Shaping Opinion Dynamics in Social Networks

Abir De*
IIT Kharagpur
abir.de@cse.iitkgp.ernet.in

Sourangshu Bhattacharya
IIT Kharagpur
sourangshu@cse.iitkgp.ernet.in

Niloy Ganguly
IIT Kharagpur
niloy@cse.iitkgp.ernet.in

ABSTRACT

A networked opinion diffusion process that usually involves extensive spontaneous discussions between connected users, is often propelled by external sources of news or feeds recommended to them. In many applications like marketing design, or product launch, etc., corporations often post curated news or feeds on social media in order to steer the users' opinions in a desired way. We call such scenarios as *opinion shaping* or *opinion control* whereby a few select users called *control users* post opinionated messages to drive the others' opinions to reach a given state. In this paper, we propose SmartShape, an opinion control package that jointly selects the control users, as well as computes the optimum rate of control messages, thereby driving the networked opinion dynamics to the desired direction. Furthermore, our proposal also includes a robust shaping suit which makes our control framework resilient to stochastic fluctuations of opinion dynamics, originating from several sources of randomness. Experiments on several synthetic and real datasets gathered from Twitter, show that SmartShape can accurately determine the quality of a set of control users as well as shape the opinion dynamics more effectively than several baselines.

ACM Reference Format:
Abir De, Sourangshu Bhattacharya, and Niloy Ganguly. 2018. Shaping Opinion Dynamics in Social Networks. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 9 pages.

1 INTRODUCTION

It is widely accepted that opinions on social networks evolve across time, often influenced by network neighbors that are represented by connecting edges. Modeling such a networked opinion dynamical process has been researched in many years, as evidenced by a plethora of works [1–8]. However, in contrast to these works that consider spontaneous opinion exchange between users, in practice, news or informations are often curated by professional journalists, market design agencies, etc. to alter the opinion of the users in a desired way. One can think of these scenarios as different *opinion shaping* or *opinion control* tasks, in which feeds or news are posted on the wall of a few people called as *control users*, in order to steer the opinions of others to a given state. In this paper, our goal is to identify appropriate control users, and devise an efficient opinion shaping mechanism, in order to curate the overall opinion dynamics in a favorable manner.

Limitations of prior works: Opinion shaping has been studied in different guises mostly by the control theorists [9–13]. Nevertheless, the traditional control approaches towards opinion shaping share several limitations. For example, they emphasize consensus-control, where the opinions of some pre-specified nodes are curated for steering others' opinion to reach consensus. Such a setting has limited applicability in most practical scenarios, as an example, in political campaigns, polarization is often the main goal. Furthermore, most of them assume control opinions as continuous signals, whereas in practice, the expressed opinions are discrete events observed only through the messages or posts. Only very recently Wang *et al.* [14] attempts to overcome these limitations by modeling control signals as discrete epochs, which, however, offers an approximate and computationally inefficient solution.

Barring the individual limitations of these existing approaches, they all have looked into the opinion control problem through the tinted glass of a naive assumption – an apriori and complete specification of the control users which, however in practice, is not known beforehand. Appropriate control users selection is the underlying premise of many information propagation models [15–18]. However, it has been ignored by the existing opinion shaping approaches, despite its compelling practical significance. In fact, the approach in [14] assigns the control signals to each and every node, which in essence means that each user is a control user who governs the opinions of others; consequently, their proposal falls wayside of any practical importance.

In practice, the utility of an opinion shaping strategy depends on two primary factors. (i) Effective selection of control users, and (ii) robustness of the shaping strategy. Control users are usually the influential nodes that are actuated through external feeds or posts which further disseminate across others and shape their opinions. However, since a user's capacity to curate the overall opinion dynamics strongly depends on both the graph structure and the opinion dynamics, computing a direct measure of influence itself is a challenging task. The traditional centrality scores may serve as proxies for influence, however they completely disregard the dynamics of opinion flow.

Besides control user selection, a robust shaping design is another critical need for effective opinion control. In general, opinion evolution follows a stochastic dynamics having multiple sources of randomness like the arrival timings, message sentiments etc. Therefore, the design of the shaping strategy must ensure that its performance is insensitive to such random opinion fluctuations. While the simplistic assumptions like deterministic dynamics, restrictive loss or consensus control are often useful notions for a shaping process to be explainable and tractable, they often hinder the control operation in a complex environment, which renders the shaping strategy fragile and practically ineffective.

Proposed approach: At the very outset, our opinion control approach uses an underlying self exciting point process model called Hawkes process [8, 19–24] which allows users' *actual or organic*

*Now affiliated to MPI for Software Systems, Germany.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

opinions to be modulated over time, by both organic and control opinions from their neighbors, expressed as sentiment messages. Using this underlying dynamics, we design SmartShape, a novel opinion shaping framework that employs a set of control users who regulate others' opinions to fulfill a desired objective, using some additional (control) messages. The objectives of SmartShape are twofolds; (i) optimal selection of control users, and (ii) computation of optimum rates for the control messages. To this aim, the proposed control suit casts the opinion shaping task as a constrained mixed integer programming (MIP) problem, where the binary variables formulate the indicator function of the control users, the objective specifies the precise control task, and the constraints capture various traits of the opinion diffusion model. In contrast to the existing works [9–13, 25] that consider the objective functions having fixed functional forms, our proposal only assumes convex shaping objectives. In a consequence, it can encompass a wide variety of opinion shaping tasks, which enhances its practical utility. Furthermore, such convexity allows the MIP to be relaxed into a convex problem that in turn plummets the computational burden of MIP, thereby providing a quick yet accurate solution. SmartShape can maintain its performance in the presence of random opinion fluctuations, as well as random posting rates. Furthermore, we derive an expression for the opinion covariance at the steady-state, which enables carrying out robust opinion shaping by introducing an additional convex constraint to the utility maximization problems.

We experiment on both synthetic and real data gathered from Twitter to evaluate the utility of our shaping strategies, which show that our framework can accurately determine the quality of a set of control users for a wide variety of opinion shaping tasks. Furthermore through detailed analysis, we point out that the influential nodes are a heterogeneous mix of users with high yet diverse centrality values. Our approach appropriately identifies these key players from different centrality measures, which renders the shaping process quick and effective.

2 MODELING OPINION DYNAMICS

Any opinion shaping strategy has to be developed on top of an opinion dynamical model. A good opinion diffusion model needs to have the following two qualities: (i) It should comply with reality, by demonstrating substantial predictive prowess against other models, and (ii) it should capture various practical facets of opinion dynamics, like consensus, polarization, convergence etc., in a unified way. Several date-driven models e.g. Biased Voter [6], AsLM [7], SLANT [8], etc. fulfil the above two criteria. Among these data-driven models, we choose SLANT [8] as the workhorse of our opinion shaping proposal. SLANT which is equipped with a point process (Hawkes) based probabilistic machinery, can accurately encapsulate the complex stochastic message dynamics. As a result, it offers a substantial performance boost beyond its competitors. In the following, we revisit the formulation of SLANT [8].

2.1 SLANT

In a nutshell, SLANT is driven by three intuitive ideas: (i) The history of messages until time t influences the arrival process of the events after time t, ii) users' opinions are hidden or latent until they decide to share it with their friends (or neighbors); and, iii) users may update their opinions about a particular topic by learning from the opinions shared by their friends,

Setup: Given a directed social network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, we denote $\mathcal{N}(u)$ as the set of users followed by a user u and each post e as $e := (u, m, t)$, where the triplet means that the user $u \in \mathcal{V}$ posted an organic message with sentiment (expressed opinion) m at time t. We denote $\mathcal{H}_u(t) = \{e_i = (u_i, m_i, t_i) | u_i = u \text{ and } t_i < t\}$ as a collection of all messages posted by user u until time t and $\mathcal{H}(t) := \cup_{u \in \mathcal{V}} \mathcal{H}_u(t)$ as the entire history of messages posted by any user until t. From now onwards, we would write $\mathcal{H}(t)$ as $\mathcal{H}_t$ to lighten the notations.

Organic dynamics of opinions and messages. At the outset, 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 until time t [26–32]. Then, we characterize the message rate of user u using the conditional intensity function $\lambda^*_u(t)$ which is associated with the conditional probability of observing an organic message event by user u in infinitesimal time interval $[t, t + dt)]$, given the history $\mathcal{H}(t)$ of posts until time t:

$$\begin{aligned} &\mathbb{P}(\ u \text{ posts a message in } [t, t+dt) | \mathcal{H}(t)) = \lambda^*_u(t)\, dt \ \ \forall u \in \mathcal{V} \\ &\text{i.e. } \mathbb{E}[d\boldsymbol{N}(t)|\mathcal{H}(t)] = \boldsymbol{\lambda}^*(t) dt \end{aligned} \tag{1}$$

where $d\boldsymbol{N}(t) := (\ dN_u(t)\ )_{u \in \mathcal{V}}$ counts the message per user in the interval $[t, t + dt)$ and $\boldsymbol{\lambda}^*(t) := (\ \lambda^*_u(t)\ )_{u \in \mathcal{V}}$ denotes the associated user intensities, which may depend on the history $\mathcal{H}(t)$ (conforming with (i)). The functional form of $\boldsymbol{\lambda}(t)$ is often designed to capture the phenomenon of interest. In this paper, we assume that the messages follow a multivariate Hawkes process where the intensity reflects the bursty nature of posts in social networks due the mutual excitation process between message events [19, 21–24, 33–41]. Here, the intensity depends on the entire history of message events $\cup_{v \in \{u \cup \mathcal{N}(u)\}} \mathcal{H}_v(t)$ before t:

$$\lambda^*_u(t) = \mu_u + b_u \sum_{e_i \in \mathcal{H}_v(t)} \kappa(t - t_i) \tag{2}$$

where the first term, $\mu_u \geqslant 0$, captures the posts by user u on her own initiative, and the second term, with $b_u \geqslant 0$, reflects the influence of previous messages posted by the users she follows have on her intensity.

We represent users' organic latent opinions as a multidimensional stochastic process $\mathbf{x}^*(t)$, where the u-th entry, $x^*_u(t) \in \mathbb{R}$, represents the opinion of user u at time t and the sign * means that it may depend on the history $\mathcal{H}(t)$. Then, every time a user u posts a message at time t, the message sentiment m reflects the expressed opinion, which is a random variable drawn from a distribution $p(m|x^*_u(t))$. The dynamics of $x^*_u(t)$ is given by,

$$x^*_u(t) = \alpha_u + \sum_{v \in \mathcal{N}(u)} a_{vu} \sum_{e_i \in \mathcal{H}_v(t)} m_i g(t - t_i) \tag{3}$$

where the first term, $\alpha_u \in \mathbb{R}$, models the original opinion a user u and the second term, with $a_{vu} \in \mathbb{R}$, models updates in user u's opinion due to the influence that previous messages of her neighbors. Such influence usually depend on node features and other factors [42]. Here, we take $\kappa(t) = e^{-\nu t}$ and $g(t) = e^{-\omega t}$ (where $\nu,\ \omega \geqslant 0$) denote an exponential triggering kernel, which models the decay of influence over time. The organic opinion and message dynamics (Eqs. (2), (3)) can be equivalently written as,

$$\boldsymbol{x}^*(t) = \boldsymbol{\alpha} + \int_0^t g(t-s)\boldsymbol{A}\boldsymbol{m}(s) \odot d\boldsymbol{N}(s) \tag{4}$$

$$\boldsymbol{\lambda}^*(t) = \boldsymbol{\mu} + \int_0^t \boldsymbol{B}\kappa(t-s)d\boldsymbol{N}(s), \tag{5}$$

with $\boldsymbol{A} = (a_{vu})_{u,v\in\mathcal{V}}$, $\boldsymbol{B} = \mathrm{diag}(b_u)$ and $\boldsymbol{x}^*(t) = (x_u^*(t))_{u\in\mathcal{V}}$.

Sentiment distribution. In our work, the sentiment distribution is assumed to be normal. That is, $m \in \mathbb{R}$, *i.e.*, $p(m|x_u(t)) = \mathcal{N}(x_u(t), \sigma_u)$. This fits well scenarios in which sentiment is extracted from text using sentiment analysis [43].

3 CONTROLLED OPINION DYNAMICS

Our opinion shaping design seeks to identify an appropriate set of control users $\mathcal{I}$ who efficiently governs $\boldsymbol{x}^*_{\mathcal{I}^c}(t)$, the opinions of other users $\mathcal{I}^c = \mathcal{V}\backslash\mathcal{I}$ using additional posts, at optimal rates. In order to steer $\boldsymbol{x}^*_{\mathcal{I}^c}(t)$, we aim to shape $\lim_{t\to\infty} \mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*_{\mathcal{I}^c}(t)]$ *i.e.* the steady state behavior of the expected opinion dynamics. Furthermore, to provide a robust shaping design that is resilient to the randomness of opinion flow, we regulate the opinion variance *i.e.* $\mathrm{Tr}(\boldsymbol{\Gamma}_{\mathcal{I}^c})$, where $\boldsymbol{\Gamma}_{\mathcal{I}^c} := \lim_{t\to\infty} \mathbb{E}_{\mathcal{H}_t}[(\boldsymbol{x}^*_{\mathcal{I}^c}(t) - \mathbb{E}_{\mathcal{H}_t}\boldsymbol{x}^*_{\mathcal{I}^c}(t))(\boldsymbol{x}^*_{\mathcal{I}^c}(t) - \mathbb{E}_{\mathcal{H}_t}\boldsymbol{x}^*_{\mathcal{I}^c}(t))^T]$. So, prior to going into the formulation of shaping tasks, it is crucial to characterize these quantities, which we describe next. We begin with formally introducing the control users, then derive the opinion dynamics under the actions of these control users. Finally, we compute differential dynamics of the expected opinion trajectory and the opinion covariance, followed by their stabilities, and the steady state expressions for stable systems under the control actions. These properties will be exploited in the subsequent section to formally devise the opinion shaping strategies.

3.1 Introducing control users

We denote the set of (non) control users as ($\mathcal{I}^c$) $\mathcal{I} \subset \mathcal{V}$. Also for tractability, we assume that these users post binary opinions, i.e. positive or negative messages (±1) with constant rates $\boldsymbol{\eta}^{\pm}_{\mathcal{I}} \in \mathbb{R}^+$. For notational consistency, assume $\boldsymbol{\eta}^{\pm} = \boldsymbol{\eta}^{\pm}_{\mathcal{V}}$. To represent the arrival process of such control messages, we introduce additional counting processes $\boldsymbol{M}^{\pm}(t)$ which regulate the rate of publications of the corresponding opinions ±1, with $\mathbb{E}[d\boldsymbol{M}^{\pm}(t)] = \boldsymbol{\eta}^{\pm}dt$.

One hot representation of $\mathcal{I}$: We define the one-hot representation of the control nodes as $\boldsymbol{\mathcal{S}} = 1[u \in \mathcal{I}]_{u\in\mathcal{V}}$. Using such representation, we can write:

$$\begin{aligned}\boldsymbol{x}^*_{\mathcal{I}^c}(t) &= ((1-\boldsymbol{\mathcal{S}}) \odot \boldsymbol{x}^*(t))_{\{u|\boldsymbol{\mathcal{S}}(u)=0\}},\\ \boldsymbol{\eta}^{\pm}_{\mathcal{I}} &= (\boldsymbol{\mathcal{S}} \odot \boldsymbol{\eta}^{\pm})_{\{u|\boldsymbol{\mathcal{S}}(u)\neq0\}}.\end{aligned} \tag{6}$$

Now, the problem of finding $\mathcal{I}$ becomes equivalent to obtaining $\boldsymbol{\mathcal{S}}$.

Characterization of $\mathcal{I}$: The theoretical characterizations of control users $\mathcal{I}$ is beyond the scope of this work. However, Section 5.1 provides a detailed empirical analysis which reveals that an optimal $\mathcal{I}$ is a set of users with high yet diverse centrality values.

3.2 Opinion dynamics under control actions

In the same sprit of the integral representation of uncontrolled opinion and message dynamics (Eq. (5)), we represent the dynamics of organic opinion and messages in the presence of control processes $\boldsymbol{M}^{\pm}$ as follows:

$$\begin{aligned}\boldsymbol{x}^*_{\mathcal{I}^c}(t) = \boldsymbol{\alpha}_{\mathcal{I}^c} + \int_0^t g(t-s)[&\boldsymbol{A}_1\boldsymbol{m}_{\mathcal{I}^c}(s) \odot d\boldsymbol{N}_{\mathcal{I}^c}(s)\\ &+ \boldsymbol{A}_2 d\boldsymbol{M}^+(s) - \boldsymbol{A}_2 d\boldsymbol{M}^-(s)]\end{aligned} \tag{7}$$

$$\boldsymbol{\lambda}^*_{\mathcal{I}^c}(t) = \boldsymbol{\mu}_{\mathcal{I}^c} + \int_0^t \kappa(t-s)\boldsymbol{B}_1 d\boldsymbol{N}_{\mathcal{I}^c}(s). \tag{8}$$

Here, $\boldsymbol{A}_1 = [a_{vu}]_{v,u\in\mathcal{I}^c}$, $\boldsymbol{A}_2 = [a_{vu}]_{v\in\mathcal{I}^c,u\in\mathcal{I}}$, and $\boldsymbol{B}_1 = \mathrm{diag}(b_u)_{u\in\mathcal{I}^c}$. Similarly one can define $\boldsymbol{\alpha}_{\mathcal{I}^c}$, and $\boldsymbol{\mu}_{\mathcal{I}^c}$. Note that in Eq. (8), the control process $\boldsymbol{M}^{\pm}$ does not influence the message intensities $\boldsymbol{\lambda}^*_{\mathcal{I}^c}(t)$ of the ordinary users. This is due to the diagonal property of $\boldsymbol{B}$ which ensures that the value of $\boldsymbol{\lambda}^*_{\mathcal{I}^c}$ only depends its *own* previous events.

SDE based representation: Given the triggering kernels to be exponential, the resulting opinion and event dynamics under control actions are Markovian, and therefore can be represented as jump stochastic differential equations. This representation will be used in subsequent sections for opinion shaping, where the messages represented by the counting process $\boldsymbol{M}^{\pm}(t)$, will act as the control signals to regulate the dynamics of the organic opinions $\boldsymbol{x}^*(t)$.

Proposition 3.1. *Given the triggering kernel $g(t) = e^{-\omega t}$ and $\kappa(t) = e^{-\nu t}$, the tuple $(\mathbf{x}^*(t), \boldsymbol{\lambda}^*(t))$ following Eqs. (4)- (5), is a Markov process, whose dynamics are defined by the following marked jumped stochastic differential equations (SDE):*

$$\begin{aligned}d\mathbf{x}^*_{\mathcal{I}^c}(t) &= \omega(\boldsymbol{\alpha}_{\mathcal{I}^c} - \mathbf{x}^*_{\mathcal{I}^c}(t))dt + [\boldsymbol{A}_1\boldsymbol{m}_{\mathcal{I}^c}(t) \odot d\boldsymbol{N}_{\mathcal{I}^c}(t)\\ &\quad + \boldsymbol{A}_2 d\boldsymbol{M}^+(t) - \boldsymbol{A}_2 d\boldsymbol{M}^-(t)]\\ d\boldsymbol{\lambda}^*_{\mathcal{I}^c}(t) &= \nu(\boldsymbol{\mu}_{\mathcal{I}^c} - \boldsymbol{\lambda}^*_{\mathcal{I}^c}(t))dt + \mathbf{B}_1\, d\boldsymbol{N}_{\mathcal{I}^c}(t).\end{aligned} \tag{9}$$

The proposition can be easily proved by differentiating Eqs. (4) and (5) respectively. A formal proof is given in [44].

Dynamics of mean and covariance of opinion dynamics: Now, in the following theorems, we provide the trajectories of mean and covariance of opinion dynamics under control actions.

Theorem 3.2 (Dynamics of $\mathbb{E}(\boldsymbol{x}^*_{\mathcal{I}^c}(t))$). *Given the message intensity of each of the non-control users $u \in \mathcal{I}^c = \mathcal{V}\backslash\mathcal{I}$ follows $\lambda_u^*(t) = \mu_u + b_{uu}\sum_{e_i\in\mathcal{H}_u(t)} e^{-\nu(t-t_i)}$, the expected opinion $\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*_{\mathcal{I}^c}(t)]$ of these non-control nodes $\mathcal{I}^c$ follows,*

$$\begin{aligned}\frac{d\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*_{\mathcal{I}^c}(t)]}{dt} = &[-\omega I + \boldsymbol{A}_1\boldsymbol{\Lambda}_{\mathcal{I}^c}(t)]\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*_{\mathcal{I}^c}(t)]\\ &+ \boldsymbol{A}_2(\boldsymbol{\eta}^+_{\mathcal{I}} - \boldsymbol{\eta}^-_{\mathcal{I}}) + \omega\boldsymbol{\alpha}_{\mathcal{I}^c},\end{aligned} \tag{10}$$

with, $\boldsymbol{\Lambda}^*_{\mathcal{I}^c}(t) = \mathrm{diag}(\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{\lambda}^*_{\mathcal{I}^c}(t)])$, *and*

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

The proof follows immediately from Proposition 3.1. A detailed proof-sketch is given in [44].

Theorem 3.3 (Dynamics of covariance matrix $\Gamma_{\mathcal{I}^c}(t)$). *Given the notations used in Theorem 3.2, the dynamics of opinion-covariance of the non-control nodes $\boldsymbol{\Gamma}_{\mathcal{I}^c}(t) := \mathbb{E}_{\mathcal{H}_t}[(\boldsymbol{x}^*_{\mathcal{I}^c}(t) - \mathbb{E}_{\mathcal{H}_t}\boldsymbol{x}^*_{\mathcal{I}^c}(t))(\boldsymbol{x}^*_{\mathcal{I}^c(t)} - \mathbb{E}_{\mathcal{H}_t}\boldsymbol{x}^*_{\mathcal{I}^c}(t))^T]$ is given by:*

$$\begin{aligned}\frac{d\boldsymbol{\Gamma}_{\mathcal{I}^c}(t)}{dt} &= -2\omega\boldsymbol{\Gamma}_{\mathcal{I}^c}(t) + \boldsymbol{\Gamma}_{\mathcal{I}^c}(t)\boldsymbol{\Lambda}_{\mathcal{I}^c}(t)\boldsymbol{A}_1^T + \boldsymbol{A}_1\boldsymbol{\Lambda}_{\mathcal{I}^c}(t)\boldsymbol{\Gamma}_{\mathcal{I}^c}(t)\\ &+ \boldsymbol{A}_1\boldsymbol{\Gamma}_{\mathcal{I}^c\,ii}(t)\boldsymbol{\Lambda}_{\mathcal{I}^c}(t)\boldsymbol{A}_1^T + \sigma^2\boldsymbol{A}_1\boldsymbol{\Lambda}_{\mathcal{I}^c}(t)\boldsymbol{A}_1^T + \boldsymbol{A}_2\,\mathrm{diag}(\boldsymbol{\eta}^+_{\mathcal{I}} + \boldsymbol{\eta}^-_{\mathcal{I}})\boldsymbol{A}_2^T\\ &+ \boldsymbol{A}_1\,\mathrm{diag}(\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*(t)])^2\boldsymbol{\Lambda}_{\mathcal{I}^c}(t)\boldsymbol{A}_1^T\end{aligned} \tag{11}$$

The theorem can be proved by computing the differential of $\Gamma_{\mathcal{I}^c}(t)$, and then using Ito calculus [45] on it. A detailed proof is given in [44].

Stability: Stability is the central challenge for any control strategy design. In the following lemmas proved in supplementary material [44], we investigate stability of mean and covariance dynamics of opinion diffusion under control actions.

Lemma 3.4 (Stability of $\mathbb{E}_{\mathcal{H}_t}(\boldsymbol{x}^*_{\mathcal{I}^c}(t))$). *Given the message dynamics follows Poisson process with $\boldsymbol{\lambda}^*(t) = \boldsymbol{\mu}$, then the expected opinion dynamics is bounded* i.e. $\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*_{\mathcal{I}^c}(t)] < \infty$ *if,* $Re[\xi(\boldsymbol{A}_1\boldsymbol{\Lambda}_{\mathcal{I}^c})] < \omega$, *where $\xi(X)$ is the eigenvalue of X.*

Lemma 3.5 (Stability of $\Gamma_{\mathcal{I}^c}(t)$). *Given the message dynamics follows Poisson process with $\boldsymbol{\lambda}^*(t) = \boldsymbol{\mu}$, then the covariance of opinion dynamics is bounded* i.e. $\Gamma_{\mathcal{I}^c}(t) < \infty$ *iff*

$$\xi[(-\omega I + \boldsymbol{A}_1\boldsymbol{\Lambda}_{\mathcal{I}^c}) \otimes I + I \otimes (-\omega I + \boldsymbol{A}_1\boldsymbol{\Lambda}_{\mathcal{I}^c}) + (\boldsymbol{A}_1 \otimes \boldsymbol{A}_1)\widehat{\boldsymbol{\Lambda}}_{\mathcal{I}^c}] < 0.$$

where $\hat{\Lambda}_{i^2,i^2} = \boldsymbol{\lambda}^*(i)$, $\Lambda := \text{diag}[\mu]$, *and $\otimes$ indicates the Kronecker product.*

Steady state behavior: Here, we discuss the asymptotic properties (expectations and variance as $t \longrightarrow \infty$) of the opinion dynamics in control environment.

Lemma 3.6 (Asymptotic for mean opinion). *If the mean opinion dynamics is stable, then the steady state average opinion of the non-control nodes is given by*

$$\lim_{t\to\infty} \mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*_{\mathcal{I}^c}(t)] = \left(I - \frac{\boldsymbol{A}_1\boldsymbol{\Lambda}_{\mathcal{I}^c}}{\omega}\right)^{-1}\left(\boldsymbol{\alpha}_{\mathcal{I}^c} + \boldsymbol{A}_2\left(\boldsymbol{\eta}^+_{\mathcal{I}} - \boldsymbol{\eta}^-_{\mathcal{I}}\right)\right), \tag{12}$$

where $\boldsymbol{\Lambda}_{\mathcal{I}^c} = \text{diag}\left[I - \frac{\boldsymbol{B}_1}{\nu}\right]^{-1}\boldsymbol{\eta}_{\mathcal{I}^c}$ *and* $\boldsymbol{B}_1 = \text{diag}([b_v]_{v\in\mathcal{I}^c})$. *Eq.* (15) *is equivalent to,*

$$(\boldsymbol{A}_1\Lambda_1 - \omega I)\mathrm{x}_{\mathcal{I}^c} + \boldsymbol{A}_2\left(\boldsymbol{\eta}^+_{\mathcal{I}} - \boldsymbol{\eta}^-_{\mathcal{I}}\right) + \omega\boldsymbol{\alpha}_{\mathcal{I}^c} = 0 \tag{13}$$

where $\boldsymbol{x}_{\mathcal{I}^c} = \lim_{t\to\infty}\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*_{\mathcal{I}^c}(t)]$.
In terms of $\boldsymbol{\mathcal{S}}$, the one-hot representation of $\mathcal{I}$, Eq. (13) *becomes,*

$$\begin{aligned}&(\boldsymbol{A}\boldsymbol{\Lambda} - \omega I)[(\boldsymbol{1} - \boldsymbol{\mathcal{S}}) \odot \boldsymbol{x}] + \boldsymbol{A}[\boldsymbol{\mathcal{S}} \odot (\boldsymbol{\eta}^+ - \boldsymbol{\eta}^-)]\\ &+ \omega(\boldsymbol{1} - \boldsymbol{\mathcal{S}}) \odot \boldsymbol{\alpha} = 0\end{aligned} \tag{14}$$

where $\boldsymbol{x}$ is the steady state organic opinion of all the users.

On steady-state for a stable system, the asymptotic value of expected opinion value does not change. So, $\frac{d\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*_{\mathcal{I}^c}(t)]}{dt} = 0$ as $t \longrightarrow \infty$ which, upon putting on Theorem 3.2, immediately proves the lemma. Now we set about computing $\Gamma_{\mathcal{I}^c}(\infty)$ *i.e.* the covariance of opinions at the steady state.

Lemma 3.7 (Asymptotic for opinion covariance). *Let the steady state solution of the covariance matrix of the non-control nodes for a stable system be* $\boldsymbol{\Gamma}_{\mathcal{I}^c} = \lim_{t\to\infty}\boldsymbol{\Gamma}_{\mathcal{I}^c}(t)$. *Then $\boldsymbol{\Gamma}_{\mathcal{I}^c}$ satisfies the following equation:*

$$\begin{aligned}&-2\omega\boldsymbol{\Gamma}_{\mathcal{I}^c} + \boldsymbol{\Gamma}_{\mathcal{I}^c}\boldsymbol{\Lambda}_{\mathcal{I}^c}\boldsymbol{A}_1^T + \boldsymbol{A}_1\boldsymbol{\Lambda}_{\mathcal{I}^c}\boldsymbol{\Gamma}_{\mathcal{I}^c} + \boldsymbol{A}_1\boldsymbol{\Gamma}_{\mathcal{I}^c ii}\boldsymbol{\Lambda}_{\mathcal{I}^c}\boldsymbol{A}^T\\ &= -\sigma^2\boldsymbol{A}_1\boldsymbol{\Lambda}_{\mathcal{I}^c}\boldsymbol{A}_1^T - \boldsymbol{A}_1\,\text{diag}(\mathrm{x}_{\mathcal{I}^c})^2\boldsymbol{\Lambda}_{\mathcal{I}^c}\boldsymbol{A}_1^T\\ &- \boldsymbol{A}_2\,\text{diag}(\boldsymbol{\eta}^+_{\mathcal{I}} + \boldsymbol{\eta}^-_{\mathcal{I}})\boldsymbol{A}_2^T,\end{aligned} \tag{15}$$

In terms, $\boldsymbol{\mathcal{S}}$, the above equation becomes

$$\begin{aligned}&-2\omega\boldsymbol{\Gamma}_{\boldsymbol{\mathcal{S}}\boldsymbol{\mathcal{S}}} + \boldsymbol{\Gamma}_{\boldsymbol{\mathcal{S}}\boldsymbol{\mathcal{S}}}\boldsymbol{\Lambda}\boldsymbol{A}^T + \boldsymbol{A}\boldsymbol{\Lambda}\boldsymbol{\Gamma}_{\boldsymbol{\mathcal{S}}\boldsymbol{\mathcal{S}}} + \boldsymbol{A}\boldsymbol{\Gamma}^{ii}_{\boldsymbol{\mathcal{S}}\boldsymbol{\mathcal{S}}}\boldsymbol{\Lambda}\boldsymbol{A}\\ &= -\sigma^2\boldsymbol{A}\boldsymbol{\Lambda}\boldsymbol{A}^T - \boldsymbol{A}\,\text{diag}((1 - \boldsymbol{\mathcal{S}})\boldsymbol{x})^2\boldsymbol{\Lambda}\boldsymbol{A}^T\\ &- \boldsymbol{A}\,\text{diag}[\boldsymbol{\mathcal{S}} \odot (\boldsymbol{\eta}^+ + \boldsymbol{\eta}^-)]\boldsymbol{A}^T\end{aligned} \tag{16}$$

where $\boldsymbol{\Gamma}_{\boldsymbol{\mathcal{S}}\boldsymbol{\mathcal{S}}} := (\boldsymbol{I} - \text{diag}[\boldsymbol{\mathcal{S}}])\boldsymbol{\Gamma}(\boldsymbol{I} - \text{diag}[\boldsymbol{\mathcal{S}}])$ *and* $\boldsymbol{\Gamma}^{ii}_{\boldsymbol{\mathcal{S}}\boldsymbol{\mathcal{S}}} = \text{diag}(\boldsymbol{\Gamma}_{\boldsymbol{\mathcal{S}}\boldsymbol{\mathcal{S}}})$.

The proof of this lemma directly comes from Theorem 3.3 by putting $d\boldsymbol{\Gamma}_{\mathcal{I}^c}/dt \longrightarrow 0$ as $t \longrightarrow \infty$. The steady state covariance matrix obtained by solving Eq (15) has a closed form solution, is given below.

Lemma 3.8. *The closed form $\boldsymbol{\Gamma}_{\mathcal{I}^c}$ is given by,*

$$\begin{aligned}vec(\boldsymbol{\Gamma}_{\mathcal{I}^c}) &= [(-\omega I + \boldsymbol{A}_1\boldsymbol{\Lambda}_{\mathcal{I}^c}) \otimes I + I \otimes (-\omega I + \boldsymbol{A}_1\boldsymbol{\Lambda}_{\mathcal{I}^c})^T\\ &+(\boldsymbol{A}_1 \otimes \boldsymbol{A}_1)\hat{\boldsymbol{\Lambda}}_{\mathcal{I}^c}]^{-1} \times vec[\sigma^2\boldsymbol{A}_1\boldsymbol{\Lambda}_{\mathcal{I}^c}\boldsymbol{A}_1^T\\ &+ \boldsymbol{A}_1\,\text{diag}(\mathbb{E}_{\mathcal{H}(\infty)\setminus\mathcal{H}(t_0)}[\boldsymbol{x}^*_{\mathcal{I}^c}(\infty)|\mathcal{H}_{t_0}])^2\boldsymbol{\Lambda}_{\mathcal{I}^c}\boldsymbol{A}_1^T]\end{aligned} \tag{17}$$

Here, $\text{vec}(\boldsymbol{X})$ *denotes the vectorization of a matrix $\boldsymbol{X}$.*

If we vectorize the matrix equation (15) using the fact $\text{vec}(AXB) = (B^T \otimes A)\text{vec}(X)$, we directly have the closed form of $\text{vec}(\boldsymbol{\Gamma}_{\mathcal{I}^c})$.

Eqs. (14) and (16) provide us the relationships between the selection vector, the organic and control message intensities of the users, the initial opinions, the steady-state opinions, and the covariance matrix. In the next section, we exploit these relationships to design a diverse range of convex programs to determine the message intensities of the control users in order to achieve a target steady-state opinion for the non-control ones, given their initial opinions.

4 EFFICIENT OPINION SHAPING

In this section, we develop the opinion control framework SmartShape, by exploiting the underlying dynamics of opinion flow in the presence of control actions. In a nutshell, SmartShape encompasses a diverse range of opinion shaping tasks as mixed integer programming problems that are further relaxed into a set of convex utility maximization problems that can be efficiently solved. Given an opinion model *i.e.* the knowledge of all the parameters $\boldsymbol{\alpha}, \boldsymbol{\mu}, \boldsymbol{A}, \boldsymbol{B}$, and the exact shaping task, we compute the selection vector $\boldsymbol{\mathcal{S}}$ (thus $\mathcal{I}$ and $\mathcal{I}^c$) and the control rates $\boldsymbol{\eta}^{\pm}$ that steer the steady-state opinions of the non-control users $\boldsymbol{x}$ by maximizing a utility function $U(.)$, a concave function of the opinions of non-control users. This utility function formally specifies the underlying shaping objective. Additionally, we associate costs with the control posts by the selected users. If the cost $\boldsymbol{c} = (c_1, c_2, \ldots, c_{|\mathcal{V}|})^T \geq 0$ is cost per control event, and C is the total budget, then we impose

$$\boldsymbol{c}^T\left(\boldsymbol{\eta}^+ + \boldsymbol{\eta}^-\right) \odot \boldsymbol{\mathcal{S}} \leq C, \text{and } \boldsymbol{\eta}^{\pm} \geq 0. \tag{18}$$

Finally, with the above impositions (Eq. (18)), we present different levels of SmartShape-Basic.

SmartShape-Basic. *Utility maximization without covariance control:* The basic form of SmartShape does not aim to regularize the variance, rather simply maximize a given utility function of opinions of organic users $\mathcal{I}^c$.

$$\underset{\boldsymbol{x},\,\boldsymbol{\mathcal{S}},\,\boldsymbol{\eta}^{\pm}}{\text{maximize}} \quad U((\boldsymbol{1} - \boldsymbol{\mathcal{S}}) \odot \boldsymbol{x}) \text{ so that, (14), and (18)} \tag{19}$$

However, this problem is MIP due to the presence of $\boldsymbol{\mathcal{S}} \in \{0, 1\}^{|\mathcal{V}|}$. We adopt the *linearization of binary variables* method described in [46–48]. In particular the linearization technique approximates each of $(1-\mathcal{S})\odot\boldsymbol{x}$, and $\mathcal{S}\odot\eta^{\pm}$ by four constraints. A cheatsheet of such techniques may also found in this blog [1].

SmartShape-Basic:

$$\underset{z,\boldsymbol{x},\boldsymbol{\eta}^{\pm},\xi^{\pm},\boldsymbol{\mathcal{S}}}{\textbf{maximize}} \quad U(z) \tag{20}$$

subject to: $(\boldsymbol{A}\boldsymbol{\Lambda}-\omega I)\boldsymbol{z}+\boldsymbol{A}\left(\xi^{+}-\xi^{-}\right)]+\omega(\boldsymbol{1}-\boldsymbol{\mathcal{S}})\odot\boldsymbol{\alpha}=0$

$\boldsymbol{c}^{T}\left(\xi^{+}+\xi^{-}\right)\leq C,\quad \xi^{+}\geq 0\quad \xi^{-}\geq 0$

% Linearization of $(1-\mathcal{S})\odot\boldsymbol{x}$	% Linearization of $\mathcal{S}\odot\eta^{\pm}$
$\underline{\boldsymbol{x}}\leq \boldsymbol{z}\leq\overline{\boldsymbol{x}}$	$\boldsymbol{c}\odot\xi^{\pm}\leq C\boldsymbol{\mathcal{S}}$
$\underline{\boldsymbol{x}}\odot(\boldsymbol{1}-\boldsymbol{\mathcal{S}})\leq \boldsymbol{z}\leq\overline{\boldsymbol{x}}\odot(\boldsymbol{1}-\boldsymbol{\mathcal{S}})$	$0\leq\xi^{\pm}\leq\boldsymbol{\eta}^{\pm}$
$\boldsymbol{x}-\boldsymbol{\mathcal{S}}\odot\overline{\boldsymbol{x}}\leq\boldsymbol{z}\leq\boldsymbol{x}-\boldsymbol{\mathcal{S}}\odot\underline{\boldsymbol{x}}$	$\boldsymbol{c}\odot\boldsymbol{\eta}^{\pm}-(\boldsymbol{1}-\boldsymbol{\mathcal{S}})C\leq\boldsymbol{c}\odot\xi^{\pm}$
$\boldsymbol{z}\leq\boldsymbol{x}+\boldsymbol{\mathcal{S}}\odot\overline{\boldsymbol{x}}$	$0\leq\boldsymbol{\mathcal{S}}\leq 1$

Here, we add an L_1 regularizer and tune the regularizer value to control the number of control nodes. The last two sets of constraints, approximate the non-convex product-terms ($(1-\boldsymbol{\mathcal{S}})\odot\boldsymbol{x}$ and $\boldsymbol{\mathcal{S}}\odot\boldsymbol{\eta}^{\pm}$). Clearly, as long as the utility function is concave, the opinion shaping problem defined by Eq. 20 is jointly convex in $\boldsymbol{z},\boldsymbol{x},\boldsymbol{\eta}^{\pm},\xi^{\pm}$, and $\boldsymbol{\mathcal{S}}$. Hence, the global optimum can be found by many well-known algorithms.

SmartShape-Centrality. *Utility Maximization with a pre-selected $\mathcal{I}$:* One can also assume $\mathcal{I}$, the set of control nodes, is already chosen using some centrality measures like PageRank, degree, closeness centralities, etc. Indeed we shall use such setting as a baseline to compare with the proposed user selection strategy of SmartShape-Basic (see Section 5.1).

$$\underset{\boldsymbol{x}_{\mathcal{I}^c},\boldsymbol{\eta}^{\pm}}{\text{maximize}} \quad U(\mathbf{x}_{\mathcal{I}^c}) \text{ subject to, (12) and (18)} \tag{21}$$

SmartShape-Robust. *Utility maximization with covariance control:* The problem formulation of SmartShape-Basic aims to steer the *average* users' opinions to a given steady-state. However, particular realizations (or trajectories) of the model may converge to a steady-state opinion *far* from the steady-state average opinion, $\lim_{t\to\infty}\mathbb{E}_{\mathcal{H}_t}[\boldsymbol{x}^*_{\mathcal{I}^c}(t)]$. Here, we reformulate the above opinion shaping problems to ensure that particular realizations of the model are typically *close* to the average, by adding a regularizer on the steady-state opinion covariance of controlled-nodes quantified as,

$$\begin{aligned}&\underset{\boldsymbol{x},\,\boldsymbol{\mathcal{S}},\,\boldsymbol{\eta}^{\pm},\,\Gamma_{\mathcal{SS}}}{\text{maximize}} && U((1-\boldsymbol{\mathcal{S}})\odot\mathbf{x})-\gamma\operatorname{Tr}(\Gamma_{\mathcal{SS}})\\ &\text{subject to} && \text{(14), (16), (18), and } \boldsymbol{\Gamma}_{\mathcal{SS}}\succeq 0.\end{aligned} \tag{22}$$

Using the linearization method as Eq. (20), we obtain :

SmartShape-Robust :

$$\underset{z,\boldsymbol{x},\boldsymbol{\eta}^{\pm},\xi^{\pm},\boldsymbol{\mathcal{S}},\Gamma_{\mathcal{SS}}}{\textbf{maximize}} \quad U(z)-\gamma\operatorname{Tr}(\Gamma_{\mathcal{SS}}) \tag{23}$$

subject to: $(\boldsymbol{A}\boldsymbol{\Lambda}-\omega I)\boldsymbol{z}+\boldsymbol{A}\left(\xi^{+}-\xi^{-}\right)]+\omega(\boldsymbol{1}-\boldsymbol{\mathcal{S}})\odot\boldsymbol{\alpha}=0$

$2\omega\Gamma_{\mathcal{SS}}+\Gamma_{\mathcal{SS}}\boldsymbol{\Lambda}\boldsymbol{A}^{T}+\boldsymbol{A}\boldsymbol{\Lambda}\Gamma_{\mathcal{SS}}+\Gamma^{ii}_{\mathcal{SS}}$
$=-\sigma^{2}\boldsymbol{A}\boldsymbol{\Lambda}\boldsymbol{A}^{T}-\boldsymbol{A}\operatorname{diag}(z)^{2}\boldsymbol{\Lambda}\boldsymbol{A}^{T}-\boldsymbol{A}\operatorname{diag}[(\xi^{+}+\xi^{-})]\boldsymbol{A}^{T}$

$\boldsymbol{c}^{T}(\xi^{+}+\xi^{-})\leq C,\quad \xi^{\pm}\geq 0,\quad \Gamma\succeq 0$

% Linearization of $(1-\mathcal{S})\odot\boldsymbol{x}$	% Linearization of $\mathcal{S}\odot\eta^{\pm}$
$\underline{\boldsymbol{x}}\leq \boldsymbol{z}\leq\overline{\boldsymbol{x}}$	$\boldsymbol{c}\odot\xi^{\pm}\leq C\boldsymbol{\mathcal{S}}$
$\underline{\boldsymbol{x}}\odot(\boldsymbol{1}-\boldsymbol{\mathcal{S}})\leq \boldsymbol{z}\leq\overline{\boldsymbol{x}}\odot(\boldsymbol{1}-\boldsymbol{\mathcal{S}})$	$0\leq\xi^{\pm}\leq\boldsymbol{\eta}^{\pm}$
$\boldsymbol{x}-\boldsymbol{\mathcal{S}}\odot\overline{\boldsymbol{x}}\leq\boldsymbol{z}\leq\boldsymbol{x}-\boldsymbol{\mathcal{S}}\odot\underline{\boldsymbol{x}}$	$\boldsymbol{c}\odot\boldsymbol{\eta}^{\pm}-(\boldsymbol{1}-\boldsymbol{\mathcal{S}})C\leq\boldsymbol{c}\odot\xi^{\pm}$
$\boldsymbol{z}\leq\boldsymbol{x}+\boldsymbol{\mathcal{S}}\odot\overline{\boldsymbol{x}}$	$0\leq\boldsymbol{\mathcal{S}}\leq 1$

[1] https://www.leandro-coelho.com/linearization-product-variables/

Since, the second equality constraint in (23) is not a convex constraint with respect to z, the optimization problem, as it is, may seem very difficult to solve efficiently. However, we can overcome this difficulty by following Lemma 3.8, and looking into the explicit expression of $\boldsymbol{\Gamma}_{\mathcal{SS}}$. Note that the objective function only depends on the trace of $\boldsymbol{\Gamma}_{\mathcal{SS}}$. To compute that, we add an additional variable t and reconstruct the second equality constraint as,

$$\begin{aligned}t\geq&-(\operatorname{vec}(\boldsymbol{I}))^{T}\boldsymbol{V}^{-1}\operatorname{vec}(\sigma^{2}\boldsymbol{A}\boldsymbol{\Lambda}\boldsymbol{A}^{T}+\boldsymbol{A}\operatorname{diag}(z)^{2}\boldsymbol{\Lambda}\boldsymbol{A}^{T}\\&+\boldsymbol{A}\operatorname{diag}\left(\xi^{+}+\xi^{-}\right)\boldsymbol{A}^{T}).\end{aligned} \tag{24}$$

$$\boldsymbol{V}=(-\omega I+\boldsymbol{A}\boldsymbol{\Lambda})\otimes I+I\otimes(-\omega I+\boldsymbol{A}\boldsymbol{\Lambda})^{T}+\sum_{i=1}^{n}(\boldsymbol{A}\otimes\boldsymbol{A}^{T})\hat{\boldsymbol{\Lambda}}$$

Such transformation (Eq. (24)) makes objective function $U(z)-\gamma t$. Note that, at the optimal condition, we have exact equality in Eq. (24). Finally we have $\boldsymbol{z},\boldsymbol{x},\boldsymbol{\eta}^{\pm},\xi^{\pm},\boldsymbol{\mathcal{S}}$ and t as the optimization variables. Note that, since all the elements in $\operatorname{diag}(z)^{2}$ are multiplied by positive constants, the inequality (24) is convex with respect to z and thus the problem is jointly convex in all variables.

4.1 Instances of utility functions

Minimax opinion shaping [MMOSH-1, MMOSH-2]: Suppose we aim to make the users with the most positive opinion as negative as possible. Then, we choose

$$\text{MMOSH-1:}\qquad U(\mathbf{x}_{\mathcal{I}^c})=-\max_{u\in[m]}x_{\mathcal{I}^c,u}. \tag{25}$$

Alternatively, in order to make the users with the most negative opinions as positive as possible, we write the utility function as

$$\text{MMOSH-2:}\qquad U(\mathbf{x}_{\mathcal{I}^c})=\min_{u\in[m]}x_{\mathcal{I}^c,u}. \tag{26}$$

Average opinion shaping [AOSH-1, AOSH-2] Suppose we aim to make the average opinions over users to be as low as possible. Then, we can choose:

$$\text{AOSH-1:}\qquad U(\mathbf{x}_{\mathcal{I}^c})=-\sum_{u\in\mathcal{I}^c}x_{u}. \tag{27}$$

On the other hand, we also take the following utility:

$$\text{AOSH-2:}\qquad U(\mathbf{x}_{\mathcal{I}^c})=\sum_{u\in\mathcal{I}^c}x_{u}. \tag{28}$$

Top-k opinion shaping [Top-k–OSH-{1,2}] Suppose we aim to make the k users with the most positive opinion as negative as possible. Then, we need to introduce an extra variable $\mathbf{y}_{\mathcal{I}^c}$, add the additional constraints

$$\text{Top-}k\text{–OSH-1:}\ \ y_{\mathcal{I}^c,u}\geq\max(x_{\mathcal{I}^c,u},0)\quad\forall u\in\mathcal{I}^{c}, \tag{29}$$

$$\text{With, Top-}k\text{–OSH-\{1,2\}:}\qquad U(\mathbf{y}_{\mathcal{I}^c})=-\sum_{i=1}^{k}|y_{[i]}|, \tag{30}$$

where $y_{[i]}$ denotes the ith largest component of y. Similarly, suppose our goal is to make the k users with the most negative opinion as positive as possible. Then, we introduce the extra variable $\mathbf{y}_{\mathcal{I}^c}$, add the additional constraints

$$\text{Top-}k\text{–OSH-2:}\qquad y_{\mathcal{I}^c,u}\leq\min(x_{\mathcal{I}^c,u},0)\quad\forall u\in\mathcal{I}^{c}, \tag{31}$$

and use the utility function given by Eq. 30.

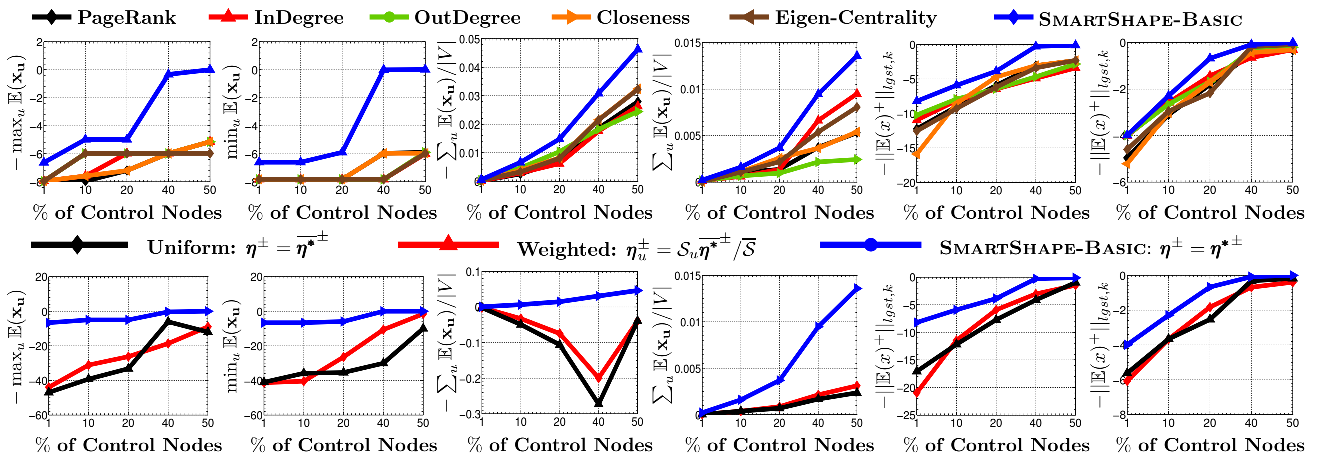

(a) MMOSH-1 (b) MMOSH-2 (c) AOSH-1 (d) AOSH-2 (e) Top-k-OSH-1 (f) Top-k-OSH-1

Figure 1: Opinion shaping on Kronecker Hierarchical (first three columns– (a) to (c)) and Politics (last three columns– (d) to (f)) networks. The top row indicates the impact of SmartShape-Basic from user selection perspective, as compared to SmartShape-Centrality where the control users are chosen based on some centrality measures like PageRank, InDegree, OutDegree, and Closeness. The bottom row indicates the efficacy of SmartShape-Basic from rate computation viewpoint, against various heuristics like "Uniform" and "Weighted".

Datasets	$\lvert\mathcal{V}\rvert$	$\lvert\mathcal{E}\rvert$	$\lvert\mathcal{H}(T)\rvert$	$\mathbb{E}[\mathbf{m}]$	**std**$[m]$
Politics	548	5271	20026	0.0169	0.1780
Movie	567	4886	14016	0.5969	0.1358
Fight	848	10118	21526	-0.0123	0.2577
Bollywood	1031	34952	46845	0.5101	0.2310
US	533	20067	18704	-0.0186	0.7135

Table 1: Real datasets statistics

5 EXPERIMENTS

We provide a comprehensive evaluation of SmartShape using both synthetic data and real data gathered from Twitter, and assess to which extent the chosen control nodes can steer the users' opinions to a given state in several types of networks with very different structures.

Synthetic datasets: We evaluate the accuracy of our shaping strategies on five types of Kronecker networks [49]: i) Homophily networks (parameter matrix $[0.96, 0.3; 0.3, 0.96]$), ii) heterophily networks ($[0.3, 0.96; 0.96, 0.3]$), iii) core-periphery networks ($[0.9, 0.5; 0.5, 0.3]$), (iv) random networks ($[0.7, 0.7; 0.7, 0.7]$); and, (v) hierarchical networks ($[0.9, 0.1; 0.1, 0.9]$). For each network, the message intensities are multivariate Hawkes, $\boldsymbol{\mu}$ and $\boldsymbol{B}$ are drawn from a uniform distribution $U(0, 1)$, and $\boldsymbol{\alpha}$ and $\boldsymbol{A}$ are drawn from a Gaussian distribution $\mathcal{N}(\mu = 0, \sigma = 1)$. Also, we use exponential kernels with parameters $\omega = 100$ and $\nu = 1$.

Real datasets: We use five real datasets [8] corresponding to various real world events, collected from Twitter, which are also summarized in Table 1. For each of these datasets, we have a social network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, and a collection of messages $\mathcal{H}(T) = \{(u_i, m_i, t_i)\}$ gathered during a time period $[0, T)$. Using this information as input, we obtain the optimal parameters $\boldsymbol{\alpha}$, $\boldsymbol{\mu}$, $\boldsymbol{A}$ and $\boldsymbol{B}$ by maximizing the corresponding likelihood function w.r.t. $\boldsymbol{\alpha}$, $\boldsymbol{\mu}$, $\boldsymbol{A}$, $\boldsymbol{B}$ [8, 50].

$$\sum_{e_i \in \mathcal{H}(T)} \log p(m_i | x^*_{u_i}(t_i)) + \sum_{e_i \in \mathcal{H}(T)} \log \lambda^*_{u_i}(t_i) - \sum_{u \in \mathcal{V}} \int_0^T \lambda^*_u(\tau)\, d\tau$$

Setup: With these parameters, we run SmartShape both in synthetic and real networks. We focus on six tasks (Eqs. (25) to (31)). For each of these tasks we find the top control nodes, and assess the performance exhibited by these nodes of different numbers. In all experiments, we set the total budget to $C = 10$, and assume all users entail the same cost. More in detail, we solve SmartShape-Basic (Eq. 20) and SmartShape-Robust (Eq. 23). We adjust the regularizer over $\boldsymbol{\mathcal{S}}$ to vary the number of control nodes over $(1, 10, 20, 40, 50)\%$ of $|\mathcal{V}|$. The performance of the control strategy is measured using the U, the value of the underlying objective function. We compare the efficacy of our user selection strategy, as well as the message rate computation against several competitors.

Baselines: To probe the utility of the proposed user selection strategy, we compare SmartShape-Basic with SmartShape-Centrality, where for the latter we fix $\mathcal{I}$ with the top $x\%$ ($x = 1, 10, 20, 40, 50$) nodes based on several centrality measures e.g. PageRank, in-degree, out-degree, closeness, and eigencentrality. On the other hand, we compare the efficacy of message rate computation with two different heuristics, e.g. *Uniform*, and *Weighted*. Here, "Uniform" indicates that the rate of each user is constant and equals to the average of the individual optimum rates, *i.e.* $\eta^{\pm}_u = \sum_{i \in \mathcal{I}^c} \eta^{\pm}_i / |\mathcal{I}^c|$. "Weighted" assumes that the message rate of each user u is proportional to her selection score $\mathcal{S}_u$, subject to the same aggregated message rate *i.e.* $\eta^{\pm}_u = \mathcal{S}_u \boldsymbol{\eta}^{*\pm}_u / \sum_u \mathcal{S}_u$.

5.1 Comparison with baselines

The primary goals of SmartShape are (i) optimal control user selection, and (ii) optimal message rate selection. We evaluate the proposed shaping strategies from these two perspectives.

Utility of optimal selection of $\mathcal{I}$. To evaluate the impact of influential user selection strategy, we compare SmartShape-Basic (Eq. (20)) with SmartShape-Centrality (Eq. (21)), where the

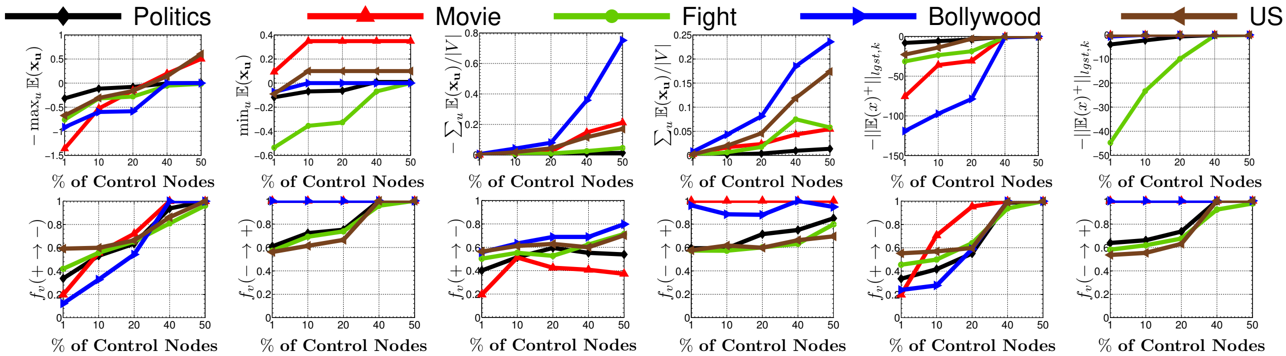

(a) MMOSH-1 (b) MMOSH-2 (c) AOSH-1 (d) AOSH-2 (e) Top-k-OSH-1 (f) Top-k-OSH-2

Figure 2: Performance variation of SmartShape-Basic with % of control nodes. The top row shows the variation of actual objective U, and the bottom row shows variation of f_v, % of users who switched their polarity in the desired direction.

latter, the control nodes are pre-selected using well-known measures of centrality such as PageRank, degree, closeness or eigencentrality. Top row of Figure 1 summarizes the results which indicate that our node-selection strategy steers the objective more efficiently than other centrality-measures. It shows that network-structure is not the only criteria for control-node selection. Rather one should also consider the utility function and the complex dynamics of opinion evolution. Our proposal combines these three factors in a unified-way. As a result, it significantly outperforms the existing baselines.

Utility of computation of optimal message rate $\eta^{*\pm}$. To understand the utility of SmartShape from control rate computation viewpoint, we compare SmartShape-Basic (Eq. (20)) with different rate selection strategies e.g. *Uniform* and *Weighted*, as described in the last subsection. The bottom row of Figure 1 summarizes the results. We observe that SmartShape steers the objective more effectively than others, indicating that it appropriately incorporates the network structure, the utility function and the complex opinion dynamics, that are not effectively combined the other two heuristics.

5.2 Characterization of control node set $\mathcal{I}^*$

The control users should have a high influence over other users to make the steering task quick and effective. Here, we aim to characterize the optimally selected user set $\mathcal{I}^*$, through various structural properties (Figure 3). In order to do that, we find out the overlap of $\mathcal{I}^*$ with top important users selected using standard centrality measures. In particular, we measure the Jaccard coefficient between $\mathcal{I}^*$ and $\mathcal{I}_{\text{Centrality}}$, $|\mathcal{I}^* \cap \mathcal{I}_{\text{Centrality}}|/|\mathcal{I}^* \cup \mathcal{I}_{\text{Centrality}}|$, with fixed number of control users. Here, $\mathcal{I}_{\text{Centrality}}$ is the set of control users chosen based on a centrality score e.g. PageRank, Degree, etc., with $|\mathcal{I}^*| = \mathcal{I}_{\text{Centrality}}|$. Figure 3 shows the variation of this overlap measure with the number of control nodes. For the bar named "All", the overlap is given by $|\mathcal{I}^* \cap \mathcal{I}_{\text{All}}|/|\mathcal{I}^* \cup \mathcal{I}_{\text{All}}|$ where $\mathcal{I}_{\text{All}} = \cup_{\text{Centrality}} \mathcal{I}_{\text{Centrality}}$. So, the bar "All" shows the overlap of the optimally chosen control users with *any* of the centrality based control user set. Figure 3 reveals three key observations. (i) $\mathcal{I}^*$ has a lot of overlaps with various centrality measures. So the control users play strong yet *different* structurally influential roles in the network. (ii) The bar named as "All" has a high value, which indicates that $\mathcal{I}^*$ is a heterogeneous blend of nodes from various centrality measures, that is accurately identified by our proposal. (iii) In most cases, the out-degree centrality has the highest overlap with $\mathcal{I}^*$, indicating that it is the key measure of influence in the context of opinion shaping.

5.3 Performance variation with $|\mathcal{I}|$

Variation of actual objective. Figure 2 shows that the performance of SmartShape-Basic for all shaping problems becomes better, as the number of control nodes increases. The shaping framework works quite effectively for US dataset (consistently second best for all shaping tasks except MMOSH-1), because it has the high average degree among all real datasets, which makes the influence of the control nodes accessible to a large number of nodes.

Change of polarity. Intuitively, by minimizing the most positive opinion across users (MMOSH-1) or the average opinions (ASOH-1) or the opinions of top-k most positive users (Top-k-OSH-1), one could expect (some) users to switch from an initial positive opinion to a steady-state negative opinion. Similarly, by maximizing the most negative opinion across users (MMOSH-2), the average opinions of the users (AOSH-2) or the opinions of the top-k most negative users (Top-k-OSH-2), one could expect (some) users to switch from an initial negative opinion to a steady-state positive opinion. Figure 2 confirms this intuition, by showing the percentage of nodes that change their opinion polarity against control node set sizes, for all the real datasets.

5.4 Analysis of SmartShape-Robust

Finally, we solve SmartShape-Robust, the same activity shaping tasks with covariance control (Eq. 23) for different penalty values γ. Figures 4, and 5 summarize the results. Figure 5 shows the variation of objective with number of nodes, where it is evidently clear that more number of control nodes help to reduce the opinion variance. Figure 4 shows the variation of objective with regularizer (γ). It shows a clear trade-off between variance and objective value. In case of MMOSH-1 and MMOSH-2 (Figures 4a and 4g), we observe the shaping performance is most robust, i.e. the performance does not change much with regularizer variation. In a few cases like AOSH-2 (Figures 4d and 4d), we observe that Bollywood suffers a higher variance. It is because, the natural activity level of the users in Bollywood is very high (See $|\mathcal{H}(t)|$ in Table 1) which is injecting an inherent high variance (through a higher $\mathbf{\Lambda}$) in the system.

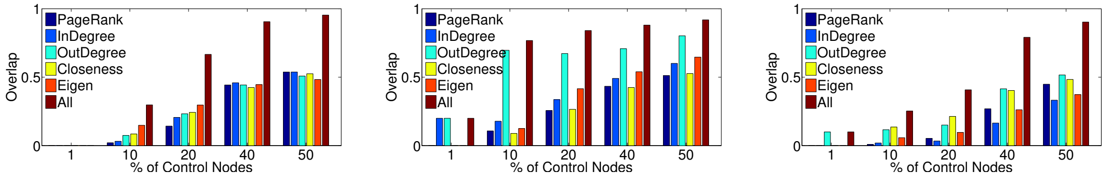

(a) Politics **(b) Movie** **(c) Bollywood**

Figure 3: Characterization of $\mathcal{I}^*$, the control users spotted by SmartShape-Basic, for three real datasets. The structural traits of $\mathcal{I}^*$ are measured using their overlap with users of top centrality scores. The overlap is measured as Jaccard coefficient between $\mathcal{I}^*$ and $\mathcal{I}_{\text{Centrality}}$, where the latter is the control users chosen based on a centrality score e.g. PageRank, Degree, etc. The bar "All" indicates the overlap of the optimal users with *any* of the centrality based user-set.

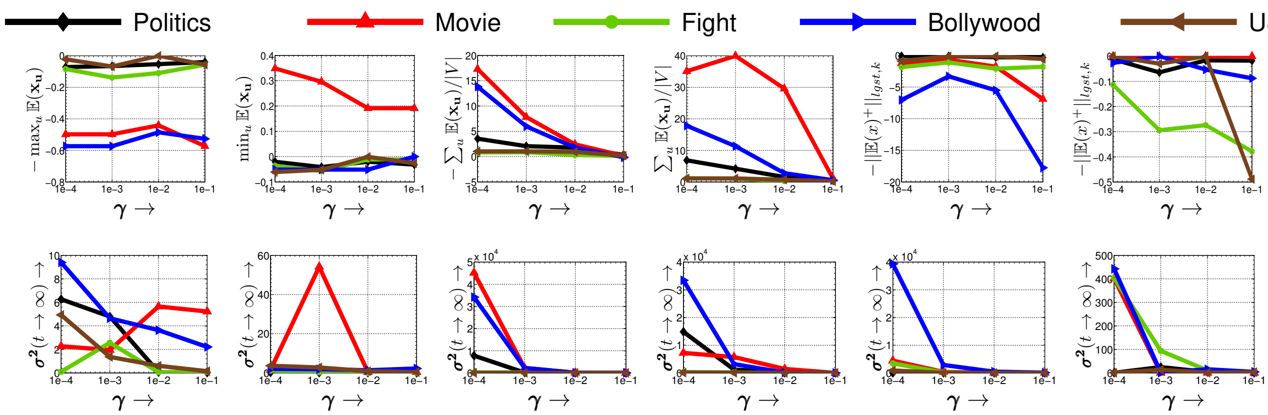

(a) MMOSH-1 **(b) MMOSH-2** **(c) AOSH-1** **(d) AOSH-2** **(e) Top-k-OSH-1** **(f) Top-k-OSH-2**

Figure 4: Performance variation of SmartShape-Robust, opinion shaping with covariance control, with respect to the regularizer γ. The top (bottom) row shows the variation of control objective (variance).

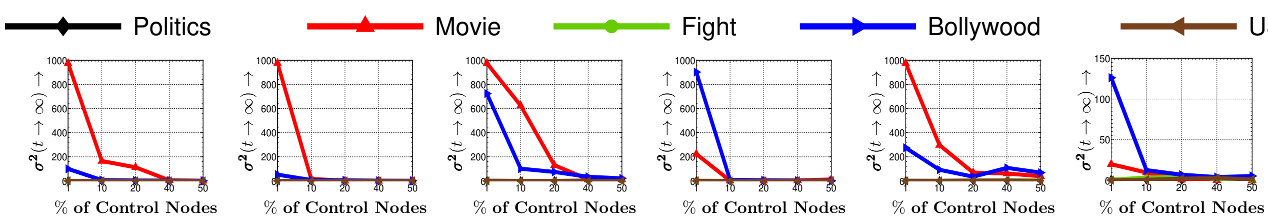

(a) MMOSH-1 **(b) MMOSH-2** **(c) AOSH-1** **(d) AOSH-2** **(e) Top-k-OSH-1** **(f) Top-k-OSH-2**

Figure 5: Variation of opinion variance for SmartShape-Robust with respect to % of control nodes.

6 CONCLUSIONS

Our principal contribution in this paper lies in designing a flexible opinion shaping framework, that employs the *best* set of k control users who can steer the opinions of others, using additional messages. To this aim, we proposed SmartShape, an novel opinion control suit which jointly selects an optimal set of influential users, as well as the rate of their messages in a unified way. SmartShape encompasses a wide variety of opinion shaping tasks addressing different shaping objectives in various practical scenarios. Furthermore, it is resilient to stochastic opinion fluctuations, as well as random message arrivals. Experiments on several synthetic and real data gathered from Twitter showed that SmartShape can identify an appropriate set of influential users who can shape the networked opinion dynamics in a desired way. Furthermore, a detailed empirical analysis revealed that these control users are a heterogeneous collection of users having high centrality values, but from different centrality measures. We believe that, apart from many immediate applications like marketing design, product campaign, etc., our proposal can also be used to understand the dynamics of fake news that are spread by many incentivized users, as well as to design counter-measures against it.

Acknowledgment: The authors acknowledge Manuel Gomez Rodriguez and Isabel Valera for useful discussions. N. Ganguly and S. Bhattacharya acknowledge an Intel grant.

REFERENCES

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

[2] 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.

[3] Frank Schweitzer and Laxmidhar Behera. Nonlinear voter models: the transition from invasion to coexistence. *The European Physical Journal B-Condensed Matter and Complex Systems*, 67(3):301–318, 2009.

[4] 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.

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

[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] Abir De, Isabel Valera, Niloy Ganguly, Sourangshu Bhattacharya, and Manuel Gomez Rodriguez. Learning and forecasting opinion dynamics in social networks. In *NIPS*. 2016.

[9] Jan Christian Dittmer. Consensus formation under bounded confidence. *Nonlinear Analysis: Theory, Methods & Applications*, 47(7):4615–4621, 2001.

[10] Seyed Rasoul Etesami and Tamer Başar. Game-theoretic analysis of the hegselmann-krause model for opinion dynamics in finite dimensions. *IEEE Transactions on Automatic Control*, 60(7):1886–1897, 2015.

[11] Vincent D Blondel, Julien M Hendrickx, and John N Tsitsiklis. On krause's multi-agent consensus model with state-dependent connectivity. *IEEE transactions on Automatic Control*, 54(11):2586–2597, 2009.

[12] Noah E Friedkin. The problem of social control and coordination of complex systems in sociology: A look at the community cleavage problem. *IEEE Control Systems*, 35(3):40–51, 2015.

[13] Irinel-Constantin Morărescu, Samuel Martin, Antoine Girard, and Aurélie Muller-Gueudin. Coordination in networks of linear impulsive agents. *IEEE Transactions on Automatic Control*, 61(9):2402–2415, 2016.

[14] Yichen Wang, Grady Williams, Evangelos Theodorou, and Le Song. Variational policy for guiding point processes. *arXiv preprint arXiv:1701.08585*, 2017.

[15] Jure Leskovec, Lada A Adamic, and Bernardo A Huberman. The dynamics of viral marketing. *ACM Transactions on the Web (TWEB)*, 1(1):5, 2007.

[16] Ali Zarezade, Abir De, Hamid Rabiee, and Manuel Gomez-Rodriguez. Cheshire: An online algorithm for activity maximization in social networks. In *arXiv preprint arXiv:1703.02059*, 2017.

[17] James Archbold and Nathan Griffiths. Limiting concept spread in environments with interacting concepts. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 1332–1340. International Foundation for Autonomous Agents and Multiagent Systems, 2017.

[18] Amulya Yadav, Hau Chan, Albert Xin Jiang, Haifeng Xu, Eric Rice, and Milind Tambe. Using social networks to aid homeless shelters: Dynamic influence maximization under uncertainty. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 740–748. International Foundation for Autonomous Agents and Multiagent Systems, 2016.

[19] Mehrdad Farajtabar, Jiachen Yang, Xiaojing Ye, Huan Xu, Rakshit Trivedi, Elias Khalil, Shuang Li, Le Song, and Hongyuan Zha. Fake news mitigation via point process based intervention. *arXiv preprint arXiv:1703.07823*, 2017.

[20] Ali Zarezade, Abir De, Utkarsh Upadhyay, Hamid R. Rabiee, and Manuel Gomez-Rodriguez. Steering social activity: A stochastic optimal control point of view. *JMLR*, 2018.

[21] Marian-Andrei Rizoiu, Lexing Xie, Scott Sanner, Manuel Cebrian, Honglin Yu, and Pascal Van Hentenryck. Expecting to be hip: Hawkes intensity processes for social media popularity. In *Proceedings of the 26th International Conference on World Wide Web*, pages 735–744. International World Wide Web Conferences Steering Committee, 2017.

[22] Seyed Abbas Hosseini, Ali Khodadadi, Ali Arabzadeh, and Hamid R Rabiee. Hnp3: A hierarchical nonparametric point process for modeling content diffusion over social media. In *Data Mining (ICDM), 2016 IEEE 16th International Conference on*, pages 943–948. IEEE, 2016.

[23] M Moore and M Davenport. Learning network structure via hawkes processes. In *Proc. Work. Signal Processing with Adaptive Sparse Structured Representations (SPARS)*, 2015.

[24] Dominik Thalmeier, Vicenç Gómez, and Hilbert J Kappen. Action selection in growing state spaces: Control of network structure growth. *Journal of Physics A: Mathematical and Theoretical*, 50(3):034006, 2016.

[25] Igor Douven and Alexander Riegler. Extending the hegselmann–krause model i. *Logic Journal of IGPL*, 18(2):323–335, 2009.

[26] Bidisha Samanta, Abir De, Abhijnan Chakraborty, and Niloy Ganguly. Lmpp: A large margin point process combining reinforcement and competition for modeling hashtag popularity.

[27] Emmanuel Bacry, Stéphane Gaïffas, Iacopo Mastromatteo, and Jean-François Muzy. Mean-field inference of hawkes point processes. *Journal of Physics A: Mathematical and Theoretical*, 49(17):174006, 2016.

[28] Nan Du, Hanjun Dai, Rakshit Trivedi, Utkarsh Upadhyay, Manuel Gomez, and Le Song. Recurrent temporal point process.

[29] Martin Jankowiak and Manuel Gomez-Rodriguez. Uncovering the spatiotemporal patterns of collective social activity. In *Proceedings of the 2017 SIAM International Conference on Data Mining*, pages 822–830. SIAM, 2017.

[30] Arun Sathanur, Mahantesh Halappanavar, Yi Shi, and Walin Sagduyu. Exploring the role of intrinsic nodal activation on the spread of influence in complex networks. *arXiv preprint arXiv:1707.05287*, 2017.

[31] Yichen Wang, Xiaojing Ye, Haomin Zhou, Hongyuan Zha, and Le Song. Linking micro event history to macro prediction in point process models. In *Artificial Intelligence and Statistics*, pages 1375–1384, 2017.

[32] Utkarsh Upadhyay, Isabel Valera, and Manuel Gomez-Rodriguez. Uncovering the dynamics of crowdlearning and the value of knowledge. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pages 61–70. ACM, 2017.

[33] Ali Zarezade, Sina Jafarzadeh, and Hamid R Rabiee. Spatio-temporal modeling of check-ins in location-based social networks. *arXiv preprint arXiv:1611.07710*, 2016.

[34] Young Lee, Kar Wai Lim, and Cheng Soon Ong. Hawkes processes with stochastic excitations. In *International Conference on Machine Learning*, pages 79–88, 2016.

[35] Behzad Tabibian, Isabel Valera, Mehrdad Farajtabar, Le Song, Bernhard Schölkopf, and Manuel Gomez-Rodriguez. Distilling information reliability and source trustworthiness from digital traces. In *Proceedings of the 26th International Conference on World Wide Web*, pages 847–855. International World Wide Web Conferences Steering Committee, 2017.

[36] Ali Zarezade, Utkarsh Upadhyay, Hamid R Rabiee, and Manuel Gomez-Rodriguez. Redqueen: An online algorithm for smart broadcasting in social networks. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pages 51–60. ACM, 2017.

[37] Mehrdad Farajtabar, Manuel Gomez-Rodriguez, Yichen Wang, Shuang Li, Hongyuan Zha, and Le Song. Co-evolutionary dynamics of information diffusion and network structure. In *Proceedings of the 24th International Conference on World Wide Web*, pages 619–620. ACM, 2015.

[38] Edward Choi, Nan Du, Robert Chen, Le Song, and Jimeng Sun. Constructing disease network and temporal progression model via context-sensitive hawkes process. In *Data Mining (ICDM), 2015 IEEE International Conference on*, pages 721–726. IEEE, 2015.

[39] Julio Cesar Louzada Pinto, Tijani Chahed, and Eitan Altman. Trend detection in social networks using hawkes processes. In *Advances in Social Networks Analysis and Mining (ASONAM), 2015 IEEE/ACM International Conference on*, pages 1441–1448. IEEE, 2015.

[40] Ali Zarezade, Ali Khodadadi, Mehrdad Farajtabar, Hamid R Rabiee, and Hongyuan Zha. Correlated cascades: Compete or cooperate. In *AAAI*, pages 238–244, 2017.

[41] Mohammad Reza Karimi, Erfan Tavakoli, Mehrdad Farajtabar, Le Song, and Manuel Gomez-Rodriguez. Smart broadcasting: Do you want to be seen? *arXiv preprint arXiv:1605.06855*, 2016.

[42] Abir De, Sourangshu Bhattacharya, Sourav Sarkar, Niloy Ganguly, and Soumen Chakrabarti. Discriminative link prediction using local, community, and global signals. *TKDE*, 2016.

[43] Aniko Hannak, Eric Anderson, Lisa Feldman Barrett, Sune Lehmann, Alan Mislove, and Mirek Riedewald. Tweetin'in the rain: Exploring societal-scale effects of weather on mood. In *ICWSM*, 2012.

[44] Supplementary material. http://www.cnergres.iitkgp.ac.in/suppmat.pdf .

[45] F. B. Hanson. *Applied stochastic processes and control for Jump-diffusions: modeling, analysis, and computation*, volume 13. Siam, 2007.

[46] Fred Glover. Improved linear integer programming formulations of nonlinear integer problems. *Management Science*, 22(4):455–460, 1975.

[47] Francisco E Torres. Linearization of mixed-integer products. *Mathematical programming*, 49(1):427–428, 1990.

[48] Hanif D Sherali and Warren P Adams. *A reformulation-linearization technique for solving discrete and continuous nonconvex problems*, volume 31. Springer Science & Business Media, 2013.

[49] Jure Leskovec, Deepayan Chakrabarti, Jon Kleinberg, Christos Faloutsos, and Zoubin Ghahramani. Kronecker graphs: An approach to modeling networks. *The Journal of Machine Learning Research*, 11:985–1042, 2010.

[50] Abir De, Sourangshu Bhattacharya, and Niloy Ganguly. Demarcating endogenous and exogenous opinion diffusion process on social networks. In *WWW*, 2018.