{I}^2), \dots, F_K(\mathbf{I}^K)), \quad (7)$$

where $F_k(\mathbf{I}^k)$ is the CDF of the influence for k^{th} contagion. The copula density function $c(\cdot)$ takes all marginal CDFs $\{F_k(\mathbf{I}^k)\}_{k=1}^K$ as its arguments, and maintains the output correlations in a parametric form.

Gaussian copula: There are a finite number of well defined copula families that can characterize several dependence structures. Though, we can investigate the choice of an appropriate copula, we consider the Gaussian copula for its favorable analytical properties. A Gaussian copula can be constructed from the multivariate Gaussian CDF, and the resulting prior on \mathbf{I} is given by a multivariate Gaussian distribution as

$$\mathbf{I} \sim \mathcal{MN}_{LN \times K}(\mathbf{m}, \mathbf{U}, \mathbf{\Omega}) = \frac{\exp\left(-\frac{1}{2}\text{tr}\left(\mathbf{U}^{-1}(\mathbf{I} - \mathbf{m})\mathbf{\Omega}^{-1}(\mathbf{I} - \mathbf{m})^T\right)\right)}{(2\pi)^{\frac{LNK}{2}} |\mathbf{\Omega}|^{\frac{LN}{2}} |\mathbf{U}|^{\frac{K}{2}}} \quad (8)$$

where $\mathbf{U} \in \mathbb{R}^{LN \times LN}$ is the row covariance matrix modeling the correlation between the influence of different nodes, $\mathbf{\Omega} \in \mathbb{R}^{K \times K}$ is the column covariance matrix modeling the correlation between the influence for different contagions, and $\mathbf{m} \in \mathbb{R}^{LN \times K}$ is the mean matrix of \mathbf{I} . The two covariances can be computed as $E[(\mathbf{I} - \mathbf{m})(\mathbf{I} - \mathbf{m})^T] = \mathbf{U}\text{tr}(\mathbf{\Omega})$ and $E[(\mathbf{I} - \mathbf{m})^T(\mathbf{I} - \mathbf{m})] = \mathbf{\Omega}\text{tr}(\mathbf{U})$ respectively. We assume that N individual nodes are spreading the contagions and influencing others independently, and thus the row covariance matrix is diagonal and can be expressed as $\mathbf{U} = \text{diag}(e_1^2, e_2^2, \dots, e_N^2) \otimes \mathcal{I}_{L \times L}$, where $e_n^2, n \in \{1, \dots, N\}$ are scalars. The posterior distribution for \mathbf{I} , which is proportional to the product of the prior in Eq. 4 and the likelihood function in Eq. 8, is given as

$$\begin{aligned} p(\mathbf{I}|\mathbf{M}, \mathbf{V}, \mathbf{\Sigma}, \mathbf{U}, \mathbf{\Omega}) &\propto p(\mathbf{V}|\mathbf{M}, \mathbf{I}, \mathbf{\Sigma}) p(\mathbf{I}|\mathbf{m}, \mathbf{U}, \mathbf{\Omega}) \\ &= \left(\prod_{k=1}^K \mathcal{N}(\mathbf{M}_k \mathbf{I}^k, \mathbf{\Sigma}_k) \right) \mathcal{MN}_{LN \times K}(\mathbf{I}|\mathbf{m}, \mathbf{U}, \mathbf{\Omega}), \end{aligned} \quad (9)$$

where $\mathbf{M} = (\mathbf{M}_1, \dots, \mathbf{M}_K) \in \mathbb{R}^{T \times LNK}$, $\mathbf{V} = (\mathbf{V}_1, \dots, \mathbf{V}_K) \in \mathbb{R}^{T \times K}$, $\mathbf{\Sigma}$ is the corresponding covariance matrix of $\mathbf{n} = (\mathbf{n}_1, \dots, \mathbf{n}_K) \in \mathbb{R}^{T \times K}$. We assume $\mathbf{\Sigma}_k \triangleq \sigma^2 \mathcal{I}_{T \times T}$ and also an identical value of $e_n^2 = e^2, \forall k = 1, \dots, K, \forall n = 1, \dots, N$. We employ maximum a posteriori (MAP) and maximum likelihood estimation (MLE), and obtain \mathbf{I} , \mathbf{m} , and $\mathbf{\Omega}$ by

$$\min_{\mathbf{I}, \mathbf{m}, \mathbf{\Omega}} \frac{1}{\sigma^2} \sum_{k=1}^K \|\mathbf{V}_k - \mathbf{M}_k \mathbf{I}^k\|_2^2 + \frac{1}{e^2} \text{tr}((\mathbf{I} - \mathbf{m})\mathbf{\Omega}^{-1}(\mathbf{I} - \mathbf{m})^T) + LN \ln |\mathbf{\Omega}| + \mathbb{1}(\mathbf{I}).$$

However, if we assume $\mathbf{\Omega}^{-1}$ to be non-sparse, the solution to $\mathbf{\Omega}^{-1}$ will not be defined (when $K > LN$) or will overfit (when K is of the same order as LN) (Rai et al., 2012). In fact, some contagions in the network can be uncorrelated, which makes the corresponding entry values in $\mathbf{\Omega}^{-1}$ zero. Hence, we add a l_1 penalty to promote sparsity of matrix $\mathbf{\Omega}^{-1}$ to obtain

$$\min_{\mathbf{I}, \mathbf{m}, \mathbf{\Omega}} \sum_{k=1}^K \|\mathbf{V}_k - \mathbf{M}_k \mathbf{I}^k\|_2^2 + \lambda_1 \text{tr}((\mathbf{I} - \mathbf{m})\mathbf{\Omega}^{-1}(\mathbf{I} - \mathbf{m})^T) - \lambda_2 \ln |\mathbf{\Omega}^{-1}| + \lambda_3 \|\mathbf{\Omega}\|_1 + \mathbb{1}(\mathbf{I}).$$

3.2 Modeling Structure of Influence Matrix \mathbf{I}

In order to better characterize the influence matrix, we propose to impose a low rank structure on the influence matrix \mathbf{I} . The nodes or the contagions in the influence network are known to form communities (or clustering structures), which may be captured using the low-rank property of the influence matrix. Note that, the sparse structure in the influence matrix implies that most individuals only influence a small fraction of contagions in the network while there can be a few nodes with wide-spread influence. We incorporate this into our formulation by using a sparsity promoting regularizer over $\mathbf{I}_{u,k}$.

$$\begin{aligned} \min_{\mathbf{I}, \mathbf{m}, \mathbf{\Omega}} \quad & \sum_{k=1}^K \|\mathbf{V}_k - \mathbf{M}_k \mathbf{I}^k\|_2^2 + \lambda_1 \text{tr}((\mathbf{I} - \mathbf{m})\mathbf{\Omega}^{-1}(\mathbf{I} - \mathbf{m})^T) \\ & - \lambda_2 \ln |\mathbf{\Omega}^{-1}| + \lambda_3 \|\mathbf{\Omega}\|_1 + \lambda_4 \|\mathbf{I}\|_* + \lambda_5 \sum_{u=1}^N \sum_{k=1}^K \|\mathbf{I}_{uk}\|_2 + \mathbb{1}(\mathbf{I}), \end{aligned} \quad (10)$$

where $\|\cdot\|_*$ denotes the nuclear norm, and $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ and λ_5 are the regularization parameters. With the estimated $\{\hat{\mathbf{I}}_{uk}\}$, one can predict the total volume of the contagion k at $T+1$ by $\hat{V}_k(T+1) = \sum_{u=1}^N \sum_{l=0}^{L-1} M_{uk}(T-l)I_{uk}(l+1)$.

4. Algorithm

We adopt an alternating optimization approach to solve the problem in Eq. 10.

Optimization w.r.t. \mathbf{m} : Given \mathbf{I} and $\mathbf{\Omega}^{-1}$, the mean matrix \mathbf{m} can be obtained by solving the following problem

$$\min_{\mathbf{m}} \quad \text{tr}((\mathbf{I} - \mathbf{m})\mathbf{\Omega}^{-1}(\mathbf{I} - \mathbf{m})^T).$$

The estimate $\hat{\mathbf{m}}$ can be analytically obtained as $\hat{\mathbf{m}} = \frac{1}{L} \mathbf{Q} \mathbf{Q}^T \mathbf{I}$.

Optimization w.r.t. $\mathbf{\Omega}$: Given \mathbf{I} and \mathbf{m} , the contagion inverse covariance matrix $\mathbf{\Omega}^{-1}$ can be estimated by solving the following optimization problem

$$\min_{\mathbf{\Omega}} \quad \lambda_1 \text{tr}((\mathbf{I} - \mathbf{m})\mathbf{\Omega}^{-1}(\mathbf{I} - \mathbf{m})^T) - \lambda_2 \ln |\mathbf{\Omega}^{-1}| + \lambda_3 \|\mathbf{\Omega}\|_1$$

The above is an instance of the standard inverse covariance estimation problem with sample covariance $\frac{\lambda_1}{\lambda_2}(\mathbf{I} - \mathbf{m})^T(\mathbf{I} - \mathbf{m})$, which can be solved using standard tools. In particular, we use the graphical Lasso procedure in (Friedman et al., 2008)

$$\hat{\mathbf{\Omega}}^{-1} = gLasso(\lambda_1/\lambda_2(\mathbf{I} - \mathbf{m})^T(\mathbf{I} - \mathbf{m}), \lambda_3). \quad (11)$$

Optimization w.r.t. \mathbf{I} : The corresponding optimization problem becomes

$$\min_{\mathbf{I}} \quad \sum_{k=1}^K \|\mathbf{V}_k - \mathbf{M}_k \mathbf{I}^k\|_2^2 + \lambda_1 \text{tr}((\mathbf{I} - \mathbf{m})\mathbf{\Omega}^{-1}(\mathbf{I} - \mathbf{m})^T) + \lambda_4 \|\mathbf{I}\|_* + \lambda_5 \sum_{u=1}^N \sum_{k=1}^K \|\mathbf{I}_{uk}\|_2 + \mathbb{1}(\mathbf{I}).$$

We rewrite the problem as

$$\min_{\mathbf{I}} \quad \ell(\mathbf{I}) + \lambda_4 \|\mathbf{I}\|_* + \mathbb{1}(\mathbf{I}). \quad (12)$$

Algorithm 1 Incremental Proximal Descent

```
1: Initialize  $\mathbf{I} = \mathbf{A}$ 
2: repeat
3:   Set  $\mathbf{I} = \mathbf{I} - \theta \nabla_{\mathbf{I}} \ell(\mathbf{I})$ 
4:   Set  $\mathbf{I} = \text{prox}_{\theta \lambda_4 \|\cdot\|_*}(\mathbf{I})$ 
5:   Set  $\mathbf{I} = P_{\mathbb{I}}(\mathbf{I})$ 
6: until convergence
7: return  $\mathbf{I}$ 
```

where $\ell(\mathbf{I}) = \sum_{k=1}^K \|\mathbf{V}_k - \mathbf{M}_k \mathbf{I}^k\|_2^2 + \lambda_1 \text{tr}((\mathbf{I} - \mathbf{m})\mathbf{\Omega}^{-1}(\mathbf{I} - \mathbf{m})^T) + \lambda_5 \sum_{u=1}^N \sum_{k=1}^K \|\mathbf{I}_{uk}\|_2$. This formulation involves a sum of a convex differentiable loss and convex non-differentiable regularizers which renders the problem non-trivial. A string of algorithms have been developed for the case where the optimal solution is easy to compute when each regularizer is considered in isolation. This corresponds to the case where the proximal operator defined for a convex regularizer $R : \mathbb{R}^{LN \times K} \rightarrow \mathbb{R}$ at a point \mathbf{Z} by $\text{prox}_R(\mathbf{Z}) = \arg \min_{\frac{1}{2}\|\mathbf{I} - \mathbf{Z}\|_F^2 + R(\mathbf{I})}$, is easy to compute for each regularizer taken separately. See (Combettes and Pesquet, 2011) for a broad overview of proximal methods. The proximal operator for the nuclear norm is given by the shrinkage operation as follows (Beck and Teboulle, 2009). If $U \text{diag}(\sigma_1, \dots, \sigma_n) V^T$ is the singular value decomposition of \mathbf{Z} , then $\text{prox}_{\lambda_4 \|\cdot\|_*}(\mathbf{Z}) = U \text{diag}((\sigma_i - \lambda_4)_+) V^T$. The proximal operator of the indicator function $\mathbb{I}(\mathbf{I})$ is simply the projection onto $I_{u,k}(l) \geq 0$, which is denoted by $P_{\mathbb{I}}(\mathbf{I})$. Next, we mention a matching serial algorithm introduced in (Bertsekas, 2011). We present here a version where updates are performed according to a cyclic order (Richard et al., 2012). Note that one can also randomly select the order of the updates. We use the optimization algorithm 1 to solve the optimization problem in Eq. 12.

5. Experiments

We compare the performance of the proposed approach to MSLIM by applying it to both synthetic and real datasets. Since the volume of a contagion over time $V_k(t)$ can be viewed as a time series, we set up this problem as a time series prediction task and evaluate the performance using the prediction mean-squared error (MSE). Furthermore, for the synthetic data set, where we have access to the true influence matrix \mathbf{I} , we also evaluate the performance of the influence matrix prediction task using the metric $\|\hat{\mathbf{I}} - \mathbf{I}\|_F$. We determined the regularization parameters for the proposed model using cross validation. In particular, we split the first 60% of the time instances as the training set and the rest for validation. Following (Wang et al., 2013), we combine the training and validation sets to re-train the model with the best selected regularization parameters and estimate the influence matrix.

5.1 Synthetic Data

We created a synthetic dataset with the number of nodes fixed at $N = 100$ and the number of contagions at $K = 20$. In addition, we assumed that $L = 10$ and $T = 20$. A rank 5 (low-rank) influence matrix \mathbf{I} was generated randomly with uniformly distributed entries. The matrix \mathbf{M} was generated with uniformly distributed random integers $\{0, 1\}$. Following our

Table 1: Prediction performance for different information diffusion models on synthetic data.

Approach	MSLIM	Proposed
Volume Prediction MSE	0.834	0.007
Influence Matrix Estimation Error	0.7681	0.62

model assumption, the volume for each \mathbf{V}_k was calculated as follows $\mathbf{V}_k = \mathbf{M}_k \times \mathbf{I}^k + \mathcal{N}(\mathbf{0}, \Delta)$ where $\mathcal{N}(\mathbf{0}, \Delta)$ is a multivariate normal distribution with covariance matrix Δ . In Table 1, we present the results obtained using the proposed approach and its comparison to MSLIM. As can be observed, for both volume prediction and influence matrix estimation tasks, the proposed approach achieves highly accurate estimates.

5.2 ISIS Twitter Data

In this section, we demonstrate the application of the proposed approach to a real-word analysis task. We begin by describing the twitter dataset used for analysis and the procedure adopted to extract the set of contagions. Following this, we discuss the problem setup and present comparisons to MSLIM on predicting the time-varying tweet volume. Finally, we present a qualitative analysis of the inferred influence structure for different contagions.

The ISIS dataset from Kaggle² is comprised of over 17,000 tweets from 112 users posted between January 2015 and May 2016. In addition to the actual tweets, meta-information such as the user name and the timestamp for each tweet are included. We performed a standard pre-processing by removing a variety of stop words, e.g. URLs, symbols. After preprocessing, we converted each tweet into a bag-of-words representation and extracted the term frequency-inverse document frequency (tf-idf) feature.

Topic Modeling: When applying our approach, the first step is to define semantically meaningful contagions. A simple way of defining topics is to directly use words as topics (e.g., ISIS). However, a single word may not be rich enough to represent a broad topic (e.g., social network sites). Hence, we propose to perform topic modeling on the tweets based on the tf-idf features. In our experiment, we obtained the topics using Non-negative Matrix Factorization (NMF), which is a popular scheme for topic discovery, with the number of topics K set at 10. Table 2 lists the top 10 words for each of the topics learned using NMF.

Volume Time Series Prediction: In our experiment, we set one day as the discrete time step for aggregating the tweet volume. The parameter L denotes the number of time steps it takes for the influence of a user to decay to zero. We set the parameter L equal to 5 since we observed that beyond $L = 5$, there is hardly any improvement in performance. The MSE on the predicted volume is computed over the entire period of observation. The comparison of the prediction MSE is presented in Table 3. It can be seen that the proposed approach significantly outperforms MSLIM in predicting the time-varying volume.

Influential Node Detection: For a contagion k , we identify the most influential nodes with respect to this contagion as nodes having high $\|\mathbf{I}_{\mathbf{u}, \mathbf{k}}\|_2$ values. First, in Figure 2(a), we

2. ISIS dataset from Kaggle is available at <https://www.kaggle.com/kzaman/how-isis-uses-twitter>.

Table 2: Top words for each topic learned using NMF with the ISIS twitter dataset.

Topic 1	isis ramiallolah iraq attack libya warreporter1 saa aamaq usa abu
Topic 2	killed soldiers today airstrikes injured wounded civilians militants iraqi attack
Topic 3	syria russia ramiallolah turkey ypg breakingnews usa group saa terror
Topic 4	state islamic fighters fighting group saudi new http wilaya control
Topic 5	aleppo nid gazau rebels north today northern syrian ypg turkish
Topic 6	assad regime myra forces rebels fsa pro islam syrian jaysh
Topic 7	al qaeda nusra abu sham ahrar islam jabhat http warreporter1
Topic 8	army iraq near ramiallolah iraqi lujah turkey ramadi west sinai
Topic 9	allah people muslims abu accept muslim make know don islam
Topic 10	breaking islamicstate forces amaqagency city fighters iraqi near area syrian




Figure 1: Comparing statistics from the estimated influence matrix with the volume of tweets corresponding to each of the users to identify influential users. In both cases, the users with a large influence score are marked in red.

plot the correlation among 10 topics learned by NMF. More specifically, we plot the pairwise correlation structure learned by our approach. It can be seen that, a strong positive correlation structure exists, which enabled the improved prediction in Table 3. Following this, we use the predicted influence matrix to select a set of highly influential nodes from the dataset. A simple approach to select the influential users can be to select the ones with a large number of tweets. However, we argue that the influence predicted in an information diffusion model can be vastly different. Consequently, we consider a user to be influential if she has a high influence score for at least one of the topics, or if she can be influential for multiple topics. For example, in Figure 1(a), we plot average influence scores of the users (averaged over all the topics) against the total number of tweets. Similarly, in Figure 1(b),




Figure 2: (a) Correlation Structure among the topics (non-black color represents positive correlation), (b) Top 9 influential users and their tweet distributions.

Table 3: Volume prediction performance on the ISIS twitter dataset.

Approach	MSLIM	Proposed
Volume Prediction MSE	2.7	0.329

we plot influence scores of the users (maximum over all the topics) against the total number of tweets. The first striking observation is that the users with high influence scores are not necessarily the ones with the most number of tweets. Instead, their impact on the information diffusion relies heavily on the complex dynamics of the implicit network.

Finally, in Figure 2(b) we plot the percentage of tweets regarding each of the topics for top 9 influential nodes. Influential nodes are obtained as a union of nodes identified based on both average and maximum influence scores. More specifically, we select the union of users with average influence score greater than 1.3 and maximum influence score greater than 1.8. In addition to displaying the distribution across topics, for each influential user, we show the total number of tweets posted by that user. It can be seen that the total number of tweets of these users vary a lot and, therefore, is not a good indication of their influence.

6. Conclusion

In this paper, we considered the problem of influential node detection and volume time series prediction. We proposed a descriptive diffusion model to take dependencies among the topics into account. We also proposed an efficient algorithm based on alternating methods to perform inference and learning on the model. It was shown that the proposed technique outperforms existing influential node detection techniques. Furthermore, the proposed model was validated both on a synthetic and a real (ISIS) dataset. We showed that the proposed approach can efficiently select the most influential users for specific contagions. We also presented several interesting patterns of the selected influential users for the ISIS dataset.

References

- Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009.
- Dimitri P Bertsekas. Incremental gradient, subgradient, and proximal methods for convex optimization: A survey. *Optimization for Machine Learning*, 2010:1–38, 2011.
- Patrick L Combettes and Jean-Christophe Pesquet. Proximal splitting methods in signal processing. In *Fixed-point algorithms for inverse problems in science and engineering*, pages 185–212. Springer, 2011.
- Nan Du, Le Song, Manuel Gomez-Rodriguez, and Hongyuan Zha. Scalable influence estimation in continuous-time diffusion networks. In *Advances in neural information processing systems*, pages 3147–3155, 2013.
- Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.
- Adrien Guille and Hakim Hacid. A predictive model for the temporal dynamics of information diffusion in online social networks. In *Proceedings of the 21st international conference on World Wide Web*, pages 1145–1152. ACM, 2012.
- Adrien Guille, Hakim Hacid, Cecile Favre, and Djamel A Zighed. Information diffusion in online social networks: A survey. *ACM SIGMOD Record*, 42(2):17–28, 2013.
- Dunia López-Pintado. Diffusion in complex social networks. *Games and Economic Behavior*, 62(2):573–590, 2008.
- Piyush Rai, Abhishek Kumar, and Hal Daume. Simultaneously leveraging output and task structures for multiple-output regression. In *Advances in Neural Information Processing Systems (NIPS)*, pages 3185–3193, 2012.
- Emile Richard, Pierre-andre Savalle, and Nicolas Vayatis. Estimation of simultaneously sparse and low rank matrices. In *Proceedings of the 29th International Conference on Machine Learning (ICML)*, pages 1351–1358, 2012.
- Yingze Wang, Guang Xiang, and Shi-Kuo Chang. Sparse multi-task learning for detecting influential nodes in an implicit diffusion network. In *AAAI*, 2013.
- Jaewon Yang and Jure Leskovec. Modeling information diffusion in implicit networks. In *2010 IEEE International Conference on Data Mining*, pages 599–608. IEEE, 2010.
- Peng Zhang, Jing He, Guodong Long, Guangyan Huang, and Chengqi Zhang. Towards anomalous diffusion sources detection in a large network. *ACM Transactions on Internet Technology (TOIT)*, 16(1):2, 2016.