

Modeling Information Cascades with Self-exciting Processes via Generalized Epidemic Models

Quyú Kong
Australian National University &
UTS & Data61, CSIRO
Canberra, Australia
quyu.kong@anu.edu.au

Marian-Andrei Rizoiu
University of Technology Sydney &
Data61, CSIRO
Sydney, Australia
marian-andrei.rizoiu@uts.edu.au

Lexing Xie
Australian National University &
Data61, CSIRO
Canberra, Australia
lexing.xie@anu.edu.au

ABSTRACT

Epidemic models and self-exciting processes are two types of models used to describe diffusion phenomena online and offline. These models were originally developed in different scientific communities, and their commonalities are under-explored. This work establishes, for the first time, a general connection between the two model classes via three new mathematical components. The first is a generalized version of stochastic Susceptible-Infected-Recovered (SIR) model with arbitrary recovery time distributions; the second is the relationship between the (latent and arbitrary) recovery time distribution, recovery hazard function, and the infection kernel of self-exciting processes; the third includes methods for simulating, fitting, evaluating and predicting the generalized process. On three large Twitter diffusion datasets, we conduct goodness-of-fit tests and holdout log-likelihood evaluation of self-exciting processes with three infection kernels – exponential, power-law and Tsallis Q-exponential. We show that the modeling performance of the infection kernels varies with respect to the temporal structures of diffusions, and also with respect to user behavior, such as the likelihood of being bots. We further improve the prediction of popularity by combining two models that are identified as complementary by the goodness-of-fit tests.

KEYWORDS

Information Diffusion, Hawkes Processes, Epidemic Models

ACM Reference Format:

Quyú Kong, Marian-Andrei Rizoiu, and Lexing Xie. 2020. Modeling Information Cascades with Self-exciting Processes via Generalized Epidemic Models. In *The Thirteenth ACM International Conference on Web Search and Data Mining (WSDM '20)*, February 3–7, 2020, Houston, TX, USA. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3336191.3371821>

1 INTRODUCTION

Epidemic models and self-exciting processes are two classes of mathematical models that have evolved separately and been applied in distinct problem domains, one in epidemiology [17] and the other in seismology [12, 29], finance [2], and neural science [15]. *Epidemic*

models typically divide the population into compartments, such as Susceptible, Infected and Recovered for the Susceptible-Infected-Recovered (SIR) model, and describe the transitions between compartments as deterministic or stochastic processes. *Self-exciting point processes* are a class of processes in which the occurrence of each event increases the likelihood of future events using time-decaying kernel functions. Both models have been used to describe events in the physical world, as well as online information diffusions [21, 23, 38, 39]. This paper aims to establish a mathematical connection between these two model classes. By achieving this, this work contributes: 1) new expressive models for self-exciting processes in finite populations; 2) methods that account for unobserved recovery events, which are common in real-world epidemiological data; 3) new tools and insights into online information diffusion.

The Hawkes process with exponential kernels and stochastic SIR process have been recently shown [31] to share a connection via the infection intensity function when the recovery time in the SIR model is *latent*. However, this result is restricted to one particular parametric family of self-exciting processes, whereas Hawkes processes allow a richer set of kernel functions, and an inequality of the connection has been overlooked. These observations lead to the question: **How to both broaden and deepen the connection between epidemic models and Hawkes processes?** The broadening is with respect to arbitrary recovery time distributions and kernel functions, while the deepening is with respect to the mathematical relationships between two model classes. To address these, we propose a generalized stochastic SIR process in which infected individuals recover independently following an arbitrary distribution of recovery times. Next, we link this process to a finite-population Hawkes process (dubbed *HawkesN* [31]) by showing that the Complementary Cumulative Distribution Function (CCDF) of the recovery time (in SIR), given the infection event history, is an upper bound of the *HawkesN* kernel. We derive relationships among three key functions: the kernel function in *HawkesN*, the SIR recovery time distribution, and the recovery hazard function. We empirically evaluate the accuracy of recovering original parameters of stochastic SIR models from a *HawkesN* model fitted only on infection events.

Connecting the two model classes will enrich the computational tools of both. One challenge emerges – **what tools can be developed and applied through the generalized connection to both classes of models?** We first enrich the generalized SIR with concepts from Hawkes processes including *event marks* (features associated with events) and *branching factors* (expected number of future events generated by a new event). We then show a simulation algorithm for the generalized SIR process by paired-sampling

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

WSDM '20, February 3–7, 2020, Houston, TX, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-6822-3/20/02...\$15.00

<https://doi.org/10.1145/3336191.3371821>

of infection and recovery times. We also present maximum log-likelihood procedures for estimating the parameters, 2 metrics for measuring goodness-of-fit and approaches for predicting final diffusion popularity, for SIR and HawkesN processes with general kernels.

While generalized models allow flexibility in the choice of parametric forms, it is important to understand **how the performances of different model formulations vary on diffusions?** On three large Twitter diffusion datasets, we show that the HawkesN model with different kernels demonstrates diverse modeling capability on diffusions with distinct temporal dynamics. For instance, on one of the datasets, *NEWS*, the HawkesN model with an exponential kernel tends to fit diffusions that are larger in event counts and shorter in time frames. We show that this can result from the participation of automated bots to the online diffusions. These observations lead us to combine models for predicting diffusion final popularity, which outperforms all other models.

The main contributions of this work include:

- A generalized stochastic SIR process with arbitrary recovery time distributions and their connection to HawkesN processes with monotonically time-decaying kernels. The generalized model is equipped with concepts from Hawkes processes including event marks and branching factors.
- A set of tools including simulation, parameter estimation, evaluation and popularity prediction algorithms for SIR processes with general recovery time distributions.
- A series of fitting, model comparison and prediction results on real-world Twitter diffusion data. We observe that the performances of general SIR processes with different recovery distributions vary with respect to diffusion dynamics. In prediction experiments, a combined model performs the best.

Related work. Effort has been put into generalizing epidemic models. Keeling and Grenfell [16] reformulate the deterministic epidemic model as integro-differential equations, and impose a Gaussian distribution on the recovery times. Streftaris and Gibson [34] specify the recovery times following a Weibull distribution and Routledge et al. [33] model them using a Rayleigh distribution. On the Hawkes processes front, a rich set of kernel functions are available including power-law [26], piece-wise linear [41], Tsallis Q-Exponential [22], and general function approximators such as neural networks [9, 14, 25, 27]. Our work links the developments from both model classes via the proposed generalized connection.

In terms of the study of information diffusion using epidemic models, Kimura et al. [18] first apply the SIS model, which allows nodes to be activated multiple times, to study information diffusion in a network. Jin et al. [13] use an enhanced SEIZ, which introduces an extra *Exposed* state (E) to the SIR model for capturing an incubation period, to detect rumors from Twitter cascades. When studying online diffusion using self-exciting processes, Zhao et al. [40] and Mishra et al. [26] both employ power-law kernel functions with Hawkes processes, which achieve state-of-art performance in popularity prediction. Rizoïu et al. [31] apply HawkesN with an exponential kernel that outperforms the Hawkes counterpart in terms of holdout log-likelihood values. Different from these works which show superior performance for a specific form in one or two




Figure 1: A sample stochastic SIR process including an infection event history until time t , i.e., $\mathcal{H}_t^C = \{t_1^I, \dots, t_4^I\}$, and recovery events $\{t_2^R, t_1^R, t_3^R\}$. Infected individuals recover at time intervals τ following a distribution $f(t)$. The bottom plot presents a corresponding realization of HawkesN events. HawkesN events generate descendants with the intensity rate $\phi(t)$. A connection between $f(t)$ and $\phi(t)$ is explored when $f(t)$ is assumed of arbitrary parametric forms. The red color marks an invalid recovery event given upcoming infections.

evaluation tasks, our analysis corroborates several aspects of tests including goodness-of-fit, holdout log-likelihood and prediction.

2 PRELIMINARIES

In this section, we discuss two classes of stochastic event models, and highlight the missing link between them.

SIR models, originally proposed by Kermack and McKendrick [17], describe the number of people infected by an epidemic in a fixed population over time. The name stands for the three possible states for individuals – those in a Susceptible state can get Infected, and those infected will eventually Recover or be Removed, and they are no longer prone to the infection. The *stochastic* variant of the SIR model [3] is concerned with individual state changes, rather than expected volumes of individuals in each state. The transition of individuals from susceptible to infected is described by the *infection process*, and that from infected to recovered by the *recovery process*.

One can represent the stochastic SIR in a fixed population of size N as two sets of random event times, for the infections and recoveries, respectively. Let \mathcal{H}_t^C denote the set of infection event times that happened before time t , and $C_t = |\mathcal{H}_t^C|$ is the number of infection events up to time t . Let i index individuals in accordance with their infection time sequence, then $\mathcal{H}_t^C = \{t_i^I \mid t_1^I = 0, t_1^I < \dots < t_{C_t}^I < t\}$. The short hand C stands for *cumulative*, i.e., \mathcal{H}_t^C and C_t are not affected by the random events of individuals recovering. Similarly, let $\mathcal{H}_t^R = \{t_j^R \mid 0 < t_j^R < t, t_j^R > t_j^I\}$ denote a set of recovery event times before time t , and let $R_t = |\mathcal{H}_t^R|$ be the number of individuals recovered by time t . We use $U()$ to denote the (index) set of individuals in an event history \mathcal{H} . It follows from the sequential indexing that $U(\mathcal{H}_t^C) = \{1, 2, \dots, C_t\}$, and that the set of recovered individuals is a subset of those infected $U(\mathcal{H}_t^R) \subset U(\mathcal{H}_t^C)$. We use \mathcal{H}_t^I to express the set of infection event times of infected individuals who have *not* recovered by time t ,

i.e., $\mathcal{H}_t^I = \{t_j^I \mid t_j^I < t, t_j^R > t\}$, and $I_t = |\mathcal{H}_t^I|$. It is easy to see that the still infected set complements the recovered set $U(\mathcal{H}_t^C) = U(\mathcal{H}_t^R) \cup U(\mathcal{H}_t^I)$, and $C_t = R_t + I_t$. Fig. 1 shows an example of a stochastic SIR process. Based on the definitions above we have: $\mathcal{H}_t^C = \{t_1^I, t_2^I, t_3^I, t_4^I\}$, $\mathcal{H}_t^R = \{t_2^R, t_1^R, t_3^R\}$, $\mathcal{H}_t^I = \{t_4^I\}$, $C_t = 4$, $R_t = 3$, $I_t = 1$, at time $t = t_3^R + \epsilon$. The *susceptible* individuals are the ones who have never been *infected*, namely are currently neither *infected* nor *recovered*: $S_t = N - C_t = N - I_t - R_t$.

The stochastic SIR process is defined by an infection event intensity function $\lambda^I(t)$ and a recovery event intensity function $\lambda^R(t)$ [36]

$$\lambda^I(t) = \beta \frac{S_t}{N} I_t; \quad \lambda^R(t) = \gamma I_t \quad (1)$$

where β and γ are known as the infection rate and the recovery rate in SIR terminology. The total infection rate is proportional to the susceptible population S_t and the infected population I_t . Each infected individual recovers independently with the same recovery rate γ , hence the total recovery rate is proportional to the size of the infected population. It is also assumed that the recovery process is *simple* [7], i.e., only one infection or recovery event can happen in any infinitesimal time interval.

Consider the random variable *recovery time* – the elapsed time between an individual’s infection and recovery. Eq. (1) implies that the *recovery time* is exponentially distributed $f(t) = \gamma e^{-\gamma t}$ [36].

Hawkes processes are a type of self-exciting point processes, i.e. processes in which the occurrence of events increases the likelihood of future events [12]. This property is modeled via the intensity function:

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

where μ is the background intensity, and $\phi : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is known as the triggering kernel – the rate of new events generated by event t_i – and the summation aggregates the influences of all past events.

HawkesN process is a finite-population variant of the Hawkes process [31]. Assuming the diffusion occurs in a fixed population of size N , the event intensity is modulated by the proportion of remaining population:

$$\lambda^H(t) = \frac{N - N_t}{N} \sum_{t_i < t} \phi(t - t_i) \quad (3)$$

N_t is the number of events up to time t , the background intensity μ is set to zero and the first event happens at time 0, i.e., $N_0 = 1$.

The stochastic SIR and Hawkes processes have been developed by separate scientific communities for modeling different natural phenomena (epidemics and financial transactions/earthquakes, respectively). It is desirable to connect these apparently disparate tools using the common language of stochastic point processes.

3 LINKING SIR AND HAWKESN

First, we present a generalized stochastic SIR model with an arbitrary recovery time distribution, and next we reveal the connection between the general stochastic SIR and HawkesN. Finally, we extend the generalized SIR model with concepts from the Hawkes models.

3.1 SIR with general recovery distributions

As discussed in Section 2, the stochastic SIR process implicitly assumes that recovery times of infected individuals are exponentially distributed. Here we relax this assumption by letting recovery times follow an arbitrary distribution $f(t)$. The recovery intensity for each individual is given by the hazard function $h(t)$ [6], i.e., the recovery time distribution conditioned on recovering after time t :

$$h(t) = \frac{f(t)}{\int_t^\infty f(\tau) d\tau} \quad (4)$$

Considering that individuals recover independently, the overall recovery event intensity is the superposition of recovery intensities of the individuals still infected at time t :

$$\lambda^R(t) = \sum_{t_i^I \in \mathcal{H}_t^I} h(t - t_i^I) = \sum_{t_i^I \in \mathcal{H}_t^I} \frac{f(t - t_i^I)}{\int_{t-t_i^I}^\infty f(\tau) d\tau} \quad (5)$$

The overall infection event intensity remains unchanged as in Eq. (1). Note that, when $f(t)$ is the exponential distribution, Eq. (5) simplifies to the infection intensity of the classic SIR in Eq. (1).

Despite being rather straightforward, to the best of our knowledge, this is the first work presenting this generalized SIR with arbitrary recovery distributions.

3.2 Marginalizing over recovery events

One of the challenges for using the SIR model for social media diffusions is that the definitions of infection and recovery are not straightforward. Infection events can be interpreted as posting, sharing or retweeting, and they are usually recorded in data traces; recovery events can be the times when these posts or discussion topics lose traction, which are rarely directly observable. This observation implies that one may treat recovery events as *latent*, and examine the expected process after marginalizing over them.

We use $\mathbb{E}_{\{t_i^R | \mathcal{H}_t^C\}} [\lambda^I(t)]$ to denote the expected infection intensity over all recovery event times up to time t :

$$\begin{aligned} \mathbb{E}_{\{t_i^R | \mathcal{H}_t^C\}} [\lambda^I(t)] &\stackrel{(a)}{=} \beta \frac{S_t}{N} \sum_{t_i^I \in \mathcal{H}_t^I} \int_{t-t_i^I}^\infty f(\tau_i | \mathcal{H}_t^C) d\tau_i \\ &\stackrel{(b)}{\geq} \beta \frac{S_t}{N} \sum_{t_i^I \in \mathcal{H}_t^I} \int_{t-t_i^I}^\infty f(\tau_i) d\tau_i \end{aligned} \quad (6)$$

Eq. (6a) follows from Rizoiu et al. [31]. Step (b) is because, given \mathcal{H}_t^C an infection history observed up to time t , the recovery event time of the i^{th} individual t_i^R ($i \in U(\mathcal{H}_t^C)$) is dependent on the entire \mathcal{H}_t^C . Fig. 1 illustrates this dependence with the red recovery event being an invalid candidate for t_1^R given $\mathcal{H}_t^C = \{t_1^I, t_2^I, t_3^I, t_4^I\}$. Intuitively, if the first individual recovers at the time of the red event, there will be zero infected individuals afterwards, rendering impossible the rest of the diffusion. We simplify the dependence using the inequality in Eq. (6b) to the recovery time distribution $f(t)$. We show that

$$\int_{t-t_i^I}^\infty f(\tau_i | \mathcal{H}_t^C) d\tau_i \geq \int_{t-t_i^I}^\infty f(\tau_i) d\tau_i \quad (7)$$

with the left and right terms being equal when $t_i^I = \max\{\mathcal{H}_t^C\}$. The proof is detailed in the online supplement [35, appendix A].

Table 1: Examples of HawkesN kernel functions $\phi(t)$, the corresponding SIR recovery time distributions $f(t)$ and hazard functions $h(t)$ following Eqs. (8)(9)(10). Parameter ranges: $\theta > 1$ for Tsallis Q-Exponential kernel, $\kappa > 0, \theta > 0, c > 0$ for all others.

HawkesN Kernel Name	HawkesN Kernel Function $\phi(t)$	SIR Recovery Time Distribution $f(t)$	SIR Recovery Hazard $h(t)$	Time Constraint t
Linear	$-\kappa\theta t + \kappa$	θ	$\frac{\theta}{-\theta t + 1}$	$(0, \frac{\kappa}{\theta})$
Quadratic	$\kappa\frac{\theta^2}{4}t^2 - \kappa\theta t + \kappa$	$-\frac{\theta^2}{2}t + \theta$	$\frac{\theta^2 t - 2\theta}{\theta^2 t^2 - 4\theta t + 4}$	$(0, \frac{2}{\theta})$
Gaussian	$\kappa e^{-\frac{t^2}{2\theta^2}}$	$\frac{t}{\theta^2} e^{-\frac{t^2}{2\theta}}$	$\frac{1}{\theta^2} t$	$(0, \infty)$
Tsallis Q-Exponential [22]	$\kappa [1 + (\theta - 1)t]^{-\frac{1}{\theta}}$	$[1 + (\theta - 1)t]^{-\frac{\theta}{\theta}}$	$1 + (\theta - 1)t$	$(0, \infty)$
Exponential [12]	$\kappa\theta e^{-\theta t}$	$\theta e^{-\theta t}$	θ	$(0, \infty)$
Power-law [26]	$\kappa(t + c)^{-(1+\theta)}$	$c^{1+\theta}(1 + \theta)(t + c)^{-(2+\theta)}$	$\frac{1 + \theta}{t + c}$	$(0, \infty)$

Comparing Eq. (6b) and Eq. (3), both $N - N_t$ (for HawkesN) and $S_t = N - C_t$ (for SIR) stand for the size of remaining susceptible population – hence the scaling factors S_t/N and $(N - N_t)/N$ are equivalent. Also, both Eq. (6b) and Eq. (3) sum over the infected population, and the integral in Eq. (6a) is a function of time since infection $t - t_i^I$. Therefore, marginalizing the recovery events reduces the infection intensity of the stochastic SIR to a lower bound – the HawkesN intensity – as long as the following relationship between the HawkesN kernel and recover time distribution holds:

$$\phi(t) = \beta \int_t^\infty f(\tau) d\tau \quad (8)$$

We can express $f(t)$ in terms of $\phi(t)$. $f(t)$ is a probability density function which implies $f(t) \geq 0$ and $\int_0^\infty f(\tau) d\tau = 1$, leading to $\phi(0) = \beta$:

$$f(t) = -\frac{\phi'(t)}{\phi(t)} \quad (9)$$

where we assume $\lim_{t \rightarrow \infty} f(t) = 0$. Eq. (8) and Eq. (9) spell out the closed-form relationship between the recovery time distribution $f(t)$ of the stochastic SIR and the kernel function $\phi(t)$ of the HawkesN process. From Eq. (8), we note that this relationship only holds when $\phi(t)$ is a monotonically decreasing function. Incorporating Eq. (9) into Eq. (4), we can express the recovery hazard function in terms of the HawkesN kernel:

$$h(t) = -\frac{\phi'(t)}{\phi(t)} \quad (10)$$

Given that $\phi(t)$ is monotonically decreasing, $-\phi'(t)$ and $h(t)$ are non-negative.

Table 1 lists six examples of HawkesN kernels, with their corresponding recovery time distributions and recovery hazard functions. The first three rows show the linear, quadratic, and Gaussian kernels, followed by the Tsallis Q-Exponential kernel used in quantum optics and atomic physics [22]. The last two examples are the exponential kernel function and the power-law kernel function, widely used for financial data, geophysics, and information diffusion [2, 12, 26].

Relation to prior work. The relationship presented by Rizoiv et al. [31] omits the inequality shown in Eq. (6b), and it is a special case of the result in this work. Their reasoning is limited to the constant recovery hazard functions and the exponentially distributed recovery times, with $f(t) = \gamma e^{-\gamma t}$ and $\phi(t) = \kappa\theta e^{-\theta t}$. The main modeling contribution compared to [31, 36] is a new set of analytical relationships among general recovery time distributions, kernel and hazard functions, in Eqs. (8)(9)(10).

3.3 Marked stochastic SIR

In real data and apart from event times, additional information about individual events is available, such as the user profile of a retweet event or patient characteristics in epidemics. Mathematically, the event history $\mathcal{H}_m^C = \{(t_1^I, m_1), \dots, (t_n^I, m_n)\}$ is a sequence of pairs of event times and extra event information also known as *event marks*. To leverage this information, marked variations of Hawkes process models are proposed to incorporate event marks as a scaling factor of kernel functions [12]. This idea leads to a marked variation of the HawkesN model, with the intensity function as:

$$\lambda_m^H(t) = \frac{N - N_t}{N} \sum_{(t_i^I, m_i) \in \mathcal{H}_m^I(t)} m_i^\rho \phi(t - t_i^I) \quad (11)$$

where ρ controls a warping effect for the mark. Using the generalized connection introduced in Section 3.2, we are able to obtain a marked stochastic SIR model, whose infection intensity function is

$$\lambda_m^I(t) = \beta \frac{S_t}{N} \sum_{(t_i^I, m_i) \in \mathcal{H}_m^I(t)} m_i^\rho \quad (12)$$

where, comparing to Eq. (1), I_t was decomposed to $\sum_{(t_i^I, m_i) \in \mathcal{H}_m^I(t)} m_i^\rho$ to account for the individual mark information. The recovery intensity $\lambda_m^R(t)$ is identical to its unmarked counterpart in Eq. (5).

3.4 Branching factor for SIR

The basic reproduction number R_0 is an important quantity in epidemic models for determining whether an epidemic is likely to occur [1]. This quantity conceptually connects to the branching factor n^* from Hawkes processes which is defined as the expected

Algorithm 1 Simulating generalized stochastic SIR

Input: Recovery time distribution $f(t)$, parameters $\{N, \beta\}$

Output: Infection event times \mathcal{H}^C and recovery event times \mathcal{H}^R

```
1: Set current time  $T = 0$ .
2: Initialize  $\mathcal{H}^C = \{0\}$  with one initial infection at time 0.
3: Initialize  $\mathcal{H}^R = \{\eta\}$  where  $\eta \sim f(t)$  and  $t_1^R = \eta$ .
4: while  $|\mathcal{H}^C| < N$  do
5:    $s = -\frac{\log(u)}{\lambda^*}$  where  $u \sim U(0, 1)$ 
6:   Compute  $\Lambda^I(t) = \int_0^t \lambda^I(\eta) d\eta$  from  $\mathcal{H}^C, \mathcal{H}^R$ 
7:    $T = T + (\Lambda^I)^{-1}(s)$ 
8:   if  $T = \infty$  then
9:     break // No infection will occur
10:  else
11:     $\eta \sim f(t)$  // Draw recovery time, update histories
12:     $\mathcal{H}^R = \mathcal{H}^R \cup \{T + \eta\}, \mathcal{H}^C = \mathcal{H}^C \cup \{T\}$ 
13: return  $\mathcal{H}^C, \mathcal{H}^R$ 
```

number of events generated by a single infection event [31], i.e., $n^* = \int_0^\infty \phi(\tau) d\tau$. Building upon this observation and Eq. (8), we define R_0 for stochastic SIR with a general recovery time distribution as

$$R_0 = n^* = \beta \int_0^\infty \int_\eta^\infty f(\tau) d\tau d\eta \quad (13)$$

Based on [28], one can also generalize R_0 to $\beta \int_0^\infty \tau f(\tau) d\tau$, but we show in [35, appendix A] that this definition is equivalent to Eq. (13).

For marked variations, this quantity is computed by taking expectation over the distribution of event marks. Particularly, for retweet cascades where the event marks are the count of user followers, a power law distribution $P(m) = (\alpha - 1)m^{-\alpha}$ of exponent $\alpha = 2.016$ is determined by Mishra et al. [26]. We obtain

$$R_0 = n^* = \beta \frac{\alpha - 1}{\alpha - 1 - \rho} \int_0^\infty \int_\eta^\infty f(\tau) d\tau d\eta \quad (14)$$

We refer to this quantity as just the branching factor n^* in the following sections to avoid confusion.

4 A SET OF TOOLS FOR STOCHASTIC SIR

In this section, we introduce a set of tools for the stochastic SIR with general recovery time distributions and HawkesN, enabling one to simulate event realizations, estimate model parameters, assess fitted results and predict final diffusion sizes.

Generalized SIR simulation. The generalized SIR proposed in Eq. (5) cannot be simulated using the approach described by Allen [1] as the recovery event rate is no longer piece-wise constant. We show a procedure of sampling general stochastic SIR processes, by sampling each infection event and its corresponding recovery time.

Starting from the first infection event at $t = 0$, Algorithm 1 iterates between two steps. Step one is to sample the recovery event time according to $f(t)$ (line 11-12), step two is to sample the next infection time by the random time change theorem [20] (line 5-7). Specifically, because future recovery times have been sampled for existing infection events, the infection event intensity can be then derived from Eq. (1) as a piece-wise constant function. The infection intensity leads to analytical forms of the cumulative

infection intensity $\Lambda^I(t) = \int_0^t \lambda^I(s) ds$ and its inverse $(\Lambda^I)^{-1}(\cdot)$. It is presented that $(\Lambda^I)^{-1}(\cdot)$ can convert a time interval sampled from a Poisson process with unit rate (line 5) to an interval generated by the intensity function $\lambda^I(t)$ (line 7) [20]. The process terminates when all N individuals have been infected (line 4), or when the infection rate falls to zero (line 8).

Parameter estimation. We use maximum likelihood to estimate model parameters given event history via standard optimization packages. The likelihood functions of stochastic SIR and HawkesN can be derived from the general likelihoods for point processes [7]. Details are in the online supplement [35, appendix B].

Suppose events are generated with an underlying stochastic SIR model. To estimate its parameters, when both infection and recovery events are observed, the stochastic SIR likelihood is maximized; when only infection events are observed, we estimate with the HawkesN likelihood to account for their latent recovery information. Due to the inequality in Eq. (6), the HawkesN likelihood is a biased estimator for stochastic SIR process parameters. We study this bias in Section 5.1 and we show empirically that it reduces as the branching factor increases.

Goodness-of-fit assessment. Given that the generalized SIR model can accommodate a wide range of recovery distribution functions, one natural question is how to assess the fitness of fitted models to observed events, choose between different parametric families and provide a guide to predict future events [5]. Due to the aforementioned random time change theorem [20], for observed infection events $t_i^I \in \mathcal{H}_t^C$ correctly described by an infection intensity function $\lambda^I(t)$, the cumulative infection intensities between infection events are time intervals generated from a Poisson process with unit rate or, equivalently, follow a unit rate exponential distribution:

$$\mathcal{T}_i = \int_{t_{i-1}^I}^{t_i^I} \lambda^I(\tau) d\tau, \quad \mathcal{T}_i \sim e^{-t} \quad (15)$$

Three statistical tests are applied: the Kolmogorov-Smirnov (KS) test and the Excess Dispersion (ED) test to measure the significance of the proposition $\{\mathcal{T}_i\} \sim e^{-t}$; the Ljung-Box (LB) test to determine the independence among $\{\mathcal{T}_i\}$.

Lallouache and Challet [19] note that the KS test is a more demanding test than the ED test. Specifically, the KS test evaluates the empirical cumulative density function (CDF) of $\{\mathcal{T}_i\}$ against the theoretical CDF of the unit rate exponential distribution (i.e., $1 - e^{-t}$) producing two values: a p-value, indicating the significant level of $\{\mathcal{T}_i\}$ not being drawn from the nominated theoretical CDF, and a distance D between the empirical CDF and the theoretical CDF [24]. As models presented in this paper are evaluated against the same theoretical CDF, we employ this distance measure D as a fitting performance metric for model comparison.

Diffusion final size prediction Point processes are generally applied for event history explanation and not optimized for prediction. To predict the diffusion final size (a.k.a the popularity for a Twitter cascade), we follow [26] by using a regression layer on top of the proposed models. We predict a quantity σ which can be interpreted as the proportion of remaining population that will be involved in the diffusion, i.e.,

$$\hat{C}_\infty = C_t + \sigma(N - C_t) \quad (16)$$

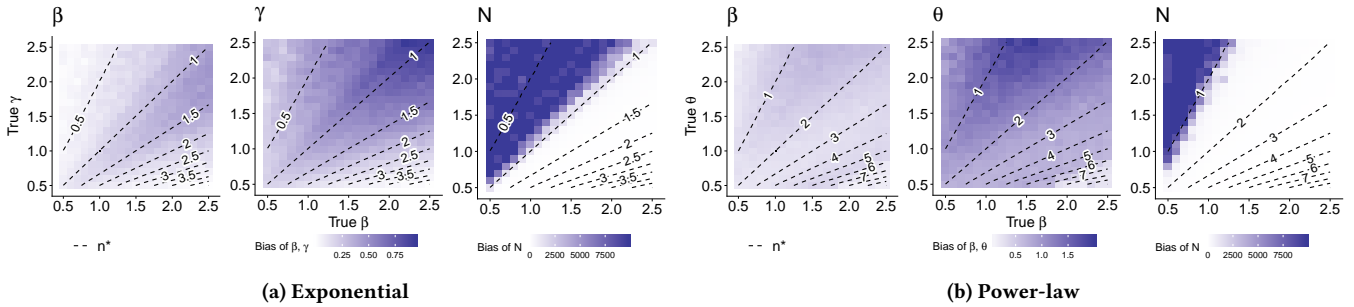


Figure 2: Bias of estimating parameters with HawkesN likelihood functions on simulated stochastic SIR infection events. Stochastic SIR with an exponential (a) and a power-law (b) recovery time distributions are evaluated. Chosen parameters are: (a) $N = 200$, $\beta = \gamma = \{0.5, 0.6, \dots, 2.5\}$; (b) $N = 200$, $c = 2$, $\beta = \theta = \{0.5, 0.6, \dots, 2.5\}$. Estimation bias is computed with *absolute errors* — lighter colors indicate lower bias. The dotted contour lines are the branching factors given the parameter sets.

where t is the observation time, C_t is the number of cumulative infection events, N is the fitted population size and \hat{C}_∞ is the predicted diffusion final size. We note that $\sigma > 1$ is possible due to the underestimation of N given observed events or the growth of population as diffusion unfolds. We use the fitted parameters and the derived branching factor (e.g., $\{\beta, \gamma, \rho, N, n^*\}$) for exponentially recovered stochastic SIR as features to train a *sigma* predictor. This setup can also be applied to HawkesN given N_t and its fitted parameters. The prediction experiment is set up to reproduce experiments in [26], and we further detail it in Section 5.

5 EXPERIMENTS

We first study the fitting of SIR parameters when the recovery times are not observed, and we design an empirical validation of the connection between the stochastic SIR and the HawkesN models through simulation and parameter estimation (in Section 5.1). Next, we investigate the performance of HawkesN models on three large Twitter cascade datasets in terms of goodness-of-fit, holdout log-likelihood and final diffusion size prediction (in Section 5.2).

Models and fitting. We use the following abbreviations when presenting our results: *EXP*, *PL* and *QEXP*, stand for Hawkes models with the exponential [12], power-law [26] and Tsallis Q-Exponential kernel functions, respectively; *EXPN*, *PLN* and *QEXPN*, refer to HawkesN models with corresponding kernel functions; *SI* is the stochastic Susceptible-Infected model as an epidemic model benchmark for comparison. The estimation of Hawkes models is performed as described by Mishra et al. [26], i.e., the model parameters are fitted on an initial training part of a cascade through maximizing the log-likelihood functions. The log-likelihood functions can be found in the online supplement [35, appendix B] for HawkesN, and in [26] for Hawkes. The parameter learning and simulation of the stochastic SI model can be adopted from the stochastic SIR model with $\gamma = 0$.

5.1 Fitting SIR parameters with latent recoveries

In many applications, including in epidemiology, the recovery events are unobserved. It is therefore desirable to be able to fit the SIR model using infections events only. In this section we show how to achieve this, and we empirically validate the connection

Table 2: Statistics of the three social media datasets.

	#cascades	#tweets	Min.	Mean	Median
<i>ActiveRT</i>	39,970	7,873,733	20	197	41
<i>Seismic</i>	166,076	34,784,488	50	209	111
<i>NEWS</i>	20,093	3,252,549	50	162	90

shown in Section 3.2 by simulating stochastic SIR and retrieving SIR parameters with the HawkesN log-likelihood functions with corresponding kernel functions. We construct a rich set of parameters for stochastic SIR with the exponential (Fig. 2a) and power-law (Fig. 2b) recovery time distributions. For each parameter set shown in Fig. 2 (each grid cell), we simulate 1000 stochastic SIR realizations (using Algorithm 1). We hide the recovery events \mathcal{H}_t^R of these realizations and we fit HawkesN processes on infection event times \mathcal{H}_t^C . We jointly fit 100 realizations at a time by summing their log-likelihood functions.

In each grid cell in Fig. 2, the colors shows the fitting bias — i.e., the *absolute error* between simulation parameters and the median of fitted parameters. Note that, for ease of comparison, we have transformed the fitted HawkesN parameters into SIR parameters (using Eqs. (8) and (9), and Table 1). Also we notice in experiments that the power-law kernel as defined in [26] is over-determined, and we fix $c = 2$ both in simulation and in fitting.

Visibly, the bias of β is relatively small due to its direct presence in the infection intensity function $\lambda^I(t)$. For the other parameters, their bias starts relatively high for low values of the branching factor (upper-left corners in Fig. 2, shown as contour lines) and gradually diminishes as the branching factor grows. When the branching factor is large (bottom-right corners in Fig. 2), the fitted parameters match closely with the simulation parameters. Processes with large branching factors are commonly of interest (e.g., $R_0 = 18$ for measles in epidemiology [4]). For this reason, this evaluation supports the application of HawkesN log-likelihood functions to retrieve SIR parameters when recovery event times are missing and high branching factors are observed.

5.2 Modeling diffusions on Twitter

Datasets. We use three publicly available Twitter datasets containing retweet cascade — individual sequences of retweet events




Figure 3: Comparing model goodness-of-fit using KS test values. (a) The most distinct model pair, EXPN and PL, on *NEWS* where colors of bins represent proportions of cascades that are better fitted by the model. **Yellow** means there is no cascade in the bins. **(b)** EXPN compares to PL on *ActiveRT* in terms of the percentage of bots involved in observed retweet cascades. The upper-panel histogram counts the number of cascades at different bot percentages; the lower-panel plot depicts proportions of cascades better fitted by EXPN or PL at given bot percentages. Two high bot-percentage cascade examples (cascade A and cascade B) better fitted by EXPN and PL, respectively, are shown on the right-hand side.

Table 3: Goodness-of-fit assessments on three datasets. Models are fitted on initial 40% of each cascade event history with marks. The numbers in each cell indicate the percentages of cascades for which each model passes the nominated statistical tests (in Section 4) at the 0.01 significance level. Darker colors signify a larger fraction of cascades passing.

	Test	EXP	EXPN	PL	PLN	QEXP	QEXPN	SI
<i>ActiveRT</i>	KS	90.6%	90.1%	92.3%	84.0%	90.5%	90.7%	86.7%
	ED	91.0%	90.3%	94.8%	86.1%	93.9%	94.3%	82.3%
	LB	96.7%	96.8%	97.9%	97.6%	97.2%	97.3%	95.6%
<i>Seismic</i>	KS	81.0%	81.3%	84.6%	79.4%	79.1%	80.8%	73.3%
	ED	86.0%	85.4%	95.0%	90.8%	97.3%	96.7%	73.0%
	LB	95.3%	95.2%	98.7%	98.6%	97.1%	97.5%	89.9%
<i>NEWS</i>	KS	94.4%	94.7%	94.5%	92.3%	92.1%	93.4%	89.3%
	ED	97.1%	96.6%	99.4%	97.4%	98.9%	99.0%	81.6%
	LB	98.4%	98.1%	99.3%	99.3%	97.0%	98.3%	90.6%

following a single initial tweet. Each tweet in the cascade is considered to be an infection event, i.e., a cascade is the collection $\mathcal{H}^C = \{(t_1^I, m_1), (t_2^I, m_2), \dots\}$ where $t_i^I \in \mathcal{H}^C$ is the time stamp of the i^{th} retweet in the cascade and m_i is its associated mark information, namely the number of followers of the user. The *Seismic* dataset was constructed by Zhao et al. [40], and it contains a subset of all tweets in a month. The *NEWS* dataset was collected by Mishra et al. [26] by crawling all tweets that contain links to popular news sites, such as New York Times and CNN, for four months in 2015. The *ActiveRT* dataset¹ was collected by Rizoiu et al. [32] over 6 months in 2014, by capturing all tweets containing links to Youtube videos. Table 2 summarizes the three datasets.

Goodness-of-fit tests. We first conduct the goodness-of-fit tests described in Section 4 on all three datasets. The first 40% of event history of each cascade is used for model fitting. Table 3 shows the percentages of cascades for which each model passes the tests at

a 0.01 significance level. We note that the passing rates of the SI model is consistently worse than other models due to the model simplicity, so we focus our following experiments and discussions only on other models.

First, we see that the statistical test on the independence of transformed event times (LB test on $\{\mathcal{T}_i\}$ in Eq. (15)) presents high passing percentages ($97.57\% \pm 1.16\%$) across all models (except SI) and datasets. The other two tests (KS test and ED test) mostly agree on the performance of models with respect to each other, despite KS being a more demanding test.

When comparing Hawkes and HawkesN, we observe an increase in performance for the Tsallis Q-Exponential kernel (from QEXP to QEXPN), and a decrease from PL to PLN. EXP and EXPN, on the other hand, share similar performance. This indicates that the effect of modulating the Hawkes intensity by a finite population for modeling retweet cascades is dependent on the choice of kernel.

Model goodness-of-fit comparison. By leveraging distances produced in the KS tests, we explore the modeling performance differences for every given dataset. Given two models M_1 and M_2 that pass KS test on a cascade \mathcal{H}_i^C , we assume M_1 fits \mathcal{H}_i^C better if it has a lower KS test distance than M_2 , denoted $D_{M_1}(\mathcal{H}_i^C) < D_{M_2}(\mathcal{H}_i^C)$. Next, we tabulate the cascades in each dataset against two dimensions: cascade duration (the time of the last event) and cascade size (number of events), both in percentiles. Fig. 3a compares the two models with the higher KS passing rate (EXPN and PL), on the *NEWS* dataset (refer to [35, appendix C] for other model pairs and datasets). Grid cells depict the proportions of cascades that are better fitted by one model or the other. Visibly, EXPN fits better cascades with larger diffusion sizes and shorter diffusion durations, whereas PL performs better on less popular cascades with longer durations. This indicates that PL and EXPN are two complementary models on *NEWS*, capturing different diffusion dynamics.

Linking modeling to botness. Here, we investigate a possible factor that induces the retweet dynamics that are better captured by EXPN compared to PL: non-human participation in cascades. We choose to analyze *ActiveRT* where the user information of individual events is available. We use the Botometer API [37] to

¹The total number of cascades in *ActiveRT* is 39,970 rather than 41,411 reported in [31] after we filtered out 1,441 duplicate cascades.




Figure 4: Fig. (a)-(b) depict holdout negative log-likelihood per event of six models on *NEWS*, with and without additional mark information. Fig. (c) shows diffusion final popularity prediction performance on *NEWS*. The red diamond shows the mean value in each boxplot – lower is better.

identify Twitter bots and we collect data for 1, 174, 248 unique users involved in the first 40% event history of the 39, 549 cascades in *ActiveRT*. Due to the rate limit of the API, we only crawled cascades that have less than 2, 500 events. Given a user i , there are three possible outcomes from the API: a botness score $b_i \in [0, 1]$ of the user; when the user has a private profile; or when the user was suspended by Twitter. As this data was collected 5 years after the creation of *ActiveRT* in 2019, we assume users suspended by Twitter are bots. Eventually, we classify users who have been suspended or who have $b_i \geq 0.6$ as bots [30].

First, we group the cascades based on the proportion of bots that participate in each of them (in percentiles). For each percentile bin, Fig. 3b displays the proportions of cascades better fitted by EXPN and PL. We only keep cascades that satisfy $|D_{EXPN}(\cdot) - D_{PL}(\cdot)| \geq 0.05$ to identify the cascades significantly better fitted by each model. This condition filters out 72.96% of cascades suggesting that EXPN and PL show similar performance on most cascades. We find that, while for most of the remaining cascades PL fits better than EXPN, when more than 90% of bots involve in a cascade, around 60% of the cascades are better fitted by EXPN. In Fig. 3b, we denote A, the bot dominated cascades better fitted by EXPN, and B, those better fitted by PL, and we show one typical example from each. Cascades in A exhibit densely clustered events, with large intervals of no activity between them, whereas the cascades in B have the events more evenly spread. Intuitively, the temporal behavior of cascades in A tends to be more bot-driven, as bots retweet each other in rapid sequences and with small delays.

Generalization to unseen data. On each of the three dataset, we fit the parameters of all six models. We follow the experimental setup in [31]: 40% of the tweets in each cascade are used to fit model parameters, and we report the negative log-likelihood on the remaining 60% of the events normalized by the event count.

Fig. 4a and b show as boxplots the generalization performance on *NEWS*, without and with marks respectively. Two conclusions emerge. First, the power-law kernel for both Hawkes and HawkesN consistently outperforms other kernel functions. This emphasizes the importance of developing the generalized SIR model, as different types of parametric kernel function might fit different types of data better. Second, HawkesN outperforms Hawkes confirming results reported in [31]. The results on the other datasets depict a very similar conclusion, and they are shown in the online supplement [35, appendix C].

Popularity prediction. We predict final retweet cascade popularity following the setup described in [26]. We observe each cascade for one hour and we fit model parameters; we predict final diffusion sizes (popularity) and test against the observed final cascade size. We measure performance using the Absolute Relative Error (ARE):

$$ARE = \frac{|\hat{N}_\infty - N_\infty|}{N_\infty}$$

where \hat{N}_∞ and N_∞ are the predicted size and the true size, respectively. We compare HawkesN models to the Hawkes models (EXP, PL), and to SEISMIC [40]. We use the GBM package in R [11] to train the σ predictor described in Section 4. Furthermore, we adopt the observation that EXPN and PL are two complementary models on *NEWS* to introduce a combined model by averaging EXPN and PL prediction outcomes [42]. Results are reported with 10-fold cross validation where 6 folds are used for testing after trained on 4 folds during each iteration.

Fig. 4c shows the prediction performances on the *NEWS*. The performances on *Seismic* (also employed by Mishra et al. [26], where it is called *TWEET-IMO*) and on *ActiveRT* are shown in the online supplement [35, appendix C]. Among all the Hawkes and HawkesN models, PL delivers the best prediction performance, and EXPN predicts better than EXP. These observations align with analyses in the previous sections. Overall, the combined model, EXPN+PL, consistently outperforms all other models, on all datasets. It provides a choice to deal with complementary modeling power of kernel functions on different cascades. This only reinforces the conclusion that there may exist more than one cascade dynamics, and that each model captures the best one of them.

6 DISCUSSION AND CONCLUSION

In this work, we introduce a connection between generalized stochastic SIR models and self-exciting point processes in a finite population. The connection stems from the relationship between the recovery time distributions in SIR and the kernel functions in HawkesN processes. In addition, we develop algorithms for simulation, parameter estimation and evaluation for the stochastic SIR processes and the corresponding HawkesN processes. The modeling insights and the computational tools describe a rich set of self-exciting kernel functions, and they are more general than traditional stochastic SIR with piece-wise constant rates. In fact, it describes SIR with arbitrary recovery time distributions, and

monotonically decreasing Hawkes kernels. We compare models with three kernel functions – an exponential, a power-law and a Tsallis Q-Exponential – on three large Twitter retweet cascade datasets. We observe differences in model performance in terms of goodness-of-fit tests. Final popularity prediction is improved by combining two complementary models.

Limitations and future work Non-monotonically decreasing kernel functions, such as the Rayleigh function, have been used in the point process literature [8, 10, 26]. Although it cannot be linked to the CCDF of recovery events in epidemics, the intuition of the Rayleigh function stems from the concept of the disease incubation period in epidemiology. We plan to broaden the connection, e.g., between HawkesN and variants in the epidemic models family.

In general, this newly established bridge between distinct classes of stochastic point processes opens up many research topics such as using modern machine learning tools to design objectives and estimation procedures, causal inference in epidemic models, and novel applications of either model in new data domains.

SOFTWARE AND RUNTIME INFORMATION

The simulation and fitting algorithms described in Section 3 have been implemented as an R package available on <http://bit.ly/34qiDTK> and the code for reproducing experiments in Section 5 can be found on <https://bit.ly/3697ojk>. Fitting HawkesN models on a 500 event cascade takes 4.8 minutes in average with 10 parallel random parameter initializations on 2.2GHz cpus where the likelihood function evaluation complexity is quadratic in the number of events.

Acknowledgments

This research is supported by Asian Office of Aerospace Research and Development (AOARD) Grant 19IOA08, Australian Research Council Discovery Project DP180101985 and the Data61, CSIRO PhD scholarship. We thank our reviewers, Swapnil Mishra and other lab members at ANU CMLab for their helpful comments. We also thank the National Computational Infrastructure (NCI) for providing computational resources, supported by the Australian Government.

REFERENCES

- [1] Linda J. S. Allen. 2008. An Introduction to Stochastic Epidemic Models. In *Mathematical Epidemiology*. Springer, Chapter 3.
- [2] Emmanuel Bacry, Iacopo Mastromatteo, and Jean-François Muzy. 2015. Hawkes processes in finance. *Market Microstructure and Liquidity* (2015).
- [3] MS Bartlett. 1949. Some evolutionary stochastic processes. *Journal of the Royal Statistical Society. Series B (Methodological)* (1949).
- [4] Fred Brauer. 2008. An Introduction to Stochastic Epidemic Models. In *Mathematical Epidemiology*. Springer, Chapter 2.
- [5] Feng Chen and Wai Hong Tan. 2018. Marked self-exciting point process modelling of information diffusion on Twitter. *The Annals of Applied Statistics* (2018).
- [6] D. R. Cox and D Oakes. 1984. *Analysis of survival data*. Routledge.
- [7] Daryl J Daley and David Vere-Jones. 2008. Conditional Intensities and Likelihoods. In *An introduction to the theory of point processes*. Vol. I. Springer, Chapter 7.2.
- [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 *CIKM*.
- [9] Nan Du, Hanjun Dai, Rakshit Trivedi, Utkarsh Upadhyay, Manuel Gomez-Rodriguez, and Le Song. 2016. Recurrent marked temporal point processes: Embedding event history to vector. In *KDD*. ACM.
- [10] Manuel Gomez-Rodriguez, David Balduzzi, and Bernhard Schölkopf. 2011. Uncovering the temporal dynamics of diffusion networks. (2011).
- [11] Brandon Greenwell, Bradley Boehmke, Jay Cunningham, and GBM Developers. 2019. *gbm: Generalized Boosted Regression Models*. R package version 2.1.5.
- [12] Alan G Hawkes. 1971. Spectra of some self-exciting and mutually exciting point processes. *Biometrika* (1971).
- [13] Fang Jin, Edward Dougherty, Parang Saraf, Yang Cao, and Naren Ramakrishnan. 2013. Epidemiological modeling of news and rumors on twitter. In *SNA-KDD Workshop*. ACM.
- [14] How Jing and Alexander J Smola. 2017. Neural survival recommender. In *WSDM*.
- [15] Don H Johnson. 1996. Point process models of single-neuron discharges. *Journal of computational neuroscience* (1996).
- [16] Matthew J Keeling and BT Grenfell. 1997. Disease extinction and community size: modeling the persistence of measles. *Science* (1997).
- [17] William Ogilvy Kermack and Anderson G McKendrick. 1927. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* (1927).
- [18] Masahiro Kimura, Kazumi Saito, and Hiroshi Motoda. 2009. Efficient Estimation of Influence Functions for SIS Model on Social Networks.. In *IJCAI*.
- [19] Mehdi Lallouache and Damien Challet. 2016. The limits of statistical significance of Hawkes processes fitted to financial data. *Quantitative Finance* (2016).
- [20] Patrick J Laub, Thomas Taimre, and Philip K Pollett. 2015. Hawkes processes. *arXiv preprint arXiv:1507.02822* (2015).
- [21] Liangda Li, Hongbo Deng, Jianhui Chen, and Yi Chang. 2017. Learning parametric models for context-aware query auto-completion via hawkes processes. In *WSDM*.
- [22] Rafael Lima and Jaesik Choi. 2018. Hawkes Process Kernel Structure Parametric Search with Renormalization Factors. *arXiv preprint arXiv:1805.09570* (2018).
- [23] Travis Martin, Jake M Hofman, Amit Sharma, Ashton Anderson, and Duncan J Watts. 2016. Exploring limits to prediction in complex social systems. In *WWW*.
- [24] Frank J Massey Jr. 1951. The Kolmogorov-Smirnov test for goodness of fit. *Journal of the American statistical Association* (1951).
- [25] Hongyuan Mei and Jason M Eisner. 2017. The neural hawkes process: A neurally self-modulating multivariate point process. In *NeurIPS*.
- [26] Swapnil Mishra, Marian-Andrei Rizoïu, and Lexing Xie. 2016. Feature Driven and Point Process Approaches for Popularity Prediction. In *CIKM*.
- [27] Swapnil Mishra, Marian-Andrei Rizoïu, and Lexing Xie. 2018. Modeling Popularity in Asynchronous Social Media Streams with Recurrent Neural Networks. In *ICWSM*.
- [28] Mark Newman. 2018. Epidemics on networks. In *Networks*. Oxford university press, Chapter 17.
- [29] Yoshiko Ogata. 1978. The asymptotic behaviour of maximum likelihood estimators for stationary point processes. *Annals of the Institute of Statistical Mathematics* (1978).
- [30] Marian-Andrei Rizoïu, Timothy Graham, Rui Zhang, Yifei Zhang, Robert Ackland, and Lexing Xie. 2018. # DebateNight: The Role and Influence of Socialbots on Twitter During the 1st 2016 US Presidential Debate. In *ICWSM*.
- [31] Marian-Andrei Rizoïu, Swapnil Mishra, Quyu Kong, Mark Carman, and Lexing Xie. 2018. SIR-Hawkes: on the Relationship Between Epidemic Models and Hawkes Point Processes. In *WWW*.
- [32] Marian-Andrei Rizoïu, 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 *WWW*.
- [33] Isobel Routledge, José Eduardo Romero Chevéz, Zulma M. Cucunubá, Manuel Gomez Rodriguez, Caterina Guinovart, Kyle B. Gustafson, Kammerle Schneider, Patrick G.T. Walker, Azra C. Ghani, and Samir Bhatt. 2018. Estimating spatiotemporally varying malaria reproduction numbers in a near elimination setting. *Nature Communications* (2018).
- [34] George Streftaris and Gavin J Gibson. 2012. Non-exponential tolerance to infection in epidemic systems—modeling, inference, and assessment. *Biostatistics* (2012).
- [35] Online supplement. 2019. Appendix: Modeling Information Cascades with Self-exciting Processes via Generalized Epidemic Models. <http://bit.ly/35APWE0>.
- [36] Ping Yan. 2008. Distribution Theory, Stochastic Processes and Infectious Disease Modelling. In *Mathematical Epidemiology*. Springer, Chapter 10.
- [37] Kai-Cheng Yang, Onur Varol, Clayton A Davis, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. 2019. Arming the public with artificial intelligence to counter social bots. *Human Behavior and Emerging Technologies* (2019).
- [38] Ali Zarezade, Utkarsh Upadhyay, Hamid R Rabiee, and Manuel Gomez-Rodriguez. 2017. Redqueen: An online algorithm for smart broadcasting in social networks. In *WSDM*.
- [39] Laijun Zhao, Hongxin Cui, Xiaoyan Qiu, Xiaoli Wang, and Jiajia Wang. 2013. SIR rumor spreading model in the new media age. *Physica A: Statistical Mechanics and its Applications* (2013).
- [40] 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 *KDD*.
- [41] Ke Zhou, Hongyuan Zha, and Le Song. 2013. Learning triggering kernels for multi-dimensional hawkes processes. In *ICML*.
- [42] Zhi-Hua Zhou. 2012. Combination Methods. In *Ensemble methods: foundations and algorithms*. Chapman and Hall/CRC, Chapter 4.

A LINKING SIR TO HAWKES

A.1 Detailed derivation of the inequality between generalized stochastic SIR and HawkesN

We use $\mathbb{E}_{\{t_i^R | \mathcal{H}_t^C\}} [\lambda^I(t)]$ to denote the expected infection intensity over all recovery event times of infected individuals up to time t :

$$\begin{aligned}
 & \mathbb{E}_{\{t_i^R | \mathcal{H}_t^C\}} [\lambda^I(t)] \\
 & \stackrel{(a)}{=} \beta \frac{S_t}{N} \mathbb{E}_{\{t_i^R | \mathcal{H}_t^C\}} [I_t] \\
 & \stackrel{(b)}{=} \beta \frac{S_t}{N} \mathbb{E}_{\{t_i^R | \mathcal{H}_t^C\}} \left[\sum_{t_i^I \in \mathcal{H}_t^C} \mathbb{1}(t_i^R - t_i^I > t - t_i^I) \right] \\
 & \stackrel{(c)}{=} \beta \frac{S_t}{N} \sum_{t_i^I \in \mathcal{H}_t^C} \mathbb{E}_{\{t_i^R | \mathcal{H}_t^C\}} [\mathbb{1}(t_i^R - t_i^I > t - t_i^I)] \\
 & \stackrel{(d)}{=} \beta \frac{S_t}{N} \sum_{t_i^I \in \mathcal{H}_t^C} \int_0^\infty \mathbb{1}(\tau_i > t - t_i^I) f(\tau_i | \mathcal{H}_t^C) d\tau_i \\
 & \stackrel{(e)}{=} \beta \frac{S_t}{N} \sum_{t_i^I \in \mathcal{H}_t^C} \int_{t-t_i^I}^\infty f(\tau_i | \mathcal{H}_t^C) d\tau_i \\
 & \stackrel{(f)}{\geq} \beta \frac{S_t}{N} \sum_{t_i^I \in \mathcal{H}_t^C} \int_{t-t_i^I}^\infty f(\tau_i) d\tau_i
 \end{aligned} \tag{17}$$

Here Eq. (17a) is due to the independence between S_t and t_i^R given \mathcal{H}_t^C . Eq. (17b) follows from decomposing the step-wise stochastic process I_t — that the recovery time of each individual therein is greater than time t , i.e., $t_i^R > t$. By definition $t_i^R > t_i^I$, we can subtract infection times t_i^I on both sides and preserve the sign of the inequality — leading to $t_i^R - t_i^I > t - t_i^I$, which is easier to model since the left hand side correspond to the recovery time of the i -th infection. $\mathbb{1}(x)$ is an indicator function that takes value 1 if the proposition x is true, 0 otherwise. Eq. (17c) pushes the expectation into the summation due to known infection events. Eq. (17d) expands the expectation for each recovery time, and uses $\tau_i = t_i^R - t_i^I \sim f(\tau_i | \mathcal{H}_t^C)$ where the i^{th} individual's recovery time distribution is conditional. Eq. (17e) uses the definition of the indicator function to change the lower bound of integration.

To show the inequality in Eq. (17f), we reduce it down to proving

$$\int_{t-t_i^I}^\infty f(\tau_i | \mathcal{H}_t^C) d\tau_i \geq \int_{t-t_i^I}^\infty f(\tau_i) d\tau_i \implies \int_0^{t-t_i^I} f(\tau_i | \mathcal{H}_t^C) d\tau_i \leq \int_0^{t-t_i^I} f(\tau_i) d\tau_i \tag{18}$$

which is equivalent to $\mathbb{P}[t_i^R < t | \mathcal{H}_t^C] \leq \mathbb{P}[t_i^R < t]$. To proof this, we reason it from the perspective of a branching structure, namely any future event is a descendent event triggered by past event. We denote the infection events triggered by i^{th} individual as $\mathcal{H}_i^C = \{t_j^I | t_j^I \in \mathcal{H}_t^C, t_i^I < t_j^I < t\}$ and use $t_m^I = \max\{\mathcal{H}_i^C\}$. Then we can see that $\mathbb{P}[t_i^R < t | \mathcal{H}_t^C] = \mathbb{P}[t_i^R < t | \mathcal{H}_i^C]$. Two possible cases emerge

- $\mathcal{H}_i^C = \emptyset$: we have $\mathbb{P}[t_i^R < t | \mathcal{H}_i^C] = \mathbb{P}[t_i^R < t]$ where the equality holds.
- $|\mathcal{H}_i^C| > 0$: we can also reduce the dependency with $\mathbb{P}[t_i^R < t | \mathcal{H}_i^C] = \mathbb{P}[t_m^I < t_i^R < t]$. We then compare $\mathbb{P}[t_m^I < t_i^R < t] = \int_{t_m^I - t_i^I}^{t-t_i^I} f(\tau) d\tau$ against $\mathbb{P}[t_i^R < t] = \int_0^{t-t_i^I} f(\tau) d\tau$ given recovery time distribution $f(t)$. As $f(t) \geq 0, \forall t \in \mathbb{R}$, we conclude $\int_0^{t-t_i^I} f(\tau) d\tau > \int_{t_m^I - t_i^I}^{t-t_i^I} f(\tau) d\tau$

Overall, that proofs $\int_{t-t_i^I}^\infty f(\tau_i | \mathcal{H}_t^C) d\tau_i \geq \int_{t-t_i^I}^\infty f(\tau_i) d\tau_i$.


A.2 Empirical analysis of the intensity difference

In this section, we explore the difference between the expected infection intensity values of generalized stochastic SIR and the corresponding HawkesN intensity. Fig. 5a presents an example of this intensity difference given a specific parameter set. For a given stochastic SIR process, we approximate its expected infection intensity by fixing the infection events and simulating recovery events via standard rejection-sampling technique [7]. HawkesN intensity values, on the other hand, can be computed from Eq. (3). Fig. 5b explores the relative intensity difference between the two at various parameter combinations.


A.3 Branching factor

For the classic SIR, the reproduction number R_0 is defined as [28]

$$R_0 = \int_0^\infty \beta \tau \gamma e^{-\gamma \tau} d\tau \tag{19}$$



(a) Intensity difference between the expected infection intensity and HawkesN intensity (an exponential recovery time distribution: $N = 50$, $\beta = 1$, $\gamma = 0.5$). The expected infection intensity is approximated with 2000 simulations.



(b) Given the infection events in (a), the relative difference between the expected infection intensity values and the HawkesN intensity values given various parameters.

Figure 5

where $\beta\tau$ is the expected number of individuals contacted by an infected individual and integrating it with the recovery time distribution leads to R_0 . We can then express it with a general recovery time distribution as $R_0 = \int_0^\infty \beta\tau f(\tau)d\tau$. We show its equivalence to Eq. (13) as following

$$\begin{aligned}
R_0 &= \int_0^\infty \beta\tau f(\tau)d\tau \\
&= \beta \left[\tau \int_0^\tau f(\eta)d\eta \right]_0^\infty - \beta \int_0^\infty \int_0^\tau f(\eta)d\eta d\tau && \text{(Integration by parts)} \\
&= \beta \left[\tau \int_0^\tau f(\eta)d\eta \right]_0^\infty - \beta \int_0^\infty (1 - \int_\tau^\infty f(\eta)d\eta)d\tau && \text{(due to } \int_0^\infty f(\eta)d\eta = 1) \\
&= \beta \left[\tau \int_0^\tau f(\eta)d\eta \right]_0^\infty - \beta [\tau]_0^\infty + \beta \int_0^\infty \int_\tau^\infty f(\eta)d\eta d\tau \\
&= \beta \left[\tau \int_0^\tau f(\eta)d\eta \right]_{\tau=\infty} - \beta [\tau]_{\tau=\infty} + \beta \int_0^\infty \int_\tau^\infty f(\eta)d\eta d\tau && \text{(due to } \int_0^0 f(\eta)d\eta = 0) \\
&= \beta \int_0^\infty \int_\tau^\infty f(\eta)d\eta d\tau && \text{(due to } \int_0^\infty f(\eta)d\eta = 1)
\end{aligned}$$

B LIKELIHOOD AND PARAMETER ESTIMATION

We conduct maximum likelihood estimation for parameter inference. For total population size N , we adopt the practice from both Jin et al. [13] and Rizoio et al. [31]: fitting N as an unknown parameter. Let Θ^E denote the set of all parameters in the stochastic SIR models, e.g., $\Theta^E = \{\beta, \gamma, N\}$ for stochastic SIR described in Section 2. To estimate Θ^E from a given stochastic SIR process until time t ($\mathcal{H}_t^C, \mathcal{H}_t^R$) with a recovery time distribution $f(t)$, the likelihood function of Θ^E can be expressed based on the log-likelihood estimator for point processes [7] as

$$\mathcal{L}(\Theta^E; \mathcal{H}_t^C, \mathcal{H}_t^R) = \sum_{t_i^I \in \mathcal{H}_t^C} \log \lambda^I(t_i^I) - \int_0^t \lambda^I(\tau)d\tau + \sum_{i \in U(\mathcal{H}_t^R)} \log f(t_i^R - t_i^I) \quad (20)$$

The first two terms of RHS of Eq. (20) comes from

$$\mathcal{L}(\Theta^E; \mathcal{H}_t^C, \mathcal{H}_t^R) = \log \prod_{t_i^I \in \mathcal{H}_t^C} \lambda^I(t_i^I) e^{-\int_0^{t_i^I} \lambda^I(u) du} = \log \prod_{t_i^I \in \mathcal{H}_t^C} \lambda^I(t_i^I) e^{-\int_{t_{i-1}^I}^{t_i^I} \lambda^I(u) du} \quad (21)$$

$$= \log e^{-\int_0^{\max(\mathcal{H}_t^C)} \lambda^I(u) du} \prod_{t_i^I \in \mathcal{H}_t^C} \lambda^I(t_i^I) \quad (22)$$

$$= -\int_0^t \lambda^I(\tau) d\tau + \sum_{t_i^I \in \mathcal{H}_t^C} \log \lambda^I(t_i^I) \quad (23)$$

When recovery events are not observed, i.e., only \mathcal{H}^C is presented, we take expectation over the recovery event history \mathcal{H}^R on Eq. (20): HawkesN log-likelihood functions after the solving integral part in Eq. (21) with different kernel functions are listed as following:

- Exponential

$$\mathcal{L}_{EXP}(\Theta^H; \mathcal{H}^C) = \sum_{j=2}^n \log \left(\lambda^H(t_j^-) \right) - \kappa \sum_{j=1}^{n-1} \sum_{l=j}^{n-1} \frac{N-l}{N} \left[e^{-\theta(t_l - t_j)} - e^{-\theta(t_{l+1} - t_j)} \right]$$

- Power-law

$$\begin{aligned} \mathcal{L}_{PL}(\Theta^H; \mathcal{H}^C) &= \sum_{t_i \in \mathcal{H}^C} \log \frac{N-i}{N} \kappa \sum_{t_j \in \mathcal{H}^C, t_j < t_i} (t_i - t_j + c)^{-(1+\theta)} \\ &\quad - \kappa \theta \sum_{t_i \in \mathcal{H}^C, t_i < t_n} \sum_{t_j \in \mathcal{H}^C, t_i \leq t_j < t_n} \frac{N-j}{N} \left[(t_j - t_i + c)^{-\theta} - (t_{j+1} - t_i + c)^{-\theta} \right] \end{aligned}$$

- Tsallis Q-EXP

$$\begin{aligned} \mathcal{L}_{Q-EXP}(\Theta^H; \mathcal{H}^C) &= \sum_{t_i \in \mathcal{H}^C} \log \frac{N-i}{N} \kappa \sum_{t_j \in \mathcal{H}^C, t_j < t_i} \left[1 + (\theta - 1)(t_i - t_j) \right]^{\frac{1}{1-\theta}} \\ &\quad - \frac{\kappa}{2-\theta} \sum_{t_i \in \mathcal{H}^C, t_i < t_n} \sum_{t_j \in \mathcal{H}^C, t_i \leq t_j < t_n} \frac{N-j}{N} \left[\left[1 + (\theta - 1)(t_j - t_i) \right]^{\frac{2-\theta}{1-\theta}} - \left[1 + (\theta - 1)(t_{j+1} - t_i) \right]^{\frac{2-\theta}{1-\theta}} \right] \end{aligned}$$

Some natural constraints are applied on $N \geq C_t$ and on other parameters as in Table 1. Eqs. (20) (21) are non-linear functions and we use an optimization tool AMPL [44] bridged with a non-linear solver Ipopt [49] for maximizing them and estimating model parameters.

After obtaining Θ^H , Eq. (9) leads us to corresponding stochastic SIR parameters Θ^E . Similarly, Eq. (8) links Θ^E inferred by Eq. (20) to Θ^H . This helps one reveal the underlying recovery processes when missing recovery event data or concentrate on infection process yet leveraging both infection and recovery events in data.

C EXTRA EXPERIMENT RESULTS

Fig. 6 presents the comparison of a most distinct model pair on each dataset. Fig. 7 shows the holdout log-likelihood values and popularity prediction performance of models fitted with event marks on all three retweet cascade datasets.

D COMPARE SIZE DISTRIBUTIONS OF STOCHASTIC SIR AND HAWKESN

In this section, we empirically show the difference of stochastic SIR and HawkesN in terms of their process final size distributions given same set of parameters. We approximate their final size distributions (empirical cumulative density functions) via simulation with 1000 simulated processes for each model and parameter set. Fig. 8 shows the results given different sets of parameters. From this plot, we notice that stochastic SIR processes consistently present higher probabilities of smaller cascade sizes and HawkesN models tend to generate larger size cascades for high branching factors.

E LIKELIHOOD OF CASCADE SIZES.

We study the probability distribution of final size for Twitter cascades. Rizoïu et al. [31] have proposed a Markov chain method to estimate the distribution of final cascade size, based on SIR's memory-less property. However, stochastic SIR with non-exponentially distributed recovery times do not have this property. To overcome this problem, we employ a simulation-based computation of the size distribution. Given a parameter set Θ^E — the parameters of HawkesN fitted on k events — we approximate the size distributions $\mathbb{P} \left[|\mathcal{H}^C| = j \mid \Theta^E, j \geq k \right]$ by converting the HawkesN parameters to stochastic SIR parameters, and applying Algorithm 1 to simulate 5000 realizations for each cascade. We construct the empirical size distribution by aggregating the sizes of the realizations and smoothing the obtained distribution. Given n the observed final size of a cascade, its likelihood under the constructed distribution is $\mathbb{P} \left[|\mathcal{H}^C| = n \mid \Theta^E \right]$.

We employ the above methodology to compute the likelihood of the observed final size for three samples — one for each dataset — each sample containing 1000 cascades. For every cascade we construct two size distributions, using HawkesN with an exponential and a power-law kernel respectively. Figure 9 aggregates the computed likelihoods, per dataset and per HawkesN kernel type. Each boxplot contains 1000 datapoints. We observe that for *ActiveRT* and for *Seismic*, the observed final size is more likely under the distribution constructed using power-law kernel than under the exponential kernel. This is likely due to power-law kernel being long-tailed, and able to explain the minority of very large cascades occurring naturally [45]. For *NEWS*, the similar performances of the two kernels are likely due to news being time-sensitive content, and on average having smaller cascades.

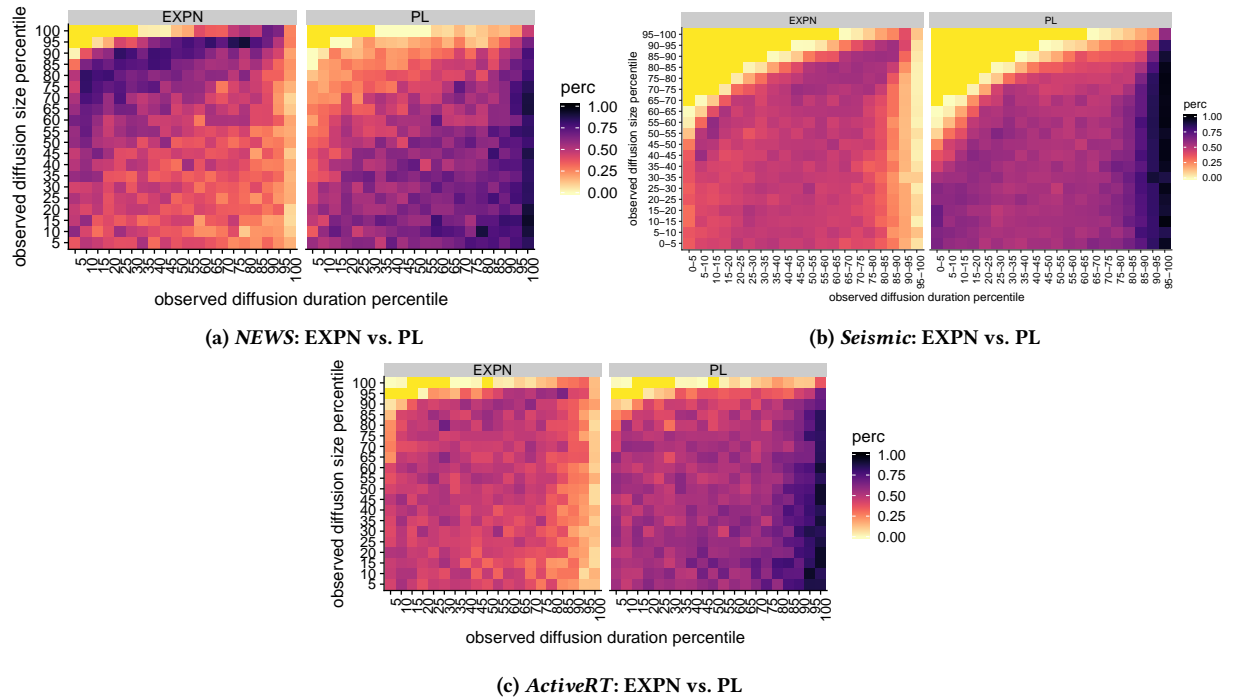


Figure 6

APPENDIX REFERENCES

- [7] Daryl J Daley and David Vere-Jones. 2008. Conditional Intensities and Likelihoods. In *An introduction to the theory of point processes*. Vol. I. Springer, Chapter 7.2.
- [44] Robert Fourer, David M Gay, and Brian W Kernighan. 1987. *AMPL: A mathematical programming language*.
- [45] Sharad Goel, Ashton Anderson, Jake Hofman, and Duncan J Watts. 2015. The structural virality of online diffusion. *Management Science* (2015).
- [13] Fang Jin, Edward Dougherty, Parang Saraf, Yang Cao, and Naren Ramakrishnan. 2013. Epidemiological modeling of news and rumors on twitter. In *SNA-KDD Workshop*. ACM.
- [28] Mark Newman. 2018. Epidemics on networks. In *Networks*. Oxford university press, Chapter 17.
- [31] Marian-Andrei Rizoiu, Swapnil Mishra, Quyu Kong, Mark Carman, and Lexing Xie. 2018. SIR-Hawkes: on the Relationship Between Epidemic Models and Hawkes Point Processes. In *WWW*.
- [49] 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* (2006).






Figure 7: Fig. (a)-(f) are holdout negative log-likelihood of models on three datasets. Fig. (g)-(i) are popularity prediction performance on three datasets



(a) EXPN vs stochastic SIR with an exponential recovery distribution



(b) PLN vs stochastic SIR with a power-law recovery distribution

Figure 8: Size distribution comparison between HawkesN models and stochastic SIR processes. The x axis and y axis are different parameter values (β, γ for (a) and c, θ for (b), respectively). For each small plot, the x axes are the cascade sizes while the y axes are the corresponding empirical cumulative densities.




Figure 9: The likelihood of observed cascade sizes on *ActiveRT*, *NEWS* and *Seismic*, using EXPN and PLN. The parameters of each model are fitted using 40% of the events in each cascade. The distributions of cascade size are approximated using 5000 simulations for each set of parameters — higher is better.