

A Related Work

Exposure shaping problems are significantly more challenging than traditional influence maximization problems which aim to identify the set of most influential users that generates spread or cascade of largest size under certain budget constraint [11]. First, in influence maximization, the state of each user is often assumed to be binary, *i.e.*, active (infected) or inactive (healthy, uninfected) which 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.

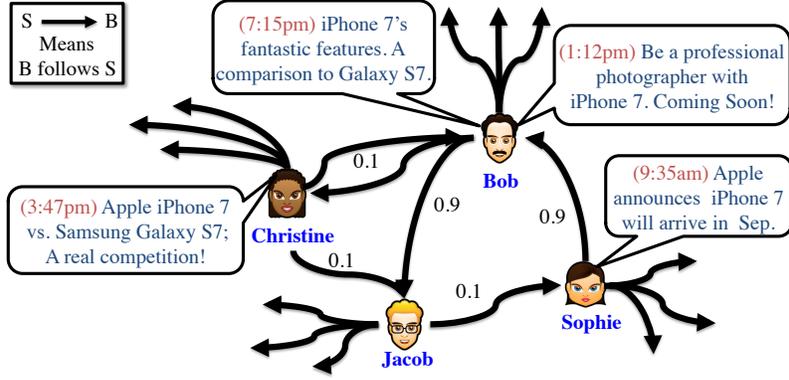
Existing work in stochastic optimal control includes jump diffusion stochastic differential equations (SDE) [12] which focuses on controlling the SDEs with the jump term driven by Poisson processes. Inspired by the opinion dynamics model proposed in [13], a multivariate jump diffusion process framework for modeling opinion dynamics over networks and determining the control over such networks is proposed [14]. In [15], 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 [16]. In [17], 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.

In contrast, our work is built on the theory of multivariate point processes [6, 4]. Their usage in modeling activity in social network is becoming increasingly popular [18, 19, 20]. More specifically, we utilize the multivariate Hawkes process [7] since its mutual and self excitation property has been proved suitable for modeling the dynamics on the network [21, 22, 23, 5, 10]. Our work is closely related to [8], which is extend 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.

B Illustration of Multi-stage Campaigning

Fig. 3-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. 3-b.



a) The social network and events

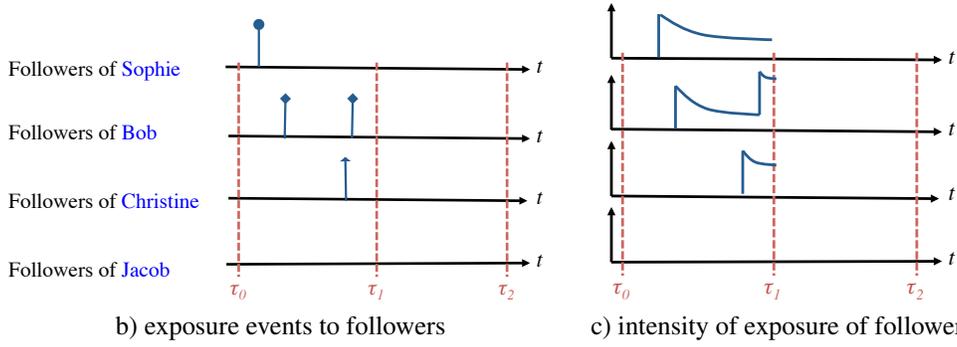


Figure 3: A social network and multi-stage campaigning to maximize the minimum exposure.

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. 3-b. This is the point where multi-stage dynamic campaigning helps. Considering the exposure intensity of followers in Fig. 3-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.

C 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

$$\eta(t) = c + \int_0^t \Phi(t-s)\eta(s)ds, \quad \forall t \in [0, T], \quad (27)$$

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

$$\Psi(t) = I + \int_0^t \Phi(t-s)\Psi(s)ds, \quad \forall t \in [0, T]. \quad (28)$$

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

$$\Psi(t) = e^{(A-\omega I)t} + \omega(A-\omega I)^{-1}(e^{(A-\omega I)t} - I) \quad (29)$$

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 (27) for every c , then substituting $\eta(t)$ by $\eta_c(t) := \Psi(t)c$ in (27) 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 (28) follows. The converse is trivial to verify. Furthermore, one can readily check that (29) satisfies (28) for $\Phi(t) = Ae^{-\omega t}\mathbf{1}_{\geq 0}(t)$. \square

Theorem 2. *Let $\Psi(t)$ satisfy (28) and $\mu(t)$ be a right-continuous piecewise constant intensity function of form (9), then the rate function $\eta(t)$ is given by*

$$\eta(t) = \sum_{k=0}^m \Psi(t - \tau_k)(c_k - c_{k-1}), \quad (30)$$

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 (30) for constant $\mu(t) = c\mathbf{1}_{[0, T]}(t)$ (i.e., $M = 1$). Suppose (30) is true for any given piecewise constant $\mu(t)$ of form (9) 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

$$\hat{\eta}(t) = \eta(t) + \Psi(t - \tau)(c - c_M)\mathbf{1}_{(\tau, T]}(t), \quad (31)$$

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

$$\begin{aligned} & \hat{\mu}(t) + \int_0^t \Phi(t-s)\hat{\eta}(s)ds \\ &= \mu(t) + (c - c_M) + \int_0^t \Phi(t-s)[\eta(s) + \Psi(s - \tau)(c - c_M)\mathbf{1}_{(\tau, T]}(s)]ds \\ &= \eta(t) + (c - c_M) + \int_0^t \Phi(t-s)\Psi(s - \tau)(c - c_M)\mathbf{1}_{(\tau, T]}(s)ds \\ &= \eta(t) + \left[I + \int_0^{t-\tau} \Phi(t - \tau - u)\Psi(u)du \right] (c - c_M) \\ &= \eta(t) + \Psi(t - \tau)(c - c_M) = \hat{\eta}(t), \end{aligned} \quad (32)$$

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 (28) of $\Psi(t)$ to get the fourth equality. This implies that the rate function is $\hat{\eta}(t)$ given in (31) 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 (28), and exogenous intensity μ is bounded and piecewise absolutely continuous on $[0, T]$ where $\mu(t+) = \mu(t)$ at all discontinuous points t , then μ is differentiable almost everywhere, and the semi-indefinite integral*

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

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

$$\eta(t) = \int_0^t \Psi(t-s)d\mu(s). \quad (34)$$

Proof. It suffices to show (34) 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., $\lim_{k \rightarrow \infty} \sup_{0 \leq t \leq T} |\mu_k(t) - \mu(t)| = 0$ [24, Thm. 2.3.6] This also implies

that $\{\mu_k\}$ is uniformly bounded. For every k , piecewise constant function μ_k has bounded variation, therefore we have by [25, Thm. 3.36] that

$$\eta_k(t) := \int_0^t \Psi(t-s) d\mu_k(s) = \int_0^t \Psi'(t-s) \mu_k(s) ds + \Psi(0) \mu_k(t) - \Psi(t) \mu_k(0), \quad (35)$$

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

$$\int_0^t \Psi'(t-s) \mu_k(s) ds \rightarrow \int_0^t \Psi'(t-s) \mu(s) ds. \quad (36)$$

Furthermore, using the uniform convergence of $\{\mu_k\}$ to μ , we know the right hand side (35) 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 μ 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 Ψ and μ 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 ψ and μ have supports in \mathbb{R}_+ , therefore integration by parts and the property of derivative of convolution [26, 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 (29), and hence $\psi(t) = Ae^{(A-\omega I)t}$ and the closed form of $\eta(t)$ follows as in the claim. \square

D Optimization Details

To find the linear and convex form of the open-loop optimization problems for the campaigning objective we define

$$\Upsilon(t) = B \int_0^t e^{(A-\omega I)s} ds = B(A-\omega I)^{-1} (e^{(A-\omega I)t} - I) \quad (37)$$

$$\Gamma(t) = B \int_0^t \Psi(s) ds = B\Upsilon(t) + B(A-\omega I)^{-1} (\Upsilon(t) - It)/\omega; \quad (38)$$

Let $\Gamma(k\Delta_M) = \Gamma_k$ and $\Upsilon(k\Delta_M) = \Upsilon_k$. Then for every $m \geq l$, Eq. (22) is rewritten as:

$$\bar{\mathcal{E}}_m(x_m, u_m) = \sum_{k=l}^{m-1} (\Gamma_{m-k+1} - \Gamma_{m-k}) u_k + \Gamma_1 u_m + \Gamma_{m-l+1} \mu + \Upsilon_{m-l+1} x_l. \quad (39)$$

Aggregating these linear forms for all $l \geq m$ yields to the following matrix equation for finding $\hat{\mathcal{E}}_l$:

$$\underbrace{\begin{bmatrix} \bar{\mathcal{E}}_l(x_l, u_l) \\ \bar{\mathcal{E}}_{l+1}(x_{l+1}, u_{l+1}) \\ \bar{\mathcal{E}}_{l+2}(x_{l+2}, u_{l+2}) \\ \vdots \\ \bar{\mathcal{E}}_{M-1}(x_{M-1}, u_{M-1}) \end{bmatrix}}_{\hat{\mathcal{E}}_l} = \underbrace{\begin{bmatrix} \Gamma_1 & 0 & 0 & \dots & 0 \\ \Gamma_2 - \Gamma_1 & \Gamma_1 & 0 & \dots & 0 \\ \Gamma_3 - \Gamma_2 & \Gamma_2 - \Gamma_1 & \Gamma_1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \Gamma_{M-l} - \Gamma_{M-l-1} & \Gamma_{M-l-1} - \Gamma_{M-l-2} & \dots & \Gamma_2 - \Gamma_1 & \Gamma_1 \end{bmatrix}}_{X_l} \underbrace{\begin{bmatrix} u_l \\ u_{l+1} \\ u_{l+2} \\ \vdots \\ u_{M-1} \end{bmatrix}}_{\hat{u}_l} + \underbrace{\begin{bmatrix} \Gamma_1 \\ \Gamma_2 \\ \Gamma_3 \\ \vdots \\ \Gamma_{M-l} \end{bmatrix}}_{Y_l} \mu + \underbrace{\begin{bmatrix} \Upsilon_1 \\ \Upsilon_2 \\ \Upsilon_3 \\ \vdots \\ \Upsilon_{M-l} \end{bmatrix}}_{W_l} x_l \quad (40)$$

Furthermore, aggregating all constraints of the form $u_m \in \mathcal{U}_m$ for $l \leq m \leq M$ into a matrix form leads to:

$$\underbrace{\begin{bmatrix} c_l^\top & \dots & 0^\top \\ \vdots & \ddots & \vdots \\ 0^\top & \dots & c_{M-1}^\top \\ I & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & I \\ -I & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & -I \end{bmatrix}}_{Z_l} \underbrace{\begin{bmatrix} u_l \\ u_{l+1} \\ u_{l+2} \\ \vdots \\ u_{M-1} \end{bmatrix}}_{\hat{u}_l} \leq \underbrace{\begin{bmatrix} C_l \\ \vdots \\ C_{M-1} \\ \alpha_l \\ \vdots \\ \alpha_{M-1} \\ 0 \\ \vdots \\ 0 \end{bmatrix}}_z. \quad (41)$$

Then, we can proceed with the open-loop control problem benefiting from the linear form of $X_l \hat{u}_l + Y_l \mu + W_l x_l = \hat{\mathcal{E}}_l$ such that $Z_l \hat{u}_l \leq z_l$.

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 Υ_m , and Γ_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 Υ_m , and Γ_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. [27], 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 μ .

E Temporal Properties

In this section we empirically study the theoretical results of section 3. 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. 4 demonstrates the activity profile of 3 random users picked in a network of 300 ones simulated 100 times to investigate Thm. 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 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. 5 demonstrate the general case where the exogenous intensity is a time-varying function in Thm. 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.

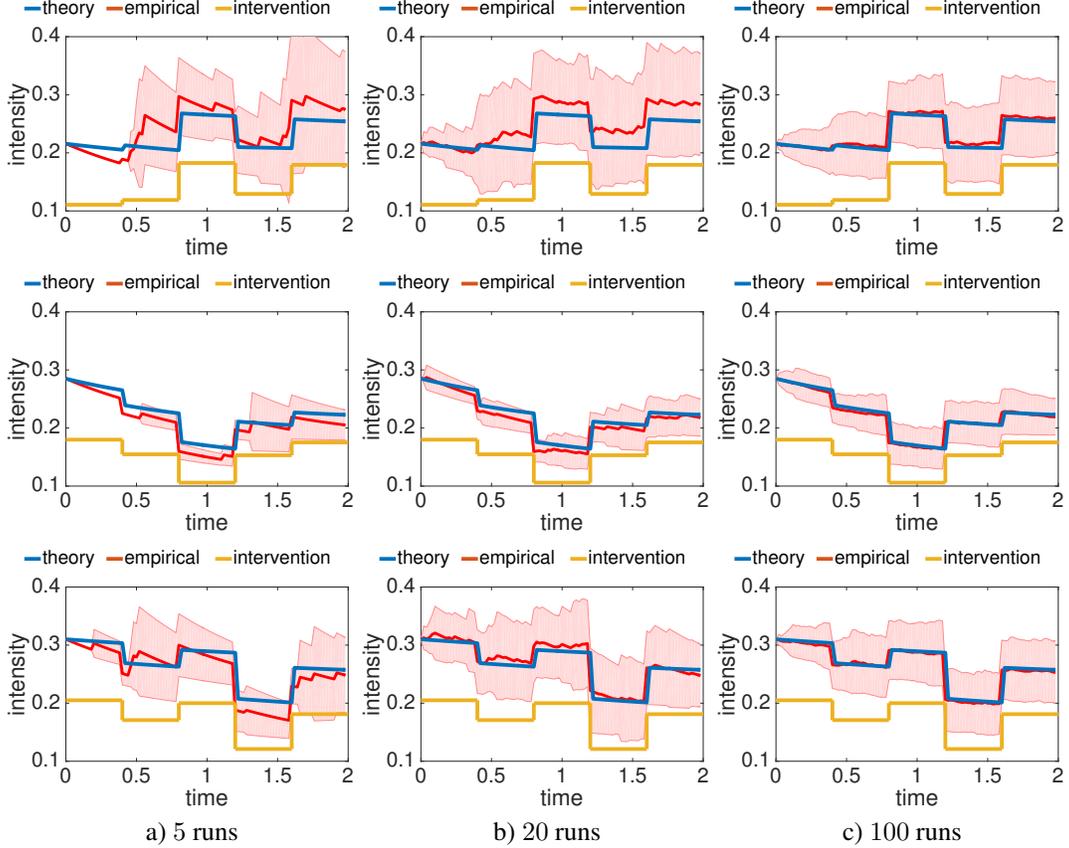


Figure 4: Empirical investigation of theoretical results in Thm. 2. blue: theoretical average intensity; red: empirical average intensity and sample standard deviation; orange: piecewise-constant exogenous intensity (interventions)

F More on experimental setup

In this section we first introduce the baselines to which we have compared the proposed algorithm.

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 (α_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.

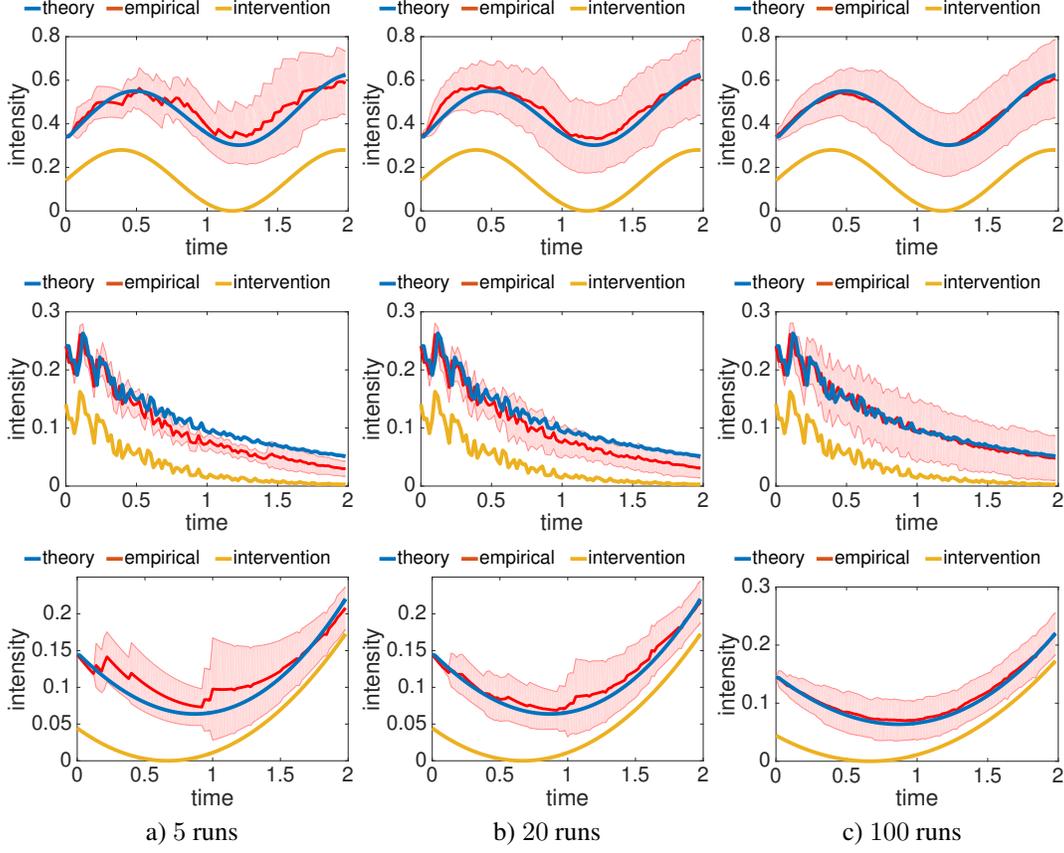


Figure 5: Empirical investigation of theoretical results in Thm. 3. blue: theoretical average intensity; red: empirical average intensity and sample standard deviation; orange: general time-varying intensity (interventions)

- **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.
- **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.

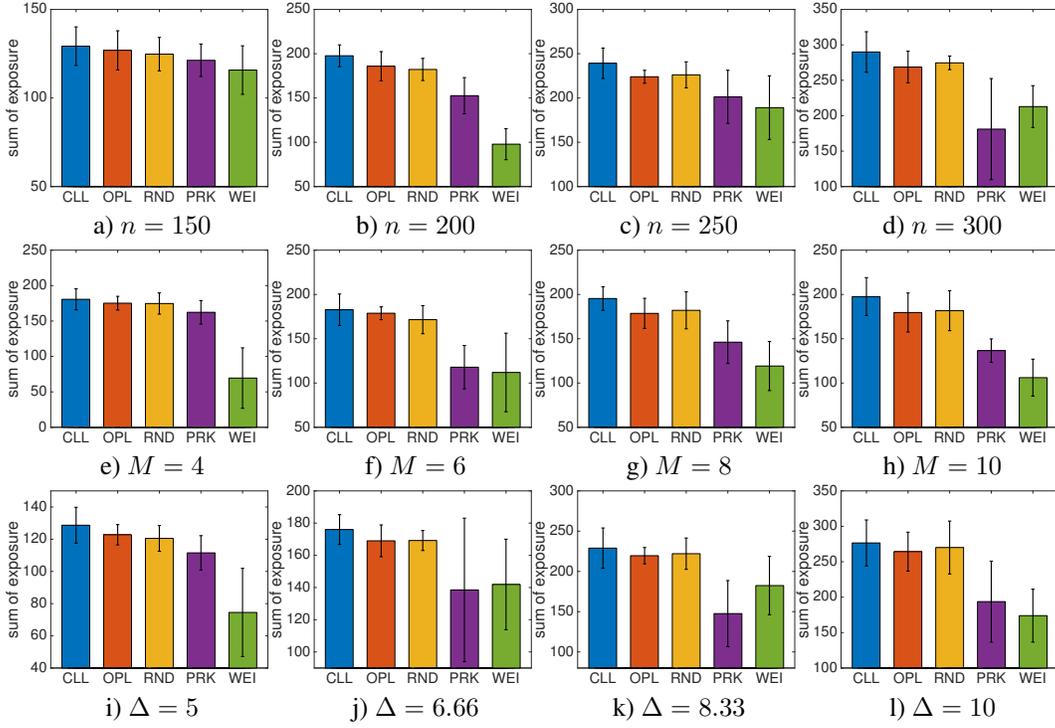


Figure 6: Caped exposure maximization results on synthetic data; top row: n varies, $M = 6$, $T = 40$; middle row: M varies, $T = 40$, $n = 200$; bottom row: T varies, $n = 200$, $M = 6$

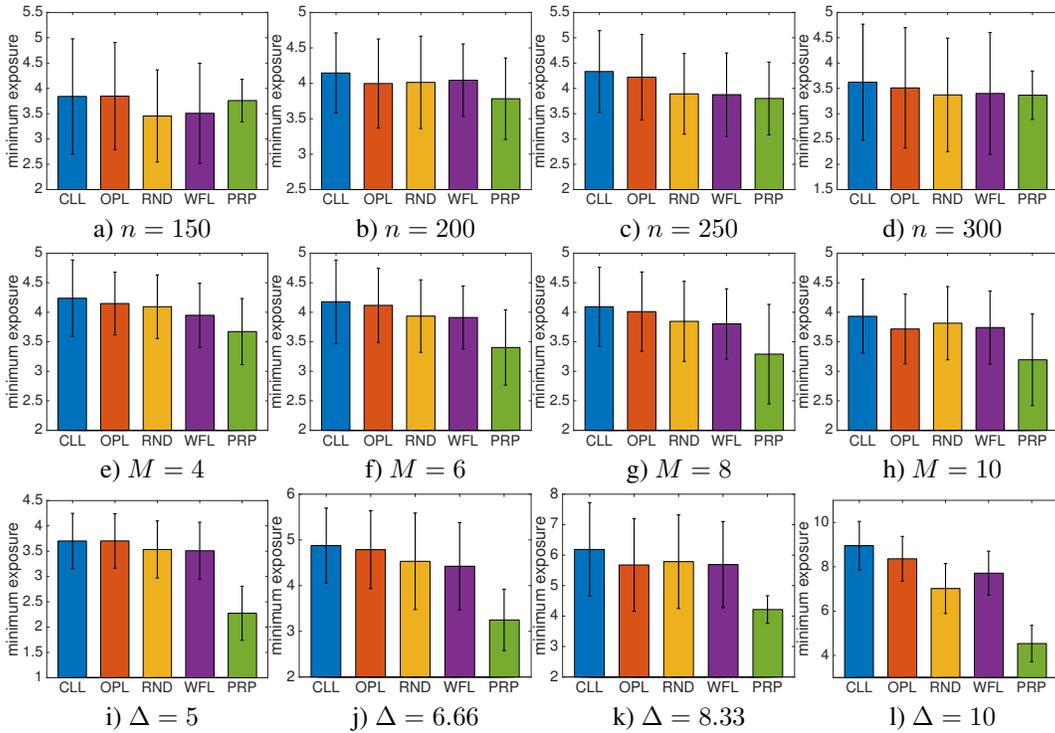


Figure 7: Minimum exposure maximization results on synthetic data; top row: n varies, $M = 6$, $T = 40$; middle row: M varies, $T = 40$, $n = 200$; bottom row: T varies, $n = 200$, $M = 6$

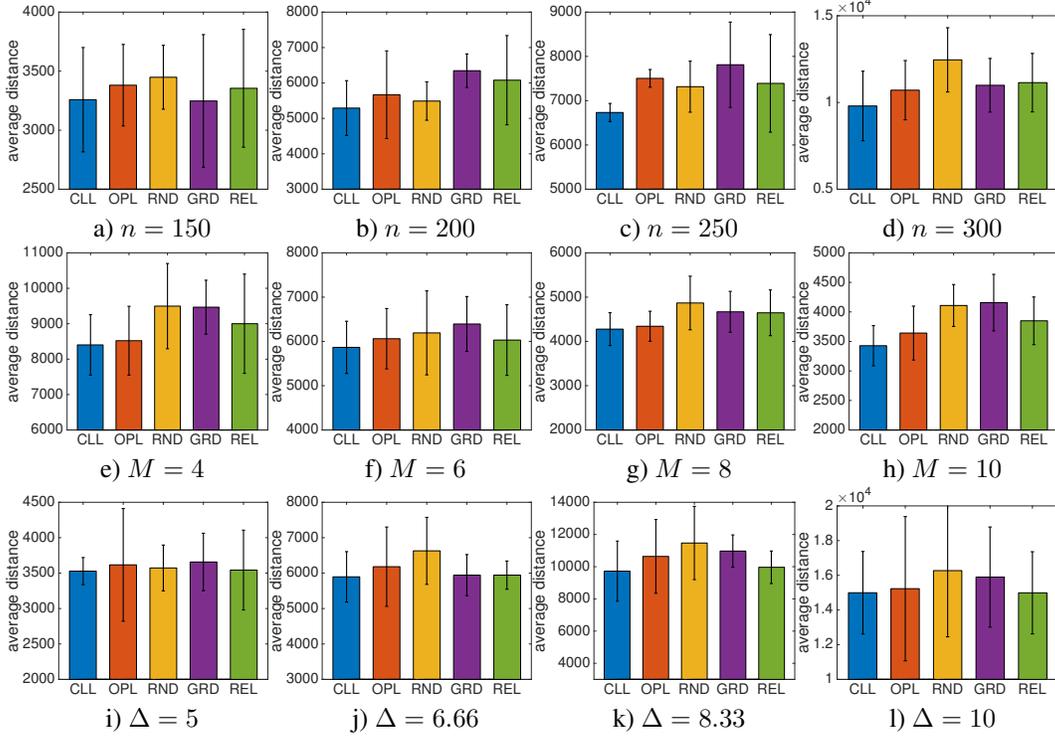


Figure 8: Least-squares exposure shaping results on synthetic data; top row: n varies, $M = 6$, $T = 40$; middle row: M varies, $T = 40$, $n = 200$; bottom row: T varies, $n = 200$, $M = 6$

- **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.

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 loosing 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 weighted exposure count. In all the synthetic and real experiments the above settings are assumed unless it is explicitly mentioned.

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.

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 loosing 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 μ , 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.

G Discussion and Future work

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.