Feature Driven and Point Process Approaches for Popularity Prediction

Swapnil Mishra

Marian-Andrei RizoIU

Lexing Xie

The Australian National University, Data 61, CSIRO, Australia
{swapnil.mishra,marian-andrei.rizoIU,lexing.xie}@anu.edu.au

ABSTRACT

Predicting popularity, or the total volume of information outbreaks, is an important subproblem for understanding collective behavior in networks. Each of the two main types of recent approaches to the problem, feature-driven and generative models, have desired qualities and clear limitations. This paper bridges the gap between these solutions with a new hybrid approach and a new performance benchmark. We model each social cascade with a marked Hawkes self-exciting point process, and estimate the content virality, memory decay, and user influence. We then learn a predictive layer for popularity prediction using a collection of cascade history. To our surprise, Hawkes process with a predictive overlay outperform recent feature-driven and generative approaches on existing tweet data [44] and a new public benchmark on news tweets. We also found that a basic set of user features and event time summary statistics performs competitively in both classification and regression tasks, and that adding point process information to the feature set further improves predictions. From these observations, we argue that future work on popularity prediction should compare across feature-driven and generative modeling approaches in both classification and regression tasks.

Keywords social media; self-exciting point process; information diffusion; cascade prediction.

1. INTRODUCTION

Popularity prediction, especially in the context of online media, has seen increased attention, as successful solutions would help both content consumers to cope with information overload, and content producers to disseminate information within limited resources. Accurate models will help producers to identify the trends, discover useful content and decimate it faster to improve the content delivery of the network. Furthermore, this would provide with more insights into understanding collective behaviour by compiling

seemingly incoherent individual responses to a single piece of information.

When presented with the task of predicting popularity, a clear gap appears between the two main classes of approaches: the feature-driven approaches and the generative approaches. *Feature-driven approaches* [3, 7, 25, 29] summarize network, user, and event history information into an extensive set of features, and they use machine learning approaches to predict future popularity. *Generative methods* [12, 34, 43, 44] leverage fine-grained timing information in the event series, however they make stronger assumptions about the diffusion process. A thorough understanding of strengths and weaknesses of the two classes is currently missing, especially since the generative methods are usually designed for explaining the mechanisms that generate attention, and not optimized for prediction. Another gap arises between the different problem settings. Prior work has addressed problems like predicting the volume of attention [34, 44], predicting if a cascade will double in size [7], whether an item will have 10 million views [33], or be among the top 5% most popular [42]. However, the understanding of how the different approaches and the different features generalize over the various problem settings is incomplete. In addition, a practical bottleneck is the availability of certain types of these features, such as network data and corresponding features – for example the number of users exposed in the diffusion up to a time t or potential number of reachable users. A third gap in the current available work is the consensus and understanding about which features are informative for feature-driven models and the lack of benchmark datasets on which new methods can be devised.

In this paper, we address the above challenges in the context of predicting the final size of retweet cascades. We build two predictive approaches, one generative and one feature-driven, and we show how to combine them into a hybrid predictor to further increase performances. The generative method is a two-layered approach, built on the intuitive *Hawkes self-exciting process* [15] model. Self-exciting point processes are a class of stochastic processes, in which the occurrence of past events makes the occurrence of future events more probable. Three key factors in information diffusion are built into the proposed model: the social influence of users, the length of “social memory” and the inherent tweet quality. We use a predictive layer on top of the generative model to make final predictions. This helps us to take into account other cascades and mitigate limitations of model assumptions and parameter estimation. Our second proposed approach, the feature-based method, uses only

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.

CIKM'16, October 24-28, 2016, Indianapolis, IN, USA

© 2016 ACM. ISBN 978-1-4503-4073-1/16/10...\$15.00

DOI: <http://dx.doi.org/10.1145/2983323.2983812>

Figure 1: Two retweet cascades announcing the death of “Mr. Spock”. The underlying diffusion processes of tw_1 and tw_2 are very different: tw_1 lasts for 2.5 hours. It attracts the attention of highly influential users, especially music related channels, however the generative kernel of tw_2 amplifies more than that of tw_1 . tw_2 diffuses faster and ends in just 12 minutes, but the final sizes for the two cascades are similar (224 retweets for tw_1 and 219 for tw_2). We model and interpret these cascades using our generative model in Sec. 3.4. x-axis: time in seconds; y-axis: number of followers each tweet can reach; $\phi_m(\tau)$: memory kernel of Eq. (2) plotted with $m = 1000$.

features which can be computed on data containing solely the message content and basic user profiles, and which were highlighted as most predictive by recent studies: user related features [25] and temporal activity features [7]. This results in a competitive feature-based predictor, which consistently outperforms the current state-of-the-art popularity prediction model [44] – reducing the mean absolute relative error by more than 200% when compared to the latter, after observing the retweet series for 10 minutes. We show that the same set of features can be employed in both regression and classification tasks. We further experiment with a hybrid predictor, which uses both data features and measures issued from the generative model, and we show additional performance improvement. We conduct experiments on a large dataset, containing 49.7 million tweets containing links to top 10 news sites in English, gathered using the public API over a period of 4 months, in 2015. Fig. 1 illustrates two of the retweet cascades in this dataset, relating the passing of the actor for “Mr. Spock”. Even though the two cascades achieve similar popularities, our generative model is capable of differentiating between their very different diffusion dynamics. For example, the memory kernel $\phi_m(\tau)$ of tw_2 – i.e. the contribution of each event to the spawning of future events – amplifies more (shown by a larger area under the $\phi_m(\tau)$), but is forgotten faster (the entire diffusion ends in just 12 minutes).

There are three main contributions of this paper:

- **Bridging the popularity prediction problem space:** we address three types of gaps: problem settings and task (e.g. regression vs. classification), type of approach

(feature-based vs. generative) and the generalization of the features across feature-driven approaches.

- **Comparative understanding of feature-driven and generative models:** we propose a generative method, based on Hawkes self-exciting processes, which features the advantages of generative methods – e.g. interpretable results – and which is adapted for predicting the final cascade popularity by using an additional predictive layer; we compare approaches to popularity prediction from the two main classes, i.e. feature-based and generative, and we show that combining them into a hybrid approach increases performances.
- **Dataset:** we identify a set of features that generalize over different settings and can be constructed with only message content and basic user features; we curate a large, domain-specific dataset of retweet diffusion cascades, which can be used by both feature-driven and generative models for prediction. Dataset along with code is publicly available for download.¹ This new dataset allows comprehensive performance benchmarks on features about users and temporal activity.

The rest of this paper is structured as follows. Sec. 2 presents related work. Sec. 3 details the proposed generative model and fitting methods, while Sec. 4 presents how the popularity prediction is performed, based on the fitted model parameters. The dataset and feature construction for the feature-driven approach are detailed in Sec. 5 and the experiments are outlined in Sec. 6. We conclude in Sec. 7.

2. RELATED WORK

We structure the discussion of related work onto the two broad previously mentioned categories: *feature-driven approaches* and *generative approaches*. For generative models, being somewhat recent, we discuss their application in popularity prediction as well as their use in other social media settings.

Predicting popularity with feature-driven models.

Data-driven approaches [11, 19] treat popularity as a non-decomposable process and take a bottom-up approach. They rely on machine learning algorithms to make the connection between item popularity and an extensive set of features. Most often the main contributions of such approaches are collecting large-scale datasets and constructing informative features. Bakshy *et al.* [4] and later Martin *et al.* [25] predict tweet popularity, defined as the total size of its retweet cascade. Both studies construct features relating to the user having posted the tweet and posted content. User features, particularly past user success – the ability of a user to start large cascades – are shown more informative, however prediction performances are limited. Martin *et al.* make the hypothesis of an inherent popularity randomness, which sets limits to its predictability. The situation is different when part of the diffusion cascade is observable, as early activity has been shown predictive for total popularity [29, 37]. When predicting whether a Facebook cascade will double in size, Cheng *et al.* [7] found that temporal features – like the resharing rate or the acquired volume of views – are the most informative features. Prediction performances are further improved when adding features which account for specific contexts of the information diffusion. For example, Li *et al.* [22] add propagation network measurements,

¹Download data and code here: <https://git.io/v6rIN>.

Chang *et al.* [6] add sophisticated features relating to audience behavior when predicting the popularity of serials and Vallet *et al.* [38] construct features relating to two social networks: Youtube and Twitter. Many proposed features cannot be used outside of the proposed context, being either data-source specific or relying on information that is not publicly accessible. The feature-driven methods that we propose use only the best-performing features shown in previous studies and which can be extracted from public data.

Generative models in social media. Generative models start by modeling a set of phenomena into a parametric model and fitting the parameters of the model to observed data. Such approaches have been successfully used in neuro-biology to model neuronal spiking activity [20], in economics to model the evolution of stock prices [14], in physics to model earthquake aftershocks occurrence [35, 17, 36] and in criminology to model crime patterns [27]. In social media such approaches have been used extensively [18, 30, 31], mainly due to their capability to model phenomena like user interactions and diffusion susceptibility. Hidden network properties, such as connexions between individuals [23, 24] or the evolution of the diffusion network topology [13], have been uncovered using generative methods, such as self-exciting point processes. The main difficulty with such approaches is scalability, since each social phenomenon must be accounted manually in the model.

Modeling and predicting popularity using generative models. Popularity modeling [9, 12, 43] and prediction [34, 44] has been a particularly fertile field for point-process based generative models. In their seminal work, Crane and Sornette [9] showed how a Hawkes point-process can account for popularity bursts and decays. Afterwards, more sophisticated models have been proposed to model and simulate popularity in microblogs [43] and videos [12]. 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.* [34] employ reinforced Poisson processes, modeling three phenomena: fitness of an item, a temporal relaxation function and a reinforcement mechanism. Zhao *et al.* [44] propose SEISMIC, which employs a double stochastic process, one accounting for infectiousness and the other one for the arrival time of events. It is the current state of the art in popularity prediction and, throughout this paper, we compare against it. Our proposed generative model uses the same information as SEISMIC – tweet timestamps and the number of followers of the user posting the tweet –, but uses a simpler, more intuitive modeling of the diffusion cascade.

Generative models are typically designed to explain popularity, not predict it. Most of the methods presented earlier introduce multiple regularization and correction factors in order to make meaningful predictions. This is sub-optimal both for prediction and for interpretation. By contrast, our Hawkes predictive model separates modeling from prediction accounting directly for the most important three factors at individual cascade level: user influence, decaying social attention and content quality, followed by a predictive layer tuned from historical data.

3. RETWEETING AS A POINT PROCESS

We first introduce the Hawkes self-exciting point process, a stochastic event model for describing retweeting cascades.

We then describe a procedure for maximum-likelihood estimation of model parameters from an observed cascade. We present observations of the Hawkes model on individual and collections cascades, illustrating the interplay of individual memory kernel and user influence in diffusion dynamics.

3.1 A self-exciting point processes

In self-exciting processes the occurrences of past events make future events more probable. In particular, we model each retweet as a stochastic event, and model three key intuitions in a social network: *magnitude of influence*, tweets by users with many followers tend to get retweeted more; *memory over time*, that most retweeting happens when the content is *fresh* [41]; and *content quality*.

The marked Hawkes process [15] is a well-known self-exciting process that captures these properties. Its event rate is written as:

$$\lambda(t) = \sum_{t_i < t} \phi_{m_i}(t - t_i) , \quad (1)$$

where $\lambda(t)$ is the arrival rate of new events. It depends on each previously occurred event (m_i, t_i) of magnitude m_i at time t_i through the triggering kernel $\phi_m(\tau)$. For the problem of estimating the size of a retweet cascade, there are no exogenously generated events (as in a general Hawkes process [16]) apart from the initial tweet, which is accounted for with magnitude m_0 at time $t_0 = 0$. The effects of m_i and t_i are separable in our triggering kernel:

$$\phi_m(\tau) = \kappa m^\beta (\tau + c)^{-(1+\theta)} . \quad (2)$$

κ describes the *virality* – or quality – of the tweet content and it scales the subsequent retweet rate; β introduces a warping effect for the user influence and it is related to the observed long-tail distribution of user influence in social networks; and $1 + \theta$ ($\theta > 0$) is the power-law exponent, describing how fast an event is *forgotten*, parameter $c > 0$ is a temporal shift term to keep $\phi_m(\tau)$ bounded when $\tau \simeq 0$; Here κm^β accounts for the magnitude of influence, and the power-law kernel $(\tau + c)^{-(1+\theta)}$ models the memory over time. We approximate user influence m using the number of followers obtained from Twitter API. Table 1 summarizes the parameters of the model, along with their interpretations.

Throughout literature, three families of functions have been mainly used to model temporal decay with Hawkes point-processes [32]: power-law functions $\phi^p(\tau) = (\tau + c)^{-(1+\theta)}$, used in geophysics [17] and social networks [9]; exponential functions $\phi^e(\tau) = e^{-\theta\tau}$, used in financial data [14]; Reyleigh functions $\phi^r(\tau) = e^{-\frac{1}{2}\theta\tau^2}$, used in epidemiology [40]. We experimented with all three kernels with the setting in Sec. 6, and found that the power-law kernel consistently outperform the other two. Consequently, we only discuss detailed results from the power law kernel in the rest of this paper.

3.2 The branching factor

One of the key quantities that describe a Hawkes process is the branching factor n^* , defined as the expected number of children events directly spawned by an event. Note that the total number of expected events from one initial event is the sum of a geometric series $\sum_{k=1}^{\infty} (n^*)^k$. For $n^* < 1$, the process is in a *subcritical regime*: the total number of retweets is bounded, the event rate $\lambda(t)$ decays to zero over time and the retweet cascade dies out. When $n^* > 1$, the process is in a so-called *supercritical regime* with the total number

Table 1: Summary of notations. Top: Quantities for estimating popularity using point processes. Bottom: Parameters of the marked Hawkes process.

Notation	Interpretation
m_i	event magnitude. Social influence of the user emitting the tweet.
t_i	event time. Tweet timestamp.
$\lambda(t)$	conditional intensity – or event rate – of a non-homogeneous Poisson Process.
$\phi_m(\tau)$	triggering kernel. Contribution of event the (m, t) to the total event rate, calculated at time $t + \tau$.
n^*	branching factor, mean number of children events spawned by a parent event. For $n^* < 1$ the cascade dies out (subcritical regime), for $n^* > 1$ the cascade explodes exponentially (supercritical regime).
T	time extent of the observed series of events in the beginning of a retweet cascade. $T = \max(t_i)$.
n	number of events in the observed series of events in the beginning of a retweet cascade. $n = \max(i)$.
N_∞	expected total number of events in a given retweet cascade, derived from the generative model.
N_{real}	the real number of events in a given retweet cascade, or total popularity.
ω	predictive factor learned using a Random Forest regressor.
$\hat{N}_\infty$	total popularity prediction made using the output of the predictive layer.

Parameter	Interpretation
θ	the power-law exponent, describing how fast an event is <i>forgotten</i> . $\theta > 0$.
κ	tweet quality. High quality tweets are more likely to generate more retweets. $\kappa > 0$.
c	temporal shift cutoff term so that $\phi_m(\tau)$ stays bounded when $\tau \simeq 0$. $c > 0$.
β	user influence power-law warping effect in $b(m)$. $\beta > 0$. $0 < \beta < \alpha - 1$ for having a defined branching factor n^* .
α	exponent of the user influence power-law distribution. Fitted to the <code>#followers</code> distribution: $\alpha = 2.016$.

of retweets being unbounded. We compute the branching factor by integrating over time and event magnitudes.

$$n^* = \int_1^\infty \int_0^\infty P(m) \phi_m(\tau) d\tau dm . \quad (3)$$

Here we assume m_i are *i.i.d.* samples from a power law distribution of social influence [21]: $P(m) = (\alpha - 1) m^{-\alpha}$. Here α is an exponent which controls the heavy tail of the distribution. We extract the number of followers m for a large sample of users from our dataset described in Sec. 5.1 and fit it to a power-law distribution following the method detailed in [8]. We obtain $\alpha = 2.016$ and we use it throughout the experiments. Substituting Eq. (2) and $P(m)$ into (3), we obtain the closed-form expression of the branching factor:

$$n^* = \kappa \frac{\alpha - 1}{\alpha - \beta - 1} \frac{1}{\theta c^\theta}, \text{ for } \beta < \alpha - 1 \text{ and } \theta > 0 . \quad (4)$$

Not only the branching factor n^* is an intuitive descriptor of the *virality* of a cascade, it is also used for predicting the total popularity in Sec 4 and as a constraint in model estimation in Sec 3.3.

3.3 Estimating the Hawkes process

The log-likelihood function. Our marked Hawkes process is completely defined by the four parameters $\{\kappa, \beta, c, \theta\}$. We now describe a maximum-likelihood estimation procedure of the model from a set of observed events, consisting of $\{(m_i, t_i), i = 0, \dots, n\}$ for time $t < T$. The log-likelihood function of a point-process described by Eq. (1) and (2) is written as ([10], Ch. 7.2):

$$\begin{aligned} \mathcal{L}(\kappa, \beta, c, \theta) &= \log P(\{(m_i, t_i), i = 1, \dots, n\}) \\ &= \sum_{i=1}^n \log(\lambda(t_i)) - \int_0^T \lambda(\tau) d\tau \\ \text{cf. (1)\&(2)} &= \sum_{i=2}^n \log \kappa + \sum_{i=2}^n \log \left(\sum_{t_j < t_i} \frac{(m_j)^\beta}{(t_i - t_j + c)^{1+\theta}} \right) \\ &\quad - \kappa \sum_{i=1}^n (m_i)^\beta \left[\frac{1}{\theta c^\theta} - \frac{(T + c - t_i)^{-\theta}}{\theta} \right] . \quad (5) \end{aligned}$$

The first two terms in Eq. 5 are from the likelihood computed using the event rate $\lambda(t)$, the last term is a normalizing factor from integrating the event rate over the observation window $[0, T]$.

Finding the maximum-likelihood solution. Eq. 5 is a non-linear objective that we optimize to find the set of parameters that maximizes data likelihood. Directly maximizing this objective is more efficient than expectation-maximization approaches [12] that try to estimate the branching structure (of which event triggers which) and that have a quadratic number of latent variables. There are a few natural constraints for each of model parameter, namely: $\theta > 0$, $\kappa > 0$, $c > 0$, and $0 < \beta < \alpha - 1$ for the branching factor to be meaningful (and positive). Furthermore, while the supercritical regimes $n^* > 1$ are mathematically valid, they will not lead to a finite cascade size prediction. We note that in reality no retweet cascade can grow indefinitely, due to finite network sizes. We further incorporate $n^* < 1$ as a non-linear constraint for the maximum likelihood estimation. We use the Ipopt [39], the large-scale interior point solver that handles both non-linear objectives and non-linear constraints, with pre-programmed gradient functions to find the solution. Details of the gradient computation and optimization can be found in an online supplement [1].

3.4 Interpreting the generative model

Individual cascades: fast vs slow. We first examine the model parameters for the two cascades illustrated in Fig. 1. These two cascades are about the same event and have similar observed popularities (224 vs. 219) through very different diffusion speeds at a complex interplay between power-law memory $(\tau + c)^{-(1+\theta)}$ and user influences. For tw_1 , the maximum-likelihood estimate of parameters are $\{\kappa = 0.75, \beta = 0.27, c = 58.46, \theta = 0.64\}$ with a corresponding $n^* = 0.12$; for tw_2 , $\{\kappa = 1, \beta = 0.42, c = 27.77, \theta = 0.77\}$, and $n^* = 0.16$ (also shown in Fig. 2). The two diffusions unfold in different ways: tw_1 has lower value of κ and higher waiting time c , hence its memory kernel $\phi(\tau)$ has lower values than that of tw_2 and a slower decay. However, tw_1 attracts the attention of some very well-followed music accounts. The original poster (`@screencrushnews`) has 12,122 followers, and among the retweeters `@TasteOfCountry` (country music) has 193,081 followers, `@Loudwire`

Figure 2: Density distribution of model parameters κ (a) and θ (b), on a sample of 17,146 cascades from the NEWS dataset. The dashed vertical lines show fitted parameter values for the two retweet cascades illustrated in Fig. 1: tw_1 in red and tw_2 in blue. We note that SEISMIC [44] uses a fixed value of $\theta = 0.242$ for all cascades (denoted in green). The distributions for c, β and n^* are shown in the supplement [1].

(rock) had 110,824 followers, @UltClassicRock (classic rock) has 99,074 followers and @PopCrush (pop music) has 114,050 followers. The resulting cascade reached 1/4 its size after half an hour, and the final tweet was sent after 4 days. In contrast, tw_2 's memory kernel has higher values, but faster decay. The most influential user in tw_2 has 412 followers, and the entire diffusion lasted for only 12.5 minutes.

Interpreting a collection of cascades. We construct a large subset of 17,146 cascades of the NEWS dataset, described in Sec. 5.1. For each cascade we fit the parameters of its generative model $\{\kappa, \beta, c, \theta\}$. Fig. 2 shows the density distribution of parameters θ and κ . As one would expect, most cascades have a low value of κ , and reflect the long-tailed distribution of content quality. n^* – descriptor of the virality of a cascade – is distributed very similarly to κ , with most cascades showing low n^* values (shown in the supplement [1]). Parameter θ – which controls for the speed of the decay of the memory kernel $\phi_m(\tau)$ – has a more surprising distribution. Recent work [44] on point processes with power-law kernels argued that a unique value of $\theta = 0.242$ is sufficient to explain the general behavior cascades. We find θ to vary in the interval $(0, 2.8]$ and showing two maxima at 0.01 and 0.34. This observation indicates that learning individual memory exponents θ allows us to better describe the diffusion dynamics of each cascade. The distributions of the other parameters are shown in the supplement: c follows a long-tail distribution; β shows a maximum of density around 0.08 and a smaller peak around $\beta \sim \alpha - 1$ resulting from the optimizer terminating at the boundary of the non-linear constraint $n^* = 1$.

4. PREDICTING POPULARITY

In this section, we first derive the number of expected future events in a cascade, we then describe a predictive layer tuned for estimating the total size, using information from other cascades in history.

4.1 The expected number of future events

Having observed a retweet cascade until time T and fitted the parameters of the point process, one can simulate a possible continuation of the cascade, until it dies out – assuming $n^* < 1$. In this section, we derive the expected number of events in the cascade, over all possible continuations. We group events based on the generation of their parents (i.e. the preceding event they are triggered by), and

Figure 3: Illustration of the rationale behind popularity prediction. The model parameters are estimated starting from a series of observed events (m_i, t_i) . Part of one possible unfolding of the diffusion cascade is simulated, using the event rate defined by Eq. (2) and simulated by thinning [28]. The event generations are shown: events in $Generation_1$ are shown in red color, $Generation_2$ in green and $Generation_k$ in blue. Note there is no theoretical limit to the number of generations or to the extent of time until the cascade dies out. Some of the parent-child relations between events in consecutive generations are shown.

estimate the number of events in each generation. We denote as $Generation_1$ the set of simulated events spawned by the observed events (shown in red in Fig. 3). Similarly, $Generation_2$ (in green in Fig. 3) groups events spawned by events in $Generation_1$, and so on. Let A_i be the expected number of events in $Generation_i$. The expected number of total events in the cascade, N_∞ , is defined as:

$$N_\infty = n + \sum_{i=1}^{\infty} A_i, \quad (6)$$

where n is the number of observed events. For generations after the first one, the best estimate about $A_i, i \geq 2$ is constructed using the average number of children events n^* and the number of events in the previous generation, i.e. $A_i = A_{i-1}n^*$. Assuming A_1 known, we derive:

$$A_i = A_{i-1}n^* = A_{i-2}(n^*)^2 = \dots = A_1(n^*)^{i-1}, i > 1 \quad (7)$$

Therefore, the second term in (6) is the sum of a converging geometric progression (assuming $n^* < 1$):

$$\sum_{i=1}^{\infty} A_i = \frac{A_1}{1 - n^*} \text{ where } n^* < 1 \quad (8)$$

A_1 could also be calculated in a fashion similar to Eq. (7). However, a more precise estimation can be obtained, given that magnitudes of the observed events – parents for events in $Generation_1$ – are known:

$$\begin{aligned} A_1 &= \int_T^\infty \lambda(\tau) d\tau = \int_T^\infty \sum_{t > t_i} \phi_{m_i}(t - t_i) dt \\ &= \kappa \sum_{i=1}^n \frac{m_i^\beta}{\theta (T + c - t_i)^\theta} \end{aligned} \quad (9)$$

We obtain the estimate of the total number of events in the cascade by introducing (9) and (8) into (6):

$$N_\infty = n + \frac{\kappa}{(1 - n^*)} \left(\sum_{i=1}^n \frac{m_i^\beta}{\theta (T + c - t_i)^\theta} \right), n^* < 1 \quad (10)$$

Figure 4: Reduction of prediction error for a subset of cascades. For each cascade, the error made when predicting final popularity using the theoretical N_∞ is shown using red circles, and with blue squares the error made when using the predictive layer $\hat{N}_\infty$. Each gray arrow shows the reduction of error for a single cascade, and pairs together a red circle with a blue square. The error reduction is shown in relation to A_1 (a) and n^* (b).

4.2 Predicting total popularity

Using generative models for prediction is often sub-optimal, often due to simplifying assumptions of the generative process. Point processes in general, are optimized for explaining observed event history and not optimized for prediction. Zhao *et al.* [44], for example, recently deployed a number of heuristic correction factors to discount the initial burst and account for a long-term decay. We consider that the Hawkes point process can benefit from systematically learning a predictor, to better account for data noise and limiting model assumptions. First, it is well known that using the number of followers as user influences m_i is at best an approximation [5] and often does not scale with retweeting propensity. Second, the network in which cascades happen can be inhomogeneous over the lifetime of the diffusion, e.g. users who respond early may have an inherently shorter memory than those responding late. Last, point processes cannot generate a meaningful prediction for supercritical cascades, and the model estimation can be prone to local minima. In this section, we leverage a collection of previously observed cascade histories to fine tune popularity prediction. This predictive layer accounts for the limitations above, and places point process estimates in a comparable framework as the feature-driven approaches.

Prediction setup. Each retweet cascade is described using four features $\{c, \theta, A_1, n^*\}$. These four features were found to be the most correlated with $\epsilon = (N_{real} - N_\infty)^2$, the error made when predicting popularity using N_∞ . Furthermore, n^* and θ are expressed as percentiles over the range observed in the training set, to account for their non-linear relation with ϵ . Note that parameters κ and β are not used: the information in κ is already included in A_1 , whereas β was found non-correlated with ϵ .

In the classification task detailed in Sec. 6.3, the Random Forest is trained to output the selected class. In the popularity prediction task – Sec. 6.1 and 6.2 – we train the Random Forest to output a scaling factor ω for the expected number of future events in a given retweet cascade. The final corrected popularity prediction is computed as:

$$\hat{N}_\infty = n + \omega \left(\frac{A_1}{1 - n^*} \right). \quad (11)$$

Effects of the predictive layer. The immediate effect of introducing the predictive layer is the reduction of

the ARE prediction error, measured as detailed in Sec. 6.1. Fig 4 illustrates this reduction on a subset of the cascades in the NEWS dataset – described in Sec. 5.1 – and its relation with A_1 and n^* . The length of the arrows is proportional with the error reduction. The range the errors made without the predictive layer increases linearly with A_1 . Consequently the error reduction when using $\hat{N}_\infty$ is generally higher when A_1 is higher (Fig. 4a). The relation between n^* and the prediction error is non-linear, therefore in Fig. 4b the horizontal axis shows the percentile values of n^* . The variation of the error reduction has two maxima, one around 40% and another one around 95%. Fig. 4 can be summarized as follows: the error reduction increases linearly with the expected number of events in *Generation*₁, but exhibits two maxima in relation with the expected number of events in *Generation* _{i} , $i \geq 2$. Another interesting observation is that N_∞ tends to over-estimate cascade sizes, in other words cascades tend to achieve less popularity than expected when observed in isolation. This aligns with the hypothesis that content items compete with each other for a finite amount of human attention [26].

5. DATA AND FEATURES FOR PREDICTION

In this section we describe the two datasets, TWEET-1MO and NEWS, used in our experiments. Next, we describe the set of features constructed for our feature-driven approach.

5.1 Datasets

TWEET-1MO: Zhao *et al.* [44] made publicly available the dataset they used to predict retweet cascade sizes. 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. For each cascade, the dataset contains information about the original tweet id and start time and, for each of the retweets in the cascade, the time offset of the retweet and the number of followers of the user posting the retweet. We randomly sample the cascades in the original dataset and we construct TWEET-1MO to contain 30,463 cascades of length between 50 and 5000. The resulting dataset has a mean cascade length of 160 and a median of 95.

NEWS: We construct a domain specific dataset of tweets dealing with news articles, by collecting data using the *free* Twitter Streaming API², over a period of four months from April 2015 to July 2015. We track the official twitter handles of ten popular news outlets, namely NewYork Times, Associated Press, CNN, Washington Post, Reuters, Yahoo News, The Guardian, BBC, Huffington Post and Google News. We also capture all tweets containing either i) user mentions of the above users or ii) a commonly used shortened url towards any of the ten news websites. This leads to a collection of 49,735,271 tweets. Most of the cascades in the NEWS dataset are of length one or two and, for compatibility with the TWEET-1MO, we present the statistics for cascades of length over 50: 110,954 cascades, mean cascade length is 158 and median length is 90.

Difference between datasets. The NEWS dataset contains several additional information, when compared to TWEET-1MO: for each user in a retweet cascade, we record her number of friends, number of posted statuses and the account creation time (these information are recorded in every tweet

²<https://dev.twitter.com/streaming/overview>

we capture through the API). This additional information is required by the feature-driven method described in Section 5.2, therefore in the experiments in Sec. 6, feature-based methods cannot run on TWEET-1Mo.

An open-access, domain-specific dataset. Many of the datasets used by prior work for popularity prediction were not made public because of user agreement terms and they are difficult to replicate. The barriers for having an open benchmark include proprietary insider information (e.g. complete Facebook network measurements [7]), privileged API access (e.g., Twitter firehose [44]) or contain restricted data from non-English sources [12, 43]. Furthermore, we argue that predicting popularity on a topic- or domain-specific dataset is as relevant as that on the firehose feed – content producers and consumers are often interested in content within a source (e.g. NYTimes) or a topic (e.g. technology). Therefore, we chose to construct NEWS, our domain-specific dataset, using only free access APIs. We present a strong basis for future benchmarking by combining the descriptive power of FEATURE-DRIVEN and modeling capabilities of HAWKES in Sec. 6.

5.2 Features for popularity prediction

Our feature-driven predictor describes each cascade using four types of features: a) basic user features, b) temporal features, c) volume features and d) past user success. They are chosen as being the most informative in recent literature [4, 7, 25, 29, 37] and require only free access to the Twitter data source. In particular, Martin *et al.* [25] compare a wide range of features, including user, content and user-content interplay features, and they show that basic user features together with past user success features account for 87% of the prediction performance. Cheng *et al.* [7] showed that temporal features relating to the unfolding of the observed part of the diffusion account for the bulk of the total classification performances. They report only 0.025 additional accuracy when using all the other features combined in addition to temporal features. Lastly, Szabo *et al.* [37] and later Pinto *et al.* [29] show that early popularity is predictive for total popularity.

Basic User Features [4, 7, 25]. Basic user statistics capture the social influence of a user. When observing the prefix of a cascade, we use the five point summary (min, median, max, 25-th and 75-th percentile) to represent the distribution for a feature:

- Number of Followers: count of user followers;
- Number of Friends: count of friends for a user;
- Number of Status: count of statuses posted by user;
- User Time: Time when user account was created.

Temporal Features [7]. We capture the temporal dynamics of a cascade using the following features:

- First Half Rate: mean waiting time between retweets, during the first half of the observed cascade;
- Second Half Rate: mean waiting time between retweets, during the second half of the observed cascade;
- Waiting Time Distribution: five point summary of waiting times between retweets, in the observed cascade;
- Exposure Distribution: five point summary of the distribution of number of followers for all users before the last retweet, in the observed cascade. Used only in the classification task in Sec. 6.3.

Figure 5: Distribution of Absolute Relative Error (ARE) on the TWEET-1Mo dataset (top row) and on the NEWS (bottom row), for SEISMIC and HAWKES. A part of the diffusion cascade is observed before making the predictions: 5 min (left column), 10 min (middle column) and 1 hour (right column). The red line and the numeric annotations denotes the median values of the distributions.

Volume [29, 37]. Number of retweets in the observed part of the cascade, which captures early popularity.

Past User Success [4, 25]. Average size of the cascades started by a given user in past. For an observed cascade, we use the five point summary of the distribution of past user success for all previously known users who participate in the cascade. This feature requires a large volume of historical data and can be computed only for users that have started cascades in the past. We detail its construction in Sec. 6.1.

6. EXPERIMENTS

We first present in Sec. 6.1 a thorough comparison between the Hawkes model and SEISMIC, the state-of-the-art generative model on popularity prediction, using two datasets. In Sec. 6.2 and Sec. 6.3 we compare the feature-driven classifier to generative approaches. Sec. 6.1 and Sec. 6.2 tackle regression tasks, i.e., predicting total cascade size after observing it for a certain time, while Sec. 6.3 tackles classification tasks, i.e., whether or not the cascade size will double from what has been observed.

6.1 Cascade size: two generative methods

We compare our point processes model, HAWKES, with SEISMIC on two datasets, TWEET-1Mo and NEWS. Note that TWEET-1Mo is the dataset on which Zhao *et al.* [44] reported results, while NEWS dataset supports extracting a richer set of features that is not available in TWEET-1Mo. In addition, Zhao *et al.* [44] showed that SEISMIC compared favorably against a number of baselines: from auto-regressive and generative models such as dynamic Poisson model [2, 9] and reinforced Poisson model [34], to linear regression and its variants [37].

Experimental setup. We run our experiments on the TWEET-1Mo dataset and a subset of NEWS: 20,093 cascades from the month of July, which have the total length of at least 50. Both algorithms observe the initial part of each cascade for a limited extent of time and predict its final popularity. We estimate separate sets of parameters for HAWKES and SEISMIC after observing the cascades for 5 minutes, 10 minutes and 1 hour, respectively. After the fitting is finished, we train the Random Forest regressor in the predictive layer of the HAWKES model. Performance of this regression layer

Table 2: (left hand side) Mean Average Relative Error (ARE) $\pm$ standard deviation obtained using HAWKES and SEISMIC on the TWEET-1Mo and NEWS datasets. (right hand side) Number of cascades for which prediction failed.

Dataset	Approach	Prediction error: different time lengths			Number of failed cascades		
		5 minutes	10 minutes	1 hour	5 minutes	10 minutes	1 hour
TWEET-1Mo	SEISMIC	2.61 $\pm$ 55.80	0.70 $\pm$ 15.58	0.51 $\pm$ 10.81	507	164	71
	HAWKES	0.36 $\pm$ 0.52	0.33 $\pm$ 0.41	0.30 $\pm$ 0.38	302	105	58
NEWS	SEISMIC	11.13 $\pm$ 282.96	0.84 $\pm$ 11.77	0.33 $\pm$ 0.92	1022	155	123
	HAWKES	0.42 $\pm$ 6.83	0.25 $\pm$ 0.60	0.22 $\pm$ 1.16	149	45	37

Table 3: Mean ARE $\pm$ standard deviation for SEISMIC, HAWKES, FEATURE-DRIVEN and HYBRID models for the NEWS July’15 dataset.

Approach	5 minutes	10 minutes	1 hour
SEISMIC	15.16 $\pm$ 375.08	0.71 $\pm$ 4.89	0.32 $\pm$ 0.40
FEAT.-DRIVEN	0.25 $\pm$ 0.18	0.22 $\pm$ 0.17	0.17 $\pm$ 0.14
HAWKES	0.27 $\pm$ 1.83	0.22 $\pm$ 0.80	0.17 $\pm$ 0.36
HYBRID	0.17 $\pm$ 0.16	0.15 $\pm$ 0.14	0.11 $\pm$ 0.12

is reported using ten fold cross-validation. 40% of cascades are used for training, 60% for testing. All experimental protocol choices – minimum cascade length threshold, random training/testing split and length of time prefixes – are made to mimic closely the original experimental setup [44].

We use the Absolute Relative Error (ARE) to measure prediction performance on a given cascade w :

$$ARE^w = \frac{|\hat{N}_\infty^w - N_{real}^w|}{N_{real}^w},$$

where $\hat{N}_\infty$ and N_{real} are the predicted and observed popularities of cascade w .

Results. The prediction results are summarized in Fig. 5 as boxplots of ARE for each combination (dataset, approach, observed time), while Table 2 shows the mean ARE $\pm$ standard deviation. Note that generative models do not output a prediction for all cascades. For example, when branching factor $n^* \geq 1$ for HAWKES the estimate for cascade size is infinite, similar failures also occur with SEISMIC when the infectiousness parameter $p \geq \frac{1}{n^*}$ ([44] Eq(7)). Here we report the number of “failed” cascades in Table 2, and only report the average ARE on cascades that both approach produce a prediction. As shown in right hand side of the Table 2, we observe that HAWKES is able to produce meaningful predictions for more cascades than SEISMIC. For example, on NEWS data with 5 minutes, we output prediction for 873 (or 0.7%) more cascades than SEISMIC.

We can see that HAWKES consistently outperforms SEISMIC, as it achieves lower average ARE on both datasets. Mean ARE for HAWKES is at least 40% better than SEISMIC on TWEET-1Mo and 33% on NEWS. We observe that NEWS dataset after 5 minutes of observation seems the most difficult to predict: HAWKES produces a mean ARE of 0.42, while the mean ARE for SEISMIC is 11.13. Fig. 5 (top row) shows that the median of HAWKES is at least 25% better than SEISMIC for any time prefix, suggesting that our approach predicts better for most of the cascades in TWEET-1Mo. Similar conclusions can be drawn from the bottom

Figure 6: Distribution of ARE on the NEWS dataset, split in time for July, for SEISMIC, FEATURE-DRIVEN, HAWKES and HYBRID, after observing 10 minutes (a) and 1 hour (b). The red line and the numeric annotations denote median value.

row of Fig. 5 (NEWS data): the median ARE of HAWKES is at most 0.19 at 5 minutes, while the best of SEISMIC models is a median ARE of .24 at 1 hour.

We are glad to see that the HAWKES model achieves better predictive performance than SEISMIC, by exploiting full flexibility in its non-linear parameters, and a predictive layer that systematically optimizes for future prediction using information from other cascades.

6.2 Cascade size: generative vs. feature-driven

In this section we compare the performances of feature-driven and generative approaches on the NEWS dataset only, as the TWEET-1Mo does not contain the necessary data to construct the features detailed in Sec. 5.2.

Experimental setup. Similar to the setup in the previous Sec 6.1, we observe cascades for 5 minutes, 10 minutes and 1 hour and fit the HAWKES and SEISMIC models for each cascade. The train-test split is different, in order to replicate closer the experimental setup in Martin *et al.* [25]: the data from first half of July (1-15) is used for training and the data from second half of July (16-31) for testing. We use the NEWS historical data from April to June to construct the past user success feature for the feature-driven approach. We consider only those users who initiated at least 2 cascades in the past as active users. Non-active users have a past user success of 0. We also construct a HYBRID approach, which leverages the features of the HAWKES model – i.e. $\{c, \theta, A_1, n^*\}$ – together with the features detailed in Sec. 5.2. We use a Random Forest regressor for predicting the final volume size after tuning the predictor on the training set with 10 fold cross-validation.

Results. Similar to the previous subsection, we report in Table 3 the mean ARE $\pm$ standard deviation and in Fig. 6 the box plots of ARE for 10 minutes and 1 hour. In accordance with the results of Sec. 6.1, HAWKES surpasses SEISMIC

Figure 7: ARE distribution over popularity percentiles, on NEWS dataset, after observing cascades for 10 minutes. The blue line denotes the overall median ARE shown in Figure 6a. Note that the y-axes are not on the same scale for all three graphs.

Table 4: Accuracy $\pm$ standard deviation when predicting whether a cascade will double its size or not.

Approach	25 tweets	50 tweets
Random Guess	0.52	0.53
HAWKESC	0.66 ± 0.013	0.70 ± 0.009
FEATURE-DRIVEN	0.79 ± 0.009	0.81 ± 0.011
HYBRID	0.79 ± 0.008	0.82 ± 0.013

for all considered setups. Note that these results are numerically different from those in Sec 6.1, because the underlying test dataset consist only the subset from second half (16-31) of July 2015. Surprisingly, even the FEATURE-DRIVEN model outperforms SEISMIC for all time prefixes, by at least 40%. In both cases, the differences are statistically significant ($p - val < 0.001$) for 10 minutes and 1 hour. The HAWKES model exhibits similar mean ARE as FEATURE-DRIVEN and higher standard deviation. However, its boxplot in Fig. 6 shows lower median and quartiles. This indicates that HAWKES predict better than FEATURE-DRIVEN for most cascades, but the error distribution of HAWKES is skewed by outliers. HYBRID performs the best, with similar boxplot summarization as HAWKES, but with a 42% better mean APE – 0.17 to 0.11 for 1 hour. The differences between HYBRID and all the other approaches are statistically significant (details in the supplement [1]). This result shows that HAWKES and FEATURE-DRIVEN approaches are complementary, likely because they summarize the cascade information differently, and the errors are uncorrelated.

ARE distribution over cascade popularity. Figure 7 presents the ARE distribution over cascade popularity as boxplots for each popularity decile. Both HAWKES and FEATURE-DRIVEN predict better for most of the cascades with the middle popularity. For the extreme ends, performance is affected, because for lower popularity percentiles even a small error is amplified to a large ARE and cascades in higher popularity bins are inherently more difficult to predict. The prediction performance of SEISMIC seems consistent across popularity bins, but always worse than HAWKES or FEATURE-DRIVEN, as already showed by the aggregated results in Fig. 6a. We also tried to explain prediction performance breaking down by NEWS sources, but we did not see any notable patterns.

6.3 Will this cascade double

The learning problem in this section is to classify whether an observed cascade will double its volume or not. The experimental setup follows that of Cheng *et al.* [7], in which cascades are observed for a fixed number of retweets, instead of for a fixed extend of time in Sec. 6.1 and 6.2. The proposed HAWKES model outputs directly the total popularity

as a cascade size and uses the predictive layer to correct the theoretical estimate. As a consequence, it cannot be used in a classification setup. Instead, we construct a classifier HAWKESC, in which a Random Forest Classifier is trained on the same features as the predictive layer of HAWKES and used to output a binary decision.

Experimental setup. We use the July subset of the NEWS dataset, filtering to only cascades of length greater than or equal to 25. We observe the cascades for two 25 retweets and 50 retweets, and we predict whether they will reach 50 and 100 retweets, respectively. We perform a 40%-60% stratified train-test split. FEATURE-DRIVEN approach uses the features mentioned in Section 5.2, except for *volume* which is constant for all cascades. HYBRID approach combines features from both HAWKES and FEATURE-DRIVEN. The random train-test split is repeated 10 times and we report mean accuracy and standard deviation.

Results. Table 4 summarizes the classification performances. A random guess (majority class) would output an accuracy of 52% and 53% for an observed length of 25 and 50 respectively. The generative-based classifier HAWKESC improves substantially over the baseline of random guess, however FEATURE-DRIVEN has the best prediction accuracy. Interestingly, combining the generative features with the data-driven features did not lead to a significant improvement, likely due to FEATURE-DRIVEN already has stronger results than HAWKESC. Admittedly, predicting whether a cascade will double in size is an easier problem than forecasting its final volume. Overall, the results suggests that the predictions are robust, and the performance are comparable to earlier results reported by Cheng *et al.* [7] (0.81 accuracy after seeing the first 50 shares).

7. CONCLUSION

This paper establishes a common benchmark that allows researchers and practitioners to compare feature-driven and generative approaches on different variants of popularity prediction problems. Existing feature-based approaches employ either proprietary insider information, or global network information that is impossible to retrieve with the limited access that are openly available. In this work, we bridge the problem space by first selecting a small set of features which can be obtained by using only free access services and constructing the FEATURE-DRIVEN method. We further propose a Hawkes self-exciting model, which intuitively aligns with the social factors responsible for diffusion of cascades: social influence of users, social memory and inherent content quality. We systematically construct a predictive layer, which helps optimize predict from other cascades. We perform extensive evaluation on two large datasets: a benchmark dataset constructed by Zhao *et al.* [44] and a domain-

specific dataset curated on news tweets. We compare the feature-driven approaches and generative approaches in two tasks: i) predicting the total size of retweeting cascades and ii) predicting whether cascades will double their size. Both our proposed generative HAWKES method and our FEATURE-DRIVEN method outperform the current state of the art predictor. Performances are further improved when combining both approaches into HYBRID, which makes us argue that popularity modeling should leverage the best that both worlds have to provide.

We plan to extend this work by including additional factors such as, different network behaviour for particular types of content, user influence distributions that are varying over time, and explicit modeling of the interplay between the unfolding of related cascades.

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.

8. REFERENCES

- [1] Supplemental material – feature driven and point process approaches for popularity prediction, 2016. <http://arxiv.org/pdf/1608.04862.pdf#page=11>.
- [2] D. Agarwal, B.-C. Chen, and P. Elango. Spatio-temporal models for estimating click-through rate. In *WWW '09*, pages 21–30, 2009.
- [3] S. Asur and B. A. Huberman. Predicting the future with social media. In *WI-IAT '10*, volume 1, pages 492–499, 2010.
- [4] E. Bakshy, J. M. Hofman, W. A. Mason, and D. J. Watts. Everyone’s an influencer. In *WSDM '11*, page 65, feb 2011.
- [5] M. Cha, H. Haddadi, F. Benevenuto, and K. P. Gummadi. Measuring user influence in twitter: The million follower fallacy. In *ICWSM '10*, May 2010.
- [6] B. Chang, H. Zhu, Y. Ge, E. Chen, H. Xiong, and C. Tan. Predicting the Popularity of Online Serials with Autoregressive Models. In *CIKM '14*, pages 1339–1348, nov 2014.
- [7] J. Cheng, L. Adamic, P. A. Dow, J. M. Kleinberg, and J. Leskovec. Can cascades be predicted? In *WWW '14*, pages 925–936, 2014.
- [8] A. Clauset, C. R. Shalizi, and M. E. J. Newman. Power-Law Distributions in Empirical Data. *SIAM Review*, 51(4):661–703, nov 2009.
- [9] R. Crane and D. Sornette. Robust dynamic classes revealed by measuring the response function of a social system. *PNAS '08*, 105(41):15649–15653, oct 2008.
- [10] D. J. Daley and D. Vere-Jones. *An introduction to the theory of point processes. Vol. I*, volume I. 2008.
- [11] F. Davletov, A. S. Aydin, and A. Cakmak. High Impact Academic Paper Prediction Using Temporal and Topological Features. In *CIKM '14*, pages 491–498, nov 2014.
- [12] W. Ding, Y. Shang, L. Guo, X. Hu, R. Yan, and T. He. Video Popularity Prediction by Sentiment Propagation via Implicit Network. In *CIKM '15*, pages 1621–1630, oct 2015.
- [13] M. Farajtabar, Y. Wang, M. Rodriguez, S. Li, H. Zha, and L. Song. Coevolve: A joint point process model for information diffusion and network co-evolution. In *NIPS '15*, pages 1945–1953, 2015.
- [14] V. Filimonov and D. Sornette. Apparent criticality and calibration issues in the Hawkes self-excited point process model: application to high-frequency financial data. *arXiv preprint*, aug 2013.
- [15] A. G. Hawkes. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1):83–90, apr 1971.
- [16] A. G. Hawkes and D. Oakes. A Cluster Process Representation of a Self-Exciting Process. *J. Appl. Probab.* '74, 11(3):493, sep 1974.
- [17] A. Helmstetter and D. Sornette. Subcritical and supercritical regimes in epidemic models of earthquake aftershocks. *J GEOPHYS RES-SOL EA '02*, 107(B10):ESE 10–1—ESE 10–21, 2002.
- [18] A. Johansen and D. Sornette. Download relaxation dynamics on the WWW following newspaper publication of URL. *PhysA '00*, 276(1-2):338–345, 2000.
- [19] K. Y. Kamath and J. Caverlee. Spatio-temporal meme prediction. In *CIKM '13*, pages 1341–1350, oct 2013.
- [20] S. Kim, D. Putrino, S. Ghosh, and E. N. Brown. A Granger causality measure for point process models of ensemble neural spiking activity. *PLoS Comput. Biol.* '11, 7(3):e1001110, mar 2011.
- [21] H. Kwak, C. Lee, H. Park, and S. Moon. What is twitter, a social network or a news media? In *WWW '10*, pages 591–600, 2010.
- [22] H. Li, X. Ma, F. Wang, J. Liu, and K. Xu. On popularity prediction of videos shared in online social networks. In *CIKM '13*, pages 169–178, oct 2013.
- [23] L. Li and H. Zha. Dyadic event attribution in social networks with mixtures of Hawkes processes. In *CIKM '13*, pages 1667–1672, 2013.
- [24] L. Li and H. Zha. Learning parametric models for social infectivity in multi-dimensional hawkes processes. In *AAAI '14*, 2014.
- [25] T. Martin, J. M. Hofman, A. Sharma, A. Anderson, and D. J. Watts. Exploring limits to prediction in complex social systems. In *WWW '16*, pages 683–694, 2016.
- [26] G. Miritello, R. Lara, M. Cebrian, and E. Moro. Limited communication capacity unveils strategies for human interaction. *Sci. Rep.* '13, 3, 2013.
- [27] G. O. Mohler, M. B. Short, P. J. Brantingham, F. P. Schoenberg, and G. E. Tita. Self-Exciting Point Process Modeling of Crime. *J. Amer. Statist. Assoc.* '11, 106(493):100–108, mar 2011.
- [28] Y. Ogata. Seismicity Analysis through Point-process Modeling: A Review. *PURE APPL GEOPHYS '99*, 155(2-4):471–507, 1999.
- [29] H. Pinto, J. M. Almeida, and M. A. Gonçalves. Using early view patterns to predict the popularity of youtube videos. In *WSDM '13*, page 365. ACM Press, feb 2013.
- [30] J. C. L. Pinto, T. Chahed, and E. Altman. Trend detection in social networks using hawkes processes. In *ASONAM '15*, pages 1441–1448, 2015.
- [31] V. Raghavan, G. V. Steeg, A. Galstyan, and A. G. Tartakovsky. Modeling temporal activity patterns in dynamic social networks. *IEEE CSS '14*, 1(1):89–107, 2014.
- [32] M. G. Rodriguez, D. Balduzzi, and B. Schölkopf. Uncovering the Temporal Dynamics of Diffusion Networks. In *ICML '11*, pages 561–568, 2011.
- [33] D. A. Shamma, J. Yew, L. Kennedy, and E. F. Churchill. Viral actions: Predicting video view counts using synchronous sharing behaviors. In *ICWSM '11*, 2011.
- [34] H.-W. Shen, D. Wang, C. Song, and A.-L. Barabási. Modeling and Predicting Popularity Dynamics via Reinforced Poisson Processes. In *AAAI '14*, pages 291–297, Canada, 2014.
- [35] A. Sornette and D. Sornette. Renormalization of earthquake aftershocks. *Geophys. Res. Lett.* '99, 26(13):1981–1984, 1999.
- [36] D. Sornette and A. Helmstetter. Endogenous versus exogenous shocks in systems with memory. *PhysA '03*, 318(3-4):577–591, 2003.
- [37] G. Szabo and B. A. Huberman. Predicting the popularity of online content. *Communicat. of the ACM*, 53(8):80–88, 2010.
- [38] D. Vallet, S. Berkovsky, S. Ardon, A. Mahanti, and M. A. Kafaar. Characterizing and Predicting Viral-and-Popular Video Content. In *CIKM '15*, pages 1591–1600, oct 2015.
- [39] A. Wächter and L. T. Biegler. On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106(1):25–57, 2006.
- [40] J. Wallinga and P. Teunis. Different epidemic curves for severe acute respiratory syndrome reveal similar impacts of control measures. *Am J Epidemiol* '04, 160(6):509–16, sep 2004.
- [41] F. Wu and B. A. Huberman. Novelty and collective attention. *PNAS '07*, 104(45):17599–601, nov 2007.
- [42] H. Yu, L. Xie, and S. Sanner. Twitter-driven youtube views: Beyond individual influencers. In *MM '14*, pages 869–872, 2014.
- [43] L. Yu, P. Cui, F. Wang, C. Song, and S. Yang. Uncovering and predicting the dynamic process of information cascades with survival model. *Know. and Inf. Syst.*, page 10, 2016.
- [44] Q. Zhao, M. A. Erdogdu, H. Y. He, A. Rajaraman, and J. Leskovec. SEISMIC: A Self-Exciting Point Process Model for Predicting Tweet Popularity. In *KDD '15*, 2015.