Inferring Social Influence and Meme Interaction with Hawkes Processes

Chuan Luo[1] Xiaolong Zheng[1] Daniel Zeng[1,2]
[1]State Key Laboratory of Management and Control for Complex Systems,
Institute of Automation, Chinese Academy of Sciences, Beijing, China
[2]Department of Management Information Systems, University of Arizona, Tucson, USA
{chuan.luo, xiaolong.zheng}@ia.ac.cn, zeng@email.arizona.edu

***Abstract*—Revealing underlying social influence among users in social media is critical to understanding how users interact, on which a lot of security intelligence applications can be built. Existing methods fail to take into account the interaction relationships among memes. In this paper, we propose to simultaneously model social influence and meme interaction in information diffusion with novel multi-dimensional Hawkes processes. Experimental results on both synthetic and real world social media data show the efficacy of the proposed approach.**

Keywords—social influence; information diffusion; Hawkes process; social media

I. INTRODUCTION

Recent years have witnessed an explosive growth of various social media sites such as online social networks, blogs, microblogs, social news websites and virtual social worlds. This gives researchers a great opportunity to study social interactions on an unprecedented scale. Since individual behaviors have an effect on the decisions of friends in a social network, the knowledge of who-influences-whom has enormous implications in security informatics [1]. Unfortunately, the influence structure among users is usually unknown and unobserved. For instance, it is very common that we only observe the times when particular users adopted a new emerging meme but we do not observe who infected them. In addition to social influence, we observe that interactions among memes may also contribute to the understanding of user behaviors. For example, an event (i.e., Bill posted a meme relevant to Microsoft yesterday) could trigger another event (i.e., Steve posted a meme relevant to Apple today). Thus, how to model both social influence and meme interaction in the information diffusion process is needed to be addressed.

In this paper, we propose a novel model based on multi-dimensional Hawkes processes to tackle this problem. The rest of this paper is organized as follows. Section II briefly surveys the related work on influence inference in social media. The proposed model is given in Section III. Section IV introduces the methods to infer model parameters. Experimental results on both synthetic and real world data will be presented in Section V. Finally, we will give a summary in Section VI.

II. RELATED WORK

The diffusion link inference problem was firstly studied as a classification problem solved by SVM classifiers. Later, researchers presented probabilistic models based on survival and event history analysis and solved it using convex programming [2]. Another study [3] proposed to address the influence inference problem in social media from the perspective of nonlinear dynamic systems. Recent works [4] started to employ the multi-dimensional Hawkes process [5] to model recurrent time stamped events occurred in the process of information diffusion to uncover the underlying influence network among users. The novelty of our paper lies in the simultaneously modeling of social influence and meme interaction in information diffusion with multi-dimensional Hawkes processes.

III. MODELS

We consider a typical scenario in information diffusion, where a set of memes $M = \{\mathrm{m} \mid \mathrm{m} = 1, 2, ..., \mathrm{M}\}$ propagate among a set of users $I = \{\mathrm{i} \mid \mathrm{i} = 1, 2, ..., \mathrm{I}\}$ through a hidden social network. The observations are a sequence of events $\{\mathrm{E_n} \mid \mathrm{n} = 1, 2, ..., \mathrm{N}\}$ ordered by time, where an event is denoted as $\mathrm{E_n} = (\mathrm{t_n}, \mathrm{i_n}, \mathrm{m_n})$, i.e., a user *i* posted a meme *m* at time *t*. We propose a novel multi-dimensional Hawkes process to infer the hidden social network and interactions among memes.

A. One-dimensional Hawkes Processes

The Hawkes process is a class of self or mutually exciting point process models. A one-dimensional Hawkes process is characterized by its conditional intensity function defined as follows:

$$\lambda(t) = \mu + \sum_{t_l < t} \kappa(\mathrm{t} - \mathrm{t}_l)$$

where μ is the base intensity which is affected by the activity of the user, t_l are the time of events occurred before time *t*, and $\kappa(\mathrm{t})$ is the time decay kernel, i.e., $\kappa(t) = \omega e^{-\omega t}$. The self-exciting property is captured by the summing over l with $t_l < t$.

B. Multi-dimensional Hawkes Processes

In order to model social influence between users, the above one-dimensional Hawkes process needs to be extended to the multi-dimensional case. Specifically, we have one Hawkes process for each user and these Hawkes processes are coupled with each other. The conditional intensity function of user *i* is expressed as follows:

$$\lambda_i(t) = \mu_i + \sum_{t_l < t} a_{i_l i}\kappa(\mathrm{t} - \mathrm{t}_l)$$

where μ_i is the base intensity of user i, and $a_{i_l i}$ captures the mutually-exciting property between user i_l and i. In other words, $a_{i_l i}$ specifies the influence of user i_l on i.

C. Proposed model

In order to take into account the effect of interaction among memes, we propose to extend the above multi-dimensional Hawkes process by associating one coefficient for each meme pair and model the diffusion process of each meme as a multi-dimensional Hawkes process.

Formally, for user i and meme m, the conditional intensity function can be described as follows:

$$\lambda_{i,m}(t) = \mu_i \gamma_m + \sum_{t_l < t} a_{i_l i} b_{m_l m}\kappa(\mathrm{t} - \mathrm{t}_l)$$

where γ_m is the base rate of meme m's propagation which is affected by the popularity of the meme. The coefficient $b_{m_l m}$ captures the interaction effect between meme m_l and m. Larger value of $b_{m_l m}$ indicates events involving meme m_l are more likely to trigger a future event involving meme m. We collect the parameters into matrix-vector forms, $\boldsymbol{\mu} = (\mu_i)$ for the user base intensity, $\boldsymbol{\gamma} = (\gamma_m)$ for the meme base intensity, $\boldsymbol{A} = (\mathrm{a}_{i_l i})$ for the social influence matrix, $\boldsymbol{B} = (\mathrm{b}_{m_l m})$ for the meme interaction matrix.

Given the time interval $[0, \mathrm{T}]$ and a sequence of events $\{(\mathrm{t}_n, \mathrm{i}_n, \mathrm{m}_n) \mid \mathrm{n} = 1, 2, ..., \mathrm{N}\}$ ordered by time, which are generated from the above conditional intensity functions, the log-likelihood function can be expressed as follows:

$$L(\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{\mu}, \boldsymbol{\gamma}) = \sum_{n=1}^{N} \ln \lambda_{i_n, m_n}(t_n) - \sum_{i=1}^{I}\sum_{m=1}^{M}\int_0^T \lambda_{i,m}(t)dt$$

An important property of social influence is sparsity [4], meaning most users only have influence on a small fraction of users in the social network, while only a small number of hub users have wide spread influence on many others. This property can be represented in the sparsity pattern of matrix $\boldsymbol{A}$. In addition, it is reasonable to assume that Matrix $\boldsymbol{B}$ may also have a similar pattern.

From the observed data, as a result, we can infer the model parameters by maximizing the regularized log-likelihood defined as follows:

$$\min_{\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{\mu}, \boldsymbol{\gamma} \ge 0} - L(\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{\mu}, \boldsymbol{\gamma}) + \lambda_1 \|\boldsymbol{A}\|_1 + \lambda_2 \|\boldsymbol{B}\|_1$$

where $\|\boldsymbol{A}\|_1, \|\boldsymbol{B}\|_1$ is the $l1$ norm of matrix $\boldsymbol{A}$ and $\boldsymbol{B}$, respectively. The parameter λ_1 and λ_2 control the effect of the sparsity regularization.

IV. INFERENCE

Because the above objective function is non-differentiable and thus difficult to optimize in general, we turn to the idea of alternating direction method of multipliers (ADMM) [6] to convert the original problem to several easier sub-problems. Specifically, the ADMM algorithm involves the following key iterative steps:

$$\boldsymbol{A}^{s+1}, \boldsymbol{B}^{s+1}, \boldsymbol{\mu}^{s+1}, \boldsymbol{\gamma}^{s+1} = \underset{A,B,\mu,\gamma \ge 0}{\arg\min}\, L_\rho(\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{\mu}, \boldsymbol{\gamma}, \boldsymbol{Z}_1^s, \boldsymbol{Z}_2^s, \boldsymbol{U}_1^s, \boldsymbol{U}_2^s)$$

$$\boldsymbol{Z}_1^{s+1} = \underset{\boldsymbol{Z}_1}{\arg\min}\, L_\rho(\boldsymbol{A}^{s+1}, \boldsymbol{B}^{s+1}, \boldsymbol{\mu}^{s+1}, \boldsymbol{\gamma}^{s+1}, \boldsymbol{Z}_1, \boldsymbol{Z}_2^s, \boldsymbol{U}_1^s, \boldsymbol{U}_2^s)$$

$$\boldsymbol{Z}_2^{s+1} = \underset{\boldsymbol{Z}_2}{\arg\min}\, L_\rho(\boldsymbol{A}^{s+1}, \boldsymbol{B}^{s+1}, \boldsymbol{\mu}^{s+1}, \boldsymbol{\gamma}^{s+1}, \boldsymbol{Z}_1^{s+1}, \boldsymbol{Z}_2, \boldsymbol{U}_1^s, \boldsymbol{U}_2^s)$$

$$\boldsymbol{U}_1^{s+1} = \boldsymbol{U}_1^s + (\boldsymbol{A}^{s+1} - \boldsymbol{Z}_1^{s+1})$$

$$\boldsymbol{U}_2^{s+1} = \boldsymbol{U}_2^s + (\boldsymbol{B}^{s+1} - \boldsymbol{Z}_2^{s+1})$$

Matrix $\boldsymbol{Z}_1$ and $\boldsymbol{Z}_2$ are two auxiliary variables, $\boldsymbol{U}_1$ and $\boldsymbol{U}_2$ are the dual variables associated with constraints $\boldsymbol{A} = \boldsymbol{Z}_1$ and $\boldsymbol{B} = \boldsymbol{Z}_2$, ρ is the penalty parameter and L_ρ is defined as follows:

$$\begin{aligned} L_\rho &= -L(\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{\mu}, \boldsymbol{\gamma}) + \lambda_1 \|\boldsymbol{Z}_1\|_1 + \lambda_2 \|\boldsymbol{Z}_2\|_1 \\ &+ \rho \mathrm{tr}(\boldsymbol{U}_1^T(\boldsymbol{A} - \boldsymbol{Z}_1)) + \rho \mathrm{tr}(\boldsymbol{U}_2^T(\boldsymbol{B} - \boldsymbol{Z}_2)) \\ &+ \frac{\rho}{2}(\|\boldsymbol{A} - \boldsymbol{Z}_1\|_F^2 + \|\boldsymbol{B} - \boldsymbol{Z}_2\|_F^2) \end{aligned}$$

A. Solving for $\boldsymbol{Z}_1$ and $\boldsymbol{Z}_2$

The parameter $\boldsymbol{Z}_1$ and $\boldsymbol{Z}_2$ can be solved with $dL_\rho / d\boldsymbol{Z}_1 = 0$ and $dL_\rho / d\boldsymbol{Z}_2 = 0$, respectively. Specifically, we have:

$$(\boldsymbol{Z}_1^{s+1})_{ij} = \begin{cases} (\boldsymbol{A}^{s+1} + \boldsymbol{U}_1^s)_{ij} - \lambda_1/\rho, if\, (\boldsymbol{A}^{s+1} + \boldsymbol{U}_1^s)_{ij} \ge \lambda_1/\rho \\ (\boldsymbol{A}^{s+1} + \boldsymbol{U}_1^s)_{ij} + \lambda_1/\rho, if\, (\boldsymbol{A}^{s+1} + \boldsymbol{U}_1^s)_{ij} \le -\lambda_1/\rho \\ 0, if\, \left|(\boldsymbol{A}^{s+1} + \boldsymbol{U}_1^s)_{ij}\right| < \lambda_1/\rho \end{cases}$$

$$(\boldsymbol{Z}_2^{s+1})_{ij} = \begin{cases} (\boldsymbol{B}^{s+1} + \boldsymbol{U}_2^s)_{ij} - \lambda_2/\rho, if\, (\boldsymbol{B}^{s+1} + \boldsymbol{U}_2^s)_{ij} \ge \lambda_2/\rho \\ (\boldsymbol{B}^{s+1} + \boldsymbol{U}_2^s)_{ij} + \lambda_2/\rho, if\, (\boldsymbol{B}^{s+1} + \boldsymbol{U}_2^s)_{ij} \le -\lambda_2/\rho \\ 0, if\, \left|(\boldsymbol{B}^{s+1} + \boldsymbol{U}_2^s)_{ij}\right| < \lambda_2/\rho \end{cases}$$

B. Solving for $\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{\mu}, \boldsymbol{\gamma}$

We rewrite the optimization problem for $\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{\mu}, \boldsymbol{\gamma}$ defined above to an equivalent form:

$$\boldsymbol{A}^{s+1}, \boldsymbol{B}^{s+1}, \boldsymbol{\mu}^{s+1}, \boldsymbol{\gamma}^{s+1} = \underset{\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{\mu}, \boldsymbol{\gamma} \ge 0}{\arg\min}\, f(\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{\mu}, \boldsymbol{\gamma})$$

which can be solved by a generalization of the EM algorithm, namely Majorize-Minimization algorithm [7]. Given any estimation $\boldsymbol{A}^{(\mathrm{d})}, \boldsymbol{B}^{(\mathrm{d})}, \boldsymbol{\mu}^{(\mathrm{d})}, \boldsymbol{\gamma}^{(\mathrm{d})}$ of $\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{\mu}, \boldsymbol{\gamma}$, we minimize a tight upper bound of $f(\boldsymbol{A},\boldsymbol{B},\boldsymbol{\mu},\boldsymbol{\gamma})$, namely a surrogate function $Q(\boldsymbol{A},\boldsymbol{B},\boldsymbol{\mu},\boldsymbol{\gamma} \mid \boldsymbol{A}^{(d)},\boldsymbol{B}^{(d)},\boldsymbol{\mu}^{(d)},\boldsymbol{\gamma}^{(d)})$. Thus, new estimation of $\boldsymbol{A},\boldsymbol{B},\boldsymbol{\mu},\boldsymbol{\gamma}$ can be solved with $dQ/d\boldsymbol{A}=0, dQ/d\boldsymbol{B}=0, dQ/d\boldsymbol{\mu}=0, dQ/d\boldsymbol{\gamma}=0$ respectively.

V. EXPERIMENT

We evaluate the proposed model by conducting experiments on both synthetic and real world data.

A. *On synthetic data*

We drew random model parameters ($\boldsymbol{A},\boldsymbol{B},\boldsymbol{\mu},\boldsymbol{\gamma}$) and simulated information diffusion events by running Ogata's thinning algorithm. Experiments were performed on 5000-10000 events involving 30 users and 30 memes. Two baselines were also evaluated, which included 1) the meme interaction oblivious Benchmark-Hawkes model and 2) the recently proposed NetRate [2] algorithm. We examine how well these models can estimate the parameters. Note that Benchmark-Hawkes model can only infer $\boldsymbol{A},\boldsymbol{\mu}$ and NetRate can only infer $\boldsymbol{A}$. Root mean square error (RMSE) was used as the evaluation metric.

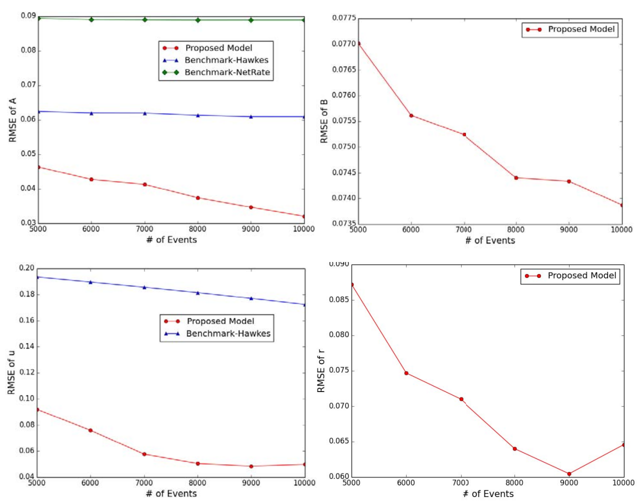

Figure 1. RMSE of parameters $\boldsymbol{A},\boldsymbol{B},\boldsymbol{\mu},\boldsymbol{\gamma}$ estimation.

The results shown in Figure 1 demonstrate the proposed model outperforms other baselines and RMSE decreases when we use more events to estimate the parameters.

B. *On real world data*

We further applied the proposed model to 15M dataset [8], in which Twitter messages were collected in the period April-May 2011, related to the political events (15M movement) occurred at that time in Spain.

The results shown in Figure 2 demonstrate the proposed model can be used to uncover the hidden user influence network and the strongest interaction relationships among different memes.

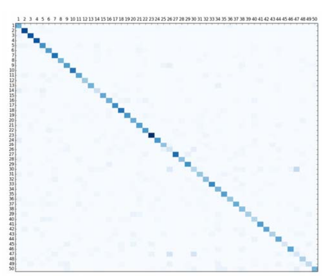

From	To
#reflexion	#europeanrevolution
#3puntosbasicos	#acampadavalencia
#elecciones22m	#22m
#yovoto	#nonosvamos
#consensodeminimos	#worldrevolution
#sinbanderas	#intereconomia
#spahishrevolution	#acampadazgz
#spahishrevolution	#spanishrevolution
#acampadapalma	#estoesreflexion
#acampadapalma	#nonosmovemos

Figure 2. Estimated user influence network represented with heatmap and the top-10 interaction relationships among different memes.

VI. CONCLUSIONS

In this paper, we have proposed a novel multi-dimensional Hawkes process to simultaneously model social influence and meme interaction in information diffusion. The effective model parameter inference method is also presented. Experimental results on both synthetic and real world data show that this model can successfully uncover the hidden social influence network and interaction relationships among memes.

ACKNOWLEDGMENT

This work was supported in part by the following grants: The National Natural Science Foundation of China under Grant No. 71025001, 71472175, 71103180, 61175040, and 61172106; The National Institutes of Health (NIH) of USA under Grant No. (Grant No.1R01DA037378-01) and the Ministry of Health under Grant No. 2012ZX10004801 and 2013ZX10004218, and by the Grant No. 2013A127.

REFERENCES

[1] C. Luo, K. Cui, X. Zheng, and D. Zeng, "Time Critical Disinformation Influence Minimization in Online Social Networks," in *Intelligence and Security Informatics Conference (JISIC), 2014 IEEE Joint*, 2014, pp. 68–74.

[2] Manuel Gomez Rodriguez, D. Balduzzi, and Bernhard Schölkopf, "Uncovering the Temporal Dynamics of Diffusion Networks," in *28th International Conference on Machine Learning*, Bellevue, Washington, 2011.

[3] C. Luo, X. Zheng, and D. Zeng, "Causal Inference in Social Media Using Convergent Cross Mapping," in *Intelligence and Security Informatics Conference (JISIC), 2014 IEEE Joint*, 2014, pp. 260–263.

[4] Ke Zhou, Hongyuan Zha, and Le Song, "Learning Social Infectivity in Sparse Low-rank Networks Using Multi-dimensional Hawkes Processes," in *AISTATS*, 2013, vol. 31, pp. 641–649.

[5] A. G. Hawkes, "Spectra of Some Self-Exciting and Mutually Exciting Point Processes," *Biometrika*, vol. 58, no. 1, p. 83, Apr. 1971.

[6] D. Gabay and B. Mercier, "A dual algorithm for the solution of nonlinear variational problems via finite element approximation," *Comput. Math. Appl.*, vol. 2, no. 1, pp. 17–40, 1976.

[7] H. D. R. and L. K., "A Tutorial on MM Algorithms," *Am. Stat.*, vol. 58, pp. 30–37, 2004.

[8] S. González-Bailón, J. Borge-Holthoefer, A. Rivero, and Y. Moreno, "The Dynamics of Protest Recruitment through an Online Network," *Sci. Rep.*, vol. 1, Dec. 2011.