

# Multistage Campaigning in Social Networks

Mehrdad Farajtabar<sup>1</sup>, Xiaojing Ye<sup>2</sup>, Sahar Harati<sup>3</sup>, Le Song<sup>1</sup>, and Hongyuan Zha<sup>1</sup>

<sup>1</sup>Georgia Institute of Technology, mehrdad@gatech.edu, {lsong,zha}@cc.gatech.edu

<sup>2</sup>Georgia State University, xye@gsu.edu

<sup>3</sup>Emory University, sahar.harati@emory.edu

## Abstract

We consider the problem of how to optimize multi-stage campaigning over social networks. The dynamic programming framework is employed to balance the high present reward and large penalty on low future outcome in the presence of extensive uncertainties. In particular, we establish theoretical foundations of optimal campaigning over social networks where the user activities are modeled as a multivariate Hawkes process, and we derive a time dependent linear relation between the intensity of exogenous events and several commonly used objective functions of campaigning. We further develop a convex dynamic programming framework for determining the optimal intervention policy that prescribes the required level of external drive at each stage for the desired campaigning result. Experiments on both synthetic data and the real-world MemeTracker dataset show that our algorithm can steer the user activities for optimal campaigning much more accurately than baselines.

## 1 Introduction

Obama was the first US president in history who successfully leveraged online social media in presidential campaigning, which has been popularized and become a ubiquitous approach to electoral politics (such as in the on-going 2016 US presidential election) in contrast to the decreasing relevance of traditional media such as TV and newspapers [1, 2]. The power of campaigning via social media in modern politics is a consequence of online social networking being an important part of people’s regular daily social lives. It has been quite common that individuals use social network sites to share their ideas and comment on other people’s opinions. In recent years, large organizations, such as governments, public media, and business corporations, also start to announce news, spread ideas, and/or post advertisements in order to steer the public opinion through social media platform. There has been extensive interest for these entities to influence the public’s view and manipulate the trend by incentivizing influential users to endorse their ideas/merits/opinions at certain monetary expenses or credits. To obtain most cost-effective trend manipulations, one needs to design an optimal campaigning strategy or policy such that quantities of interests, such as influence of opinions, exposure of a campaign, adoption of new products, can be maximized or steered towards the target amount given realistic budget constraints.

The key factor differentiating social networks from traditional media is *peer influence*. In fact, events in an online social network can be categorized roughly into two types: endogenous events where users just respond to the actions of their neighbors within the network, and exogenous events where users take actions due to drives external to the network. Then it is natural to raise the following fundamental questions regarding optimal campaigning over social networks: can we model and exploit those event data to steer the online community to a desired exposure level? More specifically, can we drive the overall exposure to a campaign to a certain level (e.g., at least twice per week per user) by incentivizing a small number of users to take more initiatives? What about maximizing the overall exposure for a target group of people?

More importantly, those exposure shaping tasks are more effective when the interventions are implemented in multiple stages. Due to the inherent uncertainty in social behavior, the outcome of each intervention may




Figure 1: A social network and multi-stage campaigning to maximize the minimum exposure.

not be fully predictable but can be anticipated to some extent before the next intervention happens. A key aspect of such situations is that interventions can't be viewed in isolation since one must balance the desire for high present reward with the penalty of low future outcome.

In this paper, the *dynamic programming* framework [3] is employed to tackle the aforementioned issues. In particular, we first establish the fundamental theory of optimal campaigning over social networks where the user activities are modeled as a multivariate Hawkes process (MHP) [4, 5] since MHP can capture both endogenous and exogenous event intensities. We also derive a time dependent linear relation between the intensity of exogenous events and the overall exposure to the campaign. Exploiting this connection, we develop a convex dynamic programming framework for determining the optimal intervention policy that prescribes the required level of external drive at each stage in order for the campaign to reach a desired exposure profile. We propose several objective functions that are commonly considered as campaigning criteria in social networks. Experiments on both synthetic data and real world network of news websites in the MemeTracker dataset show that our algorithms can shape the exposure of campaigns much more accurately than baselines.

## 2 Illustrative Example

Next sections will define the campaigning problem formally. However, we found it beneficial to formulate it more intuitively. Next section will define all the concepts appeared here rigorously.

Fig. 1-a shows a hypothetical social network highlighting 4 users and their influence on each other. They have their own set of followers which for simplicity are considered disjoint and being influenced equally. Our objective in this toy example is to maximize the minimum exposure on the followers. For the ease of exposition, we only consider 2 stages which in the beginning we can intervene. Hypothetically, we are going to advertise iPhone 7 on behalf of Apple Inc.

Now, we intuitively proceed to find the optimum intervention. At the first stage, where the state is zero (no exposure from the past), intervention through Jacob and Christine would not be part of the solution (or their share would be negligible), because the outcome would be small due to the low number of total followers and the small amount of influence. Between Bob and Sophie, we would select her. Note that Sophie has a high influence on Bob and if she is incentivized to make a blog post for the campaign, she can make Bob interested too. Therefore, the procedure will assign most budget to Sophie. Furthermore, Bob has a high influence on Jacob and will probably inspire him to do an activity. Let's follow a hypothetical scenario. Sophie makes her post at 9:35am about apple's new phone. Bob is notified about her post and will make an inspiring post about the phone's camera. However, Jacob will not respond to this campaign well, e.g., he may be traveling that day. This phenomenon is perfectly captured via the probabilistic framework for event analysis [6, 4]. On the other hand, Christine who is not a big fan of Bob's post compared to Jacob will respond to Bob's post by writing a comparison between Apple and Samsung's product at 3:47pm. This will reinforce Bob to make a separate post explaining the features of Apple's new product. These events are illustrated in Fig. 1-b.

The cascade of events at the first stage will expose followers of Sophia, Bob, and Christine to the new product. However, all things do not go according to our expectation. Jacob happens to be out of his office that day and his followers are not aware of the news. Considering the MEM as the objective function, this is quite disappointing that there are people who are exposed to the campaign 0 times. Exposures of followers are demonstrated in Fig. 1-b. This is the point where multi-stage dynamic campaigning helps. Considering the exposure intensity of followers in Fig. 1-c, Christine's, Bob's, and even Sophia's followers are still under the influence. Furthermore, they may follow their initial posts by new ones. Then, in the next stage, it looks reasonable to invest on Jacob. At least his followers will notify about the campaign and also as Jacob has an influence on Sophie, his posts may inspire Sophie again and this may increase the exposure intensity of followers of Sophie as well. Therefore, thanks to a second chance in intervention, Jacob will get the most budget at the second stage to advertise the product.

## 3 Basics and Background

An  $n$ -dimensional temporal point process is a random process whose realization consists of a list of discrete events in time and their associated dimension,  $\{(t_k, d_k)\}$  with  $t_k \in \mathbb{R}^+$  and  $d_k \in \{1, \dots, n\}$ . Many different types of data produced in online social networks can be represented as temporal point processes, such as likes and tweets. A temporal point process can be equivalently represented as a counting process,  $\mathcal{N}(t) = (\mathcal{N}^1(t), \dots, \mathcal{N}^n(t))^T$  associated to  $n$  users in the social network. Here,  $\mathcal{N}^i(t)$  records the number of events user  $i$  performs before time  $t$  for  $1 \leq i \leq n$ . Let the history  $\mathcal{H}^i(t)$  be the list of times of events  $\{t_1, t_2, \dots, t_k\}$  of the  $i$ -th user up to time  $t$ . Then, the number of observed events in a small time window  $[t, t + dt)$  of length  $dt$  is

$$d\mathcal{N}^i(t) = \sum_{t_k \in \mathcal{H}^i(t)} \delta(t - t_k) dt, \quad (1)$$

and hence  $\mathcal{N}^i(t) = \int_0^t d\mathcal{N}^i(s)$ , where  $\delta(t)$  is a Dirac delta function. The point process representation of temporal data is fundamentally different from the discrete time representation typically used in social network analysis. It directly models the time interval between events as random variables, avoids the need

to pick a time window to aggregate events, and allows temporal events to be modeled in a fine grained fashion. Moreover, it has a remarkably rich theoretical support [6].

An important way to characterize temporal point processes is via the conditional intensity function — a stochastic model for the time of the next event given all the times of previous events. Formally, the conditional intensity function  $\lambda^i(t)$  (intensity, for short) of user  $i$  is the conditional probability of observing an event in a small window  $[t, t + dt)$  given the history  $\mathcal{H}(t) = \{\mathcal{H}^1(t), \dots, \mathcal{H}^n(t)\}$ :

$$\lambda^i(t)dt := \mathbb{P}\{\text{user } i \text{ performs event in } [t, t + dt) \mid \mathcal{H}(t)\} = \mathbb{E}[d\mathcal{N}^i(t) \mid \mathcal{H}(t)], \quad (2)$$

where one typically assumes that only one event can happen in a small window of size  $dt$ . Fig. 1-b and Fig. 1-c show the events and intensity function of the activities of the 4 users in the social network respectively. Furthermore, these two can be seen as the exposure events and the exposure intensity function to their followers since for simplicity we are just considering the activities of these 4 users in the network. The functional form of the intensity  $\lambda^i(t)$  is often designed to capture the phenomena of interests.

The Hawkes process [7] is a class of self and mutually exciting point process models,

$$\lambda^i(t) = \mu^i(t) + \sum_{k:t_k < t} \phi^{id_k}(t, t_k) = \mu^i(t) + \sum_{j=1}^n \int_0^t \phi^{ij}(t, s) d\mathcal{N}^j(s), \quad (3)$$

where the intensity is history dependent.  $\phi^{ij}(t, s)$  is the impact function capturing the temporal influence of an event by user  $j$  at time  $s$  to the future events of user  $j$  at time  $t \geq s$ . Here, the first term  $\mu^i(t)$  is the exogenous event intensity modeling drive outside the network and indecent of the history, and the second term  $\sum_{k:t_k < t} \phi^{id_k}(t, t_k)$  is the endogenous event intensity modeling interactions within the network [8]. Defining  $\Phi(t, s) = [\phi^{ij}(t, s)]_{i,j=1\dots n}$ , and  $\lambda(t) = (\lambda^1(t), \dots, \lambda^n(t))^\top$ , and  $\mu(t) = (\mu^1(t), \dots, \mu^n(t))^\top$  we can compactly rewrite Eq 3 in matrix form:

$$\lambda(t) = \mu(t) + \int_0^t \Phi(t, s) d\mathcal{N}(s). \quad (4)$$

In practice it is standard to employ shift-invariant impact function, *i.e.*,  $\Phi(t, s) = \Phi(t - s)$ . Then, by using notation of convolution  $f(t) * g(t) = \int_0^t f(t - s)g(s)ds$  we have

$$\lambda(t) = \mu(t) + \Phi(t) * d\mathcal{N}(t). \quad (5)$$

## 4 From Intensity to Average Activity

In this section we will develop a closed form relation between the expected total intensity  $\mathbb{E}[\lambda(t)]$  and the intensity  $\mu(t)$  of exogenous events. This relation establish the basis of our campaigning framework. First, define the *mean function* as  $\mathcal{M}(t) := \mathbb{E}[\mathcal{N}(t)] = \mathbb{E}_{\mathcal{H}(t)}[\mathbb{E}(\mathcal{N}(t) \mid \mathcal{H}(t))]$ . Note that  $\mathcal{M}(t)$  is history independent, and it gives the average number of events up to time  $t$  for each of the dimension. Similarly, the *rate function*  $\eta(t)$  is given by  $\eta(t)dt := d\mathcal{M}(t)$ . On the other hand,

$$d\mathcal{M}(t) = d\mathbb{E}[\mathcal{N}(t)] = \mathbb{E}_{\mathcal{H}(t)}[\mathbb{E}(d\mathcal{N}(t) \mid \mathcal{H}(t))] = \mathbb{E}_{\mathcal{H}(t)}[\lambda(t) \mid \mathcal{H}(t)]dt = \mathbb{E}[\lambda(t)]dt. \quad (6)$$

Therefore  $\eta(t) = \mathbb{E}[\lambda(t)]$  which serves as a measure of activity in the network. In what follows we will find an analytical form for the average activity. Proofs are presented in Appendix A.

**Lemma 1.** Suppose  $\Psi : [0, T] \rightarrow \mathbb{R}^{n \times n}$  is a non-increasing matrix function, then for every fixed constant intensity  $\mu(t) = c \in \mathbb{R}_+^n$ ,  $\eta_c(t) := \Psi(t)c$  solves the semi-infinite integral equation

$$\eta(t) = c + \int_0^t \Phi(t - s)\eta(s)ds, \quad \forall t \in [0, T], \quad (7)$$

if and only if  $\Psi(t)$  satisfies

$$\Psi(t) = I + \int_0^t \Phi(t-s)\Psi(s)ds, \quad \forall t \in [0, T]. \quad (8)$$

In particular, if  $\Phi(t) = Ae^{-\omega t}\mathbf{1}_{\geq 0}(t) = [a_{ij}e^{-\omega t}\mathbf{1}_{\geq 0}(t)]_{ij}$  where  $0 \leq \omega \notin \text{Spectrum}(A)$ , then

$$\Psi(t) = e^{(A-\omega I)t} + \omega(A-\omega I)^{-1}(e^{(A-\omega I)t} - I) \quad (9)$$

for  $t \in [0, T]$ , where,  $\mathbf{1}_{\geq 0}(t)$  is an indicator function for  $t \geq 0$ .

Let  $\mu : [0, T] \rightarrow \mathbb{R}_+^n$  be a right-continuous piecewise constant function

$$\mu(t) = \sum_{m=1}^M c_m \mathbf{1}_{[\tau_{m-1}, \tau_m)}(t), \quad (10)$$

where  $0 = \tau_0 < \tau_1 < \dots < \tau_M = T$  is a finite partition of time interval  $[0, T]$  and function  $\mathbf{1}_{[\tau_{m-1}, \tau_m)}(t)$  indicates  $\tau_{m-1} \leq t < \tau_m$ . The next theorem shows that if  $\Psi(t)$  satisfies (8), then one can calculate  $\eta(t)$  for piecewise constant intensity  $\mu : [0, T]$  of form (10).

**Theorem 2.** Let  $\Psi(t)$  satisfy (8) and  $\mu(t)$  be a right-continuous piecewise constant intensity function of form (10), then the rate function  $\eta(t)$  is given by

$$\eta(t) = \sum_{k=0}^m \Psi(t - \tau_k)(c_k - c_{k-1}), \quad (11)$$

for all  $t \in (\tau_{m-1}, \tau_m]$  and  $m = 1, \dots, M$ , where  $c_{-1} := 0$  by convention.

Using the above lemma, for the first time, we derive the average intensity for a general exogenous intensity. Section 7 includes a few experiments to investigate these results empirically.

**Theorem 3.** If  $\Psi \in C^1([0, T])$  and satisfies (8), and exogenous intensity  $\mu$  is bounded and piecewise absolutely continuous on  $[0, T]$  where  $\mu(t+) = \mu(t)$  at all discontinuous points  $t$ , then  $\mu$  is differentiable almost everywhere, and the semi-indefinite integral

$$\eta(t) = \mu(t) + \int_0^t \Phi(t-s)\eta(s)ds, \quad \forall t \in [0, T], \quad (12)$$

yields a rate function  $\eta : [0, T] \rightarrow \mathbb{R}_+^n$  given by

$$\eta(t) = \int_0^t \Psi(t-s)d\mu(s). \quad (13)$$

**Corollary 4.** Suppose  $\Psi$  and  $\mu$  satisfy the same conditions as in Thm. 3, and define  $\psi = \Psi'$ , then the rate function is  $\eta(t) = (\psi * \mu)(t)$ . In particular, if  $\Phi(t) = Ae^{-\omega t}\mathbf{1}_{\geq 0}(t) = [a_{ij}e^{-\omega t}\mathbf{1}_{\geq 0}(t)]_{ij}$  then the rate function  $\eta(t) = \mu(t) + A \int_0^t e^{(A-\omega I)(t-s)}\mu(s)ds$ .

## 5 Multi-stage Closed-loop Control Problem

Given the analytical relation between exogenous intensity and expected overall intensity (rate function), one can solve a single one-stage campaigning problem to find the optimal constant intervention intensity [8]. Alternatively, the time window can be partitioned into multiple stages and one can impose different levels of interventions in these stages. This yields an open-loop optimization of the cost function where one selects all the intervention actions at initial time 0. More effectively, we tackle the campaigning problem in a dynamic and adaptive manner where we can postpone deciding the intervention by observing the process until the next stage begins. This is called the *closed-loop* optimization of the objective function.

In this section, we establish the foundation to formulate the problem as a multi-stage closed-loop optimal control problem. We assume that  $n$  users are generating events according to multi-dimensional Hawkes process with exogenous intensity  $\mu(t) \in \mathbb{R}^n$  and impact function  $\Phi(t, s) \in \mathbb{R}^{n \times n}$ .

## 5.1 Event exposure

Event exposure is the quantity of major interests in campaigning. The exposure process is mathematically represented as a counting process,  $\mathcal{E}(t) = (\mathcal{E}^1(t), \dots, \mathcal{E}^n(t))^\top$ : Here,  $\mathcal{E}^i(t)$  records the number of times user  $i$  is exposed (she or one of her neighbors performs an activity) to the campaign by time  $t$ . Let  $B$  be the adjacency matrix of the user network, *i.e.*,  $b_{ij} = 1$  if user  $i$  follows user  $j$  or equivalently user  $j$  influences user  $i$ . We assume  $b_{ii} = 1$  for all  $i$ . Then the exposure process is given by  $\mathcal{E}(t) = B\mathcal{N}(t)$ .

## 5.2 Stages and interventions

Let  $[0, T]$  be the time horizon and  $0 = \tau_0 < \tau_1 < \dots < \tau_{M-1} < \tau_M = T$  be a partition into the  $M$  stages. In order to steer the activities of network towards a desired level (criteria given below) at these stages, we impose a constant intervention  $u_m \in \mathbb{R}^n$  to the existing exogenous intensity  $\mu$  during time  $[\tau_m, \tau_{m+1})$  for each stage  $m = 0, 1, \dots, M-1$ . The activity intensity at the  $m$ -th stage is  $\lambda_m(t) = \mu + u_m + \int_0^t \Phi(t, s) d\mathcal{N}(s)$  for  $\tau_m \leq t < \tau_{m+1}$  where  $\mathcal{N}(t)$  tracks the counting process of activities since  $t = 0$ . Note that the intervention itself exhibits a stochastic nature: adding  $u_m^i$  to  $\mu^i$  is equivalent to incentivizing user  $i$  to increase her activity rate but it is still uncertain when she will perform an activity, which appropriately mimics the randomness in real-world campaigning.

## 5.3 States and state evolution

Note that the Hawkes process is non-Markov and one needs complete knowledge of the history to characterize the entire process. However, the conditional intensity  $\lambda(t)$  only depends on the state of process at time  $t$  when the standard exponential kernel  $\Phi(t, s) = Ae^{-\omega(t-s)}\mathbf{1}_{\geq 0}(t-s)$  is employed. In this case, the activity rate at stage  $m$  is

$$\lambda_m(t) = \mu + u_m + \underbrace{\int_0^{\tau_m} Ae^{-\omega(t-s)} d\mathcal{N}(s)}_{\text{from previous stages}} + \underbrace{\int_{\tau_m}^t Ae^{-\omega(t-s)} d\mathcal{N}(s)}_{\text{current stage}} \quad (14)$$

Define  $x_m := \lambda_{m-1}(\tau_m) - u_{m-1} - \mu$  (and  $x_0 = 0$  by convention) then the intensity due to events of all previous  $m$  stages can be written as  $\int_0^{\tau_m} Ae^{-\omega(t-s)} d\mathcal{N}(s) = x_m e^{-\omega(t-\tau_m)}$ . In other words,  $x_m$  is sufficient to encode the information of activity in the past  $m$  stages that is relevant to future. This is in sharp contrast to the general case where the state space grows with the number of events.

## 5.4 Objective function

For a sequence of controls  $u(t) = \sum_{m=0}^{M-1} u_m \mathbf{1}_{[\tau_m, \tau_{m+1})}(t)$ , the activity counting process  $\mathcal{N}(t)$  is generated by intensity  $\lambda(t) = \mu + u(t) + \int_0^t Ae^{-\omega(t-s)} d\mathcal{N}(s)$ . For each stage  $m$  from 0 to  $M-1$ ,  $x_m$  encodes the effects from previous  $m$  stages as above and  $u_m$  is the current control imposed at this stage. Let  $\mathcal{E}_m^i(t; x_m, u_m) := B \int_{\tau_m}^t d\mathcal{N}^i(s)$  be the number of times user  $i$  is exposed to the campaign by time  $t \in [\tau_m, \tau_{m+1})$  in stage  $m$ , then the goal is to steer the expected total number of exposure  $\bar{\mathcal{E}}_m^i(x_m, u_m) := \mathbb{E}[\mathcal{E}_m^i(\tau_{m+1}; x_m, u_m)]$  to a desired level. In what follows, we introduce several instances of the objective function  $g(x_m, u_m)$  in terms of  $\{\bar{\mathcal{E}}_m^i(x_m, u_m)\}_{i=1}^n$  in each stage  $m$  that characterize different *exposure shaping* tasks. Then the overall control problem is to find  $u(t)$  that optimizes the total objective  $\sum_{m=0}^{M-1} g_m(x_m, u_m)$ .

- *Capped Exposure Maximization (CEM)*: In real networks, there is a cap on the exposure each user can tolerate due to the limited attention of a user. Suppose we know the upper bound  $\beta_m^i$ , on user  $i$ 's exposure tolerance over which the extra exposure is not counted towards the objective. Then, we can form the following *capped exposure maximization*

$$g_m(x_m, u_m) = \frac{1}{n} \sum_{i=1}^n \min \{ \bar{\mathcal{E}}_m^i(x_m, u_m), \beta_m^i \} \quad (15)$$

- *Minimum Exposure Maximization (MEM)*: Suppose our goal is instead to maintain the exposure of campaign on each user above a certain minimum level, at each stage or, alternatively to make the user with the minimum exposure as exposed as possible, we can consider the following cost function:

$$g_m(x_m, u_m) = \min_i \bar{\mathcal{E}}_m^i(x_m, u_m) \quad (16)$$

- *Least-squares Exposure Shaping (LES)*: Sometimes we want to achieve a pre-specified target exposure levels,  $\gamma_m \in \mathbb{R}^n$ , for the users. For example, we may like to divide users into groups and desire a different level of exposure in each group. To this end, we can perform least-squares campaigning task with the following cost function where  $D$  encodes potentially additional constraints (e.g., group partitions):

$$g_m(x_m, u_m) = -\frac{1}{n} \|D\bar{\mathcal{E}}_m(x_m, u_m) - \gamma_m\|^2 \quad (17)$$

## 5.5 Policy and actions

By observing the counting process in previous stages (summarized in a sequence of  $x_m$ ) and taking the future uncertainty into account, the control problem is to design a policy  $\pi = \{\pi_m : \mathbb{R}^n \rightarrow \mathbb{R}^n : m = 0, \dots, M-1\}$  such that the controls  $u_m = \pi_m(x_m)$  can maximize the total objective  $\sum_{m=0}^{M-1} g_m(x_m, u_m)$ . In addition, we may have constraints on the amount of control. For example, a budget constraint on the sum of all interventions to users at each stage, or, a cap over the amount of intensity a user can handle. A feasible set or an action space over which we find the best intervention is represented as  $\mathcal{U}_m := \{u_m \in \mathbb{R}^n | c_m^\top u_m \leq C_m, 0 \leq u_m \leq \alpha_m\}$ . Here,  $c_m \in \mathbb{R}_+^n$  contains the price of each person per unit increase of exogenous intensity and  $C_m \in \mathbb{R}_+$  is the total budget at stage  $m$ . Also,  $\alpha_m \in \mathbb{R}_+^n$  is the cap on the amount of activities of the users.

To summarize, the following problem is formulated to find the optimal control policy  $\pi$ :

$$\begin{aligned} & \text{maximize}_{\pi} \sum_{m=0}^{M-1} g_m(x_m, \pi_m(x_m)), \\ & \text{subject to } \pi_m(x_m) \in \mathcal{U}_m, \text{ for } m = 0, \dots, M-1. \end{aligned} \quad (18)$$

## 6 Closed-loop Dynamic Programming Solution

We have formulated the control problem as an optimization in (18). However, when control policy  $\pi_m$  is to be implemented, only  $x_m$  is observed and there are still uncertainties in future  $\{x_{m+1}, \dots, x_{M-1}\}$ . For instance, when  $\pi_m$  is implemented according to  $x_m$  starting from time  $\tau_m$ , the intensity  $x_{m+1} := f(x_m, \pi_m(x_m))$  at time  $\tau_{m+1}$  depends on  $x_m$  and the control  $\pi_m(x_m)$ , but is also random due to the stochasticity of the process during time  $[\tau_m, \tau_{m+1})$ . Therefore, the design of  $\pi$  needs to take future uncertainties into considerations.

Suppose we have arrived at stage  $M$  at time  $\tau_{M-1}$  with observation  $x_{M-1}$ , then the optimal policy  $\pi_{M-1}$  satisfies  $g_{M-1}(x_{M-1}, \pi_{M-1}(x_{M-1})) = \max_{u \in \mathcal{U}_{M-1}} g_{M-1}(x_{M-1}, u) =: J_{M-1}(x_{M-1})$ . We then repeat this procedure for  $m$  from  $M-1$  to 0 backward to find the sequence of controls via dynamic programming such that the control  $\pi_m(x_m) \in \mathcal{U}_m$  yields optimal objective value

$$J_m(x_m) = \max_{u_m \in \mathcal{U}_m} \mathbb{E}[g_m(x_m, u_m) + J_{m+1}(f(x_m, u_m))] \quad (19)$$

### 6.1 Approximate dynamic programming

Solving (19) for finding  $J_m(x_m)$  analytically is intractable. Therefore, we will adopt an approximate dynamic programming scheme. In fact approximate control is as essential part of dynamic programming as the optimization is usually intractable due to curse of dimensionality except a few especial cases [3]. Here we adopt a suboptimal control scheme, *certainty equivalent control* (CEC), which applies at each stage the control that would be optimal if the uncertain quantities were fixed at some typical values like the

---

**Algorithm 1:** Closed-loop Multi-stage Dynamic Programming
 

---

**Input:** Intervention constraints:  $c_0 \dots c_{M-1}$ ,  $C_0 \dots C_{M-1}$ ,  $\alpha_0 \dots \alpha_{M-1}$ ,  
**Input:** Objective-specific constraints:  $\beta_0 \dots \beta_{M-1}$  for CEM and  $\gamma_0 \dots \gamma_{M-1}$  for LES  
**Input:** Time:  $T$ , Hawkes parameters:  $A, \omega$   
**Output:** Optimal intervention  $u_0 \dots u_{M-1}$ , Optimal cost:  $Cost$   
 Set  $x_0 \leftarrow 0$   
 Set  $Cost \leftarrow 0$   
**for**  $l \leftarrow 0 : M - 1$  **do**  
    $(v_l \dots v_{M-1}) = open\_loop(x_l)$  (Problems (30), (31), (32) for CEM, MEM, LES respectively)  
   Set  $u_l \leftarrow v_l$   
   Drop  $v_{l+1} \dots v_{M-1}$   
   Update next state  $x_{l+1} \leftarrow f_l(x_l, u_l)$   
   Update  $Cost \leftarrow Cost + g_l(x_l, u_l)$

---

average behavior. It results in an optimal control sequence, the first component of which is used at the current stage, while the remaining components are discarded. The procedure is repeated for the remaining stages. Algorithm 1 summarizes the dynamic programming steps. This algorithm has two parts: (i) certainty equivalence which the random behavior is replaced by its average; and (ii) the open-loop optimization. Let's assume we are at the beginning of stage  $l$  of the Alg. 1 with state vector  $x_l$  at  $\tau_l$ .

## 6.2 Certainty equivalence

We use the machinery developed in Sec. 4 to compute the average of exposure at any stage  $m = l, l + 1, \dots, M - 1$ .

$$\bar{\mathcal{E}}_m(x_m, u_m) = B\mathbb{E}[\mathcal{N}(\tau_{m+1}) - \mathcal{N}(\tau_m)] = B\mathbb{E}\left[\int_{\tau_m}^{\tau_{m+1}} d\mathcal{N}(s)\right] = B\int_{\tau_m}^{\tau_{m+1}} \eta_m(s) ds \quad (20)$$

where  $\eta_m(t) = \mathbb{E}[\lambda_m(t)]$  and  $\lambda_m(t) = \mu + u_m + x_l e^{-\omega(t-\tau_l)} + \int_{\tau_l}^t A e^{-\omega(t-s)} d\mathcal{N}(s)$  for  $t \in [\tau_m, \tau_{m+1})$ . Now, we use the superposition property of point processes [4] to decompose the process as  $\mathcal{N}(t) = \mathcal{N}^c(t) + \mathcal{N}^v(t)$  corresponding to  $\lambda_m(t) = \lambda_m^c(t) + \lambda_m^v(t)$  where the first  $\lambda_m^c(t) = \mu + u_m + \int_{\tau_l}^t A e^{-\omega(t-s)} d\mathcal{N}^c(s)$  consists of events caused by exogenous intensity at current stage  $m$  and the second  $\lambda_m^v(t) = x_l e^{-\omega(t-\tau_l)} + \int_{\tau_l}^t A e^{-\omega(t-s)} d\mathcal{N}^v(s)$  is due to activities in previous stages. According to Thm. 2 we have

$$\eta_m^c(t) := \mathbb{E}[\lambda_m^c(t)] = \Psi(t - \tau_l)\mu + \Psi(t - \tau_l)u_l + \sum_{k=l+1}^{m-1} \Psi(t - \tau_k)(u_k - u_{k-1}), \quad (21)$$

and according to Thm. 3 we have

$$\eta_m^v(t) := \mathbb{E}[\lambda_m^v(t)] = \int_{\tau_l}^t \Psi(t - s) d(x_l e^{-\omega(s-\tau_l)} \mathbf{1}_{[\tau_l, \infty)}(s)). \quad (22)$$

From now on, for simplicity, we assume stages are based on equal partition of  $[0, T]$  to  $M$  segments where each has length  $\Delta_M$ . Combining Eq. (20) and  $\eta_m(t) = \eta_m^c(t) + \eta_m^v(t)$  yields:

$$\begin{aligned} \bar{\mathcal{E}}_m(x_m, u_m) = & \Gamma((m-l+1)\Delta_M)u_l + \Gamma((m-l)\Delta_M)(u_{l+1} - u_l) + \dots \\ & + \Gamma(\Delta_M)(u_m - u_{m-1}) + \Gamma((m-l+1)\Delta_M)\mu + \Upsilon((m-l+1)\Delta_M)x_l \end{aligned} \quad (23)$$



where  $\Gamma(t)$  and  $\Upsilon(t)$  are matrices independent of  $u_m$ 's and are defined as:

$$\Upsilon(t) = B \int_0^t e^{(A-\omega I)s} ds = B(A - \omega I)^{-1}(e^{(A-\omega I)t} - I) \quad (24)$$

$$\Gamma(t) = B \int_0^t \Psi(s) ds = B\Upsilon(t) + B(A - \omega I)^{-1}(\Upsilon(t) - It)/\omega; \quad (25)$$

Note the linear relation between average exposure  $\bar{\mathcal{E}}_m(x_m, u_m)$  and intervention values  $u_l, \dots, u_{m-1}$ . Let  $\Gamma(k\Delta_M) = \Gamma_k$  and  $\Upsilon(k\Delta_M) = \Upsilon_k$ . Then for every  $m \geq l$ , Eq. (23) is rewritten as:

$$\bar{\mathcal{E}}_m(x_m, u_m) = \sum_{k=l}^{m-1} (\Gamma_{m-k+1} - \Gamma_{m-k})u_k + \Gamma_1 u_m + \Gamma_{m-l+1}\mu + \Upsilon_{m-l+1}x_l. \quad (26)$$

### 6.3 Open-loop optimization

Having found the average exposure at stages  $m = l, \dots, M-1$  we formulate an open-loop optimization to find optimal  $u_l, u_{l+1}, \dots, u_{M-1}$ . Defining  $\hat{u}_l = (u_l; \dots; u_{M-1})$  and  $\hat{\mathcal{E}}_l = (\bar{\mathcal{E}}_l(x_l, u_l); \dots; \bar{\mathcal{E}}_{M-1}(x_{M-1}, u_{M-1}))$  we can aggregate. Aggregating these linear forms for all  $l \geq m$  yields to the following matrix equation for finding  $\hat{\mathcal{E}}_l$ :

$$\underbrace{\begin{bmatrix} \bar{\mathcal{E}}_l(x_l, u_l) \\ \bar{\mathcal{E}}_{l+1}(x_{l+1}, u_{l+1}) \\ \bar{\mathcal{E}}_{l+2}(x_{l+2}, u_{l+2}) \\ \vdots \\ \bar{\mathcal{E}}_{M-1}(x_{M-1}, u_{M-1}) \end{bmatrix}}_{\hat{\mathcal{E}}_l} = \underbrace{\begin{bmatrix} \Gamma_1 & 0 & 0 & \dots & 0 \\ \Gamma_2 - \Gamma_1 & \Gamma_1 & 0 & \dots & 0 \\ \Gamma_3 - \Gamma_2 & \Gamma_2 - \Gamma_1 & \Gamma_1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \Gamma_{M-l} - \Gamma_{M-l-1} & \Gamma_{M-l-1} - \Gamma_{M-l-2} & \dots & \Gamma_2 - \Gamma_1 & \Gamma_1 \end{bmatrix}}_{X_l} \underbrace{\begin{bmatrix} u_l \\ u_{l+1} \\ u_{l+2} \\ \vdots \\ u_{M-1} \end{bmatrix}}_{\hat{u}_l} \quad (27)$$

$$+ \underbrace{\begin{bmatrix} \Gamma_1 \\ \Gamma_2 \\ \Gamma_3 \\ \vdots \\ \Gamma_{M-l} \end{bmatrix}}_{Y_l} \mu + \underbrace{\begin{bmatrix} \Upsilon_1 \\ \Upsilon_2 \\ \Upsilon_3 \\ \vdots \\ \Upsilon_{M-l} \end{bmatrix}}_{W_l} x_l$$

Defining the expanded form of constraint variables as  $\hat{c}_l = (c_l; \dots; c_{M-1})$ ,  $\hat{C}_l = (C_l; \dots; C_{M-1})$ , and  $\hat{\alpha}_l = (\alpha_l; \dots; \alpha_{M-1})$  we have

$$\underbrace{\begin{bmatrix} c_l^\top & \dots & 0^\top \\ \vdots & \ddots & \vdots \\ 0^\top & \dots & c_{M-1}^\top \\ I & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & I \\ -I & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & -I \end{bmatrix}}_{Z_l} \underbrace{\begin{bmatrix} u_l \\ u_{l+1} \\ u_{l+2} \\ \vdots \\ u_{M-1} \end{bmatrix}}_{\hat{u}_l} \leq \underbrace{\begin{bmatrix} C_l \\ \vdots \\ C_{M-1} \\ \alpha_l \\ \vdots \\ \alpha_{M-1} \\ 0 \\ \vdots \\ 0 \end{bmatrix}}_z. \quad (28)$$

In summary the following equation states the relation between intervention intensities given the constraints:

$$X_l \hat{u}_l + Y_l \mu + W_l x_l = \hat{\mathcal{E}}_l \quad \text{where} \quad Z_l \hat{u}_l \leq z_l \quad (29)$$

Then, we provide the optimization from of the above exposure shaping tasks.

For CEM consider  $\hat{\beta}_l = (\beta_l; \dots, \beta_{M-1})$ . Then the problem

$$\begin{aligned} & \text{maximize}_{\hat{h}, \hat{u}_l} \frac{1}{n} \mathbf{1}^\top \hat{h} \\ & \text{subject to } X_l \hat{u}_l + Y_l \mu + W_l x_l \geq \hat{h}, \quad \hat{\beta}_l \geq \hat{h}, \quad Z_l \hat{u}_l \leq z_l, \end{aligned} \quad (30)$$

solves CEM where  $\hat{h}$  is an auxiliary vector of size  $n(M-l)$ .

For MEM consider the auxiliary  $h$  as a vector of size  $M-l$  and  $\hat{h}$  a vector of size  $n(M-l)$ .  $\hat{h} = (h(1); \dots; h(1); h(2); \dots, h(2); \dots, h(M-l); \dots; h(M-l))$  where each  $h(k)$  is repeated  $n$  times. Then MEM is equivalent to

$$\begin{aligned} & \text{maximize}_{\hat{h}, \hat{u}_l} \mathbf{1}^\top \hat{h} \\ & \text{subject to } X_l \hat{u}_l + Y_l \mu + W_l x_l \geq \hat{h}, \quad \hat{\beta}_l \geq \hat{h}, \quad Z_l \hat{u}_l \leq z_l \end{aligned} \quad (31)$$

For LES let  $\hat{\gamma}_l = (\gamma_l; \dots; \gamma_{M-1})$  and  $\hat{D}_l = \text{diag}(D, \dots, D)$ , then

$$\begin{aligned} & \text{minimize}_{\hat{u}_l} \frac{1}{n} \|\hat{D}_l (X_l \hat{u}_l + Y_l \mu + W_l x_l) - \hat{\gamma}_l\|^2 \\ & \text{subject to } Z_l \hat{u}_l \leq z_l \end{aligned} \quad (32)$$

All the three tasks involve convex (and linear) objective function with linear constraints which impose a convex feasible set. Therefore, one can use the rich and well-developed literature on convex optimization and linear programming to find the optimum intervention.

## 6.4 Scalable optimization

All the exposure shaping problems defined above require an efficient evaluation of average intensity  $\eta(t)$  at all stages, which entails computing matrices  $X_l$ ,  $Y_l$ ,  $W_l$ , and  $Z_l$ . This leads to work with matrix exponentials and inverse matrices to obtain  $\Upsilon_m$ , and  $\Gamma_m$  for  $m = 1, \dots, M-1$ . In small or medium networks, we can rely on well-known numerical methods to compute matrix exponentials and inverse. However, in large networks, the explicit computation of  $X_l$ ,  $Y_l$ ,  $W_l$ , and  $Z_l$  becomes intractable. Fortunately, we can exploit the following key property of our convex campaigning framework: the average intensity itself and the gradient of the objective functions only depends on  $X_l$ ,  $Y_l$ ,  $W_l$ , and  $Z_l$  (and consequently on  $\Upsilon_m$ , and  $\Gamma_m$ ) through matrix-vector product operations. Similar to [8] for the computation of the product of matrix exponential with a vector, one can use the iterative algorithm by Al-Mohy et al. [9], which combines a scaling and squaring method with a truncated Taylor series approximation to the matrix exponential. For solving the sparse linear system of equation, we use the well-known GMRES method, which is an Arnoldi process for constructing an  $l_2$ -orthogonal basis of Krylov subspaces. The method solves the linear system by iteratively minimizing the norm of the residual vector over a Krylov subspace. For details please refer to [8]. Last but not least, we don't need to explicitly build  $X_l$ ,  $Y_l$ ,  $W_l$ , and  $Z_l$ . At each step of gradient computation all the operations involving them are multiplication of  $\Upsilon_1, \dots, \Upsilon_M$ , and  $\Gamma_1, \dots, \Gamma_M$  to vectors such as  $u_0, \dots, u_{M-1}$  and  $\mu$ .

## 7 Temporal Properties

In this section we empirically study the theoretical results of section 4. The empirical mean and standard deviation of the intensity averaged over multiple number of cascades is compared theoretical mean. Besides this, the other purpose of the experiment is to advocate verification process in the synthetic experiments when we used simulation to evaluate the merits of the proposed algorithm and compare to the baselines. In other words, we show that the empirical activity (and hence the average exposure) is very close to its theoretical value and it is justifiable to be used for the comparison.

Fig. 2 demonstrates the activity profile of 3 random users picked in a network of 300 ones simulated 100 times to investigate theorem. 2. The setting is similar to synthetic experiment in the main paper. Piecewise exogenous intensity (interventions) are picked randomly in  $[0.1, 0.2]$  with a slight noise. We consider 5

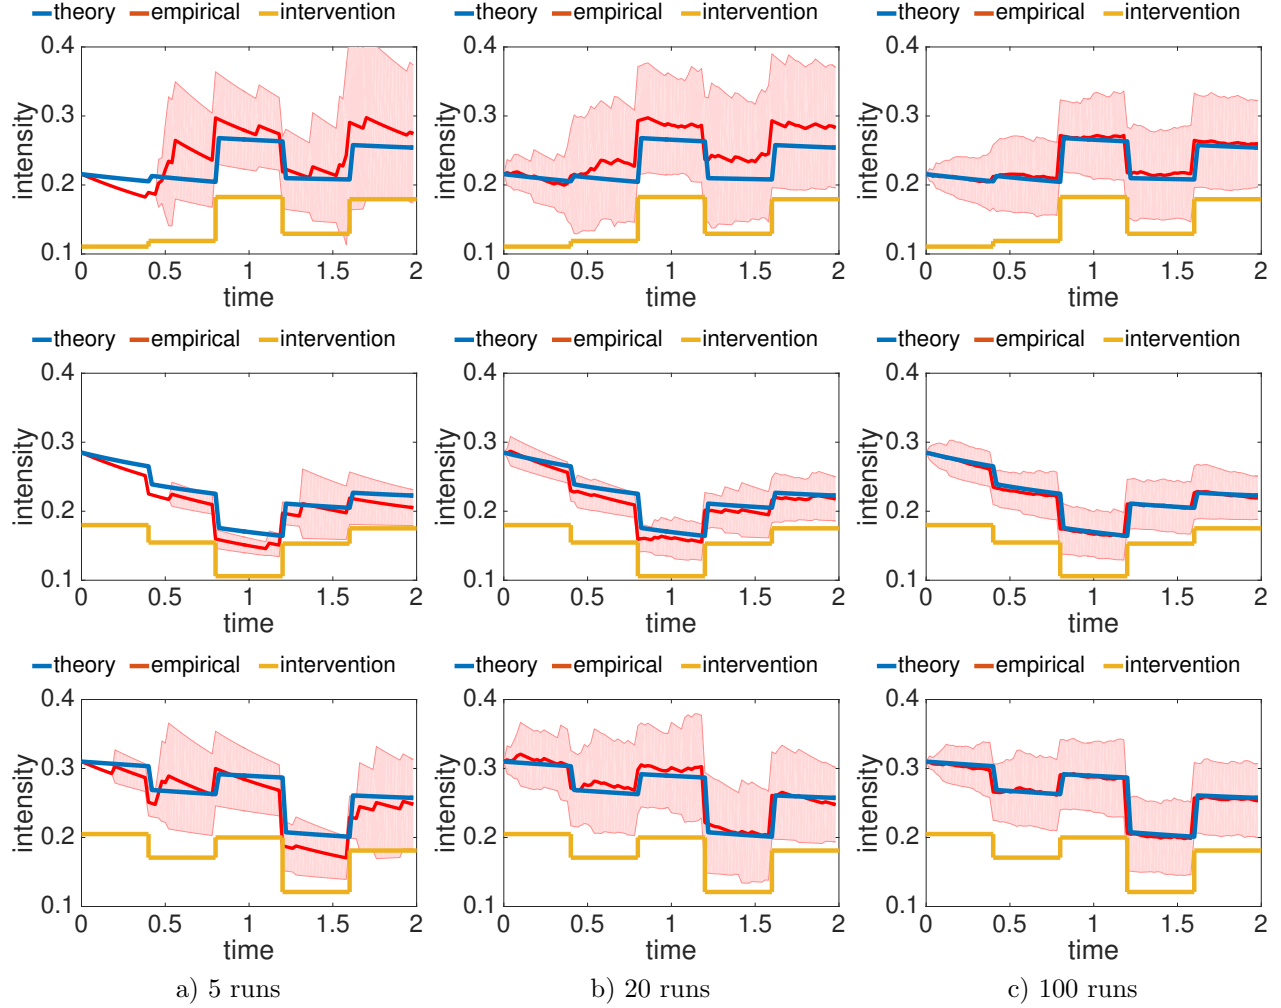


Figure 2: Empirical investigation of theoretical results in Thm. 2. blue: theoretical average intensity; red: empirical average intensity and sample standard deviation; orange: piecewise-constant exogenous intensity (interventions)

stages for changing the exogenous intensity. The empirical average and standard deviation is compared to theoretical average intensity for 3 different number of runs namely, 5, 20, and 100 times. Also, Fig. 3 demonstrate the general case where the exogenous intensity is a time-varying function in theorem. 3. We take 3 sample functions to investigate this case; a sinusoidal function; an exponential decaying function with added noise; and a quadratic function.

We observe a couple of interesting facts. Firstly, it's apparent that by increasing the number of averages the empirical intensity tends to theoretical one very fast. Secondly, as the mean becomes more accurate by increasing the number of cascades the standard deviation increases; e.g., compare the standard deviation in first and third column. Thirdly, the standard deviation is increasing with time. This is due to the fact that as time passes random elements are aggregated more and this increases variance.

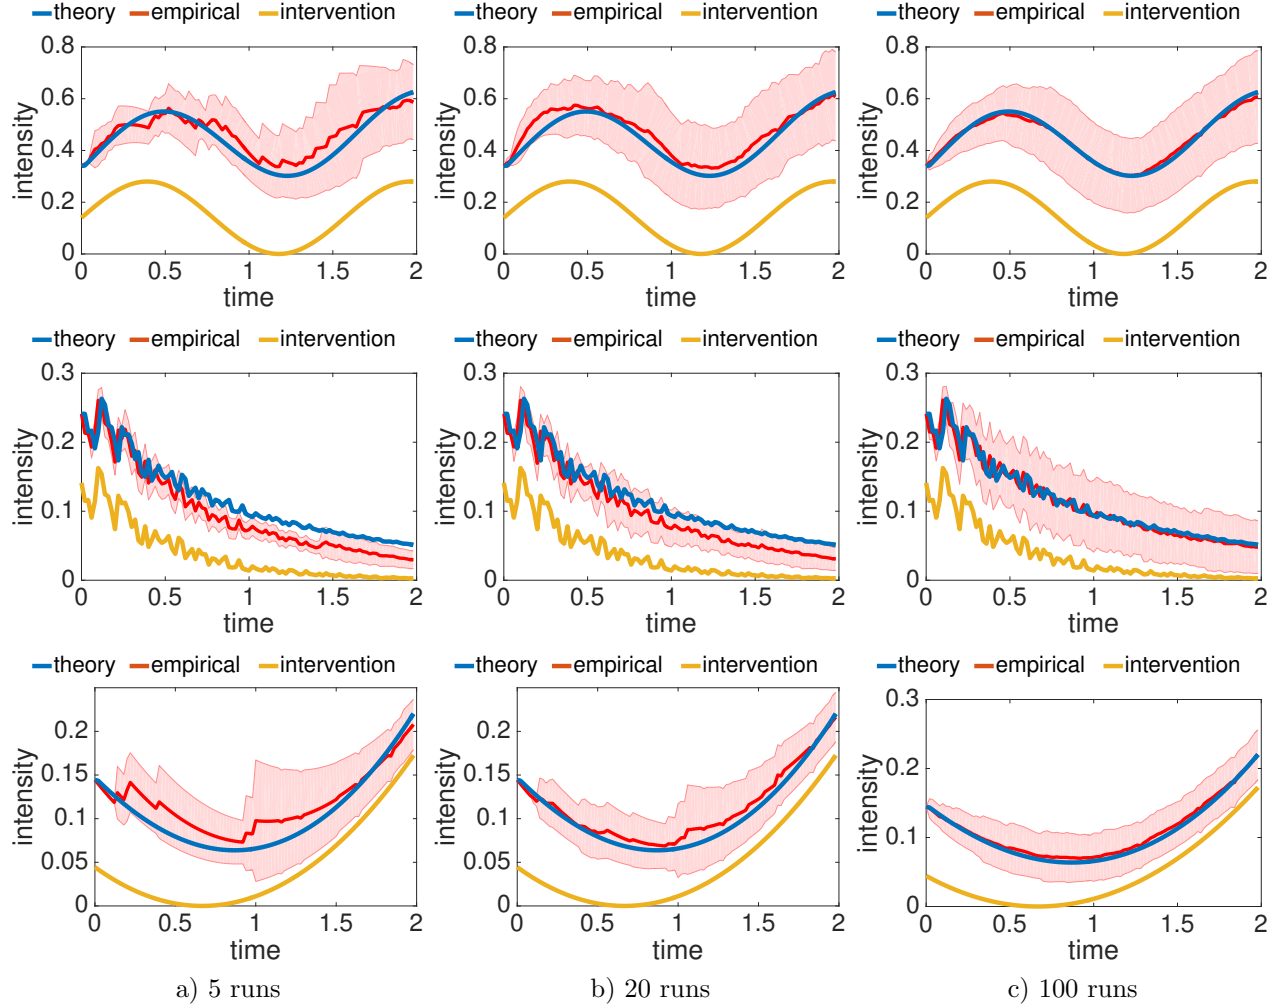


Figure 3: Empirical investigation of theoretical results in Thm. 3. blue: theoretical average intensity; red: empirical average intensity and sample standard deviation; orange: general time-varying intensity (interventions)

## 8 Experiments

We evaluate our campaigning framework using both simulated and real world data and show that our approach significantly outperforms several baselines.

### 8.1 Synthetic data generation

The network is generated synthetically using varying number of nodes. Initial exogenous intensity is set uniformly at random,  $u_m^i \sim \mathcal{U}[0, 0.1]$ . Endogenous intensity coefficients (influence matrix elements) are set similarly,  $a_{ij} \sim \mathcal{U}[0, 0.1]$ . To mimic sparse real networks half of the elements are set to 0 randomly. The matrix is scaled appropriately such that the spectral radius of the coefficient matrix is a random number smaller than one and the stability of the process is ensured.

The upper bound for the intervention intensity is set randomly from interval  $\alpha^i \sim \mathcal{U}[0, 0.1]$ . The price of each person is set  $c_m^i = 1$ , and the total budget at stage  $m$  is randomly generated as  $C_m \sim (n/10)\mathcal{U}[0, 0.1]$ . For the capped exposure maximization case the upper bound is set  $\alpha_m^i \sim \mathcal{U}[0, 1]$  and target in least-squares

exposure shaping is set similarly  $v_m^i \sim (n/10)\mathcal{U}[0, 1]$ . Furthermore, the shaping matrix  $D$  is set to  $I$ . In all the synthetic experiments  $\omega = 0.01$  which roughly means losing 63 % of influence after 100 units of time (minutes, hours, etc). Furthermore, the exposing matrix is set to the unweighted adjacency matrix *i.e.*,  $B_{ij} = 1$  if and only if  $A_{ij} \geq 10^{-4}$ . This way the results are reported in terms of the exact number of exposures and are easily interpretable. In general applications any  $B$  can be used for example using the influence matrix  $A$  yields to a wighted exposure count. In all the synthetic and real experiments the above settings are assumed unless it is explicitly mentioned.

## 8.2 Real data description and network inference

In real data, we use a temporal resolution of one hour and selected the bandwidth  $\omega = 0.001$  by cross validation. Roughly speaking, it corresponds to losing almost 50 % of the initial influence after 1 month. The upper bound for intervention intensity is set uniformly at random with mean equal to empirical intensity learned from data. The upper bound for the cap and target exposure are set similarly. For the 10 pairs of cascades we used first 3 months of data to learn the network parameters. We then drop the exogenous intensity  $\mu$ , and keep the influence network parameters  $A$ . By fixing  $A$  we use the next 6 months of data to learn the exogenous intensity of sites in the two cascades at each of the  $M$  stages and name them  $\mu_m^{c_1}$  and  $\mu_m^{c_2}$ . Given  $A$  we find the optimal intervention intensity  $u_m^{opt}$  stage by stage, for each of the three exposure shaping tasks assuming  $\mu = 0$ . Then, our prediction is: cascade  $c_1$  will reach a better objective value at stage  $m$  if  $dist(u_m^{opt}, \mu_m^{c_1}) < dist(u_m^{opt}, \mu_m^{c_2})$  and vice versa measured by cosine similarity. The prediction accuracy is then reported as a performance measure.

## 8.3 Baselines

In this section, we describe several baselines we compare our approach. Most often, these baseline methods utilize a property to prioritize users for budget assignment.

For the capped exposure maximization problem, we consider the following four baselines:

- **OPL**: It allocates the budget according to the solution to the dynamic programming in an *open loop* setting, *i.e.*, the decisions on the allocation policy are made once and for all at the initial intervention points at initial time  $t = 0$ . This is very important baseline to which comparison quantify the so called *value of information* in the context of dynamic programming and optimal control. As the name suggests it indicates how much knowing what happened so far helps making decisions for future. For the minimum and capped exposure maximization creasing  $n$  the objective function is normalized by the size of network.
- **RND**: It assigns a random point in the convex space of feasible solutions.
- **PRK**: At each stage it subtracts the previous state ( $x_m^i$ ) from the cap ( $\alpha_m^i$ ) and multiply by the page rank score of the the node ( $r^i$ ) computed with damping factor 0.85 and allocates the budget proportional to this value, *i.e.*,  $u_m^i \propto \max((\alpha_m^i - x_m^i)r^i, 0)$ . The proposed solution is then projected to the feasible set of actions in that stage and the extra amount is redistributed similarly. The process is iterated until all the budget are allocated. This baseline assumes that more central users can leverage the total activity, therefore, assigns the budget dynamically to the more connected users proportional to their page rank score.
- **WEI**: This baseline uses sum of out-going influence ( $q^i = \sum_j a_{ji}$ ) as a measure of centrality of users. Similar to the previous one it assigns budget dynamically to the users proportionally to  $u_m^i \propto \max((\alpha_m^i - x_m^i)q^i, 0)$ . This heuristic allows us to understand the effect of considering the whole network and the propagation layout with respect to only consider the direct (out-going) influence.

For the max-min exposure shaping problem, we implement the following four baselines:

- **OPL**: Similar to the previous objective it represents the open loop solution.

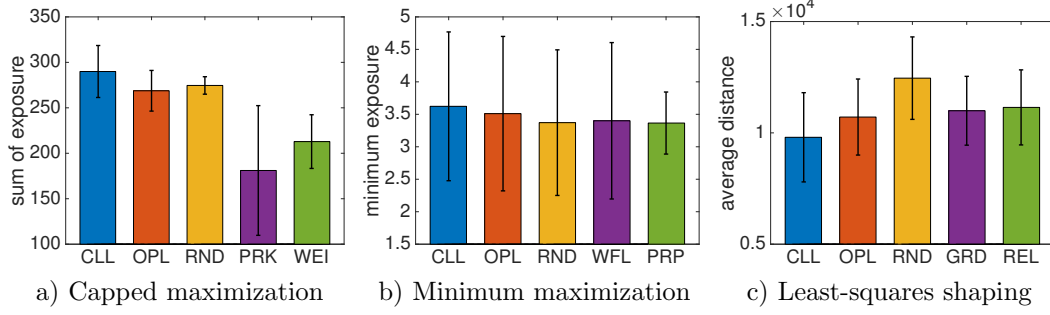


Figure 4: The objective on simulated events and synthetic network;  $n = 300$ ,  $M = 6$ ,  $T = 40$

- **RND**: Similar to the previous objective it allocates the budget randomly within the feasible set.
- **WFL**: It takes a *water filling* approach. It sorts the users in ascending order of the exposure in the previous stage. Then allocates budget to the first users until the the summation of its previous exposure and the allocated budget reaches the second lowest value or it violates a constraint. Then, assigns the budget to these two until they reach the third user with lowest exposure or a constraint is violated. This process is continued until the budget is allocated.

- **PRP**: It allocates the budget inversely proportional to the the exposure at the previous stage.

For the least-square exposure shaping problem, we compare our method with four baselines:

- **OPL**: Similar to the previous objective it represents the open loop solution.
- **RND**: Similar to the previous objective it allocates the budget randomly within the feasible set.
- **GRD**: It finds the difference between the exposure at previous stage ( $x_m^i$  and the target from  $v$  and sorts them decreasingly. Then, allocates budget one at a time until a constraint is violated. It iterates over the users until the budget is fully allocated.
- **REL**: Similar to the above finds the difference from the target but allocates the budget proportionally, *i.e.*,  $u_m^i \propto \max((v^i - x_m^i), 0)$  for all users. If one allocation violates a constraint the extra amount is reallocated in the same manner.

## 8.4 Campaigning results on synthetic networks

In this section, we experiment with a synthetic network of 300 nodes. We focus on three tasks: capped exposure maximization, minimax exposure shaping, and least square exposure shaping. To compare the methods we simulate the network with the prescribed intervention intensity and compute the objective function based on the events happened during the simulation. The mean and standard deviation of the objective function out of 10 runs are reported.

Fig. 4 summarizes the performance of the proposed algorithm (CLL) and 4 other baselines on different campaigning tasks. For **CEM**, our approach consistently outperforms the others by at least 10. This means it exposes each user to the campaign at least 10 times more than the rest consuming the same budget and within the same constraints. The extra 20 units of exposures of over OPL or value of information shows how much we gain by incorporating a dynamic closed-loop solution as opposed to open-loop one-time optimization over all stages. For **MEM**, the proposed method outperforms the others by a smaller margin, however, the 0.1 exposure difference with the second best method is not trifling. This is expected as lifting the minimum exposure is a difficult task [8]. For **LES**, results demonstrate the superiority of CLL by a large margin. The  $10^3$  difference with the second best algorithm aggregated over 6 stages roughly is translated to  $\sqrt{10^3/6} \sim 13$

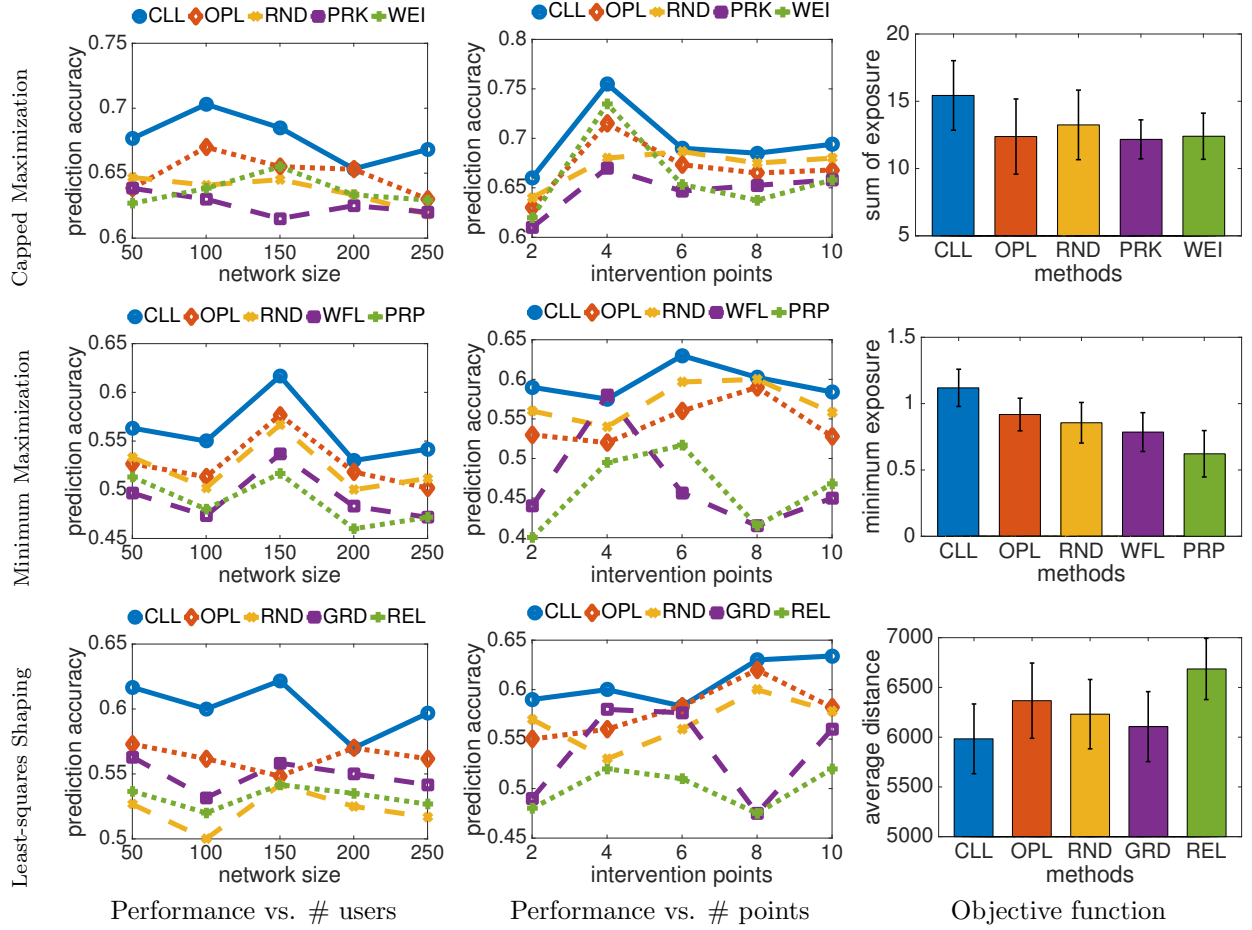


Figure 5: real world dataset results;  $n = 300$ ,  $M = 6$ ,  $T = 40$

difference in the number of exposures per user. Given the heterogeneity of the network activity and target shape, this is a significant improvement over the baselines. Appendix ?? includes further results on varying number of nodes, number of stages, and duration of each stage.

## 8.5 Campaigning results on real world networks

We also evaluate the proposed framework on real world data. To this end, we utilize the MemeTracker dataset [10] which contains the information flows captured by hyperlinks between different sites with timestamps during 9 months. This data has been previously used to validate Hawkes process models of social activity [5, 11]. For the real data, we utilize two evaluation procedures. First, similar to the synthetic case, we simulate the network, but now on a network based on the learned parameters from real data. However, the more interesting evaluation scheme would entail carrying out real intervention in a social media platform. Since this is very challenging to do, instead, in this evaluation scheme we used held-out data to mimic such procedure. Second, we form 10 pairs of clusters/cascades by selecting any 2 combinations of 5 largest clusters in the Memetracker data. Each is a cascade of events around a common subject. For any of these 10 pairs, the methods are faced to the question of predicting which cascade will reach the objective function better. They should be able to answer this by measuring how similar their prescription is to the real exogenous intensity. The key point here is that the real events happened are used to evaluate the objective function of the methods. Then the results are reported on average prediction accuracy on all stages over 10 runs of

random constraint and parameter initialization on 10 pairs of cascades.

Fig. 5, left column illustrates the performance with respect to increasing the number of users in the network. The performance drops slightly with the network size. This means that prediction becomes more difficult as more random variables are involved. The middle panel shows the performance with respect to increasing the number of intervention points. Here, a slight increase in the performance is apparent. As the number of intervention points increases the algorithm has more control over the outcome and can reach the objective function better.

Fig. 5 top row summarizes the results of **CEM**. The left panel demonstrates the predictive performance of the algorithms. CLL consistently outperforms the rest. With 65-70 % of accuracy in predicting the optimal cascade. The right panel shows the objective function simulated 10 times with the learned parameters for network of  $n = 300$  users on 6 intervention points. The extra 2.5 extra exposure per user compared to the second best method with the same budget and constraint would be a significant advertising achievement. Among the competitors OPL and RND seem to perform good. If there were no cap over the resultant exposure, all methods would perform comparably because of the linearity of sum of exposure. However, the successful method is the one who manage to maximize exposure considering the cap. Failure of PRK and WEI indicates that structural properties are not enough to capture the influence. Compared to these two, RND performs better in average, however exhibits a larger variance as expected.

Fig. 5 middle row summarizes the results for **MEM** and shows CLL outperforms others consistently. CLL still is the best algorithm and OPL and RND are the significant baselines. Failure of WFL and PRP shows the network structure plays a significant role in the activity and exposure processes.

The bottom row in Fig. 5 demonstrates the results of **LES**. CLL is still the best method. Among the competitors, OPL is still strong but RND is not performing well for this task. The objective function is summation of the square of the gap between target and current exposure. This explains why GRD is showing a comparable success, since, it starts with the highest gap in the exposure and greedily allocates the budget.

## 9 Related Work

Exposure shaping problems are significantly more challenging than traditional influence maximization problems, which aim to identify a set of users who influence others in the network and trigger a large cascade of adoptions [12, 13]. First, in influence maximization, the state of each user is often assumed to be binary. However, such assumption does not capture the recurrent nature of social activity. Second, while influence maximization methods identify a set of users to provide incentives, they do not typically provide a quantitative prescription on how much incentive should be provided to each user. Third, exposure shaping concerns about a larger variety of target states, such as minimum exposure requirement and homogeneity, not just maximization.

Existing work in stochastic optimal control includes jump diffusion stochastic differential equations (SDE) [14, 15] which focuses on controlling the SDEs with the jump term driven by Poisson processes not for Hawkes processes. Inspired by the opinion dynamics model proposed in [16], the authors in [17] proposes a multivariate jump diffusion process framework for modeling opinion dynamics over networks and determining the control over such networks.

In [18], a continuous action iterated prisoners' dilemma was used to model the interactions in a social network and extended by incorporating a mechanism for external influence on the behavior of individual nodes. Markov Decision Process (MDP) framework is proposed to develop several scheduling algorithms for optimal control of information epidemics with susceptible-infected (SI) model on Erdős-Rényi and scale-free networks [19]. In [20], the authors provided an analytically tractable model for information dissemination over networks and solved the optimal control signal distribution time for minimizing the accumulated network cost via dynamic programming. Furthermore, [21] formulated the maximization of spread of a given message in the population within the stipulated time as continuous-time deterministic optimal control problem.

In contrast, our work has been built on the well-developed theory of point processes [6, 4]. Their usage in modeling activity in social network is becoming increasingly popular [22, 23, 24, 25, 26]. More specifically, we utilizes the Hawkes process [7] which its self-exciting property has been proved to be an appropriate



choice in modeling processes of and on the networks: [27, 5, 28, 11, 29] model and infer the social activity in networks Based on Hawkes process assumption of activity in social networks [30, 31, 32, 33, 34, 35] study one or several phenomena in the social network. In [36] authors proposed a broadcasting algorithm to maximize the visibility of posts in Twitter. It only consider the direct followers and does not involve peer influence in propagation process. Our work is closely related to [8], which is extended in two significant directions here: First, we generalize their result on driving a time-dependent average intensity in the case where the exogenous intensity is not constant. Second, instead of one-shot optimization we pose the problem as a multi-stage optimal control problem which is more fit to real world applications. Then we propose a dynamic programming solution to the multi-stage optimization problem.

## 10 Conclusion and Future work

In this paper, we introduced the optimal multistage campaigning problem, which is a generalization of the activity shaping and influence maximization problems, and it allows for more elaborate goal functions. Our model of social activity is based on multivariate Hawkes process, and for the first time, we manage to derive a linear connection between a time-varying exogenous intensity (*i.e.*, the part that can be easily manipulated via incentives) and the overall network exposure of the campaign. The multistage optimal control problem is introduced and an approximate closed-loop dynamic programming approach is proposed to find the optimal interventions. This linear connection between exogenous intensity and campaign’s exposure enables developing a convex optimization framework for exposure shaping, deriving the necessary incentives to reach a global exposure pattern in the network. The method is evaluated on both synthetic and real-world held-out data and is shown to outperform several heuristics.

Experiments on synthetic and real world datasets reveal a couple of interesting facts:

- Most notable lesson is the presence of the so-called *value of information*. We have witnessed, both in synthetic and real dataset, it is possible to achieve lower cost, essentially by taking advantage of extra information. If the information was not available the controller couldn’t adapt appropriately to the unexpected behavior and consequently the cost could have been adversely affected.
- What we have empirically observed is that the performance, measured in achieving the lower cost and accurate prediction, improves with increasing the number of intervention points. The more control over social network the better one can steer the campaign towards a goal.
- The performance slightly decreases with increasing the number of nodes. That might be due to the increased dimensionality of the optimization problem.

We acknowledge that our method has indeed limitations. For the networks at the scale of web or large social networks faster and scalable methods need to be explored and developed which remains as future works. There are many other interesting venues for future work too. For example, considering competing/collaborating campaigns and their equilibria and interactions, a continuous-time intervention scheme, and exploring other approximate dynamic programming approaches remain as future work.

## References

- [1] Darrell M West. *Air Wars: Television Advertising and Social Media in Election Campaigns, 1952-2012: Television Advertising and Social Media in Election Campaigns, 1952-2012*. Sage, 2013.
- [2] Maurice Vergeer, Liesbeth Hermans, and Steven Sams. Online social networks and micro-blogging in political campaigning the exploration of a new campaign tool and a new campaign style. *Party Politics*, 19(3):477–501, 2013.
- [3] Dimitri P Bertsekas. *Dynamic programming and optimal control*, volume 1.

- [4] Daryl J Daley and David Vere-Jones. *An introduction to the theory of point processes*. Springer Science & Business Media, 2007.
- [5] Ke Zhou, Hongyuan Zha, and Le Song. Learning social infectivity in sparse low-rank networks using multi-dimensional hawkes processes. In *Artificial Intelligence and Statistics (AISTATS)*, 2013.
- [6] Odd Aalen, Ornulf Borgan, and Hakon Gjessing. *Survival and event history analysis: a process point of view*. Springer, 2008.
- [7] Alan G Hawkes. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1):83–90, 1971.
- [8] M. Farajtabar, N. Du, M. Gomez-Rodriguez, I. Valera, L. Song, and H. Zha. Shaping social activity by incentivizing users. In *Neural Information Processing Systems Conference 2014, NIPS '14*, 2014.
- [9] Awad H Al-Mohy and Nicholas J Higham. Computing the action of the matrix exponential, with an application to exponential integrators. *SIAM journal on scientific computing*, 33(2):488–511, 2011.
- [10] Jure Leskovec, Lars Backstrom, and Jon Kleinberg. Meme-tracking and the dynamics of the news cycle. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 497–506. ACM, 2009.
- [11] Shuang-Hong Yang and Hongyuan Zha. Mixture of mutually exciting processes for viral diffusion. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pages 1–9, 2013.
- [12] Matthew Richardson and Pedro Domingos. Mining knowledge-sharing sites for viral marketing. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 61–70. ACM, 2002.
- [13] David Kempe, Jon Kleinberg, and Éva 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*, pages 137–146. ACM, 2003.
- [14] Floyd B Hanson. *Applied stochastic processes and control for Jump-diffusions: modeling, analysis, and computation*, volume 13. Siam, 2007.
- [15] Yaozhong Hu and Bernt Oksendal. Partial information linear quadratic control for jump diffusions. *SIAM Journal on Control and Optimization*, 47(4):1744–1761, 2008.
- [16] Abir De, Isabel Valera, Niloy Ganguly, Sourangshu Bhattacharya, and Manuel Gomez Rodriguez. Modeling opinion dynamics in diffusion networks. *arXiv preprint arXiv:1506.05474*, 2015.
- [17] Yichen Wang, Evangelos Theodorou, Apurv Verma, and Le Song. Steering opinion dynamics in information diffusion networks. *arXiv preprint arXiv:1603.09021*, 2016.
- [18] Daan Bloembergen, Bijan Ranjbar Sahraei, Haitham Bou-Ammar, Karl Tuyls, and Gerhard Weiss. Influencing social networks: An optimal control study. In *ECAI*, pages 105–110, 2014.
- [19] Kundan Kandhway and Joy Kuri. Campaigning in heterogeneous social networks: Optimal control of si information epidemics. 2015.
- [20] Pin-Yu Chen, Shin-Ming Cheng, and Kwang-Cheng Chen. Optimal control of epidemic information dissemination over networks. *Cybernetics, IEEE Transactions on*, 44(12):2316–2328, 2014.
- [21] Aditya Karnik and Pankaj Dayama. Optimal control of information epidemics. In *Communication Systems and Networks (COMSNETS), 2012 Fourth International Conference on*, pages 1–7, 2012.

- [22] Wenzhao Lian, Ricardo Henao, Vinayak Rao, Joseph Lucas, and Lawrence Carin. A multitask point process predictive model. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 2030–2038, 2015.
- [23] Ankur P Parikh, Asela Gunawardana, and Christopher Meek. Conjoint modeling of temporal dependencies in event streams. In *UAI Bayesian Modelling Applications Workshop*. Citeseer, 2012.
- [24] Eric C Hall and Rebecca M Willett. Tracking dynamic point processes on networks. *arXiv preprint arXiv:1409.0031*, 2014.
- [25] Mehrdad Farajtabar, Manuel Gomez-Rodriguez, Nan Du, Mohammad Zamani, Hongyuan Zha, and Le Song. Back to the past: Source identification in diffusion networks from partially observed cascades. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2015.
- [26] Patrick O Perry and Patrick J Wolfe. Point process modeling for directed interaction networks. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 75(5):821–849, 2013.
- [27] Scott W Linderman and Ryan P Adams. Discovering latent network structure in point process data. In *International Conference on Machine Learning (ICML)*, 2014.
- [28] Tomoharu Iwata, Amar Shah, and Zoubin Ghahramani. Discovering latent influence in online social activities via shared cascade poisson processes. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 266–274. ACM, 2013.
- [29] Charles Blundell, Jeff Beck, and Katherine A Heller. Modelling reciprocating relationships with hawkes processes. In *NIPS*, 2012.
- [30] Mehrdad Farajtabar, Yichen Wang, Manuel Gomez-Rodriguez, Suang Li, Hongyuan Zha, and Le Song. Coevolve: A joint point process model for information diffusion and network co-evolution. In *NIPS ’15: Advances in Neural Information Processing Systems*, 2015.
- [31] Long Tran, Mehrdad Farajtabar, Le Song, and Hongyuan Zha. Netcodec: Community detection from individual activities. SIAM.
- [32] Shui-Nee Chow, Xiaojing Ye, Hongyuan Zha, and Haomin Zhou. Influence prediction for continuous-time information propagation on networks. *arXiv preprint arXiv:1512.05417*, 2015.
- [33] Hongteng Xu, Mehrdad Farajtabar, and Hongyuan Zha. Learning granger causality for hawkes processes. *arXiv preprint arXiv:1602.04511*, 2016.
- [34] Xinran He, Theodoros Rekatsinas, James Foulds, Lise Getoor, and Yan Liu. Hawkestopic: A joint model for network inference and topic modeling from text-based cascades. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 871–880, 2015.
- [35] I. Valera and M. Gomez-Rodriguez. Modeling adoption and usage of competing products. In *2015 IEEE International Conference on Data Mining*, 2015.
- [36] 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.
- [37] Omar Hijab. *Introduction to calculus and classical analysis*. Springer, 2007.
- [38] Gerald B Folland. *Real analysis: modern techniques and their applications*. John Wiley & Sons, 2013.
- [39] Ron Bracewell. The fourier transform and its applications. *New York*, 5, 1965.

## A Proofs

**Lemma 1.** Suppose  $\Psi : [0, T] \rightarrow \mathbb{R}^{n \times n}$  is a matrix function, then for every fixed constant intensity  $\mu(t) = c \in \mathbb{R}_+^n$ ,  $\eta_c(t) := \Psi(t)c$  solves the semi-infinite integral equation

$$\eta(t) = c + \int_0^t \Phi(t-s)\eta(s)ds, \quad \forall t \in [0, T], \quad (33)$$

if and only if  $\Psi(t)$  satisfies

$$\Psi(t) = I + \int_0^t \Phi(t-s)\Psi(s)ds, \quad \forall t \in [0, T]. \quad (34)$$

In particular, if  $\Phi(t) = Ae^{-\omega t}\mathbf{1}_{\geq 0}(t) = [a_{ij}e^{-\omega t}\mathbf{1}_{\geq 0}(t)]_{ij}$  where  $0 \leq \omega \notin \text{Spectrum}(A)$ , then

$$\Psi(t) = e^{(A-\omega I)t} + \omega(A-\omega I)^{-1}(e^{(A-\omega I)t} - I) \quad (35)$$

for  $t \in [0, T]$ , where,  $\mathbf{1}_{\geq 0}(t)$  is an indicator function for  $t \geq 0$ .

*Proof.* Suppose that  $\Psi(t)c$  solves (33) for every  $c$ , then substituting  $\eta(t)$  by  $\eta_c(t) := \Psi(t)c$  in (33) we obtain  $\left[\Psi(t) - I - \int_0^t \Phi(t-s)\Psi(s)ds\right]c = 0$ . Since  $c \in \mathbb{R}_+^n$  is arbitrary, we know that  $\Psi(t) - I - \int_0^t \Phi(t-s)\Psi(s)ds = 0$  for all  $t$ , and hence (34) follows. The converse is trivial to verify. Furthermore, one can readily check that (35) satisfies (34) for  $\Phi(t) = Ae^{-\omega t}\mathbf{1}_{\geq 0}(t)$ .  $\square$

**Theorem 2.** Let  $\Psi(t)$  satisfy (34) and  $\mu(t)$  be a right-continuous piecewise constant intensity function of form (10), then the rate function  $\eta(t)$  is given by

$$\eta(t) = \sum_{k=0}^m \Psi(t - \tau_k)(c_k - c_{k-1}), \quad (36)$$

for all  $t \in (\tau_{m-1}, \tau_m]$  and  $m = 1, \dots, M$ , where  $c_{-1} := 0$  by convention.

*Proof.* We prove this result by induction on partition size  $M$ . The previous lemma shows (36) for constant  $\mu(t) = c\mathbf{1}_{[0, T]}(t)$  (i.e.,  $M = 1$ ). Suppose (36) is true for any given piecewise constant  $\mu(t)$  of form (10) with  $M$  partitions. If we impose a constant control  $c \in \mathbb{R}_+^n$  (different from original  $c_M$ ) since time  $\tau \in (\tau_{M-1}, T]$ , namely the piecewise constant intensity function is updated to  $\hat{\mu}(t) := \mu(t) + (c - c_M)\mathbf{1}_{(\tau, T]}(t)$ , then we need to show that the updated rate function  $\hat{\eta}(t)$  is

$$\hat{\eta}(t) = \eta(t) + \Psi(t - \tau)(c - c_M)\mathbf{1}_{(\tau, T]}(t), \quad (37)$$

for all  $t \in [0, T]$ . This result can be verified easily for  $t \in [0, \tau]$ . If  $t \in (\tau, T]$ , then  $\hat{\mu}(t) = \mu(t) + (c - c_M)\mathbf{1}_{(\tau, T]}(t) = \mu(t) + (c - c_M)$  and

$$\begin{aligned} & \hat{\mu}(t) + \int_0^t \Phi(t-s)\hat{\eta}(s)ds \\ &= \mu(t) + (c - c_M) + \int_0^t \Phi(t-s)[\eta(s) + \Psi(s - \tau)(c - c_M)\mathbf{1}_{(\tau, T]}(s)]ds \\ &= \eta(t) + (c - c_M) + \int_0^t \Phi(t-s)\Psi(s - \tau)(c - c_M)\mathbf{1}_{(\tau, T]}(s)ds \\ &= \eta(t) + \left[ I + \int_0^{t-\tau} \Phi(t - \tau - u)\Psi(u)du \right] (c - c_M) \\ &= \eta(t) + \Psi(t - \tau)(c - c_M) = \hat{\eta}(t), \end{aligned} \quad (38)$$

where we used the fact that  $\eta(t)$  is the rate function for intensity  $\mu(t)$  to get the second equality, applied change of variables  $u = s - \tau$  to obtain the third equality, and the property (34) of  $\Psi(t)$  to get the fourth equality. This implies that the rate function is  $\hat{\eta}(t)$  given in (37) for the updated piecewise constant intensity  $\hat{\mu}(t)$  with  $M + 1$  partitions, and hence completes the proof.  $\square$

**Theorem 3.** If  $\Psi \in C^1([0, T])$  and satisfies (34), and exogenous intensity  $\mu$  is bounded and piecewise absolutely continuous on  $[0, T]$  where  $\mu(t+) = \mu(t)$  at all discontinuous points  $t$ , then  $\mu$  is differentiable almost everywhere, and the semi-indefinite integral

$$\eta(t) = \mu(t) + \int_0^t \Phi(t-s)\eta(s)ds, \quad \forall t \in [0, T], \quad (39)$$

yields a rate function  $\eta : [0, T] \rightarrow \mathbb{R}_+^n$  given by

$$\eta(t) = \int_0^t \Psi(t-s)d\mu(s). \quad (40)$$

*Proof.* It suffices to show (40) for absolutely continuous  $\mu(t)$  on  $[0, T]$  since extending the proof to piecewise absolutely continuous function is straightforward. We first define  $\mu(T) = \mu(T-)$  and obtain a continuous  $\mu(t)$  on  $[0, T]$ . Since  $[0, T]$  is compact, we know  $\mu(t)$  is uniformly continuous, and hence there exists a sequence of piecewise constant functions  $\{\mu_k\}_{k=1}^\infty$  such that  $\mu_k \rightarrow \mu$  uniformly on  $[0, T]$ , i.e.,  $\lim_{k \rightarrow \infty} \sup_{0 \leq t \leq T} |\mu_k(t) - \mu(t)| = 0$  [37, Thm. 2.3.6] This also implies that  $\{\mu_k\}$  is uniformly bounded. For every  $k$ , piecewise constant function  $\mu_k$  has bounded variation, therefore we have by [38, Thm. 3.36] that

$$\eta_k(t) := \int_0^t \Psi(t-s)d\mu_k(s) = \int_0^t \Psi'(t-s)\mu_k(s)ds + \Psi(0)\mu_k(t) - \Psi(t)\mu_k(0), \quad (41)$$

for all  $t \in [0, T]$ . Since  $\Psi \in C^1$  we know  $\Psi'$  is continuous and bounded on  $[0, T]$ . By Lebesgue's bounded convergence theorem we know

$$\int_0^t \Psi'(t-s)\mu_k(s)ds \rightarrow \int_0^t \Psi'(t-s)\mu(s)ds. \quad (42)$$

Furthermore, using the uniform convergence of  $\{\mu_k\}$  to  $\mu$ , we know the right hand side (41) converges to  $\int_0^t \Psi'(t-s)\mu(s)ds + \Psi(0)\mu(t) - \Psi(t)\mu(0)$ . Then integration by parts for piecewise absolutely continuous function  $\mu$  which has bounded variation implies that  $\eta(t) = \int_0^t \Psi(t-s)d\mu(s)$  for all  $t \in [0, T]$ .  $\square$

**Corollary 4.** Suppose  $\Psi$  and  $\mu$  satisfy the same conditions as in Thm. 3, and define  $\psi = \Psi'$ , then the rate function is  $\eta(t) = (\psi * \mu)(t)$ . In particular, if  $\Phi(t) = Ae^{-\omega t} \mathbf{1}_{\geq 0}(t) = [a_{ij}e^{-\omega t} \mathbf{1}_{\geq 0}(t)]_{ij}$  then the rate function  $\eta(t) = A \int_0^t e^{(A-\omega I)(t-s)} \mu(s)ds$ .

*Proof.* Note that both  $\psi$  and  $\mu$  have supports in  $\mathbb{R}_+$ , therefore integration by parts and the property of derivative of convolution [39, p. 126] imply that  $\eta(t) = \int_0^t \Psi'(t-s)\mu(s)ds = (\psi * \mu)(t)$ . If  $\Phi(t) = Ae^{-\omega t} \mathbf{1}_{\geq 0}(t) = [a_{ij}e^{-\omega t} \mathbf{1}_{\geq 0}(t)]_{ij}$ , then  $\Psi(t)$  is given in (35), and hence  $\psi(t) = Ae^{(A-\omega I)t}$  and the closed form of  $\eta(t)$  follows as in the claim.  $\square$

## B Extended synthetic results

For the synthetic case we can freely evaluate the properties of the proposed algorithm under several conditions. We assess the performance of the algorithm and compare to the baselines in three settings: i) increasing size of the network; ii) increasing number of intervention points; iii) increasing the time window (or equivalently the stage duration). The results are reported while keeping other parameters fixed. To compare it to the others we simulate the network with the prescribed intervention intensity and compute the objective function. The mean and standard deviation of the objective function out of 10 runs are reported.

Figures 6, 7, and 8 shows the results for CEM, MEM, and LES respectively. In each figure, the first row is for varying number of nodes, the second row is for varying number of intervention points, and the third row is for varying duration of stages. The proposed method is consistently better than the baselines. The trends and facts reported in the main paper are observed in this extended experiment. Additionally, we want to refer the high variance of baseline methods especially RND and OPL which is what we expect.

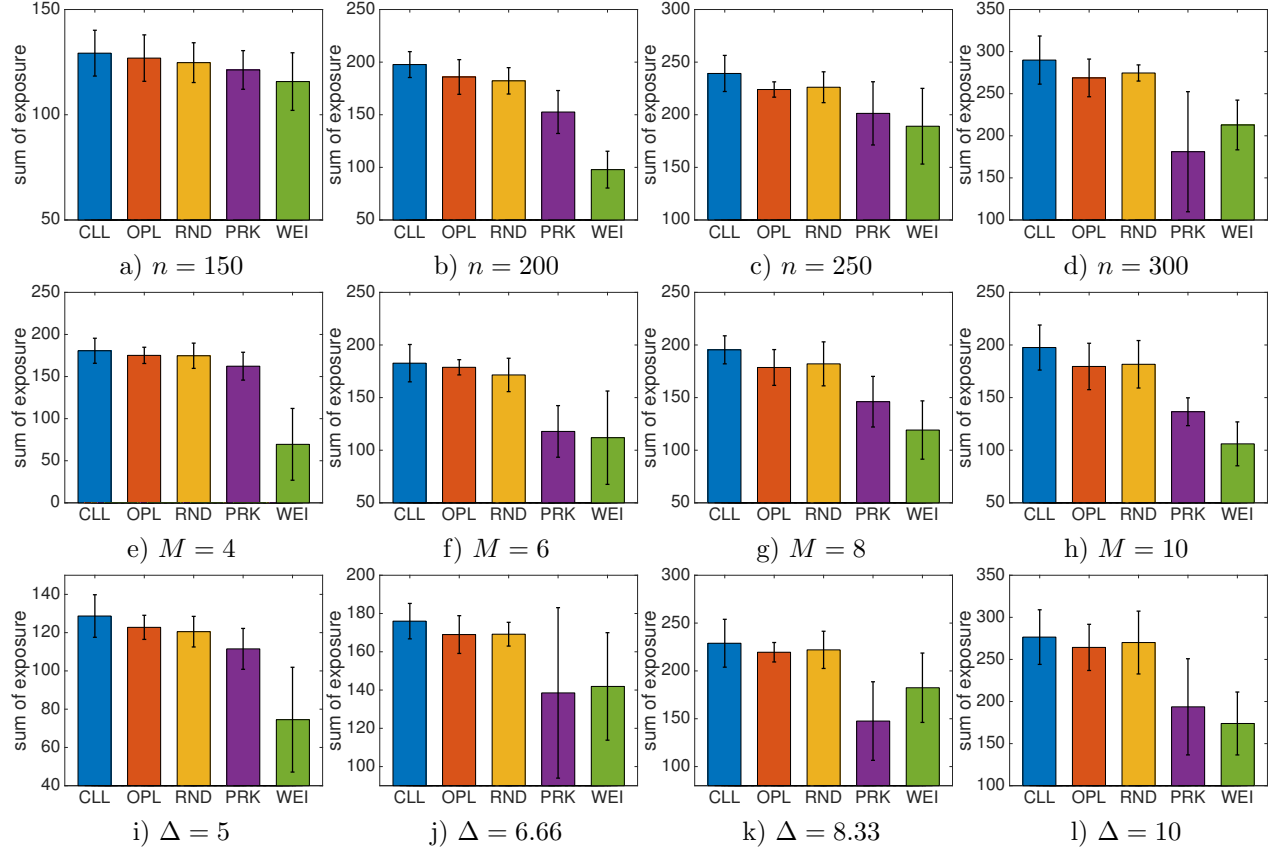


Figure 6: Caped exposure maximization results on synthetic data; top row:  $n$  varies,  $M = 6$ ,  $T = 40$ ; middle row:  $M$  varies,  $T = 40$ ,  $n = 200$ ; bottom row:  $T$  varies,  $n = 200$ ,  $M = 6$

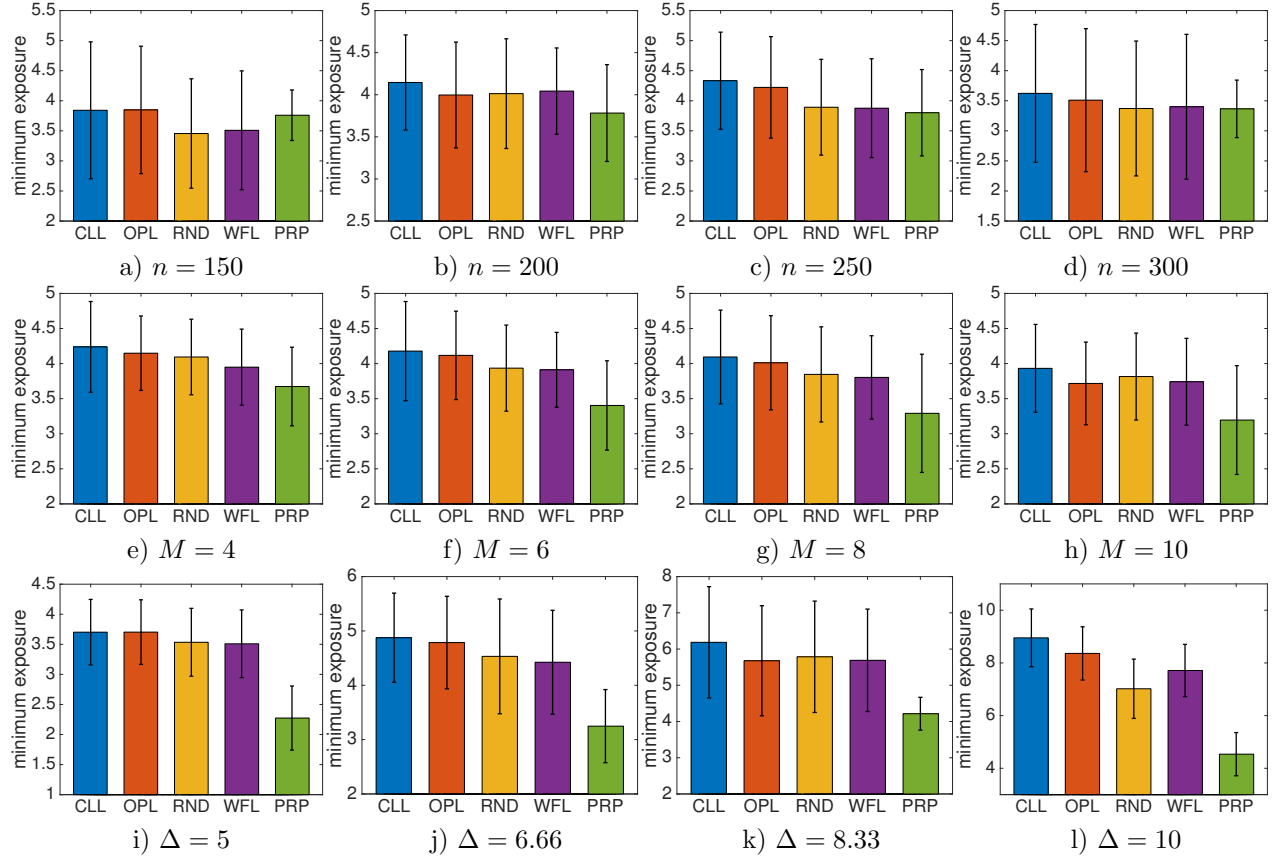


Figure 7: Minimum exposure maximization results on synthetic data; top row:  $n$  varies,  $M = 6$ ,  $T = 40$ ; middle row:  $M$  varies,  $T = 40$ ,  $n = 200$ ; bottom row:  $T$  varies,  $n = 200$ ,  $M = 6$

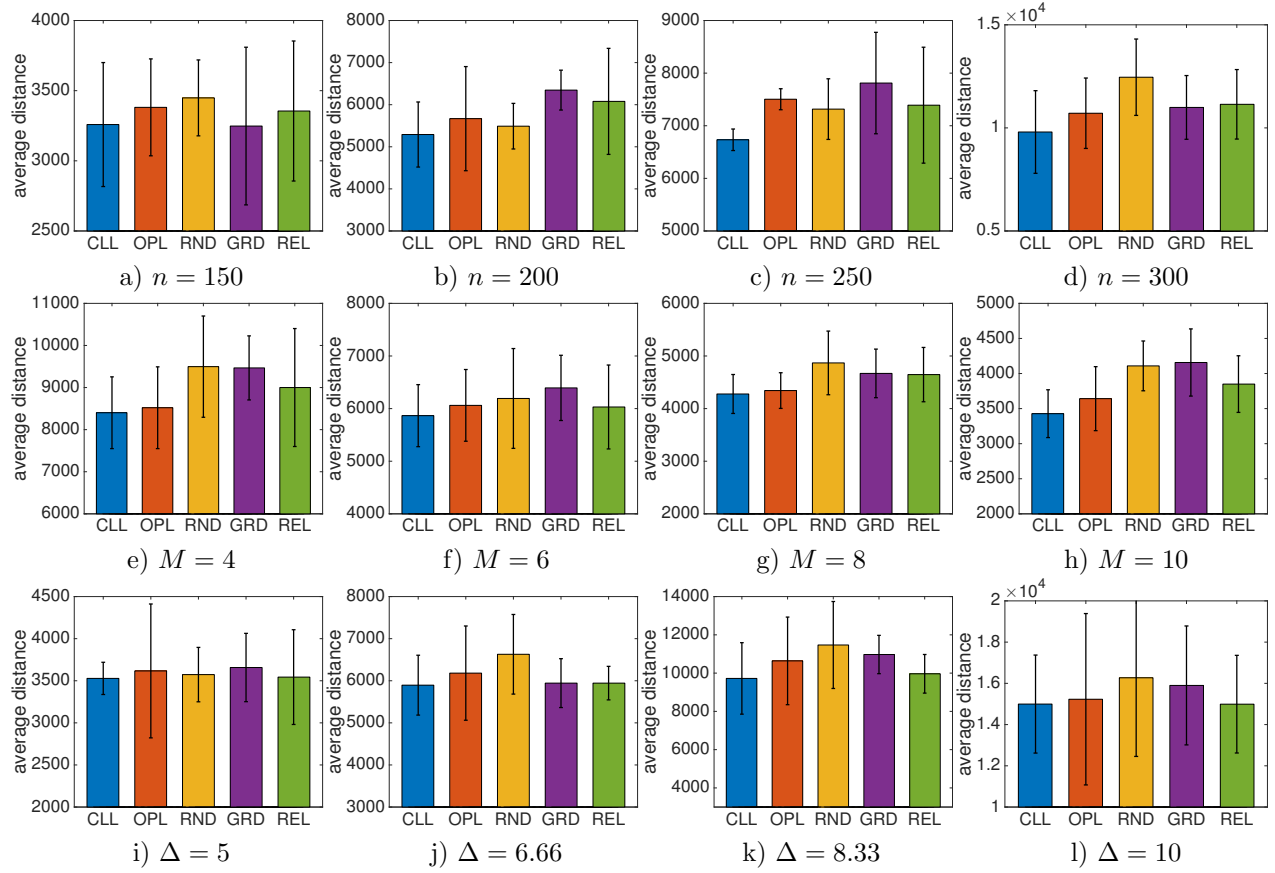


Figure 8: Least-squares exposure shaping results on synthetic data; top row:  $n$  varies,  $M = 6$ ,  $T = 40$ ; middle row:  $M$  varies,  $T = 40$ ,  $n = 200$ ; bottom row:  $T$  varies,  $n = 200$ ,  $M = 6$