SIR-Hawkes: Linking Epidemic Models and Hawkes Processes to Model Diffusions in Finite Populations

Marian-Andrei Rizoïu
ANU & Data61 CSIRO
Canberra, Australia

Swapnil Mishra
ANU & Data61 CSIRO
Canberra, Australia

Quyu Kong
ANU & Data61 CSIRO
Canberra, Australia

Mark Carman
Monash University
Melbourne, Australia

Lexing Xie
ANU & Data61 CSIRO
Canberra, Australia

ABSTRACT

Among the statistical tools for online information diffusion modeling, both epidemic models and Hawkes point processes are popular choices. The former originate from epidemiology, and consider information as a viral contagion which spreads into a population of online users. The latter have roots in geophysics and finance, view individual actions as discrete events in continuous time, and modulate the rate of events according to the self-exciting nature of event sequences. Here, we establish a novel connection between these two frameworks. Namely, the rate of events in an extended Hawkes model is identical to the rate of new infections in the Susceptible-Infected-Recovered (SIR) model after marginalizing out recovery events – which are unobserved in a Hawkes process. This result paves the way to apply tools developed for SIR to Hawkes, and vice versa. It also leads to HawkesN, a generalization of the Hawkes model which accounts for a finite population size. Finally, we derive the distribution of cascade sizes for HawkesN, inspired by methods in stochastic SIR. Such distributions provide nuanced explanations to the general unpredictability of popularity: the distribution for diffusion cascade sizes tends to have two modes, one corresponding to large cascade sizes and another one around zero.

ACM Reference Format:

Marian-Andrei Rizoïu, Swapnil Mishra, Quyu Kong, Mark Carman, Lexing Xie. 2018. SIR-Hawkes: Linking Epidemic Models and Hawkes Processes to Model Diffusions in Finite Populations. In *WWW 2018: The 2018 Web Conference, April 23–27, 2018, Lyon, France*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3178876.3186108>

1 INTRODUCTION

The research community has long been aware of the importance of the word-of-mouth phenomenon in information dissemination and in shaping user behavior in online and offline environments. In this paper, we study how information spreads online by modeling its underlying mechanism, i.e. how it passes from individual to individual. The aim is to link individual actions to collective-level measures, such as popularity or fame.

This paper is published under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW 2018, April 23–27, 2018, Lyon, France

© 2018 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC BY 4.0 License.

ACM ISBN 978-1-4503-5639-8/18/04.

<https://doi.org/10.1145/3178876.3186108>

This work addresses three open questions concerning two classes of approaches mainly used for modeling online diffusions: *epidemic models* and *Hawkes point processes*. The first open question regards the relationship between these two models. Epidemic models emerged from the field of epidemiology, and consider information as a viral contagion which spreads within a population of online users; Hawkes models have been mainly used in finance and geophysics, and view individual broadcasts of information as events in a stochastic point process. **Despite having the origins in different disciplines, these two models describe the stochastic series of discrete events; is there an inherent connection between them?** The second question is about designing more expressive diffusion models. Hawkes processes are the de facto modeling choice for social media processes, mainly because they can be easily customized to account for social factors such as the influence of users [14, 44], the length of “social memory” [23, 29] and the inherent content quality [21]. **Can we employ notions from epidemic models to design a Hawkes process more adept at describing online diffusions?** The third question concerns predicting the final size of the cascade, which intuitively reflects the popularity of the underlying message. Previous work [23, 28, 29, 44] predict a single value for the expected future popularity, however it is well known that popularity is hard to predict. There are many random factors lead to high variance in prediction [37]. **Can we compute the size distribution, to explain the high variance and hence the unpredictability?**

In this work, we address all three questions above, by drawing for the first time the connection between epidemic models and point processes, validating it both theoretically and also empirically on three large publicly available datasets of retweet cascades.

We answer the first question by studying the previously unexplored link between the Susceptible-Infected-Recovered (SIR) epidemic model [18] and the Hawkes processes. The key to the link is in the modeling of the word-of-mouth process: we regard each new each new broadcast from one user to another as an event in Hawkes, and analogous to a new infection in SIR. Starting from this observation, we show that the rate of events in an extended Hawkes model is identical to the rate of new infections in the SIR model, after taking the expectation over recovery events – which are unobserved in the Hawkes process. This is significant, as it indicates that tools developed for one approach can be applied to the other.

To answer the second question, we propose HawkesN, an extension of the Hawkes model with a finite population. The Hawkes

process [16] has no upper limit for the number of events that may occur. This is hardly a realistic assumption for social media processes such as information diffusion, which relies on a finite underlying population of humans, each broadcasting a message a finite number of times. We introduce a parameter N , denoting the finite total size of the population, and we modulate the event rate by the available population. We study the estimation of N from data and we construct a lower bound statistic to detect when parameter N does not have a valid solution. We show that the HawkesN model explains better longer event sequences.

To address the third question, we construct a probability distribution over future cascade size by applying a Markov chain technique developed for SIR to a diffusion cascade which has been partially observed and fitted using HawkesN. Based on our observations on a large sample of real diffusion cascades, we also provide a nuanced explanation for the main-stream belief that popularity is unpredictable. The distribution shows two peaks: the larger peak corresponds to the cascade extinguishing quickly after its beginning; the smaller peak corresponds to a large cascade size. At the beginning of the cascade it is impossible to distinguish between the two cases, however the posterior probability distribution after observing a prefix of the cascade can be updated to account for the observed events.

The main contributions of this work include:

- We show a previously unexplored connection between two different classes of approaches – epidemic models and Hawkes point processes – by showing that the rate of events in HawkesN is identical to the expected rate of new infections in SIR after marginalizing out recovery events.
- We introduce HawkesN – a novel class of Hawkes processes in which event intensity is modulated by the remaining population size – and we show it generalizes better to unseen data than the state-of-the-art modeling.
- We study the estimation of population size from observed data and we construct a lower bound statistic to detect when parameter N does not have a valid solution.
- We use a Markov chain tool from epidemic model theory to predict the distribution of the final size of a cascade. We provide a nuanced explanation for the main-stream belief that popularity is unpredictable.
- We construct ACTIVEERT – a new Twitter cascades benchmark dataset, publicly available (together with the HawkesN simulation and fitting R code) at: <https://github.com/computationalmedia/sir-hawkes>

2 PREREQUISITES

In this section, we briefly review a few key concepts of the Poisson and Hawkes [16] point processes (in Sec. 2.1), and of the SIR epidemic model and its bivariate process formulation (in Sec. 2.2).

2.1 Poisson and Hawkes processes

The Poisson processes. A point process is a random process whose realizations consists of event times $t_1, t_2, \dots$ [7], where t_j denotes the time of occurrence of the j -th event. In a *homogeneous* Poisson processes, the inter-arrival times $\tau_j = t_j - t_{j-1}$ are random variables *i.i.d.* exponentially distributed with parameter λ – also

called the *event rate* of the Poisson process. In *non-homogeneous* Poisson processes, the event rate is a deterministic time-continuous function $\lambda(t)$, which defines the probability of an event occurring in the infinitesimal interval around time t . Formally:

$$\begin{aligned} \mathbb{P}(N_{t+h} = n + m \mid N_t = n) &= \lambda(t)h + o(h) & \text{when } m = 1 \\ \mathbb{P}(N_{t+h} = n + m \mid N_t = n) &= o(h) & \text{when } m > 1 \\ \mathbb{P}(N_{t+h} = n + m \mid N_t = n) &= 1 - \lambda(t)h + o(h) & \text{when } m = 0 \end{aligned} \quad (1)$$

where $o(h)$ is a function so that $\lim_{h \downarrow 0} \frac{o(h)}{h} = 0$; N_t is the counting process associated with the point process, i.e. a random variable which counts the number of events up to (and including) time t .

The Hawkes process [16] is a self-exciting point process, in which each previous event occurred at the time $t_j < t$ generates new events at the rate $\phi(t - t_j)$ – also called the *kernel* of the Hawkes process. The event rate of a Hawkes process is a stochastic function dependent on previous event times, defined as:

$$\lambda(t) = \mu + \sum_{t_j < t} \phi(t - t_j) \quad (2)$$

which models the following process [20]: a new event either enters the system at the background rate μ ; or it is generated by a previous event at the rate of the corresponding kernel function.

2.2 The SIR Model

The Susceptible-Infected-Recovered (SIR) model defines three classes of individuals (also known as compartments): those *susceptible* to infection, those currently *infected* (and therefore infectious) and those *recovered* from the infection and no longer infective. SIR models the following process: when a susceptible individual meets an infectious individual, the former becomes infected at the rate β ; infected individuals recover from the infection at a constant rate γ .

Deterministic SIR. In the deterministic SIR, the individuals and their assignment to each of the three compartments are not observed. The temporal dynamics of the sizes of each of the compartments are governed by the following ordinary differential equations [1]:

$$\frac{dS(t)}{dt} = -\beta \frac{S(t)}{N} I(t) \quad (3)$$

$$\frac{dI(t)}{dt} = \beta \frac{S(t)}{N} I(t) - \gamma I(t) \quad (4)$$

$$\frac{dR(t)}{dt} = \gamma I(t). \quad (5)$$

$S(t)$, $I(t)$ and $R(t)$ are deterministic functions, denoting the sizes at time t of the susceptible, infected and recovered populations, respectively; $N = S(t) + I(t) + R(t)$ is the total population size.

There are a number of assumptions made by the SIR model. Firstly, it assumes that the population is homogeneous and individuals meet any other individual uniformly at random. Secondly, it assumes that all rates are constant: the infection rate β and the recovery rate γ . Thirdly, it assumes that the population has no births and no deaths – i.e. N is constant throughout the unfolding of the epidemic. The last assumption holds when the speed of the epidemic outpaces considerably the speed of change in the population – e.g., an average retweet diffusion only lasts minutes, compared to years of expected activity of a user on Twitter.

Stochastic SIR. Several stochastic formulations of the SIR model have been proposed [1], which model the behavior of independent and identically distributed agents. The actions of the agents are described by the same set of holistic rules defined in Eq. (3)-(5) and the same assumptions detailed above [4].

One such stochastic formulation of SIR is the bivariate point process representation [39], in which two types of events occur: *infection events* and *recovery events*. The j -th infected individual in the SIR process gets infected at time t_j^I and recovers at time t_j^R . Therefore, to each infection event corresponds a recovery event. S_t , I_t and R_t are discrete random variable taking integer values. They are the stochastic counterparts of $S(t)$, $I(t)$ and $R(t)$ respectively. Eq. (3)-(5) can be written as stochastic differential equations, but for ease of following we will keep referring to Eq. (3)-(5) in the rest of this paper. We define the *time to recovery* (i.e. the time the individual j is infectious) as: $\tau_j = t_j^R - t_j^I$. From Eq. (5) results that the times to recovery are distributed exponentially with parameter γ , and infections last on average $\frac{1}{\gamma}$ units of time.

Let C_t be the counting process of the *infection process* and R_t the counting process of the *recovery process*. Note that $C_t = N - S_t$ is total number of occurred infections (regardless if they are still infectious) and it is distinct from I_t (number of infectious at time t). Let $\mathcal{H}_t$ be the history of the bivariate epidemic process up to time t , i.e. $\mathcal{H}_t = \{t_1^I, t_2^I, \dots, t_1^R, t_2^R, \dots\}$. It can be shown that the rate of new infections $\lambda^I(t)$ and the rate of new recoveries $\lambda^R(t)$ are:

$$\lambda^I(t) = \beta \frac{S_t}{N} I_t; \quad \lambda^R(t) = \gamma I_t. \quad (6)$$

We sketch the proof for the previous statement. Yan [39] derives the probability of a new infection at time t given $\mathcal{H}_t$ as:

$$\begin{aligned} \mathbb{P}(C_{t+\delta t} - C_t = 1 | \mathcal{H}_t) &= \beta \frac{S_t}{N} I_t \delta t + o(\delta t) \\ \mathbb{P}(C_{t+\delta t} - C_t > 1 | \mathcal{H}_t) &= 0 \\ \mathbb{P}(C_{t+\delta t} - C_t = 0 | \mathcal{H}_t) &= 1 - \beta \frac{S_t}{N} I_t \delta t + o(\delta t), \end{aligned}$$

Given Eq. (1), the new infections process is a temporal point process of intensity $\beta \frac{S_t}{N} I_t$. $\lambda^R(t)$ is derived similarly.

Fig. 1 illustrates an SIR realization as a bivariate point process: five infection events occur at times $t_1^I, \dots, t_5^I$ (shown in red); five recovery events occur at $t_1^R, \dots, t_5^R$ (shown in blue). The middle panel of Fig. 1 shows the size of the infectious population I_t over time. Each new infection event increments I_t , and each new recovery decreases I_t by one. The bottom panel of Fig. 1 shows the corresponding new infection and new recovery rates. Initially, $\lambda^I(t)$ is significantly higher than $\lambda^R(t)$. As the number of susceptible individuals S_t gets depleted, the term $\frac{S_t}{N}$ in Eq. (6) inhibits $\lambda^I(t)$ which becomes zero after the fifth infection ($S_t = 0, t \geq t_5^I$). The new recovery rate also becomes zero after the last infected individual recovers ($I_t = 0, t \geq t_5^R$).

The connection between deterministic and stochastic SIR is that the mean behavior of the stochastic process asymptotically approaches that of the deterministic process [1, 39]. The connection between the deterministic and the stochastic population sizes is $S(t) = \mathbb{E}_{\mathcal{H}_t}[S_t]$, $I(t) = \mathbb{E}_{\mathcal{H}_t}[I_t]$ and $R(t) = \mathbb{E}_{\mathcal{H}_t}[R_t]$ [1]. Our own results simulating the two variants are presented in an online supplement [25].

Figure 1: An illustration of SIR as a bivariate point process: the infection process and the recovery process. (top panel) The j -th individual gets infected at time t_j^I and recovers at t_j^R . The time to recovery $\tau_j = t_j^R - t_j^I$ is the period the individual stays infectious. (middle panel) The size of the infectious population $I(t)$ over time. (lower panel) The infection rate $\lambda^I(t)$ and the recovery rate $\lambda^R(t)$ for the SIR parameters: $N = 5$, $\beta = 2$, $\gamma = 0.5$.

As will be elaborated in Sec. 3.2, the bivariate point process SIR formulation provides the link to the Hawkes point processes.

3 LINKING EPIDEMIC MODELS AND HAWKES PROCESSES

We first propose *HawkesN*, a generalization of the Hawkes process with finite population (in Sec. 3.1) and we show the connection between HawkesN and the SIR epidemic model (in Sec. 3.2).

3.1 HawkesN: a process in finite population

We generalize the Hawkes model [16] to account for finite population sizes. Intuitively cascades not only follow self-exciting word of mouth diffusions, but they are also limited by the size of the relevant community. The effect of introducing the finite population size N is that the event rate at time t is modulated by the available population. To the best of our knowledge, no prior work on modeling social processes using Hawkes models had accounted for a finite underlying population.

The event rate function in HawkesN is defined as:

$$\lambda^H(t) = \left(1 - \frac{N_t}{N}\right) \left[\mu + \sum_{t_j < t} \phi(t - t_j) \right], \quad (7)$$

where $\phi(t - t_j)$ can be the same kernel function used with Hawkes, and N_t is the counting process associated with the point process. Both $\lambda^H(t)$ and N_t are right-continuous functions. The term $1 - \frac{N_t}{N}$ scales the event rate at time t with the proportion of the events which can still occur after time t . When $t = 0$, we have $\lambda^H(t) = \mu$. When $N_t = N$, we have $\lambda^H(t) = 0$, i.e., there will be no more new events when the pool of users who can act is exhausted. When

Figure 2: An example diffusion illustrating the finite population effects in self-exciting process. (top panel) An event refers to the j^{th} user taking an action at time t_j (e.g. posting a tweet). The state of the user population is shown at each time t_j : **purple users** have performed the past observed actions; **orange users** are yet to perform any action. **(middle panel)** The counting process N_t increases by one with each event; **(lower panel)** The offspring rate $\phi_1(t)$ – the rate of events generated by this first event at time t_1 , modeled by Hawkes and by HawkesN (denoted by $\phi_1^H(t)$).

$N \rightarrow \infty$, Eq. (7) simplifies to Eq. (2). In other words, the Hawkes process is a special case of HawkesN with infinite population.

Fig. 2 illustrates the HawkesN process for an information diffusion in a population of five users ($N = 5$). Each user takes an action at most once, represented as event time $t_j, j = 1..5$. The corresponding counting process N_t is shown in the middle plot. Events $t_2..t_5$ are considered to have been triggered by event t_1 . The bottom panel compares the *offspring rates* – the rate of events generated by the first event at t_1 – for Hawkes (denoted as $\phi_1(t)$) and HawkesN (denoted as $\phi_1^H(t)$). In HawkesN, the population modulates the event rate by decreasing it after each new event and the event rate becomes zero after t_5 . The Hawkes process does not take into account the population size, i.e. it will have $\phi(t) > 0$ in Eq. 2 even after t_5 .

We use the exponential kernel function for HawkesN:

$$\phi(\tau) = \kappa \theta e^{-\theta \tau} \quad (8)$$

κ is a scaling factor and θ is the parameter of the exponential function which models the decay of *social memory*. The exponential kernel is a common choice in literature [2, 8, 11, 23, 29, 43, 44]. Other kernels have been used with Hawkes, including power-law functions [6, 17, 19, 23] and Rayleigh functions [34]. Using HawkesN with non-exponential kernel functions is left for future work.

3.2 Linking HawkesN and SIR

We now present our main result, Theorem 3.1, which links stochastic SIR and the HawkesN process.

Intuition. When modeling information diffusion, both SIR and HawkesN model the same phenomenon: users come into contact with the diffused content, which they further broadcast to other users. Each new broadcast is modeled as a new event in HawkesN, and as a new infection in SIR. The key to linking HawkesN and SIR models is the conceptual similarity between an event in HawkesN and a new infection in SIR. In HawkesN, past events generate new events at the rate $\phi(t)$, which is exponentially time-decaying in Eq. (8). In SIR, an infectious individual j infects susceptible individuals at a rate of $\frac{\beta S_t}{N}$ during the time it is infectious τ_j , which is exponentially distributed with parameter γ (discussed in Sec. 2.2).

THEOREM 3.1. *Suppose the new infections in a stochastic SIR process of parameters $\{\beta, \gamma, N\}$ follow a point process of intensity $\lambda^I(t)$. Suppose also the events in a HawkesN process with parameters $\{\mu, \kappa, \theta, N\}$ have the intensity $\lambda^H(t)$ (Eq 7). Let $\mathcal{T} = \{t_1, t_2, \dots\}$ be the set of the times to recovery of the infected individuals in SIR. The expectation of $\lambda^I(t)$ over $\mathcal{T}$ is equal $\lambda^H(t)$:*

$$\mathbb{E}_{\mathcal{T}}[\lambda^I(t)] = \lambda^H(t),$$

when $\mu = 0, \beta = \kappa \theta, \gamma = \theta$.

Note that both $\mathbb{E}_{\mathcal{T}}[\lambda^I(t)]$ and $\lambda^H(t)$ are random functions, as they depend on the random infection times t_j^I (for SIR) and the random event times t_j (for HawkesN). The expectation only removes the randomness from the recovery times t_j^R in SIR.

The rest of this section proves this theorem.

The expected new infection rate. We express S_t and I_t in Eq. (6) using indicator functions of the infection event times and the times to recovery:

$$\begin{aligned} S_t &= N - C_t = N - \sum_{j \geq 1} \mathbb{1}(t_j^I < t) \\ I_t &= C_t - R_t = \sum_{j \geq 1} \mathbb{1}(t_j^I < t, t_j^R > t) = \sum_{t_j^I < t} \mathbb{1}(t_j^I + \tau_j > t). \end{aligned} \quad (9)$$

We examine a point process consisting only of the infection events $\{t_j^I\}$. The event rate in this process is obtained by marginalizing out times of recovery:

$$\begin{aligned} \mathbb{E}_{\mathcal{T}}[\lambda^I(t)] &\stackrel{(6),(9)}{=} \mathbb{E}_{\mathcal{T}} \left[\beta \frac{S_t}{N} \sum_{t_j^I < t} \mathbb{1}(t_j^I + \tau_j > t) \right] \\ &= \sum_{t_j^I < t} \mathbb{E}_{\mathcal{T}} \left[\beta \frac{S_t}{N} \mathbb{1}(t_j^I + \tau_j > t) \right] \\ &= \sum_{t_j^I < t} \int_0^\infty \beta \frac{S_t}{N} \mathbb{1}(t_j^I + \zeta > t) r(\zeta) d\zeta \\ &= \sum_{t_j^I < t} \beta \frac{S_t}{N} \int_{t-t_j^I}^\infty r(\zeta) d\zeta, \end{aligned} \quad (10)$$

Figure 3: Visualization of the space of states of the SIR bivariate Markov Chain $\{s, i\}$. $s, i \in \mathbb{N}$ and $s + i \leq N$, in other words the space of valid states sits under the green line $s + i = N$. The initial state in an SIR epidemic is always on the green line ($S(0) + I(0) = N$). Given an initial state $\{s = S(0), i = I(0)\}$ (shown by the magenta circle), the orange area shows the space of reachable states and the gray area depicts the unreachable states. From a state $\{s, i\}$ the system can to $\{s - 1, i + 1\}$ (new infection); and to $\{s, i - 1\}$ (new recovery). The absorbing states ($\{s, 0\}$) is shown with a red line.

where $r(\zeta)$ is the exponential probability distribution function for the time to recovery (cf. Sec. 2.2). Knowing $S_t = N - C_t$, we obtain:

$$\mathbb{E}_{\mathcal{T}}[\lambda^I(t)] = \left(1 - \frac{C_t}{N}\right) \sum_{t_j^I < t} \beta e^{-\gamma(t-t_j^I)}. \quad (11)$$

We can see that Eq. (7) (with the exponential kernel in Eq. (8)) and Eq. (11) are identical when $N_t = C_t$ and under the parameter equivalence in Theorem 3.1. That is to say, the new infection point process in an SIR model and the HawkesN point-process with no background event rate are described by the same conditional intensity. This completes the proof of Theorem 3.1. We also demonstrate the equivalence empirically, through simulation and subsequent parameter fitting, in the online supplement [25].

4 DIFFUSION SIZE DISTRIBUTION

We compute the probability distribution of the final size of an information diffusion cascade which has been partially observed and fitted using HawkesN, using a Markov chain technique developed for SIR. In Sec. 4.1 we review known results on the final size distribution of an SIR epidemic. In Sec. 4.2 we employ the equivalence shown in Theorem 3.1 to compute the final size distribution of a cascade modeled with HawkesN.

4.1 Epidemic size distributions in SIR

The final size of an infection is defined as the total number of individuals that have been infected (and recovered) during the epidemic. Estimating the final size while the epidemic is in its early stages is a well-studied problem in epidemiology. In this section, we review a solution to this problem for the stochastic SIR model.

SIR as a Markov chain. The stochastic SIR introduced in Sec. 2.2 can be formulated as a bivariate continuous time-homogeneous Markov chain [1]. Each state $\sigma \in \Sigma$ is uniquely defined by the ordered pair of random variables $\{S_t, I_t\}$, denoting the sizes of the susceptible and the infected populations. Σ is the finite space of all possible states, visually represented in Fig. 3 as a triangle on the two-dimensional surface with the number of susceptible on the x-axis and the number of infected on the y-axis. From a given state $\{S_t = s, I_t = i\}$ (denoted from here on as $\{s, i\}$), there are only two other states in which the system can transition depending on the type of event that occurs: a new infection arrives, and the system transitions $\{s, i\} \rightarrow \{s - 1, i + 1\}$ with the probability $p(\{s - 1, i + 1\}|\{s, i\})$; or a new recovery arrives, and the system transitions $\{s, i\} \rightarrow \{s, i - 1\}$ with the probability $p(\{s, i - 1\}|\{s, i\})$ (shown in the figure by the red and blue arrows, respectively).

Transition matrix and probabilities. The epidemic ends when $I_t = 0$, i.e. no more infectious individuals exist to propagate the epidemic. Consequently, the states $\{s, 0\}$ are *absorbing states* – once the system arrives in one of these states, it does not transition to any other state. From a non-absorbing state $\{s, i\}$, new infections are observed at the rate $\frac{\beta}{N}si$ and new recoveries at the rate of γi (see Eq. (6)). We obtain the transition probabilities:

$$p(\{s - 1, i + 1\}|\{s, i\}) = \frac{\frac{\beta}{N}si}{\frac{\beta}{N}si + \gamma i} = \frac{\beta s}{\beta s + N\gamma}, \text{ for } i > 0$$

$$p(\{s, i - 1\}|\{s, i\}) = \frac{\gamma i}{\frac{\beta}{N}si + \gamma i} = \frac{N\gamma}{\beta s + N\gamma}, \text{ for } i > 0. \quad (12)$$

Suppose that the states in Σ are ordered from 1 to $|\Sigma|$. We define the *transition matrix* $T = [t_{kl}]$ of size $|\Sigma| \times |\Sigma|$, where $t_{kl} = p(\sigma_k = \{s_k, i_k\}|\sigma_l = \{s_l, i_l\})$ is the probability of transitioning from state σ_l to state σ_k . From Eq. (12) we obtain:

$$t_{kl} = \begin{cases} \frac{\beta s_l}{\beta s_l + N\gamma} & , s_k = s_l - 1, i_k = i_l + 1, i_l > 0 \\ \frac{N\gamma}{\beta s_l + N\gamma} & , s_k = s_l, i_k = i_l - 1, i_l > 0 \\ 1 & , s_k = s_l, i_k = i_l = 0 \\ 0 & , \text{otherwise} \end{cases}. \quad (13)$$

Note that the sum of each column j in the transition matrix M is equal to 1, as it contains the probabilities of transitioning from σ_l to another state.

Probability state vector and size distribution. Let π be the probability state vector $\pi \in \mathbb{R}^{|\Sigma| \times 1}$, with the l^{th} position of π giving the probability that the system is currently found in state σ_l . Let $\pi_0 = [0, \dots, 0, 1, 0, \dots, 0]$ be the initial probability vector, with the value of 1 corresponding to $\tilde{\sigma}$ the *initial state*, and zero everywhere else. Starting from π_0 , we compute the probability state vector after one transition as $\pi^{(1)} = M \times \pi_0$. $\pi^{(2)} = M^2 \times \pi_0$ gives the probabilities after two transitions, $\pi^{(3)} = M^3 \times \pi_0$ after three etc. Given that there are at most $N - 1$ infection events and N recovery events in an SIR realization, the system is guaranteed to converge after $2N - 1$ steps [1]. At convergence, all states except the absorbing states have a probability of zero in $\pi^{(2N-1)}$. We denote as $P(s)$ the value in $\pi^{(2N-1)}$ for the state $\{s, 0\}$, $s = 0, 1, \dots, N - I_0$.

From an initial state $\tilde{\sigma}$ the system can finish in any of the absorbing states $\{s, 0\}$ with the probability $P(s)$. The distribution of final size of the diffusion is that of the random variable $N' = N - s$.

4.2 Cascade size distribution in HawkesN

Given Theorem 3.1, computing the probability distribution over the final size of the cascade after observing an arbitrary number of events, conceptually amounts to changing the initial state $\tilde{\sigma} = \{\tilde{s}, \tilde{i}\}$ and using the method described in Sec. 4.1. Suppose we observed l events in HawkesN, we have $\tilde{s} = N - l$. The recovery events are not observed and the exact size of the infected population $\tilde{i}$ is not known. We compute its expectation over times to recovery $\mathcal{T}$ as:

$$\mathbb{E}_{\mathcal{T}}[\tilde{i}] = \mathbb{E}_{\mathcal{T}} \left[\sum_{j=1}^l \mathbb{1}(t_j^I + \tau_j > t_l) \right] \stackrel{\text{Eq. (10), (11)}}{=} \sum_{j=1}^l e^{-\gamma(t_l - t_j^I)} \quad (14)$$

and we run the method in Sec. 4.1 starting from the initial state $\tilde{\sigma} = \{N - l, \mathbb{E}_{\mathcal{T}}[\tilde{i}]\}$. $\mathbb{E}_{\mathcal{T}}[\tilde{i}]$ in Eq. (14) is a real number, that we round to the closest integer.

In our discussion in Sec. 6.3, we study two probability distributions: the *a priori* distribution is computed starting from the initial state $\tilde{\sigma} = \{N - 1, 1\}$ and it is dependent on model parameters only. The *a posteriori* distribution is the size distribution after observing l events, it is computed starting from $\tilde{\sigma} = \{N - l, \mathbb{E}_{\mathcal{T}}[\tilde{i}]\}$ and it is dependent on model parameters and the observed event times $t_1, t_2, \dots, t_l$ (cf. Eq. (14)).

5 FITTING HAWKESN TO DATA

In epidemiology, the size of population N is usually considered a fixed known parameter – e.g. the number of people in a community. For online diffusions, it could be possible to estimate N from past diffusions (as discussed in Sec. 8). However, in this section we analyze the case when N is not known beforehand and needs to be estimated from observed data.

5.1 The likelihood of HawkesN

Let $\{t_1, t_2, \dots, t_n\}$ be a set of event times assumed to have been generated from a HawkesN process described in Sec. 3.1. When modeling diffusion cascades, it is typically assumed that every event apart from t_1 is a reaction to the first event, i.e., the background intensity is zero $\mu = 0, \forall t > 0$ [23]. We estimate the remaining HawkesN parameters $\{\kappa, \theta, N\}$ by maximizing the log-likelihood function of the point process (see the online supplement [25] or Daley and Vere-Jones [7, Ch. 7.2]):

$$\mathcal{L}(\kappa, \theta, N) = \sum_{j=1}^n \log \lambda^H(t_j) - \int_0^{t_n} \lambda^H(\tau) d\tau. \quad (15)$$

We further detail the integral term:

$$\begin{aligned} \int_0^{t_n} \lambda^H(\tau) d\tau &= \int_0^{t_n} \left(1 - \frac{N_t}{N}\right) \sum_{t_j < t} \phi(t - t_j) dt \\ &= \sum_{j=0}^{n-1} \int_{t_j}^{t_n} \left(1 - \frac{N_t}{N}\right) \phi(t - t_j) dt \\ &= \sum_{j=0}^{n-1} \sum_{l=j}^{n-1} \frac{N-l}{N} \int_{t_l}^{t_{l+1}} \phi(t - t_j) dt \\ \text{cf. (8)} &= \kappa \sum_{j=0}^{n-1} \sum_{l=j}^{n-1} \frac{N-l}{N} \left[e^{-\theta(t_l - t_j)} - e^{-\theta(t_{l+1} - t_j)} \right]. \quad (16) \end{aligned}$$

Table 1: Percentage of non-valid N solutions and accuracy of detection using statistic $\mathcal{S}$.

Percentage observed	5%	10%	20%	40%	80%
Non-valid N (SIR)	4%	0%	0%	0%	0%
Non-valid N (HawkesN)	56%	42%	37%	18%	0%
valid N roots	44	58	63	82	100
found using $\mathcal{S} < 0$	40	57	63	82	100

Eq. (15) is a non-linear objective and there are a few natural constraints for each of the model parameters, namely: $\theta > 0, \kappa > 0$ and $N \geq n$. We use the mathematical modeling language AMPL [10], which provides an interface to different tools for continuous optimization, including automatic gradient computation and solvers. We choose as solver Ipopt [33], a common choice in literature for large problems with non-linear objectives. More details can be found in the online supplement [25].

5.2 Estimating population size N

Here we examine the case when the population size N is the only unknown. The purpose is to identify how difficult it is to retrieve the value of N from data. Having a value of N for which the derivative is zero is a necessary condition for a local maximum in the log-likelihood function in Eq. (15). We write the derivative of the log-likelihood with respect to N :

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial N} &= \sum_1^n \frac{(\mu(t) + \sum_{t_j < t} \phi(t - t_j)) \frac{j-1}{N^2}}{(\mu(t) + \sum_{t_j < t} \phi(t - t_j)) \frac{N-j+1}{N}} \\ &\quad - \kappa \sum_{j=0}^{n-1} \sum_{l=j}^{n-1} \frac{l}{N^2} \left[e^{-\theta(t_l - t_j)} - e^{-\theta(t_{l+1} - t_j)} \right] \\ &= \sum_1^n \frac{j-1}{N(N-j+1)} - \frac{1}{N^2} \sum_{j=0}^{n-1} \sum_{l=j}^{n-1} l \kappa \left[e^{-\theta(t_l - t_j)} - e^{-\theta(t_{l+1} - t_j)} \right]. \quad (17) \end{aligned}$$

Knowing that $N - j + 1 \leq N, \forall j = 1..n$, we construct a lower bound for $\frac{\partial \mathcal{L}}{\partial N}$:

$$\frac{\partial \mathcal{L}}{\partial N} \geq \frac{1}{N^2} \underbrace{\left(\frac{(n-1)n}{2} - \sum_{j=0}^{n-1} \sum_{l=j}^{n-1} l \kappa \left[e^{-\theta(t_l - t_j)} - e^{-\theta(t_{l+1} - t_j)} \right] \right)}_{\mathcal{S}(\kappa, \theta, \{t_1, t_2, \dots, t_n\})} \quad (18)$$

We define the right hand side of Eq. (18) $\mathcal{S}(\kappa, \theta, \{t_1, t_2, \dots, t_n\})$ as a statistic of the observed event times t_j and of fixed model parameters κ and θ . The statistic does not depend on N . Given parameters κ and θ , and a set of event times t_j , when the statistic $\mathcal{S}(\kappa, \theta, \{t_1, t_2, \dots, t_n\}) > 0$ it is guaranteed that the log-likelihood function $\mathcal{L}$ keeps monotonically increasing with $N \geq n$, and no valid solution exists for N .

Difficulty of estimating N . We illustrate the difficulty of estimating N by simulation, and we show that this is dependent on the number of observed events in the cascade. Starting from the set of

Figure 4: (top panel) Example of a simulated SIR process realization, containing 10 infection events and 10 recovery events. **(middle panel)** The log-likelihood as a function of N for HawkesN (over the infection events, in red) and for SIR (in blue). **(bottom panel)** The derivative of the log-likelihood w.r.t. N . For HawkesN, the derivative is always positive and the log-likelihood monotonically increases, and there is no valid N solution. For SIR, the log-likelihood has a maximum at 102, close to the simulated value $N = 100$. Statistic value for this realization $S = 13.77$.

Figure 5: Median and 15%/85% confidence intervals for estimating N using the maximum log-likelihood using HawkesN (red color) and SIR (blue color).

parameters $\mu = 0, \kappa = 5, \theta = 0.2, N = 100$, we simulate 100 realizations using stochastic SIR. Assuming fixed all parameters except N , we study the validity and the quality of estimating N over increasingly longer prefixes of each cascade containing a percentage of all events in the range [5%, 100%]. SIR observes both infection and recovery events in each prefix, while HawkesN observes infection events only.

We implement a numerical procedure for finding N : we divide the range $[0, 200]$ into 1000 intervals and we numerically search each interval for a root for Eq. 17 using `uniroot` in R. This is a slow procedure which provides a ground truth against which we compare the statistic S . Table 1 shows for how many cascades there

Table 2: Datasets profiling: number of cascades and number of tweets tweets; min, mean and median cascade size.

	#cascades	#tweets	Min.	Mean	Median
ACTIVE RT [28]	41,411	8,142,892	20	197	41
SEISMIC [44]	166,076	34,784,488	50	209	111
NEWS [23]	20,093	3,252,549	50	162	90

is no valid solution for N for SIR (first row) and HawkesN (second row). Five observed percentages are shown 5%, 10%, 20%, 40% and 80%. For HawkesN, 56% of cascades do not have a valid N after observing 5% of the events. We observe that, when more than 50% of the cascade is observed, a solution for N exists. For SIR however, only 10% of the cascade have a non-valid N at the beginning of the cascades, and all cascades have valid solutions to N once more than 10% is observed. This indicates that it is more difficult to estimate N in HawkesN than in SIR. The bottom two rows of Table 1 show how many cascades have valid solutions for N and for how many of these the statistic S is negative. $S < 0$ mis-identifies only 4 valid solutions (out of 44) after observing 5% of the cascade and it identifies all valid solutions for percentages greater than 10%. Note that S is a lower bound for the log-likelihood and it is guaranteed to find all non-valid solutions.

Fig. 4 shows an example of a cascade with 10 infection events and 10 recovery events. No valid solution exists for N when using HawkesN – the log-likelihood function is monotonically increasing. When the recovery events are observed in SIR, N has a feasible solution close to the ground truth. This indicates that the timing of the recovery events (not observed in HawkesN) embeds information about the size of the population. Fig. 5 confirms this conclusion, showing that SIR estimates correctly N even at the beginning of cascades, whereas HawkesN requires observing around 50% of the the cascade to estimate N correctly. For the full details of the simulation and additional analytic analysis, please consult the online supplement [25].

6 EXPERIMENTS AND RESULTS

In this section, we investigate the performances of HawkesN on three Twitter diffusion datasets (described in Sec. 6.1). We evaluate the generalization performance of HawkesN (in Sec. 6.2) and we profile cascade size distributions and we provide a new explanation for the perceived popularity unpredictability (Sec. 6.3).

6.1 Datasets

We use three datasets of retweet diffusion cascades in Twitter, used in previous work. For each tweet in each cascade, we have information about the time offset of the retweet and the number of followers of the user posting the retweet. The ACTIVE RT dataset was collected by Rizoio et al. [28] during 6 months in 2014. It contains more than 41k retweet cascades related to more than 13k Youtube videos, each cascade containing at least 20 tweets. The SEISMIC dataset was collected by Zhao et al. [44]. It contains a sample of all tweets during a month (i.e. using the firehose Twitter API restricted access), further filtered so that the length of each cascade is greater than 50. The NEWS dataset was collected by Mishra et al. [23] over a period of

Figure 6: Performances of HawkesN explaining unobserved data, using holdout negative log likelihood. The performance over all cascades in a dataset are summarized using boxplots, lower is better (a). The percentage of observed events in each cascade used to train HawkesN is varied between 10% and 95%. We use 1000 cascades randomly sampled from NEWS. (b), (c) and (d) The performances on all cascades of ACTIVE RT, SEISMIC and NEWS, for Hawkes and HawkesN, when observing 40% and 80% of each cascade.

Figure 7: Final size probability distribution for a NEWS cascade. HawkesN was fitted on the first 27 events (left) and 47 events (right). The black line shows the *a priori* distribution, after seeing a single event. the blue line shows the *a posteriori* distribution, after seeing the observed events. Shown with vertical lines: the number of observed events in gray dashed, the deterministic prediction in magenta dashed and the actual observed cascade size in red.

four moths in 2015. They selected tweets containing links to news articles, by tracking the official twitter handles of popular news outlets, such NewYork Times, or CNN. Each cascade contains at least 50 tweets. Table 2 summarizes these datasets.

6.2 Generalization to unobserved data

All three datasets described in the previous section also contain user information for each tweet. The tweets are pairs $\{m_j, t_j\}$, $t = 1, \dots, n$, where m_j is the number of followers of the user having emitted tweet j at time t_j . In this section, we choose to use the modified exponential kernel function proposed by Mishra et al. [23], which also accounts for the number of followers for a user: $\phi(\tau) = \kappa m^\eta \theta e^{-\theta \tau}$. More details about the marked HawkesN and its equivalence with SIR are found in the online supplement [25].

We empirically validate HawkesN by studying how it generalizes to unseen data. We compare HawkesN with the Hawkes model for information diffusion, proposed by Mishra et al. [23]. We adopt the setup in [2, 8, 11, 28, 29, 44]: the first few events in a diffusion are observed and used to fit the models. Hawkes is fitted as described

in [23], and HawkesN is fitted as described in Sec. 5.1. The population size N is also fitted from data. We measure the holdout likelihood, i.e. the likelihood of the events in the unobserved period. The lower the negative holdout likelihood, the better the model generalizes to unseen data. We report the *per event* holdout negative likelihood, to render the results comparable across holdout sets containing different numbers of events. Given the analysis in Sec. 5.2, we chose to observe a given proportion of each cascade, to render the results comparable across cascades of different length.

Fig. 6a shows the generalization performances of HawkesN, when varying the percentage of observed events from 10% to 95%. Consistent with the conclusions in Sec. 5.2, we observe a high variance of performance when observing less than 40% of each cascade. The basic Hawkes model shows less variance at lower percentages (shown in the online supplement [25]). Plots (b) to (d) in Fig. 6 show the generalization performance of Hawkes and HawkesN, on the three datasets, for the observed percentages of 40% and 80%. Visibly, HawkesN has a consistently lower median value for the negative log-likelihood than Hawkes for higher observed percentages. The mean negative log-likelihood values are comparable for HawkesN and Hawkes on NEWS and SEISMIC. On ACTIVE RT the mean of HawkesN is higher – likely due to Youtube videos behaving differently, with some old ones (e.g. Music) still being shared.

For higher observed percentages, the mean negative log-likelihood improves for HawkesN and it degrades for Hawkes. This indicates that the modulation factor $(1 - \frac{N}{N_t})$ in Eq. 7 helps improve likelihood, and HawkesN fits longer event sequences better.

6.3 Explaining popularity unpredictability

In this section, we study the probability distribution of population size for real-life cascades. Both left and right plots in Fig. 7 show the same cascade from the NEWS dataset, with the HawkesN parameters fit on 27 and 47 events respectively (here N is a meta-parameter fixed at $N = 80$). The *a priori* probability size distribution – the distribution after observing only the first event – shows two maxima: one around very small values of cascade size, and one around ~ 65 . This provides the following explanation for the general perceived unpredictability of online popularity. For cascades showing

a bi-modal *a priori* size distribution, there are two likely outcomes: either it dies out early or it reaches a large size compared to the maximum population N . At time $t = 0$ it is impossible to differentiate between the two outcomes. The situation is different after observing a number of events. The *a posteriori* probability distribution (shown in Fig. 7 with a blue line) reflects the information gained from the observed events and it shows a single maximum towards the higher size values. The more events we observe, the higher the likelihood of the true value of cascade size. We also observe that the size distribution gets narrower as we observe more events, i.e. there is less uncertainty in cascade size prediction. This provides another explanation to why autoregressive popularity prediction approaches [5, 27, 31] achieve higher results. Online popularity has been previously claimed to be unpredictable [21], however as far as we know this is the first explanation for it, based on analytical results on size distributions.

7 RELATED WORK

We structure some of the related work in the field of social media analysis into two broad categories, based on the used framework: point process approaches and epidemic models.

Point process approaches. Point-process based generative models are a popular choice for popularity modeling [6, 8, 40] and prediction [22, 28, 29, 44]. In their seminal work, Crane and Sornette [6] link popularity bursts and decays to the effects of a Hawkes self-exciting process. More sophisticated models have been proposed to model and simulate popularity in microblogs [40] and videos [8]. These approaches successfully account for the social phenomena which modulate online diffusion: the “rich-get-richer” phenomenon and social contagion. Certain models can output an estimate for the total size of a retweet cascade. Shen et al. [29] employ reinforced Poisson processes, modeling three phenomena: fitness of an item, a temporal relaxation function and a reinforcement mechanism; while SEISMIC [44] employs a double stochastic process, one accounting for infectiousness and the other one for the arrival time of events. Our work differs from the above in two aspects. First it proposes a generalization of the Hawkes model, which operates in a finite population – which is a more realistic assumption. Second, it outputs a size probability distribution and it explains the perceived unpredictability of online popularity [21].

A recently emerging body of work employs Stochastic Differential Equations to formulate Hawkes point processes. RedQueen [43] and Cheshire [42] are two algorithms aimed at optimizing social influence, which they formulate as a stochastic optimal control problem. Wang et al. [35] use stochastic control and reinforcement learning to address the user activity guiding problem and feedback in social systems. Later, Wang et al. [36] use the stochastic differential equation model to link the microscopic event data and macroscopic inference, and to approximate its probability distribution. The similarity between the above and our work is at the level of tools, by using stochastic calculus to link the event-level to the event rate and compute expected quantities. However, none of the above links point processes to epidemic models. The advantage of our solution is that it enables to leverage the mature tools in epidemic models to the field of information diffusion.

Epidemic model approaches. Despite being developed for the field of epidemiology, epidemic models have been applied to information diffusion problems through the analogy of information spread as a disease. Classic epidemic models were early applied in the knowledge and scientific theory diffusion study [13] and latter employed in many areas, such as economic and finance time series analysis [30]. Pastor-Satorras and Vespignani [26] applied SIS (Susceptible-Infected-Susceptible) epidemic model to simulate computer virus transmission over the Internet. A series of studies analyzed the spread of rumors in complex networks based on an epidemic model [15, 24, 32, 41]. More recently, Woo and Chen [38] modeled topic diffusion in web forums using an SIR model; Martin et al. [21] fitted an epidemic model to retweet cascades and used the computed basic reproduction number to theorize the unpredictability of online popularity. Bauckhage et al. [3] use a stochastic SIR model to characterize attention dynamics of viral videos. Goel et al. [12] apply large-scale agent based SIR simulation on a random network to study the virality on Twitter diffusions. Feng et al. [9] propose a fractional SIR model in which the infection probability of a node is proportional to its fraction of infected neighbors and apply it on Sina Weibo data. However, these work do not leverage tools specific to epidemic models (e.g. the probability distribution of size), nor do they link to point process models as our work does.

8 CONCLUSION

In this work, we present a previously unexplored connection between Hawkes point processes and SIR epidemic models. First, we establish a novel connection between these two frameworks by linking the rate of events in an extended Hawkes model to the rate of new infections in the Susceptible-Infected-Recovered (SIR) model after marginalizing out recovery events – which are unobserved in a Hawkes process. This paves the way to applying tools developed for one approach to the other approach. It also leads to HawkesN, an extension of the Hawkes process with a finite number of events. Finally, we present a novel method to compute the probability distribution of the final size of a cascade after observing its initial unfolding using HawkesN, which is based on a Markov chain tools developed for SIR. We use the probability of cascade size on a large sample of real cascades to provide a nuanced explanation for the general unpredictability of popularity.

Assumptions, limitations and future work. This work assumes a fixed population (users don’t enter, nor do they exit). A link could be drawn between evolving populations in SIR and $\mu(t) \neq 0$ in HawkesN. The current work assumes that the maximum population size N is estimated for each cascade, while observing it. Future work could use other observed similar cascades to infer the size of a “thematic neighborhood” before a cascade starts unfolding. Finally, allowing for user-specific behavior in the SIR model or kernel functions other than the exponential function requires more advanced SIR formulations, such as an agent-based formulation.

Acknowledgments. This material is based on research sponsored by the Air Force Research Laboratory, under agreement number FA2386-15-1-4018. We thank the National Computational Infrastructure (NCI) for providing computational resources, supported by the Australian Government. We thank Aditya Krishna Menon for insightful discussions.

REFERENCES

- [1] Linda J. S. Allen. 2008. An Introduction to Stochastic Epidemic Models. In *Mathematical Epidemiology*. Springer, Berlin, Heidelberg, Chapter 3, 81–130. https://doi.org/10.1007/978-3-540-78911-6_3
- [2] Peng Bao, Hua-Wei Shen, Xiaolong Jin, and Xue-Qi Cheng. 2015. Modeling and Predicting Popularity Dynamics of Microblogs using Self-Excited Hawkes Processes. In *Proceedings of the 24th International Conference on World Wide Web - WWW '15 Companion*. ACM Press, New York, New York, USA, 9–10. <https://doi.org/10.1145/2740908.2742744>
- [3] Christian Bauckhage, Fabian Hadji, and Kristian Kersting. 2015. How Viral Are Viral Videos?. In *ICWSM*. 22–30.
- [4] Georgiy V. Bobashev, D. Michael Goedecke, Feng Yu, and Joshua M. Epstein. 2007. A hybrid epidemic model: Combining the advantages of agent-based and equation-based approaches. In *Proceedings - Winter Simulation Conference*. IEEE, 1532–1537. <https://doi.org/10.1109/WSC.2007.4419767>
- [5] Biao Chang, Hengshu Zhu, Yong Ge, Enhong Chen, Hui Xiong, and Chang Tan. 2014. Predicting the Popularity of Online Serials with Autoregressive Models. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management - CIKM '14*. ACM Press, New York, New York, USA, 1339–1348. <https://doi.org/10.1145/2661829.2662055>
- [6] Riley Crane and Didier Sornette. 2008. Robust dynamic classes revealed by measuring the response function of a social system. *Proceedings of the National Academy of Sciences* 105, 41 (oct 2008), 15649–15653. <https://doi.org/10.1073/pnas.0803685105>
- [7] D J Daley and D Vere-Jones. 2008. *An introduction to the theory of point processes. [Vol. I]*. Vol. I. xviii+573 pages. <https://doi.org/10.1007/b97277>
- [8] Wanying Ding, Yue Shang, Lifan Guo, Xiaohua Hu, Rui Yan, and Tingting He. 2015. Video Popularity Prediction by Sentiment Propagation via Implicit Network. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*. ACM, 1621–1630. <https://doi.org/10.1145/2806416.2806505>
- [9] Ling Feng, Yanqing Hu, Baowen Li, H Eugene Stanley, Shlomo Havlin, and Lidia A Braunstein. 2015. Competing for attention in social media under information overload conditions. *PLoS one* 10, 7 (2015), e0126090.
- [10] Robert Fourer, David M Gay, and Brian W Kernighan. 1987. *AMPL: A mathematical programming language*. AT&T Bell Laboratories Murray Hill, NJ 07974.
- [11] Shuai Gao, Jun Ma, and Zhumin Chen. 2015. Modeling and Predicting Retweeting Dynamics on Microblogging Platforms. In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining - WSDM '15*. ACM Press, New York, New York, USA, 107–116. <https://doi.org/10.1145/2684822.2685303>
- [12] Sharad Goel, Ashton Anderson, Jake Hofman, and Duncan J Watts. 2015. The structural virality of online diffusion. *Management Science* 62, 1 (2015), 180–196.
- [13] William Goffman. 1971. A Mathematical Method for Analyzing the Growth of a Scientific Discipline. *Journal of the ACM (JACM)* 18, 2 (apr 1971), 173–185. <https://doi.org/10.1145/321637.321640>
- [14] Manuel Gomez-Rodriguez, Le Song, Nan Du, Hongyuan Zha, and Bernhard Schölkopf. 2016. Influence Estimation and Maximization in Continuous-Time Diffusion Networks. *ACM Transactions on Information Systems* 34, 2 (feb 2016), 1–33. <https://doi.org/10.1145/2824253>
- [15] Daniel Gruhl, R. Guha, David Liben-Nowell, and Andrew Tomkins. 2004. Information Diffusion Through Blogspace. In *Proceedings of the 13th International Conference on World Wide Web (WWW '04)*. ACM, New York, NY, USA, 491–501. <https://doi.org/10.1145/988672.988739>
- [16] Alan G. Hawkes. 1971. Spectra of some self-exciting and mutually exciting point processes. *Biometrika* 58, 1 (apr 1971), 83–90. <https://doi.org/10.1093/biomet/58.1.83>
- [17] Agnès Helmstetter and Didier Sornette. 2002. Subcritical and supercritical regimes in epidemic models of earthquake aftershocks. *Journal of Geophysical Research: Solid Earth* 107, B10 (2002), ESE 10–1—ESE 10–21. <https://doi.org/10.1029/2001jb001580> arXiv:cond-mat/0109318
- [18] W. O. Kermack and A. G. McKendrick. 1927. A Contribution to the Mathematical Theory of Epidemics. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 115, 772 (aug 1927), 700–721. <https://doi.org/10.1098/rspa.1927.0118>
- [19] Ryota Kobayashi and Renaud Lambiotte. 2016. TiDeH: Time-Dependent Hawkes Process for Predicting Retweet Dynamics. In *ICWSM 2016*. arXiv:1603.09449
- [20] Patrick J. Laub, Thomas Taimre, and Philip K. Pollett. 2015. Hawkes Processes. (jul 2015). arXiv:1507.02822 <http://arxiv.org/abs/1507.02822>
- [21] Travis Martin, Jake M. Hofman, Amit Sharma, Ashton Anderson, and Duncan J. Watts. 2016. Exploring Limits to Prediction in Complex Social Systems. In *Proceedings of the 25th International Conference on World Wide Web*. 683–694. <https://doi.org/10.1145/2872427.2883001> arXiv:1602.01013
- [22] Hongyuan Mei and Jason Eisner. 2017. The Neural Hawkes Process: A Neurally Self-Modulating Multivariate Point Process. In *Advances in Neural Information Processing Systems* 30. 6757–6767. arXiv:1612.09328 <https://arxiv.org/abs/1612.09328>
- [23] Swapnil Mishra, Marian-Andrei Rizoio, and Lexing Xie. 2016. Feature Driven and Point Process Approaches for Popularity Prediction. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management - CIKM '16*. ACM Press, Indianapolis, IN, USA, 1069–1078. <https://doi.org/10.1145/2983323.2983812>
- [24] Yamin Moreno, Maziar Nekovee, and Amalio F. Pacheco. 2004. Dynamics of rumor spreading in complex networks. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* 69, 6 2 (jun 2004), 066130. <https://doi.org/10.1103/PhysRevE.69.066130> arXiv:cond-mat/0312131
- [25] online supplement. 2017. Appendix: SIR Hawkes: Linking Epidemic Models and Hawkes Point Processes for Online Information Diffusion. (2017). <https://arxiv.org/pdf/1711.01679.pdf#page=11>.
- [26] Romualdo Pastor-Satorras and Alessandro Vespignani. 2001. Epidemic spreading in scale-free networks. *Physical Review Letters* 86, 14 (apr 2001), 3200–3203. <https://doi.org/10.1103/PhysRevLett.86.3200> arXiv:cond-mat/0010317
- [27] Henrique Pinto, Jussara M. Almeida, and Marcos A. Gonçalves. 2013. Using early view patterns to predict the popularity of youtube videos. In *Proceedings of the sixth ACM international conference on Web search and data mining - WSDM '13*. ACM Press, New York, New York, USA, 365. <https://doi.org/10.1145/2433396.2433443>
- [28] Marian-Andrei Rizoio, Lexing Xie, Scott Sanner, Manuel Cebrian, Honglin Yu, and Pascal Van Hentenryck. 2017. Expecting to be HIP: Hawkes Intensity Processes for Social Media Popularity. In *26th International Conference on World Wide Web - WWW '17*. ACM Press, Perth, Australia, 735–744. <https://doi.org/10.1145/3038912.3052650> arXiv:1602.06033
- [29] Hw Shen, Dashun Wang, Chaoming Song, and Al Barabási. 2014. Modeling and Predicting Popularity Dynamics via Reinforced Poisson Processes. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*. AAAI Press, Québec City, Québec, Canada, 291–297. arXiv:arXiv:1401.0778v1
- [30] Ernest S. Shtatland and Timur Shtatland. 2008. Another Look at Low-Order Autoregressive Models in Early Detection of Epidemic Outbreaks and Explosive Behaviors in Economic and Financial Time Series. In *SGF Proceedings*.
- [31] Gabor Szabo and Bernardo a. Huberman. 2010. Predicting the popularity of online content. *Commun. ACM* 53, 8 (aug 2010), 80. <https://doi.org/10.1145/1787234.1787254> arXiv:0811.0405
- [32] Daniel Trpevski, Wallace K. S. Tang, and Ljupco Kocarev. 2010. Model for rumor spreading over networks. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* 81, 5 (may 2010), 056102. <https://doi.org/10.1103/PhysRevE.81.056102>
- [33] A Wächter and L T Biegler. 2006. On the Implementation of a Primal-Dual Interior Point Filter Line Search Algorithm for Large-Scale Nonlinear Programming. *Mathematical Programming* 106, 1 (2006), 25–57.
- [34] Jacco Wallinga and Peter Teunis. 2004. Different epidemic curves for severe acute respiratory syndrome reveal similar impacts of control measures. *American journal of epidemiology* 160, 6 (sep 2004), 509–16.
- [35] Yichen Wang, Evangelos Theodorou, Apurv Verma, and Le Song. 2016. A Stochastic Differential Equation Framework for Guiding Online User Activities in Closed Loop. (mar 2016). arXiv:1603.09021 <http://arxiv.org/abs/1603.09021>
- [36] Yichen Wang, Xiaojing Ye, Haomin Zhou, Hongyuan Zha, and Le Song. 2017. Linking Micro Event History to Macro Prediction in Point Process Models. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, Vol. 54. 1375–1384. <http://proceedings.mlr.press/v54/wang17f.html>
- [37] Duncan J. Watts. 2011. *Everything is obvious : once you know the answer*. Crown Business. 335 pages.
- [38] Jiyoung Woo and Hsinchun Chen. 2016. Epidemic model for information diffusion in web forums: experiments in marketing exchange and political dialog. *SpringerPlus* 5, 1 (dec 2016), 66. <https://doi.org/10.1186/s40064-016-1675-x>
- [39] Ping Yan. 2008. Distribution Theory, Stochastic Processes and Infectious Disease Modelling. In *Mathematical Epidemiology*, Wu J. Brauer F., van den Driessche P. (Ed.). Springer, Berlin, Heidelberg, Chapter 10, 229–293. https://doi.org/10.1007/978-3-540-78911-6_10
- [40] Linyun Yu, Peng Cui, Fei Wang, Chaoming Song, and Shiqiang Yang. 2017. Uncovering and predicting the dynamic process of information cascades with survival model. *Knowledge and Information Systems* 50, 2 (feb 2017), 633–659. <https://doi.org/10.1007/s10115-016-0955-7> arXiv:1505.07193
- [41] Damián H. Zanette. 2002. Dynamics of rumor propagation on small-world networks. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* 65, 4 (mar 2002), 041908. <https://doi.org/10.1103/PhysRevE.65.041908> arXiv:0110324
- [42] Ali Zarezade, Abir De, Hamid Rabiee, and Manuel Gomez Rodriguez. 2017. Cheshire: An Online Algorithm for Activity Maximization in Social Networks. (mar 2017). arXiv:1703.02059 <http://arxiv.org/abs/1703.02059>
- [43] Ali Zarezade, Utkarsh Upadhyay, Hamid Rabiee, and Manuel Gomez Rodriguez. 2017. RedQueen: An Online Algorithm for Smart Broadcasting in Social Networks. In *10th ACM International Conference on Web Search and Data Mining*. arXiv:1610.05773 <http://arxiv.org/abs/1610.05773>
- [44] Qingyuan Zhao, Murat A Erdogdu, Hera Y He, Anand Rajaraman, and Jure Leskovec. 2015. SEISMIC: A Self-Exciting Point Process Model for Predicting Tweet Popularity. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*.