arXiv:1605.06855v1 [cs.SI] 22 May 2016

Smart broadcasting: Do you want to be seen?

Mohammad Reza Karimi*[1], Erfan Tavakoli*[1], Mehrdad Farajtabar[2],
Le Song[2], and Manuel Gomez-Rodriguez[3]

[1]Sharif University, mkarimi@ce.sharif.edu, erfan.tavakoli71@gmail.com
[2]Georgia Institute of Technology, mehrdad@gatech.edu, lsong@cc.gatech.edu
[3]Max Planck Institute for Software Systems, manuelgr@mpi-sws.org

Abstract

Many users in online social networks are constantly trying to gain attention from their followers by *broadcasting* posts to them. These broadcasters are likely to gain greater attention if their posts can remain visible for a longer period of time among their followers' most recent feeds. Then when to post? In this paper, we study the problem of smart broadcasting using the framework of temporal point processes, where we model users feeds and posts as discrete events occurring in continuous time. Based on such continuous-time model, then choosing a broadcasting strategy for a user becomes a problem of designing the conditional intensity of her posting events. We derive a novel formula which links this conditional intensity with the "visibility" of the user in her followers' feeds. Furthermore, by exploiting this formula, we develop an efficient convex optimization framework for the "when-to-post" problem. Our method can find broadcasting strategies that reach a desired "visibility" level with provable guarantees. We experimented with data gathered from Twitter, and show that our framework can consistently make broadcasters' post more visible than alternatives.

1 Introduction

The popularization of social media and online social networking has empowered political parties, small and large corporations, celebrities, as well as ordinary people, with a platform to build, reach and broadcast information to their own audience. For example, political leaders use social media to present their character and personalize their message in hopes of tapping younger voters[1]; corporations increasingly rely on social media for a variety of tasks, from brand awareness to marketing and customer service [7]; celebrities leverage social media to bring awareness to themselves and strengthen their fans' loyalty[2]; and, ordinary people post about their lives and express their opinions to gain recognition from a mix of close friends and acquaintances[3]. However, social media users often *follow* hundreds of *broadcasters*, and they often receive information at a rate far higher than their cognitive abilities to process it [13]. This also means that many broadcasters actually share quite a portion of their followers, and they are constantly competing for attention from these followers.

In this context, these followers' attention becomes a scarce commodity of great value [8], and broadcasters would like to consume a good share of it so that their posted contents are *noticed* and possibly liked or shared. As a consequence, there are myriads of articles and blog entries about the *best times* to broadcast information in social media and social networking, as well as data analytics tools to find these times[4][5]. However, the best

*Authors contributed equally. This work was done during the authors' internships at Max Planck for Software Systems.
[1]http://www.nytimes.com/2012/10/08/technology/campaigns-use-social-media-to-lure-younger-voters.html
[2]http://www.wsj.com/articles/what-celebrities-can-teach-companies-about-social-media-1444788220
[3]http://www.pewinternet.org/topics/social-networking/
[4]http://www.huffingtonpost.com/catriona-pollard/the-best-times-to-post-on_b_6990376.html
[5]http://blog.klout.com/2015/07/whens-the-best-time-to-post-on-social/

time to post on social media depends on a variety of factors, often specific to the broadcaster in question, such as their followers' daily and weekly behavior patterns, their location or timezone, and the number of broadcasters and volume of information competing for their attention in these followers' feeds (be it in the form of a Twitter user's timeline, a Facebook user's wall or an Instagram user's feed). Therefore, the problem of finding the best times to broadcast messages and elicit attention (be it views, likes or shares), in short, the *when-to-post* problem, requires careful reasoning and smart algorithms, which have been largely inexistent until very recently [19].

In this paper, we develop a novel framework for the when-to-post problem, where we measure the gained attention or *visibility* of a broadcaster as the time that at least one post from her is among the most recent k received stories in her followers' feed. A desirable property of this time based visibility measure is that it is easy to estimate from real data. In order to measure the achieved visibility for a particular deployed broadcasting strategy, one only need to use a separate held-out set of the followers' feeds, independently of the broadcasted content. This is in contrast to other measures based on, *e.g.*, the number of likes or shares caused by a broadcasting strategy. These latter measures are difficult to estimate from real data and often require actual interventions, since they depend on other confounding factors such as the follower's reaction to the post content [6], whose effect is difficult to model accurately [5].

More specifically, we will model users' feeds and posts as discrete events occurring in continuous time using the framework of temporal point processes. Our model explicitly characterize the continuous time interval between posts by means of conditional intensity functions [1]. Based on such continuous-time model, then choosing a strategy for a broadcaster becomes a problem of designing the conditional intensity of her posting events. We derive a novel formula which can link the conditional intensity of an *arbitrary* broadcaster with her visibility in her followers' feeds. Interestingly, we can show that the average visibility is concave in the space of (piece-wise) smooth intensity functions. Based on this result, we propose a convex optimization framework to address a diverse range of visibility shaping tasks given budget constraints. Our framework allows us to conduct fine-grained control of a broadcaster's visibility across her followers. For instance, our framework can steer the visibility in such a way that some time intervals are favored over others, *e.g.*, times when the broadcasters' followers are on-line. In addition to the novel framework, we develop an efficient gradient based optimization algorithm, which allows us to find optimal broadcast intensities for a variety of visibility shaping tasks in a matter of milliseconds. Finally, we experimented on a large real-world dataset gathered from Twitter dataset, and show that our framework can consistently make broadcasters' posts more visible than alternatives.

Related work. The work most closely related to ours is by Spasojevic et al. [19], who introduced the when-to-post problem. In their work, they first perform an empirical study on the best times to post in Twitter and Facebook by analyzing more than a billion messages and responses. Then, they design several heuristics to (independently) pinpoint at the times that elicited the greatest number of responses in a training set and then show that these times also lead to more responses in a held-out set. In our work, we measure attention by means of visibility, a measure that is not confounded with the message content and can be accurately evaluated on a held-out set, and then develop a convex optimization framework to design complete broadcasting strategies that are *provably* optimal.

There have been an increasing number of empirical studies on understanding attention and information overload on social and information networks [2, 14, 16, 13]. The common theme is to investigate whether there is a limit on the amount of ties (*e.g.*, friends, followees or phone contacts) people can maintain, how people distribute attention across them, and how attention influences the propagation of information. In contrast, in this work, we focus on optimizing a social media user's broadcasting strategy to capture the greatest attention from their followers.

Our work also relates to the influence maximization problem, extensively studied in recent years [18, 15, 4, 10], which aims to find a set of nodes in a social network whose initial adoptions of certain idea or product can trigger the largest expected number of follow-ups. In this line of work, the goal is finding these influential users but not to find the best times for these users to broadcast their messages, which is our goal here. Only very recently, Farajtabar et al. [11] have developed a convex optimization framework to find broadcasting strategies, however, their focus is on steering the overall activity in the network to a certain

state by incentivizing a few influential users, in contrast, we focus on maximizing visibility as measured on a broadcaster's audience's feeds.

Finally, the framework of temporal point processes, which our work builds upon, has been increasingly used to model a wide range of phenomena in social media and social networking sites, *e.g.*, from social influence [11], network evolution [12], opinion dynamics [9] or product competition [20].

2 Background on Point Processes

A temporal point process is a stochastic process whose realization consists of a list of discrete events localized in time, $\{t_i\}$ with $t_i \in \mathbb{R}^+$ and $i \in \mathbb{Z}^+$. Many different types of data produced in online social networks can be represented as temporal point processes, such as the times of tweets, retweets or likes in Twitter. A temporal point process can be equivalently represented as a counting process, $N(t)$, which records the number of events before time t. Then, in a infinitesimally small time window dt around time t, the number of observed event is

$$dN(t) = \sum_{t_i \in \mathcal{H}(t)} \delta(t - t_i)\, dt, \tag{1}$$

and hence $N(t) = \int_0^t dN(s)$, where $\delta(t)$ is a Dirac delta function. It is often assumed that only one event can happen in a small window of size dt, and hence $dN(t) \in \{0, 1\}$.

An important way to characterize temporal point processes is via the intensity function — the stochastic model for the time of the next event given all the times of previous events. The intensity function $\lambda(t)$ (intensity, for short) is the probability of observing an event in a small window $[t, t + dt)$, *i.e.*,

$$\lambda(t)dt = \mathbb{P}\left\{\text{event in } [t, t + dt)\right\}. \tag{2}$$

Based on the intensity, one can obtain the expectation of the number of events in the windows $[t, t + dt)$ and $[0, t)$ respectively as

$$\mathbb{E}[dN(t)] = \lambda(t)\, dt \text{ and } \mathbb{E}[N(t)] = \int_0^t \lambda(\tau)\, d\tau \tag{3}$$

There is a wide variety of functional forms for the intensity $\lambda(t)$ in the growing literature on social activity modeling using point processes, which are often designed to capture the phenomena of interests. For example, retweets have been modeled using multidimensional Hawkes processes [11, 22], new network links have been predicted using survival processes [21, 12], and daily and weekly variations on message broadcasting intensities have been captured using inhomogeneous Poisson processes [17].

In this work, since we are interested on optimizing message broadcasting intensities, we use inhomogeneous Poisson processes, whose intensity is a time-varying function $\lambda(t) = g(t) \geqslant 0$.

3 From Intensities to Visibility

In this section, we will present our model for the posting times of broadcasters and the feed story arrival times of followers using point processes parameterized by intensity functions. Based on these models, we will then define our visibility measure, and derive a novel link between the visibility measure and the intensity functions of a broadcaster and her followers.

Representation of broadcast and feed. Given a directed social network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with $m = |\mathcal{V}|$ users, we assume that each user can be both broadcaster and follower. Then, we will use two sets of counting processes to modeling each user's activity, the first set for the user's broadcasting activity, and the second set for the user's feed activity.

More specifically, we represent the broadcasting times of the users as a set of counting processes denoted by a vector $\boldsymbol{N}(t)$, in which the u-th dimension, $N_u(t) \in \{0\} \cup \mathbb{Z}^+$, counts the number of messages user u

broadcasted up to but not including time t. Then, we can characterize the message rate of these users using their corresponding intensities

$$\mathbb{E}[d\boldsymbol{N}(t)] = \boldsymbol{\lambda}(t)\, dt. \tag{4}$$

Furthermore, given the adjacency matrix $\boldsymbol{A} \in \{0,1\}^{m\times m}$ corresponding to the social network $\mathcal{G}$, where $A_{uv} = 1$ indicates that v follows u, and $A_{uv} = 0$ otherwise, we can represent the feed story arrival times of the users as a sum of the set of broadcasting counting processes. That is

$$\boldsymbol{M}(t) = \boldsymbol{A}^T \boldsymbol{N}(t), \tag{5}$$

which essentially aggregates for each user the counting processes of the broadcasters followed by this user. Then, we can characterize the feed rates using intensity functions

$$\mathbb{E}[d\boldsymbol{M}(t)] = \boldsymbol{\gamma}(t)\, dt, \tag{6}$$

where $\boldsymbol{\gamma}(t) := \boldsymbol{A}^T \boldsymbol{\lambda}(t) = (\gamma_1, \ldots, \gamma_m)^T$.

Finally, from the perspective of a pair of broadcaster (or user) u and her follower v, it is useful to define the feed rate of v due to other broadcasters (or users) followed by v as

$$\gamma_{v\setminus u}(t) := \gamma_v(t) - \lambda_u(t), \tag{7}$$

where we assume $\gamma_{v\setminus u}(t) := 0$ if v does not follow u, $A_{uv} = 0$.

Definition of Visibility. Consider a broadcaster u and her follower v, and we note that v may follow many other broadcasters other than u. Thus, at any time t, user v may see stories originated from multiple broadcasters. We can model the times and origins of all these stories present in v's current feed as a first-in-first-out (FIFO) queue[6] of pairs

$$\mathcal{H}_v(t) := \left\{ (t_{(i)}, u_{(i)}) : t \geqslant t_{(1)} \geqslant \ldots \geqslant t_{(I-1)} \geqslant t_{(I)},\, u_{(i)} \in \mathcal{N}^-(v) \right\},$$

where $\cdot_{(i)}$ denotes the i-th element in the queue, $t_{(i)}$ is the time when v receives a story from broadcaster $u_{(i)}$, $\mathcal{N}^-(v)$ denotes the set of broadcasters followed by v, and I is the length of the queue. The length I accounts for the fact that online social platforms typically set a maximum number of stories that can be displayed in the feed, *e.g.*, currently Twitter has $I = 20$. The FIFO queue is to model the fact that when a new story arrives, the oldest story, $(t_{(I)}, u_{(I)})$, at the bottom of the feed will be removed, and the ordering of the remaining stories will be shifted down by one slot, *i.e.*,

$$i + 1 \leftarrow i, \quad \forall i = 1, \ldots, I-1$$

and the newly arrived story will be appended to the beginning of the queue as $t_{(1)}$ and appear at the top of the feed. For simplicity, we assume that the queue is always full at the time of modeling.

In the list $\mathcal{H}_v(t)$, we keep track of the rank $r_{uv}(t)$ of the most recent story posted by the broadcaster u among all the stories received by user v by time t, *i.e.*,

$$r_{uv}(t) = \min \left\{ i \;:\; u_{(i)} = u \right\}. \tag{8}$$

Then, given an observation time window $[0, T]$, and a deterministic sequence of broadcasting events, we can define the deterministic visibility of broadcaster u at k with respect to follower v as

$$\mathcal{T}_{uv}(k) := \int_0^T \mathbb{I}[r_{uv}(t) \leqslant k]\, dt, \tag{9}$$

which is the amount of times that at least one story from broadcaster u is among the most recent k stories in user v's feed.

[6] In this work, we assume the social network sorts stories in each user's feed in inverse chronological order.

Since the sequence of broadcasting events are generated from stochastic processes, we will consider the expected value of $\mathcal{T}_{uv}(k)$ instead. If we first denote the probability that at least one story from broadcaster u is among the k most recent stories in follower v's feed as

$$f_{uv}(t,k) = \mathbb{P}\left\{r_{uv}(t) \leqslant k\right\}, \tag{10}$$

then the expected (or average) visibility $\mathcal{V}(k)$ can be defined as

$$\mathcal{V}_{uv}(k) := \mathbb{E}\left[\mathcal{T}_{uv}(k)\right] = \int_0^T f_{uv}(t,k)\,dt, \tag{11}$$

given the integral is well-defined. In some scenarios, one may like to favor some periods of times (*e.g.*, times in which the follower is online), encode such preference by means of a time significance function $s(t) \geqslant 0$ and consider $f_{uv}(t,k)s(t)$ instead of just $f_{uv}(t,k)$.

Note that the visibility $\mathcal{V}_{uv}(k)$ is defined for a pair of broadcaster u and her follower v given k. We will focus our later exposition on a particular of u and v, and omit the subscript $\cdot_{uv}$ and simply use notation such as $f_k(t)$, $\mathcal{V}(k)$. However, we note that the computation of the visibility for a pair of users u and v may depend on the broadcast and feed intensities of all users in the network.

Computation of Visibility. In this section, we derive an expression for the average visibility, given by Eq. 11, using the broadcaster posting and follower feed representation, given by Eqs. 4-7. This link is crucial for the convex visibility shaping framework in Section 5.

Given a broadcaster u with $\lambda_u(t) = \lambda(t)$ and her follower v with $\gamma_v(t) = \gamma(t)$ and $\gamma_{v\setminus u}(t) = \mu(t)$, we first compute the probability $f_1(t)$ that at least one message from the broadcaster is among the $k = 1$ most recent ones received by v at time t. By definition, one can easily realize that $f_1(t)$ satisfies the following equation:

$$f_1(t+dt) = \underbrace{f_1(t)\left(1-\mu(t)dt\right)}_{\text{1. Remains the most recent}} + \underbrace{\left(1-f_1(t)\right)\lambda(t)dt,}_{\text{2. Becomes the most recent}} \tag{12}$$

where each term models one of the two possible situations:

1. The most recent message received by follower v by time t was posted by broadcaster u (*w.p.* $f_1(t)$) and none of the other broadcasters that v follows posts a message in $[t, t+dt]$ (*w.p.* $1-\mu(t)dt$).
2. The most recent message received by follower v by time t was posted by a different broadcaster (*w.p.* $1-f_1(t)$) and broadcaster u posts a message in $[t, t+dt]$ (*w.p.* $\lambda(t)dt$) which becomes the most recent one.

Then, by rearranging terms and letting $dt \to 0$, one finds that the probability satisfies the following differential equation:

$$f_1'(t) = -(\mu(t)+\lambda(t))f_1(t) + \lambda(t). \tag{13}$$

We can proceed with the induction step for $f_k(t)$ with $k > 1$. In particular, by definition, $f_k(t)$ satisfies the following equation:

$$f_k(t+dt) = \underbrace{f_{k-1}(t)}_{\text{1. Was among } k-1} + \underbrace{\left(f_k(t)-f_{k-1}(t)\right)\left(1-\mu(t)dt\right)}_{\text{2. Remains on the } k\text{-th position}} + \underbrace{\left(1-f_k(t)\right)\lambda(t)dt,}_{\text{3. Becomes the most recent}} \tag{14}$$

where each term models one of the three possible situations:

1. The last message posted by broadcaster u by time t is among the most recent $k-1$ ones received by follower v (*w.p.* $f_{k-1}(t)$) and, independent of whether a message is posted by any other broadcaster or not, this message will remain among the most recent k at $t+dt$.
2. The last message posted by broadcaster u by time t is the k-th one (*w.p.* $(f_k(t)-f_{k-1}(t))$) and none of the other broadcasters followed by v posts a message in $[t, t+dt]$ (*w.p.* $1-\mu(t)dt$)
3. The last k messages received by follower v by time t were posted by other broadcasters (*w.p.* $1-f_k(t)$) and broadcaster u posts a message in $[t, t+dt]$ (*w.p.* $\lambda(t)dt$), becoming the most recent one.

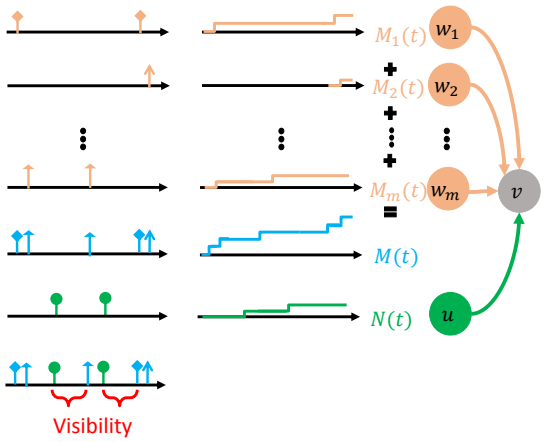

Figure 1: The visibility shaping problem. A social media user u broadcast $N_u(t)$ messages at a rate $\lambda_u(t)$. Her messages accumulate in each of her followers' feeds, which receives $M_v(t)$ messages at a rate $\gamma_v(t) = \lambda_u(t) + \gamma_{v \setminus u}(t)$, where $\gamma_{v \setminus u}(t)$ denotes the message rate due to other broadcasters v follows. For each follower, the average visibility of a user u's messages is defined as the time that a post from user u is among the last k stories the follower received. In the visibility shaping problem, the goal is to optimize $\lambda_u(t)$ to steer visibility.

By rearranging terms and letting $dt \to 0$, we uncover a recursive relationship between $f_k(t)$ and $f_{k-1}(t)$, by means of the following differential equation:

$$f_k{}'(t) = -(\mu(t) + \lambda(t)) f_k(t) + \mu(t) f_{k-1}(t) + \lambda(t), \tag{15}$$

Perhaps surprisingly, we can find a closed form expression for $f_k(t)$, given by the following Lemma (proven in the Appendix A):

Lemma 1. *Given a broadcaster with message intensity $\lambda(t)$ and one of her followers with feed message intensity due to other broadcasters $\mu(t)$. The probability $f_k(t)$ that at least one message from the broadcaster is among the k most recent ones received by the follower at time t can be uniquely computed as*

$$f_k(t) = \frac{\int_0^t \lambda(\tau) e^{-\int_\tau^t \lambda(x)dx} \Gamma[k, \int_\tau^t \mu(x)dx] \, d\tau}{(k-1)!}, \tag{16}$$

given the boundary conditions $f_1(0) = \ldots = f_k(0) = 0$ and the incomplete Gamma function defined as $\Gamma[k, x] = \int_x^\infty \tau^{k-1} e^{-\tau} d\tau$.

4 On the Concavity of Visibility

Once we have a formula that allows us to compute the average visibility given *any* arbitrary intensities for the broadcasters, we will now show that, remarkably, the average visibility is concave in the space of smooth intensity functions. Moreover, we will also show that the average visibility is concave with respect to the parameters of piecewise constant functions, which we will use in our experiments.

Smooth intensity functions. In this section, we assume that the message intensity of the broadcaster belongs to the space $\mathcal{H}$ of all smooth functions. Before we proceed, we need the following definition:

Definition 2. *Given the space $\mathcal{H}$ of all smooth functions, a functional $J : \mathcal{H} \to \mathbb{R}$ is concave if for every $g_1, g_2 \in \mathcal{H}$ and $0 < \alpha < 1$:*

$$J[\alpha g_1 + (1-\alpha) g_2] \geq \alpha J[g_1] + (1-\alpha) J[g_2]. \tag{17}$$

A functional J is convex if $-J$ is concave.

It readily follows that the probability $f_k(t)$, given by Eq. 16, is a functional with $\lambda(\cdot)$ as input. Moreover, the following two theorems, proven in Appendices B and C, establish the concavity of $f_k(t)$ and $\mathcal{V}(k)$ with respect to $\lambda(\cdot)$.

Theorem 3. *Given a broadcaster with message intensity $\lambda(t)$ and one of her followers with feed message intensity due to other broadcasters $\mu(t)$. The probability $f_k(t)$ that at least one message from the broadcaster is among the k most recent ones received by the follower at time t, given by Eq. 16, is concave with respect to $\lambda(\cdot)$.*

Theorem 4. *Given a broadcaster with message intensity $\lambda(t)$ and one of her followers with feed message intensity due to other broadcasters $\mu(t)$. The visibility $\mathcal{V}(k)$, given by Eq. 11, is concave with respect to $\lambda(\cdot)$.*

Given the above results, one could think of finding the optimal (general) message intensity $\lambda(t)$ that maximize (a function of) the average visibilities across a broadcaster's followers. However, in practical applications, this may be inefficient and undesirable, instead, one may focus on a simpler parametrized family of intensities, such as piecewise constant intensity functions, which will be easier to optimize and fit using real data. To this aim, next, we prove that the average visibility is also concave on the parameters defining piecewise constant intensity functions.

Piecewise constant intensity functions. In this section, we assume that the message intensity $\lambda(t)$ of the broadcaster belongs to the space of piecewise constant functions $\lambda : [0, T] \to \mathbb{R}$, denoted by $\mathcal{G}$, which we parametrized as follows:

$$\lambda(t) = \sum_{m=1}^{M} a_m \mathbb{I}(\tau_{m-1} \leq t < \tau_m), \tag{18}$$

where $a_m \geq 0$, M is the number of pieces, $\tau_i - \tau_{i-1} = T/M = \Delta$ and $\tau_0 = 0$.

As the reader may have noticed, the results from the previous section are not readily usable since Lemma 1 requires the intensity functions to be smooth. However, we will now show that, for every function $\lambda(t) \in \mathcal{G}$, there is a sequence of smooth functions $\lambda_n(t) \in \mathcal{H}$ such that $\lim_{n\to\infty} \lambda_n(t) = \lambda(t)$ and, this will sufficient to prove concavity. Before we proceed, we need the following definition:

Definition 5. *A functional $J : \mathcal{G} \to \mathcal{H}$ is said to be continuous at $\lambda(\cdot) \in \mathcal{H}$ if for every $\epsilon > 0$, there is a $\delta > 0$ such that*

$$|J[\lambda] - J[\hat{\lambda}]| < \epsilon \tag{19}$$

provided that $||\lambda - \hat{\lambda}|| < \delta$, where $||\cdot||$ is a norm in $\mathcal{H}$.

It readily follows that the probability f_k is a continuous functional on $\mathcal{H}$. Moreover, we need the following lemma (proven in Appendix D) to prove the concavity:

Lemma 6. *For every $\lambda(t) \in \mathcal{G}$, there is a sequence of smooth functions $\lambda_n(t) \in \mathcal{H}$ where $\lim_{n\to\infty} \lambda_n(t) = \lambda(t)$.*

Using Lemma 6, for any $\lambda(t) \in \mathcal{G}$, it follows that

$$f_k(\lambda(\cdot)) = \lim_{n\to\infty} f_k(\lambda_n(t)) \tag{20}$$

where $\lambda_n(t)$ is a sequence of smooth functions such that $\lim_{n\to\infty} \lambda_n(t) = \lambda(t)$. As a consequence, we can establish the concavity of $f_k(t)$ and $\mathcal{V}(k)$ with respect to $a_1, \ldots, a_M$ with the following Theorem (proven in Appendix E):

Theorem 7. *f_k and $\mathcal{V}(k)$ are concave functionals in the space of piecewise constant functions $\mathcal{G}$.*

Corollary 8. *If we represent $\lambda \in \mathcal{G}$ using Eq. 18, $f_k(t)$ and $\mathcal{V}(k)$ are concave with respect to $a_1, \ldots, a_m$.*

Algorithm 1: Projected Gradient Descent for Visibility Shaping

Initialize $\boldsymbol{c}$;
repeat
  1- Project $\boldsymbol{c}$ into the polytope $\boldsymbol{c} \geqslant 0$, $\boldsymbol{c}^\top \mathbf{1} \leqslant C$;
  2- Find the gradient $\boldsymbol{g}(\boldsymbol{c})$;
  3- Update $\boldsymbol{c}$ using the gradient $\boldsymbol{g}(\boldsymbol{c})$;
until *convergence*;

5 Convex Visibility Shaping Framework

Given the concavity of the average visibility, we now propose a convex optimization framework for a variety of visibility shaping tasks. In all these tasks, our goal is to find the optimal message intensity $\lambda_u(t)$ for broadcaster u that maximizes a particular nondecreasing concave utility function $U(\boldsymbol{\mathcal{V}}_u(k))$ of the average visibility of broadcaster u in all her followers within a time window $[0, T]$, *i.e.*,

$$\begin{array}{ll} \text{maximize}_{\lambda_u(t)} & U(\boldsymbol{\mathcal{V}}_u(k)) \\ \text{subject to} & \lambda_u(t) \geq 0 \quad t \in [0,T] \\ & \int_0^T \lambda(t)\, dt \leq C, \end{array} \tag{21}$$

where $\boldsymbol{\mathcal{V}}_u(k) = (\mathcal{V}_{uv}(k))_{v \in \mathcal{N}(u)^+}$, $\mathcal{N}(u)^+$ denotes the broadcaster u's followers, $\mathcal{V}_{uv}(k)$ denotes the average visibility in follower v, the first constraint asserts the intensity function remains positive, and the second limits the average number of messages broadcasted within $[0, T]$ to be no more than C.

We next discuss two instances of the general framework, which achieve different goals (their constraints remain the same and hence omitted). More generally, the flexibility of our framework allows to use any nondecreasing concave utility function.

Average Visibility Maximization (AVM). The goal here is to maximize the sum of the visibility for all the broadcaster's followers, *i.e.*,

$$\text{maximize}_{\lambda_u(t)} \quad \sum_{v \in \mathcal{N}(u)^+} \mathcal{V}_{uv}(k) \tag{22}$$

Minimax Visibility Maximization (MVM). Suppose our goal is instead to keep the visibility in the n followers with the smallest visibility value above a certain minimum level, or, alternatively make the average visibility across the n followers with the smallest visibility as high as possible. Then, we can perform the following *minimax visibility maximization* task

$$\text{maximize}_{\lambda_u(t)} \sum_{i=1}^{n} \mathcal{V}_{uv_{[i]}}(k), \tag{23}$$

where $\mathcal{V}_{uv_{[i]}}(k)$ denotes the average visibility in the follower with the i-th smallest visibility among all the broadcaster's followers.

6 Scalable Algorithm

To solve the visibility shaping problems defined above, we need to be able to (efficiently) evaluate the probability function f_k and visibility $\mathcal{V}(k)$. However, a direct evaluation by means of Eqs. 16 and 11 seem difficult. Here, we present an alternative representation of the probability function f_k and the visibility $\mathcal{V}(k)$ for piece-wise constant intensity functions, which allow us to compute these quantities very efficiently. Based on this result, we present an efficient gradient based algorithm to find the optimum intensity.

Assume the broadcaster's message intensity $\lambda(t)$ and the follower's feed message intensity due to other broadcasters $\mu(t)$ adopt the following form:

$$\lambda(t) = \sum_{m=1}^{M} c_m \mathbb{I}(\tau_{m-1} \leq t < \tau_m) \text{ and } \mu(t) = \sum_{m=1}^{M} b_m \mathbb{I}(\tau_{m-1} \leq t < \tau_m).$$

Then, each piece m in the above intensities satisfies the recurrence relation given by Eq. 15, which we rewrite as

$$f_k'(t) + (b_m + c_m) f_k(t) = c_m + b_m f_{k-1}(t),$$

and one can easily prove by induction that, in general, the solution of the above differential equation for each time interval $\tau_{m-1} \leq t < \tau_m$ is given by

$$f_k(t) = e^{-(b_m+c_m)t}(\alpha_{k-1,k} t^{k-1} + \cdots + \alpha_{0,k}) + \beta_k \tag{24}$$

where $\alpha_{i,k} = \frac{b^i}{i!}(h_{k-i} - \beta_{k-i})$, $\beta_i = 1 - \left(\frac{b_m}{b_m + c_m}\right)^i$, and h_i is the probability $f_i(\tau_{i-1})$ at the beginning of time interval. Such representation allows for an efficient evaluation of $f_k(t)$.

Next, we also need to compute the integral of $f_k(t)$ to efficiently compute the visibility $\mathcal{V}(k)$. Without loss of generality, we represent the time for each piece in a normalized time window $[0, 1]$. Then, the integral of $f_k(t)$ can be written as follows:

$$\mathcal{V}(k) = \int_0^1 f_k(t)\, dt = \beta_k + \sum_{i=0}^{k-1} \alpha_{i,k} \int_0^1 e^{-(b_m+c_m)t} t^i\, dt = \beta_k + \sum_{i=0}^{k-1} \frac{\alpha_{i,k}}{(b_m + c_m)^{i+1}} [i! - \Gamma(i+1, b_m + c_m)] \tag{25}$$

where note that the last term is efficiently computable since, for integer values of n, the incomplete Gamma function $\Gamma(n, x) = (n-1)!\, e^{-x} \sum_{i=0}^{n-1} \frac{x^i}{i!}$.

Given Eq. 25, we can now easily compute the gradient of the visibility $\mathcal{V}(k)$, which we can then use to design an efficient gradient based algorithm. For brevity, we just show the gradient for $k = 1$. Let $\boldsymbol{c} = (c_1, \ldots, c_M)$, and $\boldsymbol{y} = (y_0, \ldots, y_{M-1}, y_M)$ be the values of $f_k(t)$ at the beginning of each time interval, then,

$$\frac{\partial \mathcal{V}(1)}{\partial c_i} = \frac{1}{(b_i + c_i)^2} \left(-\frac{\partial y_i}{\partial c_i}(b_i + c_i) + (y_i - y_{i-1}) + b_i \right) + \sum_{m=i+1}^{M} \frac{1}{b_m + c_m} \left[\frac{\partial y_{m-1}}{\partial c_i} - \frac{\partial y_m}{\partial c_i} \right].$$

where we can easily compute $\partial y_j / \partial c_m$ recursively as

$$\frac{b_m}{(b_m + c_m)^2} - \left(y_{m-1} - \frac{c_m}{(b_m + c_m)} - \frac{b_m}{(b_m + c_m)^2} \right) e^{-(b_m + c_m)},$$

if $j = m$, and $e^{-(b_j + c_j)} \frac{\partial y_{j-1}}{\partial c_m}$, if $j > m$.

Once we have an efficient way to compute the visibility $\mathcal{V}(k)$ and its gradient, we can readily design a projected gradient descent algorithm to find the optimal message intensity $\lambda_u(t)$ in the visibility shaping problems described in Section 5. Note that, since our optimization problems are convex, there is a unique optimum and convergence is guaranteed. Moreover, for the projection step, we solve a quadratic program, minimizing the distance to the feasible polytope. Algorithm 1 summarizes the overall algorithm.

7 Experiments

Dataset description and experimental setup. We use data gathered from Twitter as reported in previous work [3], which comprises the following three types of information: profiles of 52 million users, 1.9 billion directed follow links among these users, and 1.7 billion public tweets posted by the collected users. The follow link information is based on a snapshot taken at the time of data collection, in September 2009.

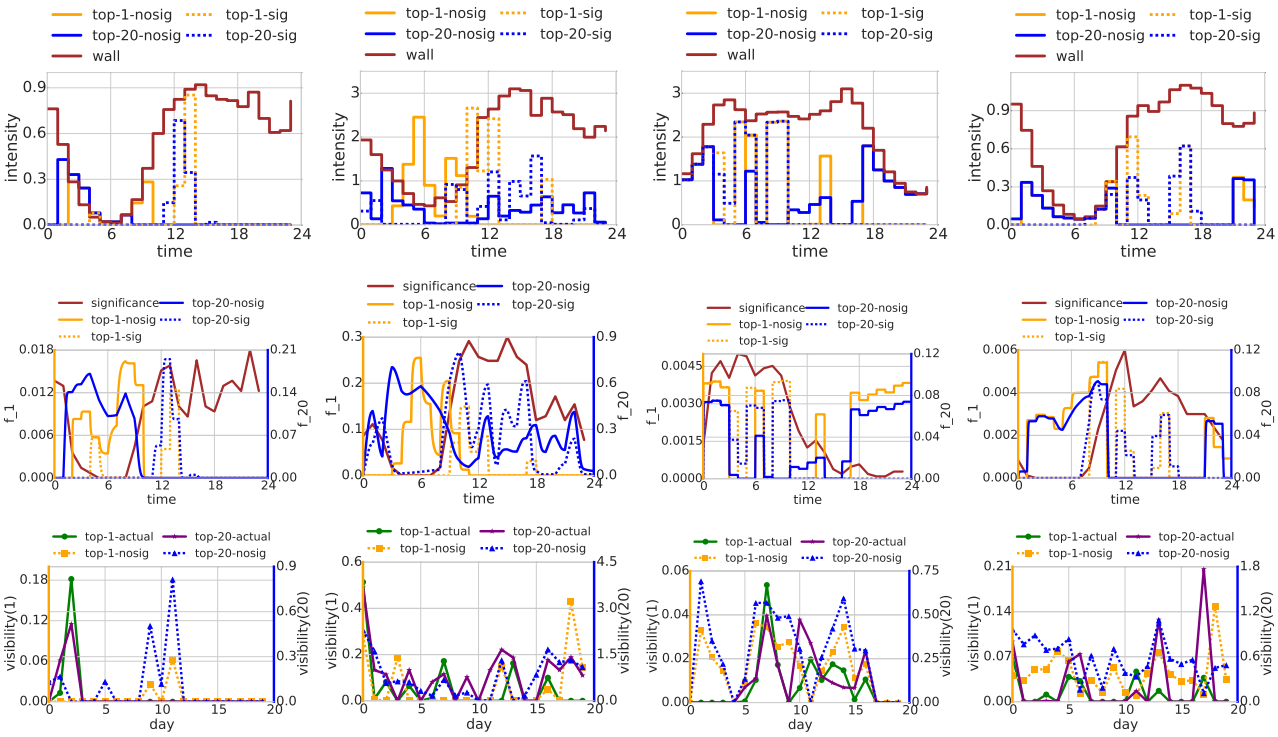

Figure 2: Intensities and top-k probabilities and visibilities. We focus on four broadcasters (one per column) and solve the AVM problem for one of their followers, picked at random. The first row shows the follower's timeline intensity ($\mu(t)$, in brown) fitted using events from the training set, and the optimized intensities, as given by our framework, that maximize visibility for k=1, 20 on the training set with and without significance ($\lambda^*(t)$, in solid and dashed yellow and blue, respectively). The second row shows the top-k probability for the optimized intensities with and without significance for k=1, 20 ($f_k^*(t)$, in solid and dashed yellow and blue) as well as the follower's significance ($s(t)$, in brown). The third row compares the average visibility achieved by the optimized intensities without significance for k=1, 20 ($\mathcal{V}^*(k)$, in yellow and blue) to the average visibility achieved by the broadcaster's posting activity ($\mathcal{V}(k)$, in green and purple) on a held-out set.

Here, we focus on the tweets published during a six and a half month period, from February 2, 2009 to August 13, 2009. In particular, we sample 10,000 users uniformly at random as broadcasters and record all the tweets they posted. Moreover, for each of these broadcasters, we track down all their followers and record all the tweets they posted as well as reconstruct their *true* timelines by collecting all the tweets published by the people they follow.

In our experiments, we use the first three and a half month period, from February 2 to May 13 to fit the piecewise constant intensities of the followers' timelines and the followers' significance, which we use in our convex visibility shaping framework. Here, the follower's significance is the probability that she is on-line, estimated as a piecewise (hourly) constant probability from the tweets-retweets the follower posted – if a follower tweeted or retweeted in an hour, we assume it was on-line during that hour. Then, we use the last three month period, from May 14 to August 13, to evaluate our framework. We refer to the former period as the training set and the latter as the test set. We experiment both with T=24 hours (M=24, Δ=1 hour) and T=7 days (M=24 $\times$ 7, Δ=1 hour), and set the budget C to be equal to the average number of tweets per T the broadcaster posted in the training period.

Evaluation schemes. Throughout this section, we use three different evaluation schemes, with an increasing resemblance to a real world scenario:

Theoretical objective: We compute the theoretical value of the utility using the broadcaster intensity under study, be it the (optimal) intensity given by our convex visibility shaping framework, the intensity given by an alternative baseline, or the the broadcaster's (true) fitted intensity.

Simulated objective: We simulate events both from the broadcaster intensity under study and each of the followers' timeline fitted intensities. Then, we estimate empirically the overall utility based on the simulated events. We perform 100 independent simulation runs and report the average and standard error (or standard deviation) of the utility.

Held-out data: We simulate events from the broadcaster intensity under study, interleave these generated events on the true followers' timelines recorded as test set, and compute the corresponding utility. We perform 10 independent simulation runs and report the average and standard error (or standard deviation) of the utility.

Intensities, top-k probabilities and visibilities. We pay attention to four broadcasters, picked at random, and solve the average visibility maximization task for one of their followers, also picked at random. Our goal here is to shed light on the influence that the follower's timeline intensity and significance have on the optimized broadcaster's intensity as well as its corresponding visibility and top-k probability for different values of k. Figure 2 summarizes the results, which show that (i) including the significance in the visibility definition shifts the optimized intensities away from the times in which the followers are not online (first row); (ii) the optimized intensities typically achieve a higher average visibility than the one achieved by the broadcaster's true posting activity on a held-out set (third row); and (iii) the optimized intensities are more concentrated in time for $k = 1$ (first row) and achieve a higher average visibility and top-k probability for $k = 20$ (second and third row).

Solution quality. In this section, we perform a large scale evaluation of our framework across all 10,000 broadcasters in terms of the three evaluation schemes described above and compare its performance against several baselines. Here, we consider the definition of visibility that incorporates significance since, as argued previously, may lead to more effective broadcasting strategies[7].

In the average visibility maximization task, we compare our framework with three heuristics, in which the broadcaster distributes the available budget uniformly at random (RAVM), proportionally to $\sum_{i=1}^{n} \mu_i(t)$ (IAVM) and proportionally to $\sum_{i=1}^{n} s_i(t)\mu_i(t)$ (PAVM), respectively. In the minimax visibility maximization task, we also compare with three heuristics. The first two heuristics are similar to two of the ones just mentioned for AVM, *i.e.*, the broadcaster distributes the available budget uniformly at random (RMVM) and proportionally to $\sum_{i=1}^{n} \mu_i(t)$ (IMVM). In the third heuristic, the broadcaster distributes its budget following a greedy procedure: at each iteration k, it first finds the user with the least visibility given $\lambda^{(k-1)}(t)$ and then solves the average visibility maximization for that user given a budget of C/n. Finally, it outputs the intensity $\lambda(t) = \sum_{k=1}^{n} \lambda^{(k)}(t)$. The greedy procedure starts with $\lambda^{(0)}(t) = C/M$. Additionally, for the held-out comparison, we also compute the actual average intensity that the broadcaster achieved in reality.

Figure 3 summarizes the results by means of a box plot, which shows the utilities achieved by our framework and the heuristics normalized with respect to the utility achieved by the broadcasters' fitted true intensity (by the posts during the test set for the third evaluation scheme). That means, if $y = 1$, the optimized intensity achieves the same utility as the broadcaster's recorded posts. For the average visibility maximization task, the intensities provided by our method achieve 1.5× higher theoretical objective and 1.3× higher utility on a held-out set, in average (black dashed line), than the broadcaster's fitted intensities. In contrast, alternatives fail at providing any gain, *i.e.*, $y \leq 1$ for a half of the broadcasters. Finally, for the minimax visibility maximization task, which is significantly harder, the intensities provided by our method achieve 1.6× higher theoretical objective and 1.4× higher average utility on a held-out set, in average (black dashed line), than the broadcaster's fitted intensities. In this case, although our method outperforms the baselines by large margins in terms of theoretical and simulated objectives, the baselines achieved almost the same average utility on the held-out set. The theoretical and simulated objective are almost equal in all cases, as one may have expected.

[7]We obtain qualitatively similar results if we omit the significance in the definition of visibility. Actually, in such case, our framework beats the baselines by a greater margin.

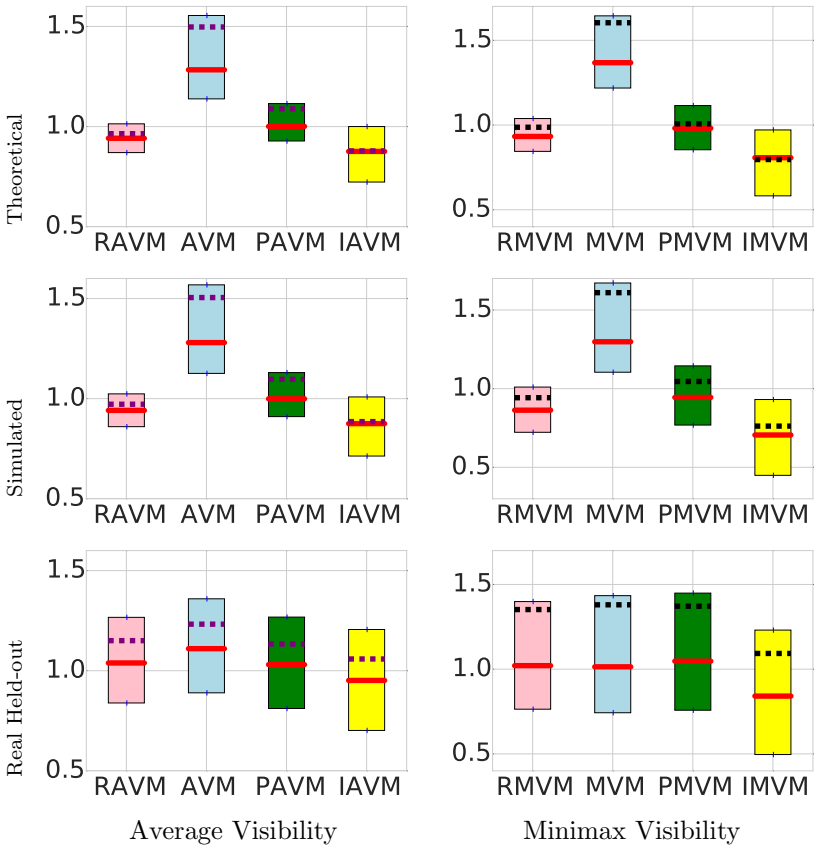

Figure 3: Visibility shaping for 10,000 broadcasters. The left (right) column corresponds to AVM (MVM), evaluated using the theoretical objective (first row), the simulated objective (second row) and the held-out data (third row). The red (blue) dashed line shows the median (mean) objective popularity and the box limits correspond to the 25%–75% percentiles.

Solution quality vs. # of followers. Figure 4(a) shows the average visibilities achieved by our optimized intensities for the AVM task, normalized by the average visibility that the corresponding broadcasters' fitted intensities achieve, against number of followers for the same 10,000 broadcasters as above. Independently of the number of followers, we find that the intensities provided by our method consistently outperform the broadcaster's fitted intensities.

Visibility vs. k. Figure 4(b) shows the average visibility achieved by our optimized intensities for the AVM task against k for the four broadcasters from Figure 2.

Scalability. Figure 4(c) shows that our convex optimization framework easily scale to broadcasters with thousands of followers. For example, given a broadcaster with 2,000 followers, our algorithm takes $\sim$250 milliseconds to find the optimal intensity for the average visibility maximization using a single machine with 64 cores and 1.5 TB RAM.

8 Conclusions

In this paper, we developed a novel framework to solve the when-to-post problem, in which we model users' feeds and posts as discrete events occurring in continuous time. Under such continuous-time model, then choosing a strategy for a broadcaster becomes a problem of designing the conditional intensity of her posting events. The key technical idea that enables our framework is a novel formula which can link the conditional

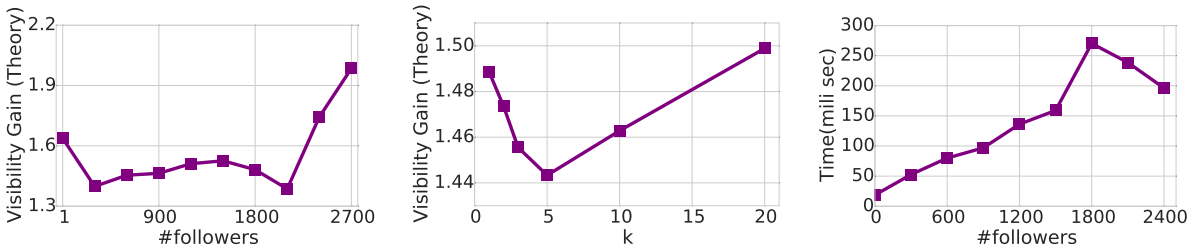

a) Followers b) k c) Running time

Figure 4: Panels show the average visibility against (a) # of followers and (b) k. Panel (c) plots running time.

intensity of an *arbitrary* broadcaster with her visibility in her followers' feeds, defined as the time that at least one post from her is among the most recent k received stories in her followers' feed. In addition to the framework, we develop an efficient gradient based optimization algorithm, which allows us to find optimal broadcast intensities for a variety of visibility shaping tasks in a matter of seconds. Experiments on large real-world data gathered from Twitter revealed that our framework can consistently make broadcasters' posts more visible than alternatives.

Our work also opens many interesting venus for future work. For example, we assume that the social network sorts stories in each user's feed in inverse chronological order. While this is a realistic assumption for some social networks (*e.g.*, Twitter), there are other social networks (*e.g.*, Facebook) where the feed is curated algorithmically. It would be very interesting to augment our framework to such cases. In this work, we model users' intensities using inhomogeneous Poisson processes, whose intensities are history independent and deterministic. Extending our framework to point processes with stochastic and history dependent intensity functions, such as Hawkes processes, would most likely provide more effective broadcasting strategies. In this work, we validate our framework on two visibility shaping tasks, average visibility maximization and minimax visibility maximization, however, there are many other useful tasks one may think of, such as visibility homogenization. Finally, it would be very interesting to investigate the scenario in which there are several smart broadcasters using our algorithm.

References

[1] O. Aalen, O. Borgan, and H. K. Gjessing. *Survival and event history analysis: a process point of view*. Springer, 2008.

[2] L. Backstrom, E. Bakshy, J. M. Kleinberg, T. M. Lento, and I. Rosenn. Center of attention: How facebook users allocate attention across friends. *ICWSM*, 2011.

[3] M. Cha, H. Haddadi, F. Benevenuto, and P. K. Gummadi. Measuring User Influence in Twitter: The Million Follower Fallacy. *ICWSM*, 2010.

[4] W. Chen, C. Wang, and Y. Wang. Scalable influence maximization for prevalent viral marketing in large-scale social networks. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2010.

[5] J. Cheng, L. Adamic, P. Dow, J. Kleinberg, and J. Leskovec. Can cascades be predicted? In *Proceedings of the 23rd international conference on World wide web*, 2014.

[6] T. Chenhao, L. Lee, and B. Pang. The effect of wording on message propagation: Topic- and author-controlled natural experiments on twitter. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, 2014.

[7] E. Constantinides. Foundations of social media marketing. *Procedia-Social and behavioral sciences*, 148:40–57, 2014.

[8] M. B. Crawford. *The world beyond your head: On becoming an individual in an age of distraction.* Macmillan, 2015.

[9] A. De, I. Valera, N. Ganguly, S. Bhattacharya, and M. Gomez-Rodriguez. Modeling opinion dynamics in diffusion networks. *arXiv preprint arXiv:1506.05474*, 2015.

[10] N. Du, L. Song, M. Gomez-Rodriguez, and H. Zha. Scalable influence estimation in continuous-time diffusion networks. In *Advances in Neural Information Processing Systems*, 2013.

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

[12] M. Farajtabar, Y. Wang, M. Gomez-Rodriguez, S. Li, H. Zha, and L. Song. Coevolve: A joint point process model for information diffusion and network co-evolution. In *Advances in Neural Information Processing Systems*, pages 1945–1953, 2015.

[13] M. Gomez-Rodriguez, K. Gummadi, and B. Schölkopf. Quantifying information overload in social media and its impact on social contagions. In *8th International AAAI Conference on Weblogs and Social Media*, 2014.

[14] N. Hodas and K. Lerman. How visibility and divided attention constrain social contagion. *SocialCom*, 2012.

[15] D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, 2003.

[16] G. Miritello, R. Lara, M. Cebrian, and E. Moro. Limited communication capacity unveils strategies for human interaction. *Scientific reports*, 3, 2013.

[17] N. Navaroli and P. Smyth. Modeling response time in digital human communication. In *Ninth International AAAI Conference on Web and Social Media*, 2015.

[18] M. Richardson and P. Domingos. Mining knowledge-sharing sites for viral marketing. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, 2002.

[19] N. Spasojevic, Z. Li, A. Rao, and P. Bhattacharyya. When-to-post on social networks. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 2127–2136. ACM, 2015.

[20] I. Valera and M. Gomez-Rodriguez. Modeling adoption and usage of competing products. In *IEEE International Conference on Data Mining*, 2015.

[21] D. Vu, D. Hunter, P. Smyth, and A. Asuncion. Continuous-time regression models for longitudinal networks. In *Advances in Neural Information Processing Systems*, 2011.

[22] Q. Zhao, M. Erdogdu, H. He, A. Rajaraman, and J. Leskovec. Seismic: A self-exciting point process model for predicting tweet popularity. In *21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2015.

A Proof of Lemma 1

We will prove this lemma by induction on k. For the case $k = 1$, $f_1(t)$ satisfies a first-order linear differential equation,

$$f_1'(t) = -(\mu(t) + \lambda(t))f_1(t) + \lambda(t), \tag{26}$$

whose unique solution is

$$f_1(t) = \int_0^t \lambda(\tau)e^{-\int_\tau^t(\lambda+\mu)(x)dx}. \tag{27}$$

as long as assuming $f_1(t) = 0$. Then, using that $\Gamma[1, \int_\tau^t \mu(x)dx] = e^{-\int_\tau^t \mu(x)dx}$, we can rewrite the solution as

$$f_1(t) = \int_0^t \lambda(\tau)e^{-\int_\tau^t \lambda(x)dx}\Gamma[1, \int_\tau^t \mu(x)dx]d\tau,$$

which proves the theorem for $k = 1$. Now, in the inductive step we assume the hypothesis is true for $1, 2, \ldots k-1$ and we prove it for k. We start by rewriting the differential given by Equation 15 as

$$f_k{}'(t) + (\mu(t) + \lambda(t))f_k(t) = \lambda(t) + \mu(t)f_{k-1}(t), \tag{28}$$

where, by assumption, $f_{k-1}(t)$ is unique and known. Then, as long as $f_k(0) = 0$, the above differential equation has a unique solution and thus we only need to find $f_k(t)$ that satisfies it. To do so, we rewrite the right hand side of the differential equation using the inductive hypothesis as

$$\lambda(t) + \mu(t)\frac{\int_0^t \lambda(\tau)e^{-\int_\tau^t \lambda(x)dx}\Gamma[k-1, \int_\tau^t \mu(x)dx]d\tau}{(k-2)!},$$

which, using $\Gamma[k-1, x] = \frac{1}{k-1}(\Gamma[k, x] + \frac{\partial\Gamma[k,x]}{\partial x})$, can be expressed as

$$\lambda(t) + \mu(t)\frac{\int_0^t \left(\lambda(\tau)e^{-\int_\tau^t \lambda(x)dx}\Gamma[k, \int_\tau^t \mu(x)dx] + \frac{\partial\Gamma[k,\int_\tau^t \mu(x)dx}{\partial \int_\tau^t \mu(x)dx}\right)d\tau}{(k-1)!} \tag{29}$$

Next, we hypothesize that

$$f_k(t) = \frac{\int_0^t \lambda(\tau)e^{-\int_\tau^t \lambda(x)dx}\,\Gamma[k, \int_\tau^t \mu(x)dx]\,d\tau)}{(k-1)!}, \tag{30}$$

and rewrite Eq. 29 as

$$\lambda(t) + \ \mu(t)f_k(t) + \frac{\mu(t)\int_0^t \lambda(\tau)e^{-\int_\tau^t \lambda(x)dx}\frac{\partial\Gamma[k,\int_\tau^t \mu(x)dx]}{\partial \int_\tau^t \mu(x)dx}d\tau}{(k-1)!}. \tag{31}$$

Then, by the fundamental theorem of calculus,

$$\frac{\partial\Gamma[k, \int_\tau^t \mu(x)dx]}{\partial \int_\tau^t \mu(x)dx} = \frac{\partial\Gamma[k, \int_\tau^t \mu(x)dx]}{\partial t} \times \frac{\partial t}{\partial \int_\tau^t \mu(x)dx} = \frac{\partial\Gamma[k, \int_\tau^t \mu(x)dx]}{\partial t} \times \frac{1}{\mu(t)}$$

and thus

$$\lambda(t) + \mu(t)f_k(t) + \frac{\int_0^t \lambda(\tau)e^{-\int_\tau^t \lambda(x)dx}\frac{\partial\Gamma[k,\int_\tau^t \mu(x)dx]}{\partial t}d\tau}{(k-1)!}. \tag{32}$$

Finally, using that for differentiable functions g and h, $gh' = (gh)' - g'h$, we have that

$$\int_0^t (\lambda(\tau)e^{-\int_\tau^t \lambda(x)dx})(\frac{\partial \Gamma[k, \int_\tau^t \mu(x)dx]}{\partial t})d\tau = \underbrace{\int_0^t \frac{\partial(\lambda(\tau)e^{-\int_\tau^t \lambda(x)dx}\Gamma[k, \int_\tau^t \mu(x)dx])}{\partial t}d\tau}_{(k-1)!(f_k'(t)-\lambda(t))} - \underbrace{\int_0^t (\frac{\partial \lambda(\tau)e^{-\int_\tau^t \lambda(x)dx}}{\partial t})(\Gamma[k, \int_\tau^t \mu(x)dx])d\tau}_{(k-1)!\lambda(t)f_k(t)}$$

and then we can rewrite Eq. 32 as

$$\lambda(t) + \mu(t)f_k(t) + \frac{(k-1)!(f_k'(t) - \lambda(t)) + (k-1)!\lambda(t)f_k(t)}{(k-1)!},$$

which simplifies to

$$f_k{}'(t) + (\mu(t) + \lambda(t))f_k(t).$$

This asserts that hypothesized solution for $f_k(t)$ in Eq. 30 satisfies Eq. 28, hence, it is the unique solution for $f_k(t)$.

B Proof of Theorem 3

From Lemma 1, we know that

$$f_k(t) = \frac{\int_0^t (\lambda(\tau)e^{-\int_\tau^t \lambda(x)dx})\Gamma[k, \int_\tau^t \mu(x)dx]d\tau}{(k-1)!}.$$

Using integration by parts, we can rewrite the above expression as

$$f_k(t) = 1 - \frac{e^{-\int_0^t \lambda(x)dx}\Gamma[k, \int_0^t \mu(x)dx]}{(k-1)!} - \frac{\int_0^t (e^{-\int_\tau^t \lambda(x)dx})\frac{\partial \Gamma[k, \int_\tau^t \mu(x)dx]}{\partial \tau}d\tau}{(k-1)!}$$

Lemma 9 tells us that $e^{-\int_0^t \lambda(x)dx}$ and $e^{-\int_\tau^t \lambda(x)dx}$ are convex with respect to $\lambda(\cdot)$. Moreover, using Lemma 10 and the fact that $\frac{\partial \Gamma[k, \int_\tau^t \mu(x)dx]}{\partial \tau} > 0$, it follows that the function $\int_0^t (e^{-\int_\tau^t \lambda(x)dx})\frac{\partial \Gamma[k, \int_\tau^t \mu(x)dx]}{\partial \tau}d\tau$ is convex. Finally, given that $\Gamma[k, \int_0^t 0] > 0$, we can conclude that $f_k(t)$ is concave with respect to $\lambda(\cdot)$.

Lemma 9. *Functional $J[\lambda] = e^{-\int_a^t \lambda(x)dx}$ is convex with respect to $\lambda(\cdot)$ for any constant $a \leq t$.*

Proof. We simply verify that $J[\lambda]$ satisfies the definition of convexity, as given by Eq. 17:

$$J[\alpha\lambda_1 + (1-\alpha)\lambda_2] = e^{-\int_a^t \alpha\lambda_1(x) + (1-\alpha)\lambda_2(x)\,dx} \leq \alpha e^{-\int_a^t \lambda_1(x)\,dx} + (1-\alpha)e^{-\int_a^t \lambda_2(x)\,dx} = \alpha J[\lambda_1] + (1-\alpha)J[\lambda_2]$$

where the inequality follows from the arithmetic-geometric mean inequality, *i.e.*, $\theta x + (1-\theta)y \geq x^\theta y^{1-\theta}$ for all positive x, y, and $0 < \theta < 1$. □

Lemma 10. *If the functional $J_t[\lambda(\cdot)]$ is convex with respect to $\lambda(\cdot)$. Then, given any arbitrary function $g(x) \geq 0$, the functional $L[\lambda] = \int_0^t J_\tau[\lambda(.)]g(\tau)d\tau$ is also convex with respect to $\lambda(\cdot)$.*

Proof. We verify that the functional $L[\lambda] = \int_0^t J_\tau[\lambda(.)]g(\tau)d\tau$ verifies the definition of convexity, as given by Eq. 17:

$$\begin{aligned} L[\alpha\lambda_1 + (1-\alpha)\lambda_2] &= \int_0^t J_\tau[\alpha\lambda_1 + (1-\alpha)\lambda_2]g(\tau)d\tau \leq \alpha \int_0^t J_\tau[\lambda_1]g(\tau)d\tau + (1-\alpha)\int_0^t J_\tau[\lambda_2]g(\tau)d\tau \\ &= \alpha L[\lambda_1] + (1-\alpha)L[\lambda_2] \end{aligned}$$

where the inequality holds using that, given any two arbitrary functions h_1 and h_2 such that $h_1(x) \geq h_2(x) \geq 0$ for all $x \in \mathcal{D}$, then $\int_{\mathcal{D}} h_1(x)g(x)dx \geq \int_{\mathcal{D}} h_2(x)g(x)dx$ given $g(x) \geq 0$ for all $x \in \mathcal{D}$. □

C Proof of Theorem 4

Theorem 3 proves the concavity of $f_k(t)$ with respect to $\lambda(\cdot)$. Therefore, $1 - f_k(t)$ is convex and, using Lemma 10, it holds that $\int_0^T (1 - f_k(t))s(t)dt = \int_0^T s(t)dt - \int_0^T f_k(t)s(t)dt$ is also convex. Then, since $\int_0^T s(t)dt$ is constant, $\int_0^T f_k(t)s(t)dt$ is concave with respect to $\lambda(\cdot)$ and the proof is complete.

D Proof of Lemma 6

Each piecewise continues function can be represented as summation of a number of *heaviside* step functions. The count is equal to the number of discontinuity points. However, each heaviside function itself is the limit of smooth *tanh* functions. Therefore, the piecewise continues function will be the limit of a finite summation of smooth *tanh* functions.

E Proof of Theorem 7

Consider two piecewise constant functions $\lambda(\cdot), \mu(\cdot) \in \mathcal{G}$. According to Lemma 6 there exist sequence of smooth functions such that $\lim_{n\to\infty} \lambda_n = \lambda$ and $\lim_{n\to\infty} \lambda'_n = \lambda'$. Because of the concavity of f_k in $\mathcal{H}$ we know for $0 < \alpha < 1$:

$$f_k[\alpha\lambda_n(\cdot) + (1-\alpha)\lambda'_n(\cdot)] \geq \alpha f_k[\lambda_n(\cdot)] + (1-\alpha) f_k[\lambda'_n(\cdot)].$$

Taking the limit and using the continuity of f_k we get:

$$f_k[\alpha\lambda(\cdot) + (1-\alpha)\lambda'(\cdot)] \geq \alpha f_k[\lambda(\cdot)] + (1-\alpha) f_k[\lambda'(\cdot)]. \tag{33}$$

Accompanied with convexity of space $\mathcal{G}$ the theorem is proved.