Modeling Opinion Dynamics in Diffusion Networks

Abir De[1], Isabel Valera[2], Niloy Ganguly[1], Sourangshu Bhattacharya[1], and Manuel Gomez Rodriguez[2]

[1]IIT Kharagpur, {abir.de, niloy, sourangshu}@cse.iitkgp.ernet.in
[2]Max Plank Institute for Software Systems, {ivalera, manuelgr}@mpi-sws.org

Abstract

Social media and social networking sites have become a global pinboard for exposition and discussion of news, topics, and ideas, where social media users increasingly form their opinion about a particular topic by learning information about it from her peers. In this context, whenever a user posts a message about a topic, we observe a noisy estimate of her current opinion about it but the influence the user may have on other users' opinions is hidden. In this paper, we introduce a probabilistic modeling framework of opinion dynamics, which allows the underlying opinion of a user to be modulated by those expressed by her neighbors over time. We then identify a set of conditions under which users' opinions converge to a steady state, find a linear relation between the initial opinions and the opinions in the steady state, and develop an efficient estimation method to fit the parameters of the model from historical fine-grained opinion and information diffusion event data. Experiments on data gathered from Twitter, Reddit and Amazon show that our model provides a good fit to the data and more accurate predictions than alternatives.

1 Introduction

Social media and social networking sites are increasingly used by people to express their opinions, give their "hot takes", on the latest breaking news, political issues, sports events, and new products. As a consequence, there has been an increasing interest on leveraging social media and social networking sites to sense and predict *opinions* as well as understand *opinion dynamics*. For example, political parties routinely use social media to sense people's opinion about their political discourse[1]; quantitative investment firms measure and trade investor sentiment using social media [17]; and, large corporations increasingly leverage brand sentiment, estimated from users' posts, likes and shares in social media and social networking sites, to design their marketing campaigns[2]. In this context, multiple methods for sensing opinions, typically based on sentiment analysis [25], have been proposed in recent years [23, 26]. However, methods for predicting opinions are still scarce [6, 7], despite the extensive literature on theoretical models of opinion dynamics [5, 8].

In this paper, we aim at developing a realistic modeling framework of opinion dynamics that does not only fit fine grained opinion data from social media and social networking sites, but also provides accurate predictions of the individual social media users' opinions over time. To this end, we propose a model that captures two intuitive key processes driving opinion dynamics: *informational influence* and *social influence*. The former process accounts for the idea that the users may form or update their opinion about a particular topic by learning from the information and opinions that their friends/neighbors share. It has been one of the main underlying premises used by many well-known theoretical models of opinion dynamics [5, 8]. The latter process, i.e., social influence, accounts for the impact that different users have on the activity level in the network. In other words, some users may either express their opinions more frequently than others, or be more influential and thus trigger a greater number of follow-ups every time they express their opinion. However, despite its relevance on closely related processes such as information diffusion and social activity [10, 11, 18], social influence has been surprisingly ignored by models of opinion dynamics [6].

[1] {http://www.nytimes.com/2012/10/08/technology/campaigns-use-social-media-to-lure-younger-voters.html}
[2] http://www.nytimes.com/2012/07/31/technology/facebook-twitter-and-foursquare-as-corporate-focus-groups.html

More in detail, we propose a novel continuous-time modeling framework of opinion dynamics in social networks, which accounts for both social influence and informational influence. The key idea behind our approach is to model each user's opinion as a continuous-time stochastic process, which is modulated over time by the opinions *expressed* asynchronously by her neighbors. The proposed formulation also captures characteristic properties of the opinion dynamics, previously studied in the literature [6], such as *stubbornness*, *conformity* and *compromise*. Furthermore, for several instances of our modeling framework, we identify the conditions under which opinions converge to a steady state of consensus or polarization. In such cases, we derive a closed-form expression for the relationship between the users' steady state opinions and the initial opinions they start with. We then develop an efficient method to find the optimal model parameters that jointly maximize the likelihood of an observed set of opinions and information diffusion events. Finally, we experiment with both synthetic and real data gathered from Twitter, Reddit and Amazon, and show that our model provides a good fit to the data and more accurate predictions than several state of the art models of opinion dynamics [6, 7, 8, 13, 28].

Further related work. There is an extensive line of work on theoretical models of opinion dynamics and opinion formation [1, 2, 5, 8, 13, 15, 28]. However, previous works typically share the following limitations: (i) they do not distinguish between opinion, which is a latent continuous quantity, and sentiment, which is a noisy (discrete, such as thumbs up/down, or continuous, such as text sentiment) observation of the opinion; (ii) they do not distinguish between informational and social influence; (iii) they focus on analyzing only the steady state of the users' opinion, neglecting the transient behavior of real opinion dynamics; (iv) their model parameters are difficult to fit from real fine-grained opinion data and instead are set arbitrarily; and, (v) they consider opinions to be updated sequentially in discrete time, however, opinions may be updated asynchronously, since opinions are expressed asynchronously [10]. More recently, there have been some efforts on designing models that overcome some of the above limitations and can be used for prediction [6, 7], however, they do not distinguish between opinion and sentiment nor between information and social influence, and still consider opinions to be updated sequentially in discrete time. In this work, we introduce a modeling framework that overcomes the above limitations and, by doing so, achieves more accurate and fine-grained predictions than alternatives.

2 Proposed Model

In a social network, whenever a user posts a message about a topic, she is revealing a noisy estimate of her current (hidden) opinion about the topic. Here, we think of users' opinions as stochastic processes that may evolve over time, and think of the sentiment users express in each message as noisy samples from these stochastic processes localized in time. Our modeling framework aims to uncover the evolution of these processes by modeling two key phenomena: informational influence, by which a user's message may modulate other users' opinions, and social influence, by which a user's message may trigger additional messages or replies by others. We illustrate both processes in Figure 1. Next, we formulate our generative model, starting from the data it is designed for.

Given a directed social network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, we record each message as $e := (u, m, t)$, where the triplet means that the user $u \in \mathcal{V}$ posted a message with sentiment m at time t. Given a collection of messages $\{e_1 = (u_1, m_1, t_1), \ldots, e_n = (u_n, m_n, t_n)\}$, the history $\mathcal{H}_u(t)$ gathers all messages posted by user u up to but not including time t, i.e.,

$$\mathcal{H}_u(t) = \{e_i = (u_i, m_i, t_i) | u_i = u \text{ and } t_i < t\}, \tag{1}$$

and $\mathcal{H}(t) := \cup_{u \in \mathcal{V}} \mathcal{H}_u(t)$ denotes the entire history of messages up to but not including time t.

We represent the users' opinions as a multidimensional (latent) stochastic process $\mathbf{x}^*(t)$, in which the u-th entry, $x_u(t) \in \mathbb{R}$, represents the opinion of user u at time t. Here, the sign * means that the opinion $x_u^*(t)$ may depend on the history $\mathcal{H}(t)$. Then, every time a user u posts a message at time t, we draw its sentiment m from a sentiment distribution $p(m|x_u^*(t))$. Further, we represent the message times by a set of counting processes. In particular, we denote the set of counting processes as a vector $\boldsymbol{N}(t)$, in which the u-th entry, $N_u(t) \in \{0\} \cup \mathbb{Z}^+$, counts the number of messages user u posted up to but not including time t. That is, $N_u(t) = |\mathcal{H}_u(t)|$ is the size of the history $\mathcal{H}_u(t)$. Then, we can characterize the message rate of the users using their corresponding intensities as

$$\mathbb{E}[d\boldsymbol{N}(t) \,|\, \mathcal{H}(t)] = \boldsymbol{\lambda}^*(t)\, dt, \tag{2}$$

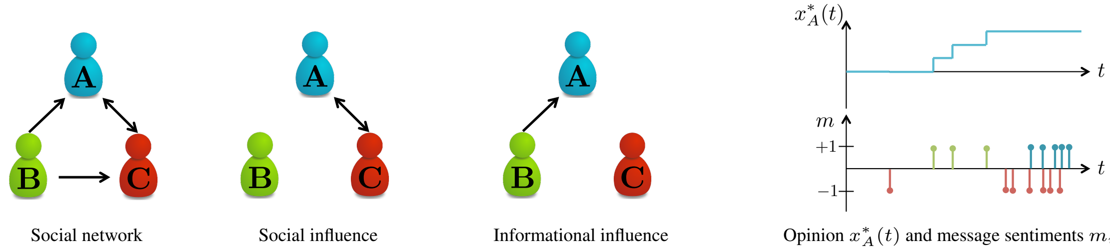

Figure 1: Our model of opinion dynamics in social networks. A user has a hidden opinion $x^*(t)$ and every time she posts a message i at time t_i, we observe the sentiment m_i expressed in the message, which depends on $x^*(t_i)$. When a user receives messages posted by her in-neighbors, she may update her opinion $x^*(t)$ by means of *informational influence*, and may be stimulated to post more messages through *social influence*. In the above example, user A updates her opinion, $x_A^*(t)$ every time she reads a message posted by user B; and A and C stimulate each other to post more messages.

where $d\boldsymbol{N}(t) := (\ dN_u(t)\ )_{u\in\mathcal{V}}$ denotes the number of message events per user in the window $[t, t+dt)$ and $\boldsymbol{\lambda}^*(t) := (\ \lambda_u^*(t)\ )_{u\in\mathcal{V}}$ denotes the vector of intensities associated to all the users, which may depend on the history $\mathcal{H}(t)$. Moreover, we denote the neighborhood of user u by $\mathcal{N}(u)$. Next, we specify the functional form of the intensity functions $\boldsymbol{\lambda}^*(t)$, the users' opinions $\mathbf{x}^*(t)$, and the sentiment distribution $p(m|x_u^*(t))$.

Intensity for messages. We leverage the multivariate version of Hawkes processes [22] to model the message intensities, as in a number of recent works [4, 9, 16, 21]. In particular, each intensity $\lambda_u^*(t)$ takes the following form:

$$\lambda_u^*(t) = \mu_u + \sum_{v\in u\cup\mathcal{N}(u)} b_{vu} \sum_{e_i\in\mathcal{H}_v(t)} \kappa(t-t_i) = \mu_u + \sum_{v\in u\cup\mathcal{N}(u)} b_{vu}\,(\kappa(t)\star dN_v(t)), \tag{3}$$

where the first term, $\mu_u \geq 0$, models the publication of messages by user u on her own initiative, and the second term, with $b_{vu} \geq 0$, models the publication of additional messages by user u due to social influence, *i.e.*, the influence that previous messages posted by users that u follows has on her intensity. Here, $\kappa(t)$ denotes a nonnegative triggering kernel, which models the decay of social influence over time, and $\star$ denotes the convolution operation.

Stochastic process for opinion. The opinion $x_u^*(t)$ of a user u at time t takes the following form:

$$x_u^*(t) = \alpha_u + \sum_{v\in\mathcal{N}(u)} a_{vu} \sum_{e_i\in\mathcal{H}_v(t)} m_i g(t-t_i) = \alpha_u + \sum_{v\in\mathcal{N}(u)} a_{vu}\,(g(t)\star(m_v(t)dN_v(t))), \tag{4}$$

where the first term, $\alpha_u \in \mathbb{R}$, models the original opinion a user u starts with, and the second term, with $a_{vu} \in \mathbb{R}$, models updates in user u's opinion due informational influence, *i.e.*, the influence that previous messages with opinions m_i posted by users u follows has on her opinion. Here, $g(t)$ denotes a nonnegative triggering kernel, which models the decay of informational influence over time, if any.

Importantly, this functional form allows for stubbornness, conformity and compromise, which previous literature has identified as characteristic properties of opinion dynamics [6]. In particular, stubborn users do not update their opinions by means of informational influence, and can be characterized by $a_{vu} = 0$ for all v; conforming users's opinion do not have an opinion at the beginning, and can be characterized by $\alpha_u = 0$; and compromised users have a starting opinion that updates over time by means of informational influence, and can be characterized by $\alpha_u \neq 0$ and $a_{vu} \neq 0$ for some v.

Sentiment distribution. The particular choice of sentiment distribution $p(m|x_u^*(t))$ depends on the recorded data. In this paper, we allow for both continuous sentiment, $m \in (-\infty, \infty)$, and discrete binary sentiment, $m \in \{-1, 1\}$, and consider:

I. **Gaussian Distribution** The sentiment is assumed to be a real random variable $m \in \mathbb{R}$, *i.e.*, $p(m|x_u(t)) = \mathcal{N}(x_u(t), \sigma_u)$.

II. **Logistic.** The sentiment is assumed to be a binary random variable $m \in \{-1, 1\}$, *i.e.*, $p(m|x_u(t)) = 1/(1 + \exp(-m \cdot x_u(t)))$.

However, our estimation method in Section 4 can be easily adapted to any log-concave sentiment distribution.

3 Model Properties

In this section we aim to identify under which conditions users' average opinion, $\mathbb{E}_{\mathcal{H}(t)}[\mathbf{x}^*(t)]$, converges to a steady state and, if so, find the steady state opinion $\lim_{t\to\infty} \mathbb{E}_{\mathcal{H}(t)}[\boldsymbol{x}^*(t)]$. Throughout this section (and in Appendices A-E), we assume the sentiment $m \in \mathbb{R}$ and the sentiment distribution satisfies that $\mathbb{E}[m|x_u^*(t)] = x_u^*(t)$ (*e.g.*, $p(m|x_u^*(t)) = \mathcal{N}(x_u^*(t), \sigma_u)$), and we write $\mathcal{H}_t = \mathcal{H}(t)$ to lighten the notation. Additionally, we denote the eigenvalues of a matrix $\boldsymbol{X}$ by $\xi(\boldsymbol{X})$.

Perhaps surprisingly, we can show that the behavior of the average opinion satisfies the following (proven in Appendix A):

Lemma 1. *The average opinion $\mathbf{x}^*(t)$ in the model of opinion dynamics defined by Eqs. 3 and 4 with exponential triggering kernels with parameters ω and ν satisfies the following differential equation:*

$$\frac{d\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)]}{dt} = \boldsymbol{A}\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t) \odot \boldsymbol{\lambda}^*(t)] - \omega \mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)] + \omega \boldsymbol{\alpha}, \tag{5}$$

where $\boldsymbol{A} = (a_{vu})_{v,u\in\mathcal{G}}$, $\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_{|\mathcal{V}|}]^\top$, and the sign $\odot$ denotes pointwise product. Moreover, if $\lambda_u^*(t) = \mu_u + \int_0^t b_{uu} k(t-\theta) dN_u(\theta)$ for all $u \in \mathcal{G}$. Then, Eq. 5 simplifies to:

$$\frac{d\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)]}{dt} = [-\omega I + \boldsymbol{A}\boldsymbol{\Lambda}(t)]\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)] + \omega\boldsymbol{\alpha}, \tag{6}$$

where

$$\boldsymbol{\Lambda}(t) = \mathrm{diag}[\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{\lambda}^*(t)]], \text{ and}$$
$$\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{\lambda}^*(t)] = \left[e^{(\boldsymbol{B}-\nu I)t} + \nu(\boldsymbol{B}-\nu I)^{-1}\left(e^{(\boldsymbol{B}-\nu \boldsymbol{I})t} - \boldsymbol{I}\right)\right]\boldsymbol{\mu},$$

where $\boldsymbol{B} = \mathrm{diag}[[b_{11}, \ldots, b_{|\mathcal{V}||\mathcal{V}|}]^\top]$, $\boldsymbol{\mu} = [\mu_1, \ldots, \mu_{|\mathcal{V}|}]^\top$, and $\mathrm{diag}[\mathbf{x}]$ is a diagonal matrix with the elements of vector $\mathbf{x}$ in the diagonal.

A general analysis of convergence using Eq. 5 is, in general, difficult. Here, we consider three particular instances of our model under which such convergence analysis is feasible and then discuss the implications of our results at the end of this section.

I. Poisson intensity. Consider $b_{vu} = 0$ for all $v, u \in \mathcal{G}$, *i.e.*, user's messages follow a Poisson process with rate $\lambda_u^*(t) = \mu_u$. Then, the average opinion $\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)]$ and the steady state average opinion $\lim_{t\to\infty} \mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)]$ are given by (proven in Appendices B and C):

Theorem 2. *Given the conditions of Lemma 1 and $\lambda_u^*(t) = \mu_u$ for all $u \in \mathcal{G}$, then,*

$$\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)] = \left[e^{(\boldsymbol{A}\boldsymbol{\Lambda}_1 - \omega I)t} + \omega(\boldsymbol{A}\boldsymbol{\Lambda}_1 - \omega I)^{-1}\left(e^{(\boldsymbol{A}\boldsymbol{\Lambda}_1 - \omega \boldsymbol{I})t} - \boldsymbol{I}\right)\right]\boldsymbol{\alpha}, \tag{7}$$

where $\boldsymbol{\Lambda}_1 = \mathrm{diag}[\boldsymbol{\mu}]$.

Theorem 3. *Given the conditions of Theorem 2, if $Re[\xi(\boldsymbol{A}\boldsymbol{\Lambda}_1)] < \omega$, then,*

$$\lim_{t\to\infty} \mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)] = \left(I - \frac{\boldsymbol{A}\boldsymbol{\Lambda}_1}{w}\right)^{-1}\boldsymbol{\alpha}. \tag{8}$$

II. Univariate Hawkes intensities. Consider $b_{vu} = 0$ for all $v, u \in \mathcal{G}, v \neq u$, *i.e.*, $\lambda_u^*(t) = \mu_u + \int_0^t b_{uu} k(t-\theta) dN_u(\theta)$. Then, the steady state average opinion $\lim_{t\to\infty} \mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)]$ is given by (proven in Appendix D):

Theorem 4. *Given the conditions of Lemma 1 and $\lambda_u^*(t) = \mu_u + \int_0^t b_{uu} k(t-\theta) dN_u(\theta)$ for all $u \in \mathcal{G}$. If the transition matrix $\Phi(t)$ associated to the time-varying linear system described by Eqs. 6 satisfies that $||\Phi(t)|| \leq \gamma e^{-ct}$ $\forall t > 0$, where $\gamma, c > 0$, then,*

$$\lim_{t\to\infty} \mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)] = \left(I - \frac{\boldsymbol{A}\boldsymbol{\Lambda}_2}{\omega} \right)^{-1} \boldsymbol{\alpha}, \tag{9}$$

where $\boldsymbol{\Lambda}_2 = \operatorname{diag}\left[I - \frac{B}{\nu} \right]^{-1} \boldsymbol{\mu}$.

III. Stubborn users. Consider $x_v^*(t) = \alpha_v$ for all $v \in \mathcal{G}\backslash\{u\}$ and $x_u^*(t)$ is given by Eq. 4, *i.e.*, all users except u are stubborn [27]. Then, the steady state average opinion of user u, $\lim_{t\to\infty} \mathbb{E}_{\mathcal{H}_t}[x_u^*(t)]$, is given by (proven in Appendix E):

Theorem 5. *Consider the model of opinion dynamics defined by Eq. 3 and 4 but assume that $x_v^*(t) = \alpha_v$ for all $v \in \mathcal{G}\backslash\{u\}$. Then, if $Re[\xi(\boldsymbol{B})] < \nu$,*

$$\lim_{t\to\infty} \mathbb{E}_{\mathcal{H}_t}[x_u^*(t)] = \boldsymbol{\alpha} + \frac{\nu}{\omega} \boldsymbol{A}_{u*}^T \operatorname{diag}[\boldsymbol{\alpha}][\nu I - \boldsymbol{B}]^{-1} \boldsymbol{\mu}, \tag{10}$$

where $\boldsymbol{A}_{u*}$ denotes the u-th row of $\boldsymbol{A}$ and $\boldsymbol{B} = (b_{vu})_{v,u\in\mathcal{G}}$.

Implications. The above results do not only identify the conditions that ensure the existence of a steady state average opinion (**I**: $Re[\xi(\boldsymbol{A}\boldsymbol{\Lambda}_1)] < \omega$; **II**: $||\Phi(t)|| \leq \gamma e^{-ct}$; and **III**: $Re[\xi(\boldsymbol{B})] < \nu$) but also provide a closed-form expression for its value (Eq. 8–10). Therefore, these results establish the foundations to investigate consensus and polarization in arbitrary networks as well as carrying out opinion shaping [9], which are very interesting venues for future work. Additionally, if the user message times are Poisson, Theorem 2 allows for a fine grained analysis of the temporal evolution of average opinions over time, as illustrated in Figure 3.

4 Model Parameter Estimation

Given a collection of messages $\mathcal{H}(T) = \{(u_i, m_i, t_i)\}$ recorded during a time period $[0, T)$ in a social network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, our goal is to find the optimal parameters $\boldsymbol{\alpha}$, $\boldsymbol{\mu}$, $\boldsymbol{A}$ and $\boldsymbol{B}$ by solving a maximum likelihood estimation (MLE) problem. To this end, it is easy to show that the log-likelihood of the messages is given by

$$\mathcal{L}(\boldsymbol{\alpha}, \boldsymbol{\mu}, \boldsymbol{A}, \boldsymbol{B}) = \underbrace{\sum_{e_i\in\mathcal{H}(T)} \log p(m_i | x_{u_i}^*(t_i))}_{\text{message sentiments}} + \underbrace{\sum_{e_i\in\mathcal{H}(T)} \log \lambda_{u_i}^*(t_i) - \sum_{u\in\mathcal{V}} \int_0^T \lambda_u^*(\tau)\, d\tau}_{\text{message times}}. \tag{11}$$

Then, we can find the optimal parameters $(\boldsymbol{\alpha}, \boldsymbol{\mu}, \boldsymbol{A}, \boldsymbol{B})$ using maximum likelihood estimation (MLE) as

$$\underset{\boldsymbol{\alpha}, \boldsymbol{\mu}\geq 0, \boldsymbol{A}, \boldsymbol{B}\geq 0}{\text{maximize}} \quad \mathcal{L}(\boldsymbol{\alpha}, \boldsymbol{\mu}, \boldsymbol{A}, \boldsymbol{B}). \tag{12}$$

Note that, as long as the sentiment distributions are log-concave, the MLE problem above is concave and thus can be solved efficiently. Moreover, the problem decomposes in $2|\mathcal{V}|$ independent problems, two per user u, since the first term in Eq. 11 only depends on $(\boldsymbol{\alpha}, \boldsymbol{A})$ whereas the last two terms only depend on $(\boldsymbol{\mu}, \boldsymbol{B})$, and thus can be readily parallelized. The global optimum can be found by many algorithms. Here, in order to find $(\mu^*, \boldsymbol{B}^*)$, we use spectral projected gradient descent [3], which we found that work well in practice. For finding $(\boldsymbol{\alpha}^*, \boldsymbol{A}^*)$, we used a standard logistic regressor solver for logistic sentiment distributions, and a standard least-square solver for Gaussian sentiment distributions.

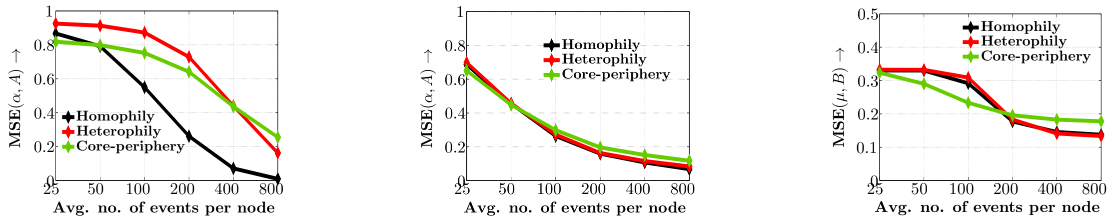

(a) Gaussian $p(m|x_u^*(t))$, $(\boldsymbol{\alpha}, \mathbf{A})$ (b) Logistic $p(m|x_u^*(t))$, $(\boldsymbol{\alpha}, \mathbf{A})$ (c) Hawkes intensity, $(\boldsymbol{\mu}, \mathbf{B})$

Figure 2: Performance of model estimation for several 512-node kronecker networks in terms of mean squared error (MSE) between estimated and true parameters. As we feed more messages into the estimation procedure, the estimation becomes more accurate.

5 Experiments

We validate our model using both synthetic and real data gathered from Twitter, Reddit and Amazon. We first use synthetic data to show that our model is able to produce opinion dynamics that converge to a steady state of consensus and polarization, and our estimation method can accurately recover the true model parameters from recorded messages using synthetic data. We then use real data to show that our model can accurately predict users' opinions, significantly outperforming five state of the art methods [6, 7, 8, 13, 28], and uncover users' opinions coevolution. Appendices F and G provide additional experimental results.

5.1 Experiments on synthetic data

Model parameters estimation. Here, we evaluate the accuracy of our model estimation procedure, described in Section 4. To this end, we generate three types of synthetic networks using a well-known model of directed social networks, the Kronecker graph model [19]: i) Homophily/assortative networks with (parameter matrix $[0.96, 0.3; 0.3, 0.96]$); ii) heterophily/dissortative networks (parameter matrix $[0.3, 0.96; 0.96, 0.3]$); and iii) core-periphery networks (parameter matrix $[0.9, 0.5; 0.5, 0.3]$). For each network, we draw the parameters $\boldsymbol{\mu}$ and $\boldsymbol{B}$ from a uniform distribution $U(0, 1)$, and the opinion parameters $\boldsymbol{\alpha}$ and $\boldsymbol{A}$ from a Gaussian distribution $\mathcal{N}(\mu = 0, \sigma = 1)$. We use exponential kernels with parameters $\omega = 100$ and $\nu = 1$, respectively, for opinions $\mathbf{x}_u^*(t)$ and intensities $\boldsymbol{\lambda}^*(t)$, and experiment both with logistic and Gaussian sentiment distributions. We simulate from the model adapting an efficient implementation of Ogata's algorithm [24]. We evaluate the accuracy of our model estimation procedure in terms of mean squared error (MSE), between the estimated parameters ($\hat{x}$) and the true parameters (x), *i.e.*, $\mathbb{E}[(x - \hat{x})^2]$. Figure 2 shows that as we feed more messages into the estimation procedure, the estimation becomes more accurate. Note that the MSE for $(\boldsymbol{\alpha}, \boldsymbol{A})$ is larger than for $(\boldsymbol{\mu}, \boldsymbol{B})$ because $a_{ji} \in (-\infty, \infty)$ while $b_{ji} \in (0, 1)$. Appendix F contains additional experiments using Forest Fire networks [20].

Consensus and polarization. In Section 3, we provided closed-form expressions of the steady state average opinion for several instances of our model. Here, we verify those results empirically by simulating our model on three different small networks using Gaussian sentiment distribution and Poisson intensities, *i.e.*, $\lambda_u^*(t) = \mu_u$, $\mu \sim U(0, 1)$. Figure 3 summarizes the results, which show that (i) the theoretical average opinions given in Eq. 7 closely match the empirical estimates and, (ii) our model is able to produce opinion dynamics that converge to negative consensus (G_1), positive consensus (G_2) and polarization (G_3).

5.2 Experiments on real data

Data description. We experimented with four Twitter datasets about current real-world events (Tw: Politics, Tw: Movie, Tw: Fight and Tw: Bollywood), which we gathered in-house, a dataset gathered from Amazon[3] (Amzn: food)

[3]https://snap.stanford.edu/data/web-FineFoods.html

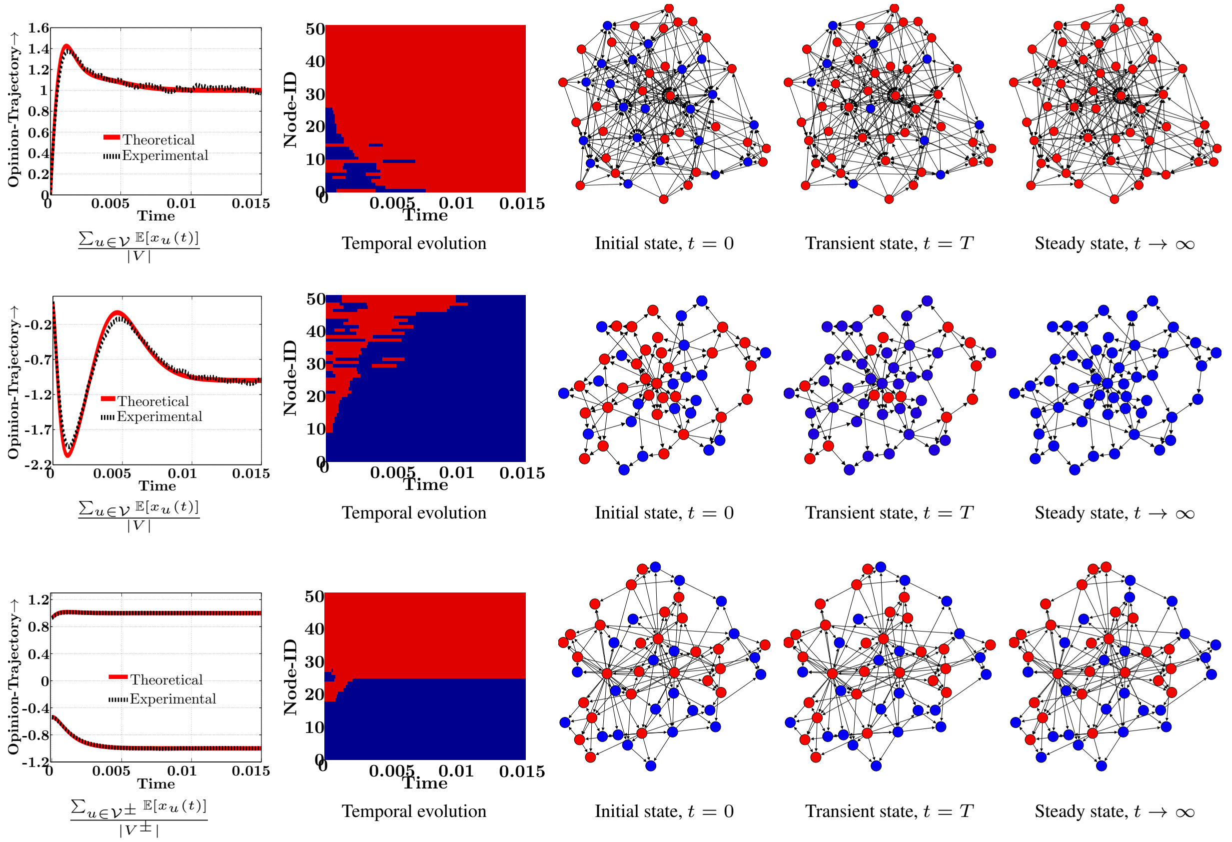

Figure 3: Opinion dynamics on three 50-node networks G_1 (top), G_2 (middle), G_3 (bottom). The first column shows the temporal evolution of the theoretical and empirical average opinions (in the third row, we compute the average separately for positive and negative opinions in the steady state). The second column shows the polarity of theoretical average opinion per user over time. The three right columns show three snapshots of the opinions in the networks at different times. The number of nodes with positive opinion (in red) and negative opinion (in blue) at time $t = 0$ is the same in all networks, however, at $t \to \infty$, G_1 reaches positive consensus, G_2 reaches negative consensus and G_3 gets uniformly polarized, due to their model parameter values and network structure. In all three networks, the sentiment distributions are Gaussian and the intensities are Poisson.

and a dataset gathered from Reddit (Reddit) [7]. Statistics about these datasets are shown in Table 1. Appendix F contains more detailed information about the way the Twitter datasets are build as well as a description of the current real-world events they correspond to.

Experimental setup. For each recorded message i, our model requires its timing t_i, which is already explicit in the above datasets, and sentiment value m_i, which we compute as follows. For the Twitter datasets, we compute each message sentiment m_i using a popular sentiment analysis toolbox, specially designed for Twitter [12], which returns values in $(-1, 1)$. For the Amazon and Reddit datasets, we compute each message sentiment m_i using a popular sentiment analysis toolbox, LIWC[4], which returns values in $(-1, 1)$. We consider the sentiment polarity to be simply $\text{sign}(m)$.

The Amazon and Reddit datasets do not explicitly provide a social network, which we build as follows. For the Amazon dataset, we assume two users *follow* each other if they have posted at least 15 reviews on the same type of food. For the Reddit dataset, we assume a user v *follows* a user u if v has posted a message within a week after u

[4] `http://www.liwc.net/`

Dataset	Nodes, $|\mathcal{V}|$	Edges, $|\mathcal{E}|$	Msgs, $|\mathcal{H}(T)|$	$\mathbb{E}[\mathbf{m}]$	std$[m]$
Tw: Politics	548	5271	20026	0.0169	0.1780
Tw: Movie	567	4886	14016	0.6969	0.1358
Tw: Fight	848	10118	21526	-0.0123	0.2577
Tw: Bollywood	1031	34952	46845	0.5101	0.2310
Amzn: Food	447	13498	24815	0.2108	0.1648
Reddit	556	4629	64,366	-0.0088	0.1333

Table 1: Real datasets statistics

Dataset	HGOM	AsLM	DeGroot	Voter	Bias. Voter	Flocking
Tw: Politics	**0.0344**	0.0557	0.1557	0.0761	0.0549	0.0782
Tw: Movie	**0.0482**	0.1123	0.2868	0.5606	0.1135	0.5702
Tw: Fight	**0.0373**	0.0549	0.0714	0.098	0.0670	1.0341
Tw: Bollywood	**0.0825**	0.1704	0.2889	0.3787	0.1470	0.3443
Amz: Food	**0.0593**	0.0791	0.1567	0.1008	0.0781	0.0968
Reddit	**0.0723**	0.1083	0.1465	0.1842	0.0869	0.1853

(a) Mean squared error (MSE) on sentiment value prediction. The sentiment value $m \in (-1, 1)$.

Dataset	HGOM	AsLM	DeGroot	Voter	Biased Voter	Flocking
Tw: Politics	**0.0539**	0.1389	0.1643	0.1938	0.1214	0.2148
Tw: Movie	0	0	0	0	0	0
Tw: Fight	**0.0507**	0.0669	0.1125	0.1552	0.1004	0.2844
Tw: Bollywood	**0.0457**	0.0470	0.0472	0.0942	0.0540	0.0914
Amz: Food	**0.0891**	0.0988	0.1142	0.1104	0.1023	0.3087
Reddit	**0.0065**	0.0720	0.0871	0.1389	0.0476	0.1418

(b) Failure rate (FR) on sentiment polarity prediction. The sentiment polarity $\text{sign}\,(m) \in \{-1, 1\}$.

Table 2: Sentiment prediction performance using a 10% held-out set for each dataset.

posted a message in the same Reddit group.

Sentiment prediction. For each dataset, we first estimate the parameters of our model, which we denote as HGOM, using messages from a training set (90% of the messages). Since the message sentiment $m \in (-1, 1)$, we use Gaussian sentiment distributions. Moreover, $\kappa(t)$ and $g(t)$ are exponential triggering kernels, with decay parameters set by cross-validation. Then, we predict the sentiment value m for each message in a held-out set (10% of the messages) using the trained HGOM model, and compare its performance with five state of the art opinion models: the asynchronous linear model (AsLM) [7], DeGroot's model [8], the voter model [28], the biased voter model [6], and the flocking model [13]. Here, we evaluate the performance by (i) computing the mean squared error (MSE) between the true (m) and the estimated ($\hat{m}$) sentiment value for all messages in the held-out set, *i.e.*, $\mathbb{E}[(m - \hat{m})^2]$, and (ii) computing the failure rate, defined as the probability that the true polarity $\text{sign}\,(m)$ and the estimated polarity $\text{sign}\,(\hat{m})$ do not coincide, *i.e.*, $\mathbb{P}(\text{sign}(m) \neq \text{sign}(\hat{m}))$. Table 2 summarizes the results, using a 10% held-out set for each dataset. Our method consistently outperforms alternatives both in predicting sentiment polarity and sentiment value.

Then, we train our model, using also Gaussian sentiment distribution, for different training (held-out) set sizes and compare its performance with DeGroot's model and AsLM. Figure 4 summarizes the results, which show that our model outperforms competing methods for any training size and always benefit from increasing the training sets, in contrast with competing methods, which sometimes degrade its performance. Here, we only compare with DeGroot's model and AsLM since the other methods do not need to be trained since they base their predictions simply on previous events.

Finally, we investigate the predictive performance of our model for discrete-valued sentiment messages. To this end, for the Twitter and Reddit datasets, we only record the polarity of the sentiment provided by [12] and LIWC respectively; and for the Amazon dataset, since the sentiment provided by LIWC has almost always positive polarity,

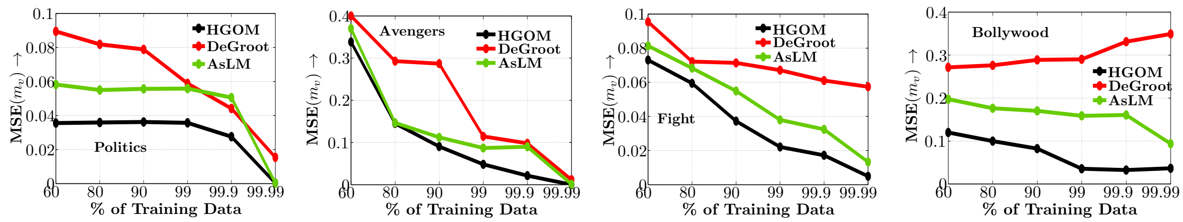

(a) Politics (b) Movie (c) Fight (d) Bollywood

Figure 4: Mean squared error (MSE) on sentiment value prediction against training set size for all datasets gathered from Twitter. Appendix G shows similar trends for Amazon Food and Reddit.

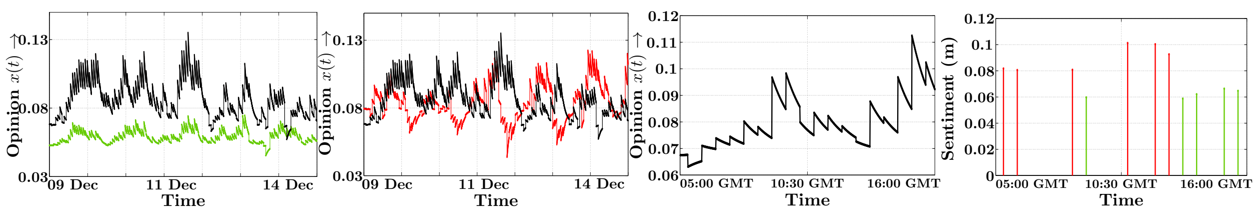

(a) $x^*_A(t)$ vs $x^*_B(t)$ (b) $x^*_A(t)$ vs $x^*_C(t)$ (c) $x^*_A(t)$ (d) $m_B(t)$ and $m_C(t)$

Figure 5: Opinion dynamics. Panels (a, b) show that user A's opinion, $x^*_A(t)$ (in black), co-evolves with user B's opinion, $x^*_B(t)$ (in green), and user C's opinion, $x_C(t)$ (in red). Panels (c, d) show that user A's opinion $x^*_A(t)$ (in black) is updated by means of informational influence every time user B (in green) and C (in red) post a message with a sentiment m.

we consider all the review (message) scores equal to 5 (52.6% of the total ratings) to be positive, while scores below 5 are considered negative.

For each dataset, we first estimate the parameters of our model using messages from a training set (90% of the messages). In this case, since we only record sentiment polarity, we use logistic sentiment distributions, and denote our model as HLOM. We then compare our model with the Voter model in terms of sentiment polarity prediction. Table 3 summarizes the results, where we observe that our model outperforms the Voter model for all the datasets except for the "Tw: Movie", in which all the sentiments are positive and both models achieve the same perfect score.

Opinions co-evolution and informational influence. Figure 5 visualizes the opinions $x^*(t)$, as inferred by our model, for a user (A, in black) and two of the users she follows (B, in green; C, in red), picked at random. Here, it is apparent that users' opinions coevolve by means of informational influence and the users in red and green exert informational influence over the user in black. In particular, every time users B and C post a message with a sentiment m, user A updates her opinion, as shown in Panels (c, d). Here, note that user A also updates her opinion when additional people she follows, which we do not show in the figure, post messages. For clarity, we show only three nodes, however, we found qualitatively similar results across nodes.

6 Conclusions

We proposed a continuous-time modeling framework of opinion dynamics, naturally designed to fit fine-grained opinion data. The key innovation of our framework is distinguishing latent opinions from expressed opinions, and accounting for both informational and social influence. As a result, it predicts expressed opinions over time more accurately than previous state-of-the-arts. Our proposed model opens up many interesting venues for future work. For example, a general analysis of convergence in our model would be of great interest, since it would allow us to investigate the effect of complex communication patterns on the emergence of consensus and polarization. Moreover, it would

Dataset	HLOM	Voter
Tw: Politics	0.4565	0.4820
Tw: Movie	0	0
Tw: Fight	0.1966	0.2558
Tw: Bollywood	0.0085	0.0602
Amz: Food	0.1092	0.1231
Reddit	0.2190	0.3770

Table 3: Failure rate (FR) on sentiment polarity prediction using a 10% held-out set for each dataset. Here, HLOM and Voter only have access to the sentiment polarities to make predictions.

be interesting to leverage our framework to design opinion shaping algorithms and reason about large-scale opinion dynamics in large datasets.

References

[1] R. Axelrod. The dissemination of culture a model with local convergence and global polarization. *Journal of conflict resolution*, 41(2):203–226, 1997.

[2] D. Bindel, J. Kleinberg, and S. Oren. How bad is forming your own opinion? In *FOCS*, 2011.

[3] E. G. Birgin, J. M. Martínez, and M. Raydan. Nonmonotone spectral projected gradient methods on convex sets. *SIAM Journal on Optimization*, 10(4), 2000.

[4] C. Blundell, J. Beck, and K. A. Heller. Modelling reciprocating relationships with hawkes processes. In *NIPS*, 2012.

[5] P. Clifford and A. Sudbury. A model for spatial conflict. *Biometrika*, 60(3):581–588, 1973.

[6] A. Das, S. Gollapudi, and K. Munagala. Modeling opinion dynamics in social networks. In *WSDM*, 2014.

[7] A. De, S. Bhattacharya, P. Bhattacharya, N. Ganguly, and S. Chakrabarti. Learning a linear influence model from transient opinion dynamics. In *CIKM*, 2014.

[8] M. H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.

[9] M. Farajtabar, N. Du, M. Gomez-Rodriguez, I. Valera, L. Song, and H. Zha. Shaping social activity by incentivizing users. In *NIPS*, 2014.

[10] M. Gomez-Rodriguez, D. Balduzzi, and B. Schölkopf. Uncovering the Temporal Dynamics of Diffusion Networks. In *ICML*, 2011.

[11] M. Gomez Rodriguez, J. Leskovec, and A. Krause. Inferring networks of diffusion and influence. In *KDD*, 2010.

[12] A. Hannak, E. Anderson, L. F. Barrett, S. Lehmann, A. Mislove, and M. Riedewald. Tweetin'in the rain: Exploring societal-scale effects of weather on mood. In *ICWSM*, 2012.

[13] R. Hegselmann and U. Krause. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3), 2002.

[14] D. Hinrichsen, A. Ilchmann, and A. Pritchard. Robustness of stability of time-varying linear systems. *Journal of Differential Equations*, 82(2):219 – 250, 1989.

[15] P. Holme and M. E. Newman. Nonequilibrium phase transition in the coevolution of networks and opinions. *Physical Review E*, 74(5):056108, 2006.

[16] T. Iwata, A. Shah, and Z. Ghahramani. Discovering latent influence in online social activities via shared cascade poisson processes. In *KDD*, 2013.

[17] T. Karppi and K. Crawford. Social media, financial algorithms and the hack crash. *Theory, Culture & Society*, 2015.

[18] D. Kempe, J. M. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. In *KDD*, 2003.

[19] J. Leskovec, D. Chakrabarti, J. M. Kleinberg, C. Faloutsos, and Z. Ghahramani. Kronecker graphs: An approach to modeling networks. *JMLR*, 2010.

[20] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graphs over time: densification laws, shrinking diameters and possible explanations. In *KDD*, 2005.

[21] S. W. Linderman and R. P. Adams. Discovering latent network structure in point process data. In *ICML*, 2014.

[22] T. Liniger. *Multivariate Hawkes Processes*. PhD thesis, ETHZ, 2009.

[23] B. O´Connor, R. Balasubramanyan, B. R. Routledge, and N. A. Smith. From tweets to polls: Linking text sentiment to public opinion time series. *ICWSM*, 2010.

[24] Y. Ogata. On lewis' simulation method for point processes. *IEEE Transactions on Information Theory*, 27(1):23–31, 1981.

[25] B. Pang and L. Lee. Opinion mining and sentiment analysis. *Foundations and trends in information retrieval*, 2(1-2):1–135, 2008.

[26] A. Tumasjan, T. O. Sprenger, P. G. Sandner, and I. M. Welpe. Predicting elections with twitter: What 140 characters reveal about political sentiment. *ICWSM*, 2010.

[27] E. Yildiz, A. Ozdaglar, D. Acemoglu, A. Saberi, and A. Scaglione. Binary opinion dynamics with stubborn agents. *ACM Transactions on Economics and Computation*, 1(4):19, 2013.

[28] M. E. Yildiz, R. Pagliari, A. Ozdaglar, and A. Scaglione. Voting models in random networks. In *Information Theory and Applications Workshop*, pages 1–7, 2010.

A Proof of Lemma 1

Using that $\mathbb{E}[m_v(\theta)|x_v^*(\theta)] = x_v^*(\theta)$, we can compute the average opinion of user u across all possible histories from Eq. 4 as

$$\begin{aligned}\mathbb{E}_{\mathcal{H}_t}[x_u^*(t)] &= \alpha_u + \sum_{v \in \mathcal{N}(u)} a_{uv} \int_0^t g(t-\theta)\mathbb{E}_{\mathcal{H}_t}[m_v(\theta) dN_v(\theta)] \\ &= \alpha_u + \sum_{v \in \mathcal{N}(u)} a_{uv} \int_0^t g(t-\theta)\mathbb{E}_{\mathcal{H}_\theta}\left[x_v^*(\theta)\lambda_v^*(\theta)\right] d\theta,\end{aligned}$$

and we can write the expectation of the opinion for all users in vectorial form as

$$\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)] = \boldsymbol{\alpha} + \boldsymbol{A} \int_0^t g(t-\theta)\mathbb{E}_{\mathcal{H}_\theta}[\boldsymbol{x}^*(\theta) \odot \boldsymbol{\lambda}^*(\theta)] d\theta, \tag{13}$$

where the sign $\odot$ denotes pointwise product. Then, by differentiating Eq 13, we obtain

$$\frac{d\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)]}{dt} = \boldsymbol{A}\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t) \odot \boldsymbol{\lambda}^*(t)] - \omega \mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)] + \omega \boldsymbol{\alpha}. \tag{14}$$

Now, assume $b_{vu} = 0$ for all $v, u \in \mathcal{G}, v \neq u$. Then, $\lambda_v^*(t)$ only depends on user v's history and, since $x_v^*(t)$ only depends on the history of the user v's neighbors $\mathcal{N}(v)$, we can write

$$\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t) \odot \boldsymbol{\lambda}^*(t)] = \mathbb{E}_{\mathcal{H}_t}\left[\boldsymbol{x}^*(t)\right] \odot \mathbb{E}_{\mathcal{H}_t}\left[\boldsymbol{\lambda}^*(t)\right],$$

and rewrite Eq. 14 as

$$\frac{d\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)]}{dt} = \boldsymbol{A}(\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)] \odot \mathbb{E}_{\mathcal{H}_t}[\boldsymbol{\lambda}^*(t)]) - \omega \mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)] + \omega \boldsymbol{\alpha}. \tag{15}$$

We can now compute $\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{\lambda}^*(\theta)]$ analytically as follows. From Eq. 3, we obtain

$$\mathbb{E}_{\mathcal{H}_t}\left[\boldsymbol{\lambda}^*(t)\right] = \boldsymbol{\mu} + \int_0^t \boldsymbol{B}\kappa(t-\theta)\mathbb{E}_{\mathcal{H}_\theta}\left[\boldsymbol{\lambda}^*(\theta)\right] d\theta, \tag{16}$$

where $\boldsymbol{\mu} = [\mu_1, \mu_2, ..., \mu_{|\mathcal{V}|}]^\top$ and $\boldsymbol{B} = (b_{vu})_{v,u \in \mathcal{V}}$, where $b_{vu} = 0$ for all $v \neq u$, by assumption. Then, we apply the Laplace transform and obtain

$$\hat{\boldsymbol{\lambda}}(s) = [I - \boldsymbol{B}\hat{\kappa}(s)]^{-1}\boldsymbol{\mu}, \tag{17}$$

where $\hat{\boldsymbol{\lambda}}(s)$ and $\hat{\kappa}(s)$ denote the Laplace transforms of $\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{\lambda}^*(t)]$ and $\kappa(t)$ respectively. Now, using that $\kappa(t) = \exp(-\nu t)$, we can write

$$\hat{\boldsymbol{\lambda}}(s) = \left[I - \frac{\boldsymbol{B}}{s+\nu}\right]^{-1}\boldsymbol{\mu},$$

and obtain $\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{\lambda}^*(t)]$ in time domain as

$$\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{\lambda}^*(t)] = \left[e^{(\boldsymbol{B}-\nu I)t} + \nu(\boldsymbol{B}-\nu I)^{-1}\left(e^{(\boldsymbol{B}-\nu \boldsymbol{I})t} - \boldsymbol{I}\right)\right]\boldsymbol{\mu}. \tag{18}$$

Finally, using Eq. 15, Eq. 18, and $\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)] \odot \mathbb{E}_{\mathcal{H}_t}[\boldsymbol{\lambda}^*(t)] = \boldsymbol{\Lambda}(t)\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)]$, where $\boldsymbol{\Lambda}(t) := \text{diag}[\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{\lambda}^*(t)]]$, we obtain

$$\frac{d\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)]}{dt} = [-\omega I + \boldsymbol{A}\boldsymbol{\Lambda}(t)]\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)] + \omega \boldsymbol{\alpha}. \tag{19}$$

B Proof of Theorem 2

Using Lemma 1 and $\lambda_u^*(t) = \mu_u$, we obtain

$$\frac{d\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)]}{dt} = [-\omega I + \boldsymbol{A}\boldsymbol{\Lambda}]\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)] + \omega\boldsymbol{\alpha}, \tag{20}$$

where $\boldsymbol{\Lambda}(t) = \mathrm{diag}[\boldsymbol{\mu}]$. Then, we apply the Laplace transform to the expression above and obtain

$$\hat{\boldsymbol{x}}(s) = \left(1 + \frac{\omega}{s}\right)[sI + \omega I - A\boldsymbol{\Lambda}]^{-1}\boldsymbol{\alpha}.$$

Finally, applying the inverse Laplace transform, we obtain the average opinion $\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)]$ in time domain as

$$\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)] = \left[e^{(\boldsymbol{A}\boldsymbol{\Lambda}-\omega I)t} + \omega(\boldsymbol{A}\boldsymbol{\Lambda} - \omega I)^{-1}\left(e^{(\boldsymbol{A}\boldsymbol{\Lambda}-\omega \boldsymbol{I})t} - \boldsymbol{I}\right)\right]\boldsymbol{\alpha}.$$

C Proof of Theorem 3

Theorem 2 states that the average users' opinion $\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)]$ in time domain is given by

$$\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)] = \left[e^{(\boldsymbol{A}\boldsymbol{\Lambda}-\omega I)t} + \omega(\boldsymbol{A}\boldsymbol{\Lambda} - \omega I)^{-1}\left(e^{(\boldsymbol{A}\boldsymbol{\Lambda}-\omega \boldsymbol{I})t} - \boldsymbol{I}\right)\right]\boldsymbol{\alpha}.$$

If $Re[\xi(\boldsymbol{A}\boldsymbol{\Lambda})] < \omega$, where $\xi(\boldsymbol{X})$ denote the eigenvalues of matrix $\boldsymbol{X}$, it easily follows that

$$\lim_{t\to\infty} \mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)] = \left(I - \frac{\boldsymbol{A}\boldsymbol{\Lambda}}{w}\right)^{-1}\boldsymbol{\alpha}. \tag{21}$$

D Proof of Theorem 4

Lemma 1 states that

$$\frac{d\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)]}{dt} = [-\omega I + \boldsymbol{A}\boldsymbol{\Lambda}(t)]\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)] + \omega\boldsymbol{\alpha}, \tag{22}$$

where $\boldsymbol{\Lambda}(t) = \left[e^{(\boldsymbol{B}-\nu I)t} + \nu(\boldsymbol{B} - \nu I)^{-1}\left(e^{(\boldsymbol{B}-\nu \boldsymbol{I})t} - \boldsymbol{I}\right)\right]\boldsymbol{\mu}$. In such systems, solutions can be written as [14]

$$\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)] = \Phi(t)\alpha + \omega\int_0^t \Phi(s)\boldsymbol{\alpha}ds, \tag{23}$$

where the transition matrix $\Phi(t)$ defines as a solution of the matrix differential equation

$$\dot{\Phi}(t) = [-\omega I + \boldsymbol{A}\boldsymbol{\Lambda}(t)]\Phi(t) \text{ with } \Phi(0) = I.$$

If $\Phi(t)$ satisfies $||\Phi(t)|| \leq \gamma e^{-ct}\ \forall t > 0$ for $\gamma, c > 0$ then the steady state solution to Eq. 23 is given by [14]

$$\lim_{t\to\infty} \mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)] = \left(I - \frac{A\boldsymbol{\Lambda}}{\omega}\right)^{-1}\boldsymbol{\alpha}.$$

where $\boldsymbol{\Lambda} = \lim_{t\to\infty}\boldsymbol{\Lambda}(t) = \mathrm{diag}\left[I - \frac{B}{\nu}\right]^{-1}\boldsymbol{\mu}$.

E Proof of Theorem 5

By assumption, $\mathbb{E}_{\mathcal{H}_t}[m_v|x_v(t)] = \alpha_v$ for all $v \in \mathcal{G}\backslash\{u\}$. Then, the expected opinion of node u at time t is given by

$$\mathbb{E}_{\mathcal{H}_t}[x_u^*(t)] = \alpha_u + \sum_{v \in V} a_{uv}\alpha_v \int_0^t g(t-\theta)\mathbb{E}_{\mathcal{H}_\theta}[\lambda_v^*(\theta)]d\theta. \tag{24}$$

Next, we will compute the Laplace transform of the above expression. We start by writing the average intensity for all users in vectorial form as

$$\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{\lambda}^*(t)] = \boldsymbol{\mu} + \int_0^t \boldsymbol{B}\kappa(t-\theta)\mathbb{E}_{\mathcal{H}_\theta}[\boldsymbol{\lambda}^*(\theta)]d\theta, \tag{25}$$

and then apply the Laplace transform and obtain,

$$\hat{\boldsymbol{\lambda}}^*(s) = [I - B\hat{\kappa}(s)]^{-1}\frac{\boldsymbol{\mu}}{s} \tag{26}$$

where $\hat{\kappa}(s)$ is the Laplace transform of the exponential kernel $\kappa(t) = \exp(-\nu t)$. Using Eq. 24 and Eq. 26, the Laplace transform of $\mathbb{E}_{\mathcal{H}_t}[x_u^*(t)]$ is given by

$$\hat{x}_u(s) = \frac{\alpha_u}{s} + \boldsymbol{A}_{u*}^T \operatorname{diag}[\boldsymbol{\alpha}]g(s)\hat{\boldsymbol{\lambda}}^*(s), \tag{27}$$

where $\boldsymbol{A}_{u*}$ is the u^{th} row of $\boldsymbol{A}$. Assuming an exponential kernel $g(t) = \exp(-\omega t)$ and substituting $\hat{\mu}(s)$, we obtain

$$\hat{x}_u(s) = \frac{\boldsymbol{\alpha}}{s} + \boldsymbol{A}_{u*}^T \operatorname{diag}[\boldsymbol{\alpha}]\frac{s+\nu}{s+\omega}[sI + \nu I - B]^{-1}\frac{\boldsymbol{\mu}}{s}. \tag{28}$$

Then, if $Re[\xi(B)] < \nu$, it easily follows that

$$\lim_{t\to\infty} \mathbb{E}_{\mathcal{H}_t}[x_u^*(t)] = \boldsymbol{\alpha} + \frac{\nu}{\omega}\boldsymbol{A}_{u*}^T \operatorname{diag}[\boldsymbol{\alpha}][\nu I - B]^{-1}\boldsymbol{\mu}. \tag{29}$$

F Additional experiments on synthetic data

Parameter estimation. We evaluate the accuracy of our model estimation procedure, described in Section 4, in the same Kronecker networks as in Section 5.1 in terms of log-likelihood over a test set of messages, disjoint from the set of messages used for model estimation. Figure 6 shows that as we feed more messages into the estimation procedure, the test log-likelihood achieved by the estimated parameters becomes closer to the test log-likelihood achieved by the true parameters.

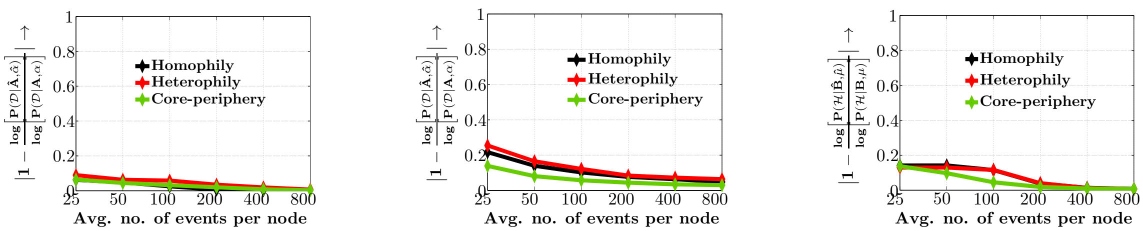

(a) Gaussian $p(m|x^*_u(t))$, msg sentiments (b) Logistic $p(m|x^*_u(t))$, msg sentiments (c) Hawkes intensity, msg times

Figure 6: Performance of model estimation for several 512-node kronecker networks in terms of test log-likelihood. As we feed more messages into the estimation procedure, the estimation becomes more accurate. In all three networks, we used $p_b = 0.3$.

Additionally, we evaluate the accuracy of our model estimation procedure in Forest-Fire networks [20]. Figure 7 summarizes the results. We also observe that as we feed more messages into the estimation procedure, the mean squared error (MSE) becomes smaller and the test log-likelihood achieved by the estimated parameters becomes closer to the test log-likelihood achieved by the true parameters.

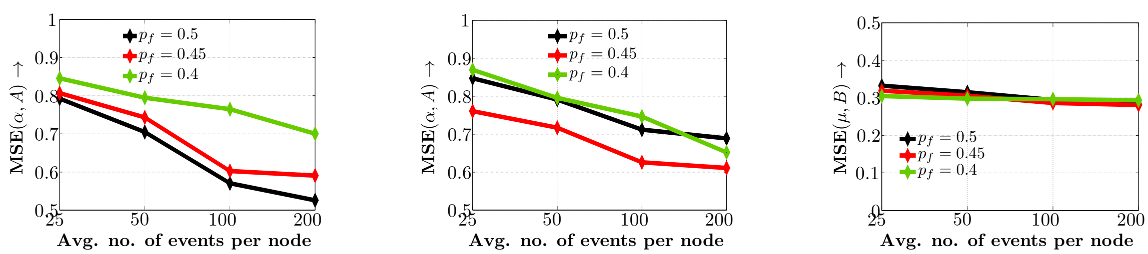

(a) Gaussian $p(m|x^*_u(t))$, $(\boldsymbol{\alpha}, \mathbf{A})$ (b) Logistic $p(m|x^*_u(t))$, $(\boldsymbol{\alpha}, \mathbf{A})$ (c) Hawkes intensity, $(\boldsymbol{\mu}, \mathbf{B})$

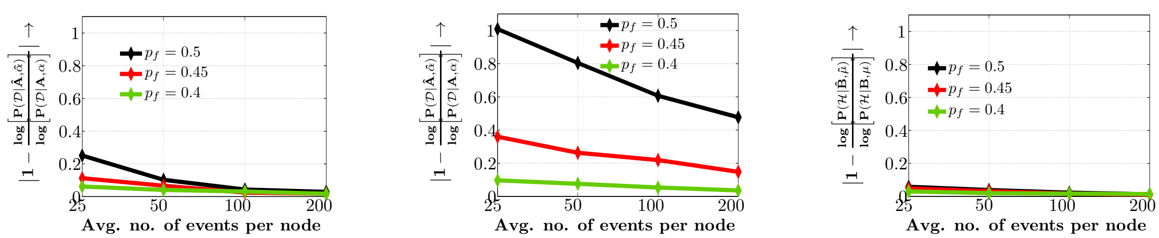

(d) Gaussian $p(m|x^*_u(t))$, msg sentiments (e) Logistic $p(m|x^*_u(t))$, msg sentiments (f) Hawkes intensity, msg times

Figure 7: Performance of model estimation for several 512-node Forest-Fire networks in terms of test log-likelihood. As we feed more messages into the estimation procedure, the estimation becomes more accurate

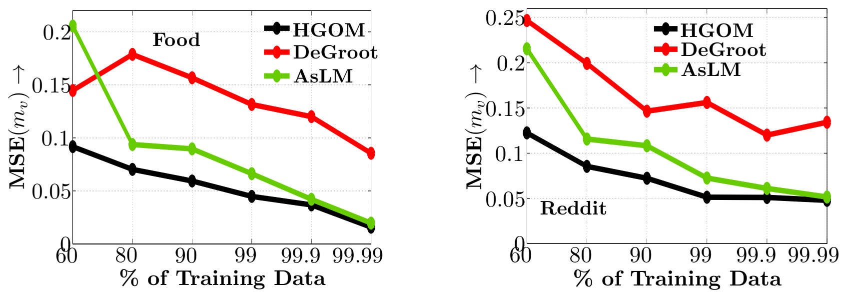

(a) Amazon: Food, MSE

(b) Reddit, MSE

Figure 8: Mean squared error (MSE) on sentiment value prediction against training set size for the Amazon and Reddit datasets.

G Additional experiments on real data

G.1 Twitter dataset description

We used the Twitter search API[5] to collect all the tweets (corresponding to a 2-3 weeks period around the event date) that contain hashtags related to the following events/topics:

- **Politics:** Similarly to [7], we collect all the tweets, from 9th to 15th of December 2013, related to the Delhi Assembly election 2013.
- **Movie:** We collect all the tweets, from April 28th to May 5th, 2015, about the discussion on the release of "*Avengers: Age of Ultron*" movie that took place on May 1, 2015.
- **Fight:** We collect all the tweets, from April 29th to May 7th, 2015, about the discussion on the professional boxing match between the eight-division world champion Manny Pacquiao and the undefeated five-division world champion Floyd Mayweather Jr., which took place on May 2, 2015 in Las Vegas, Nevada.
- **Bollywood:** We collect all the tweets, from May 5th to 16th, 2015, about the discussion on the verdict that declared guilty to Salman Khan (a popular Bollywood movie star) on May 6, 2015, for causing death of a person by rash and negligible driving

We then built the follower-followee network for the users that posted the collected tweets using the Twitter rest API[6]. Finally, we filtered out users that posted less than 200 tweets during the account lifetime, follow less than 100 users, or have less than 50 followers.

G.2 Additional results for sentiment value prediction

Here, we compare the performance of our model with DeGroot's model and AsLM for the Amazon and Reddit datasets in terms of sentiment value prediction against training set size. Figure 8 summarizes the results, which show that our model improves its performance when increasing the training set and outperforms the competing methods for any training size.

[5] `https://dev.twitter.com/rest/public/search`
[6] `https://dev.twitter.com/rest/public`