

From Micro to Macro: Uncovering and Predicting Information Cascading Process with Behavioral Dynamics

Linyun Yu¹, Peng Cui¹, Fei Wang², Chaoming Song³, Shiqiang Yang¹

¹Tsinghua National Laboratory for Information Science and Technology
Department of Computer Science and Technology, Tsinghua University

²Department of Computer Science and Engineering, University of Connecticut

³Department of Physics, University of Miami

yuly12@mails.tsinghua.edu.cn, cuip@tsinghua.edu.cn, fei_wang@uconn.edu
chaoming.song@gmail.com, yangshq@tsinghua.edu.cn

Abstract—Cascades are ubiquitous in various network environments. How to predict these cascades is highly nontrivial in several vital applications, such as viral marketing, epidemic prevention and traffic management. Most previous works mainly focus on predicting the final cascade sizes. As cascades are typical dynamic processes, it is always interesting and important to predict the cascade size at any time, or predict the time when a cascade will reach a certain size (e.g. an threshold for outbreak). In this paper, we unify all these tasks into a fundamental problem: *cascading process prediction*. That is, given the early stage of a cascade, how to predict its cumulative cascade size of any later time? For such a challenging problem, how to understand the micro mechanism that drives and generates the macro phenomena (i.e. cascading process) is essential. Here we introduce behavioral dynamics as the micro mechanism to describe the dynamic process of a node’s neighbors getting infected by a cascade after this node getting infected (i.e. one-hop subcascades). Through data-driven analysis, we find out the common principles and patterns lying in behavioral dynamics and propose a novel Networked Weibull Regression model for behavioral dynamics modeling. After that we propose a novel method for predicting cascading processes by effectively aggregating behavioral dynamics, and present a scalable solution to approximate the cascading process with a theoretical guarantee. We extensively evaluate the proposed method on a large scale social network dataset. The results demonstrate that the proposed method can significantly outperform other state-of-the-art baselines in multiple tasks including cascade size prediction, outbreak time prediction and cascading process prediction.

Keywords—Information Cascades, Social Network, Dynamic Processes Prediction.

I. INTRODUCTION

In a network environment, if decentralized nodes act on the basis of how their neighbors act at earlier time, these local actions often lead to interesting *macro* dynamics - cascades. In online social networks, the information a user can get and engage in is highly dependent on what his/her friends share, and thus information cascades naturally occur and become the major mechanism for information communication. There has been a growing body of research on these information cascades because of their big potential in various vital applications such as viral marketing, epidemic prevention, and traffic management. Most of them focus on characterizing these information

cascades and discovering their patterns in structures, contents and temporal dynamics.

Recently, predictive modeling on information cascades has aroused considerable research interests. Earlier works focus on predicting the final size of information cascades based on content, behavioral and structural features [3], [6]. As only large cascades are of interest in most real applications, Cui et al.[6] propose a data driven approach to predicting whether the final size will surpass a threshold for outbreak. More recently, Cheng et al.[3] investigate the problem of continuously predicting whether the cascade will double the current size in future. However, the previous works are all about cascade size, which do not include the whole of information cascades. Information cascade is a typical dynamic process, and the temporal scale is critical for understanding the cascading mechanism. Furthermore, it is highly nontrivial to predict when a cascade breaks out, and, more ambitiously, to predict the evolving process of a cascade (i.e. cascading process, as shown in 1 (a)). In this paper, we move one step forward to ask: Is the cascading process predictable? That is, given the early stage of an information cascade, can we predict its cumulative cascade size of any later time?

It is apparent that the targeted problem is far more challenging than those in previous works. The commonly used cascade-level macro features for size prediction, such as the content, increasing speed and structures in the early stage are not distinctive and predictive enough for the cascade sizes at any later time. A fundamental way to address this problem is to look into the *micro* mechanism of cascading processes. Intuitively, an information cascading process can be decomposed into multiple local (one-hop) subcascades. When a node involves in a cascade, one or more of its offspring nodes will also involve in the cascade with a temporal scaling. If the dynamic process of these subcascades can be accurately modeled, then the cascade process can be straightforwardly predicted by an additive function of these local subcascades.

Here we exploit *behavioral dynamics* as the micro mechanism to represent the above mentioned dynamic process of local subcascades. Given a node involving in a cascade at t_0 , its behavioral dynamic aims at capturing the changing process

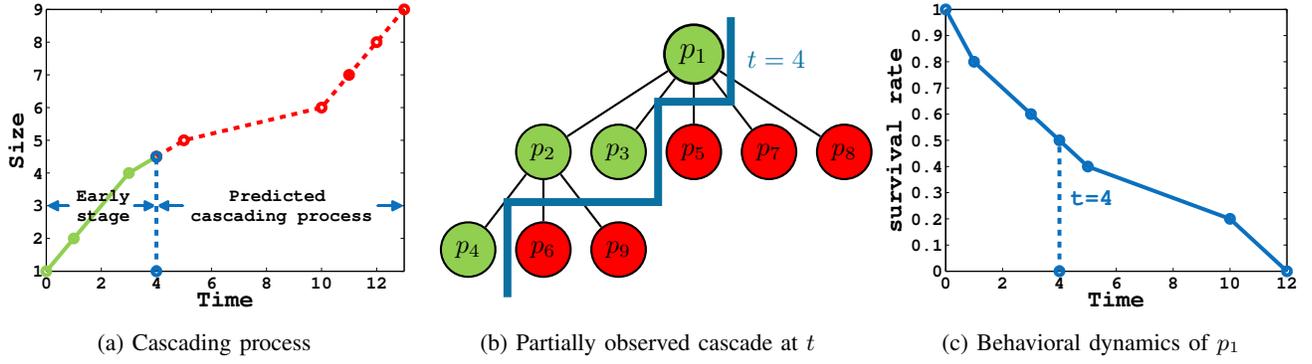


Fig. 1: Illustrations of cascading process prediction. (a) gives the early stage of a cascading process before t . (b) shows the partially observed cascade, where nodes in green (red) represent the observed (unobserved) nodes involved before (after) t . (c) portrays the cascading retweeting dynamics of the followers of p_1 with respect of time.

of the cumulative number of its offspring nodes that involve in the cascade with time evolving. By definition, this is a non-decreasing counting process and can be well represented by survival model [17]. A paucity of recent research works have exploited the survival theory to model how the occurrence of event at a node affects the time for its occurrence at other nodes (i.e. diffusion rate), and their results demonstrate the superiority of continuous-time survival model to uncover temporal processes[11], [18]. However, their targeted problem is to uncover the hidden diffusion networks, and thus suppose the parameters of the survival function on each edge to be fixed. This will cause the unexpected result that all the cascades with the same root node (or early involved nodes) will be anticipated to have the same cascading processes, which makes these models inapplicable in our problem.

In this paper, we propose a novel method for cascading process prediction, as shown in Figure 1. Given the early stage of a cascading process before t in Figure 1 (a), we illustrate the partially observed cascade as shown in Figure 1 (b), where nodes in green (red) represent the observed (unobserved) nodes involved before (after) t . Given the behavioral dynamics of node p_1 represented by its survival rates, and the number of its offspring nodes that have involved before t , we can predict the cumulative number of its offspring nodes that involve in the cascade at any time $t' > t$. After conducting similar predictions on all the observed nodes, the cascading process after t can be predicted by an additive function over all local predictions from behavioral dynamics.

More specifically, how to model behavioral dynamics and further predict cascading process based on continuous-time survival theory also entail many challenges. First, it is unclear what distribution form the behavioral dynamics follow. Although Exponential and Rayleigh distributions are commonly used to characterize the temporal scaling of pairwise interactions, behavioral dynamics in this paper are a reflection of collective behaviors and are proved to be inconsistent with these simple distributions in real data. Second, the parameters in survival models are difficult to interpret, which limits the generality of the learned model. Given the distribution form of data, the parameters of survival model can always be learned from real data in maximum likelihood manner. However, it is unsure what these parameters stands for and the learned

model cannot be generalized to out-of-sample nodes (i.e. the nodes whose behavioral dynamic data is not included in the data). Third, the predictive models based on survival theory are computationally expensive due to the continuous-time characteristic, which makes them infeasible in real applications. How to design an effective and interpretable model for behavioral dynamics modeling and a scalable solution for cascading process prediction are still open issues.

In this paper, we conduct extensive statistical analysis on large scale real data and find that the behavioral dynamics cannot be well captured by simple distributions such as Exponential and Rayleigh distribution, but the general form of Exponential and Rayleigh, Weibull distribution, can well preserve the characteristics of behavioral dynamics. Also, we discover strong correlations between the parameters of a node's behavioral dynamics and its neighbor nodes behavioral features. Enlightened by these, we propose a NETWORKED WEIBULL REGRESSION (NEWER) model for parameter learning of behavioral dynamics. In particular, besides the maximum likelihood estimation term, we also assume the parameters of a node can be regressed by the behavioral features of its neighbor nodes and thus impose networked regularizers to improve the interpretability and generality of the model. Based on the behavioral dynamics, we further propose an additive model for cascading process prediction. To make it scalable, we propose an efficient sampling strategy for approximation with a theoretical guarantee.

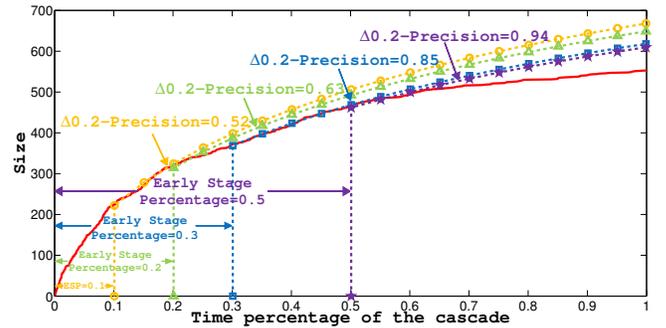


Fig. 2: Showcase of cascading process prediction for a real cascade. The red line represents the groundtruth cascading process. The others are prediction results based on different early stage information.

We extensively evaluate the proposed method in a complete dataset from a population-level social network in China, including over 320 million users, 1.2 billion edges and 340 million cascades. In all the testing scenarios, the proposed method can significantly outperform other baseline methods. Figure 2 is a showcase of cascading process prediction by the proposed method. We show that by accurately modeling behavioral dynamics of social network users, we can predict the whole cascading process with only early 30 percent cascade information, and get the average precision of 0.85 with 20% error tolerance. Also, the accurate predictions of final cascade size, cascade outbreaking time are all implied in the predicted cascading process.

The main contributions of this paper are:

(1) Enlightened by the cascading size prediction works, we move one step forward to attempt cascading process prediction problem, which implies several vital problems such as cascade size prediction, outbreaking time prediction as well as evolving process prediction.

(2) We find out the common principles and patterns lying in behavioral dynamics and propose a novel Networked Weibull Regression model for behavioral dynamics modeling accordingly, which significantly improves the interpretability and generality of traditional survival models.

(3) We propose a novel method for predicting macro cascading process by aggregating micro behavioral dynamics, and propose a scalable solution to approximate the cascading process with a theoretical guarantee.

II. RELATED WORK

Prediction on Cascades. In recent years, many methods have been proposed to make prediction on cascades. Most of them focus on predicting the future size of a cascade, and the common way is to select vital nodes and place sensors on them. For example, Cohen *et al.* [4] focus on exploring the topological characteristics of the cascade. Cui *et al.* [6] proposes to optimize the size prediction problem using dynamic information. Cheng *et al.* [3] introduces temporal feature into the problem and they predict the growing size of the cascade. Rather than attempt to predict the cascade size, we focus on predicting the cascading process which considers both time and volume information together.

Survival Model. Survival model is a method try to analysis things according to the time duration until one or more events happen. In recent years, researchers started modeling information diffusion using continuous models. Myers *et al.* [15] proposed CONNIE to infer the diffusion network base on convex programming while leaving the transmission rate to be fixed, later on Rodriguez *et al.*[18] proposed NETRATE which allowing the transmission rate to be different in different edges. Subsequently, Rodriguez *et al.* [11] give an additive model and a multiplicative model to describe information propagation base on survival theory. Most of these works focus on discovering the rules and patterns to the edges in the social network and is hard to extend to make predictions for cascades since the correlation between transmission rates on edges is little. In contrast, our work focus more on predictive modeling by grouping correlated edges together so that we can make predictions for edges base on the information of other edges.

Influence Modeling and Maximization. Influence modeling and maximization aims to evaluate users' importance in social networks. This is first proposed by Domingos *et al.* [8] to select early starters to trigger a large cascade. Then Kempe *et al.* [13] proposed Stochastic Cascade Model to formalize the problem and Chen *et al.* [2] proposed a scalable solutions. Recently the approach was extended to adding opinion effects [1], [10], topic effects [7] or time decay effects [19] on the models. Our work is distinct from existing works in the following way: Rather than quantify the influence on nodes, we will predict the cascading process.

III. PRELIMINARIES

This section presents the dataset information, discovered patterns and validated hypotheses to support the model design and solution.

A. Dataset Description

The dataset in this paper is from Tencent Weibo, one of the largest Twitter-style websites in China. We collect all the cascades in 10 days generated between Nov 15th and Nov 25th in 2011. The dataset contains in total 320 million users with their social relations, 340 million cascades¹ with their explicit cascading processes. The distribution of cascade size is shown in Figure 3. We can see that the cascade size follows Power-Law distribution, and the majority of cascades have very small size, which are not of interest for many applications. As the paper intends to predict cascading process, we filter out the cascades with the size of less than 5, and maintain the remaining 0.59 million cascades with obvious cascading process for statistical analysis and experiments.

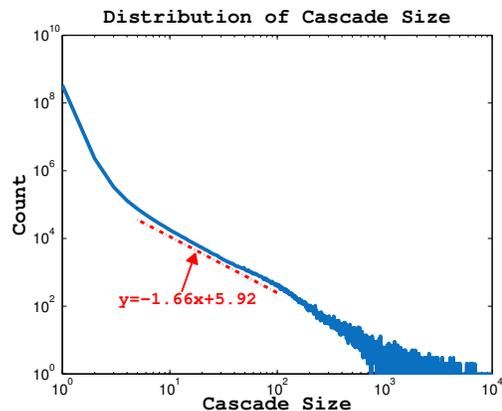


Fig. 3: Distribution of cascade size. The red straight line is the linear fitting result to the blue curve, showing the size distribution fits power-law.

B. Characteristics of Behavioral Dynamics

As mentioned before, behavioral dynamics play a central role in uncovering and predicting cascade processes. Here we investigate the characteristics of behavioral dynamics to enlighten the modeling of behavioral dynamics. By definition, the behavioral dynamics of a user capture the changing process of the cumulative number of his/her followers retweet a post after the user retweeting the post. Then the behavioral dynamics of

¹Here the cascades are information cascades. When a user retweet/generate a post, several of his/her followers will further retweet the post and so on so forth to form a information cascade.

TABLE I: Parametric Models

model	density function	survival function	hazard function	ks-static in Weibo
Exponential	$\lambda_i e^{-\lambda_i t}$	$e^{-\lambda_i t}$	λ_i	0.2741
Power Law	$\frac{\alpha_i}{\delta} \left(\frac{t}{\delta}\right)^{-\alpha_i - 1}$	$\left(\frac{t}{\delta}\right)^{-\alpha_i}$	$\frac{\alpha_i}{t}$	0.9893
Rayleigh	$\alpha_i t e^{-\alpha_i \frac{t^2}{2}}$	$e^{-\alpha_i \frac{t^2}{2}}$	$\alpha_i t$	0.7842
Weibull	$\frac{k_i}{\lambda_i} \left(\frac{t}{\lambda_i}\right)^{k_i - 1} e^{-\left(\frac{t}{\lambda_i}\right)^{k_i}}$	$e^{-\left(\frac{t}{\lambda_i}\right)^{k_i}}$	$\frac{k_i}{\lambda_i} \left(\frac{t}{\lambda_i}\right)^{k_i - 1}$	0.0738

a user can be straightforwardly represented by averaging the size growth curve of all subcascades that spread to the user and his/her followers. However, Figure 4 shows that the size growth curves vary significantly for different subcascades of the same user, which means that such a representation is not fit to characterize behavioral dynamics. Here we normalize the size growth process by the cascade final size and adopt survival function to describe the behavioral dynamics where the survival rate represents the percentage of nodes that has not been but will be infected. As shown in Figure 4, a user’s survival function is quite stable for different subcascades although their size growth patterns vary.

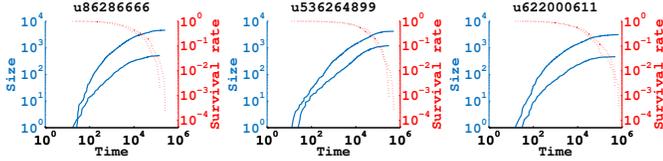


Fig. 4: The size growth curves and their corresponding survival function for 3 users.

Then can we use the behavioral dynamics represented by survival function to predict the size growth curve of a subcascade? We provide positive answer with the assistance of early stage information. For example, if we know the subcascade size at an early time t_0 , then the survival function can be straightforwardly transformed from percentage dimension into size dimension.

C. Parametrize Behavioral Dynamics

For the ease of computation and modeling, we need to parameterize the behavioral dynamics in our case. In state-of-the-art, Exponential and Rayleigh distributions are often used to describe the dynamics of user behaviors in different settings [9], [12]. Here we testify these distribution hypothesis on our real data and find that these distributions cannot well capture both the shape and scale characteristics of behavioral dynamics. Thus, we turn to the general form of Exponential and Rayleigh distributions, the Weibull distribution [16], and find it adequate for parameterizing behavioral dynamics. In order to quantify the effect of parametrization, we calculate KS-Statistic for the three candidate distributions as shown in Table I. It displays that Weibull distribution performs much better than Exponential and Rayleigh distribution. The improvement is attributed to the high degree of freedom of Weibull distribution as it has two parameters λ and k to respectively control the scale and shape of the behavioral dynamics.

D. Covariates of Behavioral Dynamics

If subcascades for all users are sufficient, the parameters of behavioral dynamics can be directly learned from data. However, this suffers from several drawbacks: (1) some users may

TABLE II: Behavioral features and structural for users.

Behavioral features	
<i>inflowrate</i>	the number of the posts user received in a certain period.
<i>outflowrate</i>	the number of the posts user sent in a certain period.
$inflowrate_{F(i)}^{avg}$	average inflow rate of followers to the user, or $\frac{\sum_i retweet(i) \cdot inflow(i)}{\sum_i retweet(i)}$ where $F(u)$ is the follower to user i (and the same as following).
$retweetrate_{F(u)}^{avg}$	average retweet rate of followers to the user, or $\frac{\sum_i retweet(i) \cdot retweetrate(i)}{\sum_i retweet(i)}$.
Structural features	
<i>follower number</i>	number of the followers to the user.
<i>followee number</i>	number of users this user follows.

have no or very sparse subcascade in training dataset, which makes these users’ behavioral dynamics inaccurate or even unknown; (2) it is difficult to interpret the parameters directly learned from data, which prohibits us from getting insightful understanding of the behavioral dynamics. To address these, we investigate the covariates of behavioral dynamics here. As the behavioral dynamics of a user are to capture the collective responses of his/her followers, we assume the parameters of the user’s behavioral dynamics should be correlated with the behavioral features of his/her followers (network neighbors). Hence, we extract a set of behavioral features for each user as listed in Table II.² For each user with enough subcascades in our dataset, we learned their λ and k directly from data. And then, we calculate the correlations between the learned parameters and their followers’ collective behavioral features. The examples given in Figure 5 indicate obvious correlations between the learned parameters with these behavioral features. Therefore, we can use these behavioral features as covariates to regress the parameters of behavioral dynamics.

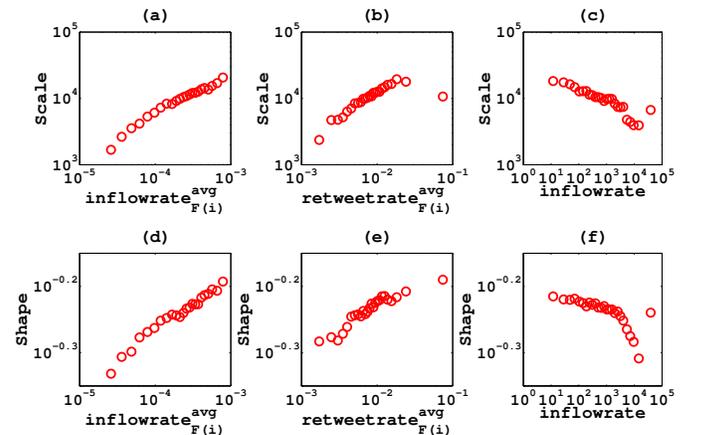


Fig. 5: Correlations between the survival function parameters and the behavioral features

²We think that follower with different retweet number will have different effects to the user, so we modify the weights on each term of *follower_avg_inflow_rate* and *follower_avg_retweet_rate*.

E. From Behavioral Dynamics to Cascades

After validating that the behavioral dynamics can potentially be accurately modeled and predicted, the key problem is whether we can derive the macro cascading process from micro behavioral dynamics. Intuitively, the cascading process cannot be perfectly predicted at early stage by behavioral dynamics. Given any time t , we can only use the behavioral dynamics of the users that involved before t to predict the cascading process after t . Consequently, the prediction coverage is restricted to all the followers of these users, while the users beyond this scope are neglected. These uncovered users may potentially affect the performance of cascading process prediction.

Fortunately, we observe two interesting phenomena in real data.

Minor dominance. Although each user has behavioral dynamics, the behavioral dynamics of different users make significantly different contributions to the cascading process. It is intuitive that the behavioral dynamics of an active user with 1 million followers contribute much more than that of an inactive user with 5 followers. The data also coincides with our intuition. According to Figure 6 (a), it can be observed that a very small number of nodes whose behavioral dynamics dominate the cascading process underpin the idea of just using the behavioral dynamics of these dominant nodes for cascading process prediction.

Early stage dominance. Enlightened by the minor dominance phenomenon, we further ask whether the dominant nodes are prone to join cascades in early stage. Here, Figure 6 (b) depicts the time distribution of these dominant nodes joining in cascades, and we can see that most of these nodes in actual join cascades in the very early stage.

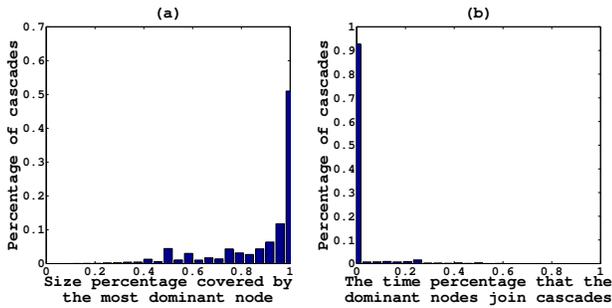


Fig. 6: **Minor dominance and early stage dominance in information cascades.**

Taking these two phenomena into account together, it is safe to design a model exploiting the behavioral dynamics of infected nodes in early stage to predict the cascading process.

IV. METHODOLOGY

This section introduces the NETworked WEibull Regression (NEWER) and cascade prediction methods in detail.

A. Problem Statement

Given a network $G = \langle U, A \rangle$, where U is a collection of nodes and A is the set of pairwise directed/undirected relationships. An event (e.g., tweet) can be originated from one node and spread (e.g., by retweeting) to its neighboring nodes. A cascade is typically formed by repeating this

process. Therefore, a cascade can be represented by a set of nodes $C = \{u_1, u_2, \dots, u_m\}$, where u_1 is the root node. In a cascade, each node will get infected by the event only once, so it is tree-structured. For every node u_i in the cascade, we denote its parent node as $rp(u_i)$. The time stamp that u_i gets infected is $t(u_i)$, and $t(u_i) \leq t(u_{i+1})$. Then the partial cascade before time t is denoted by $C_t = \{u_i | t(u_i) \leq t\}$, and its size $size(C_t) = |C_t|$ where $|\cdot|$ is the cardinality of a set. Then the cascade prediction problem can be defined as:

Cascade Prediction: Given the early stage of a cascade C_t , predict the cascade size $size(C_{t'})$ ($t' \geq t$).

B. Survival Analysis

Survival analysis is a branch of statistics that deals with analysis of time duration until one or more events happen, such as death in biological organisms and failure in mechanical systems [14]. It is a useful technique for cascade prediction. More concretely, let τ_0 be a non-negative continuous random variable representing the waiting time until the occurrence of an event with probability density function $f(t)$, the **survival function**

$$S(t) = Pr\{\tau_0 \geq t\} = \int_t^{\infty} f(t) \quad (1)$$

encodes the probability that the event occurs after t , the **hazard rate** is defined as the event rate at time t conditional on survival until time t or later ($\tau_0 \geq t$), i.e.,

$$\lambda(t) = \lim_{dt \rightarrow 0} \frac{Pr(t \leq \tau_0 < t + dt | \tau_0 \geq t)}{dt} = \frac{f(t)}{S(t)} \quad (2)$$

$S(t)$ and $\lambda(t)$ are the two core quantities in survival analysis.

C. NETworked WEibull Regression Model

The Weibull distribution is commonly used in survival analysis. In network scenario, if we think the time that an event (e.g., retweet) happened on a node as a survival process, we can fit a Weibull distribution to the survival time of node i , then its corresponding density, survival and hazard functions

$$f_i(t) = \frac{k_i}{\lambda_i} \left(\frac{t}{\lambda_i} \right)^{k_i-1} \exp^{-\left(\frac{t}{\lambda_i} \right)^{k_i}} \quad (3)$$

$$S_i(t) = \exp^{-\left(\frac{t}{\lambda_i} \right)^{k_i}} \quad (4)$$

$$h_i(t) = \frac{k_i}{\lambda_i} \left(\frac{t}{\lambda_i} \right)^{k_i-1} \quad (5)$$

where $t > 0$ is the average event happening time to node i , $\lambda_i > 0$ and $k_i > 0$ is the scale and shape parameter of the Weibull distribution. In the following we will assume the network nodes are users and the event is retweeting.

Likelihood of retweeting dynamics. Supposing there are N users in total, T_i is a set of m_i time stamps and each element $T_{i,j}$ indicates the j -th retweet time stamp to the post of the i -th user. We sort those time stamps out in increasing order so that $T_{i,j+1} \geq T_{i,j}$. We assume $T_{i,j} \geq 1$ and $T_{i,m_i} > 1$. Then the likelihood of the event data can be written as follows:

$$L(\lambda, k) = \prod_{i=1}^N \prod_{j=1}^{m_i} (h_i(T_{i,j}) \cdot S_i(T_{i,j}))$$

$$= \prod_{i=1}^N \prod_{j=1}^{m_i} \left(k_i \cdot T_{i,j}^{k_i-1} \cdot \lambda_i^{-k_i} \cdot e^{-T_{i,j} \cdot \lambda_i^{-k_i}} \right) \quad (6)$$

$$\log L(\lambda, k) = \sum_{i=1}^N l_i(\lambda_i, k_i) \quad (7)$$

where $l_i(\lambda_i, k_i) = m_i \log k_i + (k_i - 1) \sum_{j=1}^{m_i} \log T_{i,j} - m_i k_i \log \lambda_i - \lambda_i^{-k_i} \sum_{j=1}^{m_i} T_{i,j}^{-k_i}$.

As discovered in section III-D, the survival characteristics of the user is correlated with the behavioral features of him/her. Then we can parameterize those parameters in the personalized Weibull distributions using those behavioral features. More formally, let x_i be a r dimensional feature vector for user i , we parameterize λ_i and k_i with the following linear function:

$$\log \lambda_i = \log x_i * \beta \quad (8)$$

$$\log k_i = \log x_i * \gamma \quad (9)$$

where β and γ are r -dimensional parameter vector for λ and k . We attempt to find the scale and shape parameter of every user so that the likelihood of the observed data is maximized, at the same time we can also get the parameter vectors for out-of-sample extensions.

We use the Equation (8) and (9) to replace λ_i and k_i in the log likelihood function Equation (7) to solve the parameters. To further enhance the interpretability, we also add ℓ_1 sparsity regularizers on β and γ respectively to enforce model sparsity. Combining everything together, we can obtain the *NETworked WEibull Regression* (NEWER) formulation which aims to minimize the following objective:

$$F(\lambda, k, \beta, \gamma) = G_1(\lambda, k) + \mu G_2(\beta, \lambda) + \eta G_3(\gamma, k) \quad (10)$$

$$G_1(\lambda, k) = -\log L(\lambda, k) \quad (11)$$

$$G_2(\lambda, \beta) = \frac{1}{2N} \|\log \lambda - \log X \cdot \beta\|^2 + \alpha_\beta \|\beta\|_1 \quad (12)$$

$$G_3(k, \gamma) = \frac{1}{2N} \|\log k - \log X \cdot \gamma\|^2 + \alpha_\gamma \|\gamma\|_1 \quad (13)$$

Optimization. To minimize $F(\lambda, k, \beta, \gamma)$ in Equation (10), we first prove that the function is lower bounded. We have the following theorem.

Theorem 1: $F(\lambda, k, \beta, \gamma)$ has global minimum.

Proof: See the appendix. ■

With this theorem, the following coordinate descent strategy can be used to solve the problem with guaranteed convergence. At each iteration, we solve the problem with one group of variables with others fixed.

$$\left\{ \begin{array}{l} \text{For } it = 1, \dots, it_{max} \\ \lambda^{[it+1]} = \underset{\lambda}{\operatorname{argmin}} F(\lambda, k^{[it]}, \beta^{[it]}, \gamma^{[it]}) \\ k^{[it+1]} = \underset{k}{\operatorname{argmin}} F(\lambda^{[it+1]}, k, \beta^{[it]}, \gamma^{[it]}) \\ \beta^{[it+1]} = \underset{\beta}{\operatorname{argmin}} F(\lambda^{[it+1]}, k^{[it+1]}, \beta, \gamma^{[it]}) \\ \gamma^{[it+1]} = \underset{\gamma}{\operatorname{argmin}} F(\lambda^{[it+1]}, k^{[it+1]}, \beta^{[it+1]}, \gamma) \end{array} \right. \quad (14)$$

For solving the subproblem with respect to λ or k , we use Newton's Method. For subproblem with respect to β and γ , we use standard LASSO solver [20].

D. Efficient cascading process prediction

It should be born in mind that cascading prediction is intended to perform early prediction of its size at any later time. In the following we will present two models to achieve this goal.

1) *Basic Model:* The entire flow of the basic model we proposed is illustrated in Algorithm 1:

Algorithm 1 Basic Model

Input:

Set of users U involved in the cascade C before time t_{limit} , survival functions of users $S_{u_j}(t)$, predicting time t_e ;

Output:

Size of cascade $size(C_{t_e})$;

- 1: **for all** user $u_i \in U$ **do**
 - 2: creates a subcascade process with $replynum(u_i) = 0$
 - 3: **if** u_i is not root node **then**
 - 4: $replynum(rp(u_i)) = replynum(rp(u_i)) + 1$
 - 5: **end if**
 - 6: **end for**
 - 7: $sum = 1$
 - 8: **for all** user $u_i \in U$ **do**
 - 9: $deathrate(u_i) = \max\left(1 - S_{u_i}(t_{limit} - t(u_i)), \frac{1}{|V|}\right)$
 - 10: $fdrate(u_i) = \max\left(1 - S_{u_i}(t_e - t(u_i)), \frac{1}{|V|}\right)$
 - 11: $sum = sum + \frac{replynum(u_i) \cdot fdrate(u_i)}{deathrate(u_i)}$
 - 12: **end for**
 - 13: **return** $size(C_{t_e}) = sum$
-

When a new node u_i is added into the cascade at $t(u_i)$, the algorithm will launch a process to estimate the final size of the subcascade that u_i will generate, with temporal size counter $replynum(u_i)$ and survival function $S_{u_i}(t)$ starting at $t(u_i)$. If u_i is involved by others, the algorithm also increases the temporal size of the retweet set of its parent $rp(u_i)$ by one.

After all the information before the deadline is collected, the result will be finalized by aggregating all the value estimated by every subcascade process. Since the post number is at most $|V|$ (all nodes in the network are involved into the cascade), the value of death rate $deathrate(u_i)$ and final death rate $fdrate(u_i)$ (complement to their survival rates) at line 9 and line 10 is set to be $1/|V|$ when it is lower than $1/|V|$.

Complexity Analysis. Only constant time operations is involved in the two for-loops. Therefore, the complexity of the algorithm is $O(n)$ where n is the number of users in the cascade.

2) *Sampling Model:* Although the basic model solves the estimation problem, real applications often need to estimate the cascade size dynamically so that the changes can be monitored.

To make the algorithm scalable, the number of recalculations should be limited, while the estimated value of size should fall into an acceptable error scope. We can utilize the following two facts to make the estimation process more efficient: (1) For a subcascade generated by u_i , the estimation of the size will always be zero if there is no user involved into it, which means we can ignore the calculation. (2) If we

do not re-estimate the final number of a subcascade (when there is no new user involved into it), the temporal size counter $replynum(u_i)$ and final death rate $edrate(u_i)$ will not change but the death rate $deathrate_{u_i}(t)$ will increase over time. Supposing the previous time stamp of the subcascade set estimation is t_0 , it will cause a relative error rate of $\frac{deathrate_{u_i}(t_1)}{deathrate_{u_i}(t_0)} - 1$ at t_1 . Hence, the relative error rate will be at most ϵ if we re-estimate the final number of the subcascade at $S_u^{-1}(1 - (1 + \epsilon) \cdot (deathrate_{u_i}(t_0)))$. By exploring those two tricks, we propose a sampling model shown in Algorithm 2:

Algorithm 2 Sampling Model

Input:

survival functions of users $S_{u_j}(t)$, and set of users U in one cascade C (given dynamically);

Output:

for every size prediction request to t_e at t_0 , output $size(C_{t_e})$;

```

1:  $sum = 0$ ;
2: while request = model.acceptRequest do
3:   switch (request.type)
4:     case APPROXIMATION:
5:       return  $size(C_{t_e}) = sum$ 
6:     case INVOLVED_USER:
7:        $u_i = \text{request.user}$ ,  $t_0 = \text{request.time}$ 
8:       creates a subcascade process:
9:          $t(u_i) = t_0$ ,  $app(u_i) = 0$ ,  $replynum(u_i) = 0$ ,
10:         $fdrate(u_i) = \max\left(\frac{1}{|V|}, 1 - S_{rp(u_i)}(t_e - t_0)\right)$ ;
11:       if  $u_i$  is root node then
12:          $sum = 1$ ;
13:       else
14:          $t_{rep} = t_0 - t(rp(u_i))$ ;
15:          $replynum(rp(u_i)) = replynum(rp(u_i)) + 1$ ;
16:          $sum = sum - app(rp(u_i))$ ;
17:          $deathrate(rp(u_i)) = \max\left(\frac{1}{|V|}, 1 - S_{rp(u_i)}(t_{rep})\right)$ ;
18:          $app(rp(u_i)) = \frac{replynum(rp(u_i)) \cdot fdrate(rp(u_i))}{deathrate(rp(u_i))}$ ;
19:          $sum = sum + app(rp(u_i))$ ;
20:          $t_{new} = S_{rp(u_i)}^{-1}(1 - (1 + \epsilon) \cdot deathrate(rp(u_i))) + t(rp(u_i))$ ;
21:         sendRequest(THRESHOLD_CHANGE,  $rp(u_i)$ ,  $t_{new}$ );
22:       end if
23:     case THRESHOLD_CHANGE:
24:        $u_i = \text{request.user}$ ,  $t_0 = \text{request.time}$ 
25:        $sum = sum - app(u_i)$ ;
26:        $deathrate(u_i) = \max\left(\frac{1}{|V|}, 1 - S_{u_i}(t_0 - t(u_i))\right)$ ;
27:        $t_{new} = S_{u_i}^{-1}(1 - (1 + \epsilon) \cdot deathrate(u_i)) + t(u_i)$ ;
28:       sendRequest(THRESHOLD_CHANGE,  $u_i$ ,  $t_{new}$ );
29:        $app(u_i) = \frac{replynum(u_i) \cdot fdrate(u_i)}{deathrate(u_i)}$ ;
30:        $sum = sum + app(u_i)$ ;
31:     end switch
32: end while

```

Complexity Analysis. The following theorem analyzes the complexity of Algorithm 2.

Theorem 2: With an overall $O(n \log_{1+\epsilon}(|V|))$ counting to estimate the number of subcascades, the sampling model can approximate the size of the cascade estimated by the basic model at any time with an relative error rate of at most ϵ .

Proof: For each approximation request, we only need to report the number directly; for every new subcascade, the initially operation number is also constant, and we need to do at most $O(\log_{1+\epsilon}(|V|))$ times threshold adjustment for

subcascade which has users involved in, since the lowerbound of $deathrate$ is $\frac{1}{|V|}$ and the upperbound is 1(all the people are involved in the cascade). Above all, the final complexity is $O(t) + O(n) + O(n \log_{1+\epsilon}(|V|)) = O(t + n \log_{1+\epsilon}(|V|))$ for each cascade (with n users and t requests). If we put this algorithm into an online environment, the complexity will be $O(T + N \log_{1+\epsilon}(|V|)) \sim O(T)$ for all the cascades with N Users in total³ (we see $\log_{1+\epsilon}(|V|)$ as a constant with respect to T and $N \sim T$ as the number of users involved in cascades increases over time). ■

With this model, for cascade final size prediction, we just need to set the prediction time t_e to be infinite so that the $deathrate$ of all subcascades will be 1. For outbreak time prediction, we can make a binary search with respect to time t_e , checking whether the cascade size will be more or less than the size number at t_{mid} and make the decision eachtime. We can also solve the process prediction problem by asking the outbreak time of size restricted from the current size to the final size of the cascade one by one. In all cases, the sampling model will be n times faster than the base model.

V. EXPERIMENTS

In order to evaluate the performances and fully demonstrate the advantages of the proposed method, we conduct a series of experiments on the dataset introduced in Section III-A. The results of multiple tasks are reported, including cascade size prediction, outbreak time prediction and cascading process prediction. Also, it will address more insights about the proposed method.

A. Baselines and Evaluation Metrics

Since we are the first to investigate cascading process prediction problem, no previous models can be adopted as direct baselines. Here, we implemented the following methods which can be potentially applied into our targeted problem as baselines:

- Cox Proportional Hazard Regression Model (Cox) : This model assumes that the behavioral dynamics of all users have different scale parameters while sharing the same shape parameter. We use the same covariates as in our model and find the optimal scale parameters for all users and the shared shape parameter. We implement it as in [5].
- Exponential/Rayleigh Proportional Hazard Regression Model (Exponential/Rayleigh): Since the shape parameters of both Exponential and Rayleigh distributions are fixed values (1 for Exponential distribution and 2 for Rayleigh distribution), they are two special cases of Cox model.
- Log-linear Regression Model (Log-linear): We refer to [3] which extracted 4 classes features to characterize cascades, including node features, structural features of cascades, temporal features and content features. In our case, we ignore the content features which are not covered in our dataset and also reported by [3] to be unimportant for cascade prediction. Then we use log-linear regression model to predict the cascade size.

³It will be counted multiple times if a specific user involves in multiple cascades

It is noted that Log-linear can only predict cascade size but not for time-related prediction, while Cox, Exponential and Rayleigh models are applied to all prediction tasks. Also, the goal of Cox, Exponential and Rayleigh models are to elucidate the behavioral dynamics. After that, we use the same cascade prediction model as in our method to conduct cascade-level predictions.

For each cascade, our dataset includes its complete cascading process as the groundtruth. Next, we use the following metrics to evaluate the performances:

- **Root Mean Square Log Error (*RMSLE*):** In Power-Law distributed data, it is not reasonable to use standard RMSE to evaluate the prediction accuracy. For example, for a cascade with the groundtruth size of 1000, it is significantly different to predict its size to be 2000 or 0, but they have the same RMSE. Thus, we first calculate the logarithmic results for both the groundtruth and predicted value, then calculate RMSE on the logarithmic results to evaluate the accuracy of the proposed method and baselines.
- **Precision with σ -Tolerance ($\Delta\sigma$ -Precision):** In real applications, a small deviation from the groundtruth value is often acceptable. In our case, we regard the predicted value within the range of $(1 + \sigma)^{\pm 1} \text{groundtruth}$ as tolerable, and the resulted precision is $\Delta\sigma$ -Precision.

For parameter setting, there are 4 parameters in our method, including μ, η, α_β and α_γ . We tune these parameters by grid searching, and the optimal parameters used in our experiments are $\mu = 10, \eta = 10, \alpha_\beta = 6 * 10^{-5}, \alpha_\gamma = 8 * 10^{-6}$.

B. Cascade Size Prediction

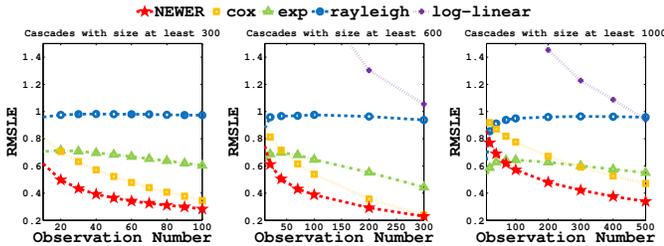


Fig. 7: *RMSLE* results of different methods with different number of observed nodes in cascades. We did not show log-linear results in the first sub-figure since their *RMSLE* value are all greater than 1.5.

We randomly separate the cascades into 10 folds, and conduct a 10-fold cross validation by using 9 of them as training data and the other one as testing data. For cascades with size over k , we use the first $s (s < k)$ nodes as observed data, and the target is to predict the final cascade sizes.

The prediction performances of all the methods are shown in Figure 7. It can be seen that the proposed method NEWER significantly outperforms other baselines in *RMSLE* value in different sized datasets. The baselines that has the closest performance with NEWER is the Cox model. We can see that the margins of improvement from Cox to NEWER are more obvious in the dataset with larger k . In a certain dataset, the margins are more evident with smaller s . These results

TABLE III: **Running time for different methods in different dataset under a server with 3.4GHZ Quad Core Intel i7-3770 CPU and 16GB memory.**

Method	Base Model	Improved Model ($\delta = 0.1$)	Directed Learning Method
Size ≥ 20	$8.47 * 10^5 s$	10.73s	899s
Size ≥ 50	$7.61 * 10^5 s$	8.62s	899s
Size ≥ 100	$6.65 * 10^5 s$	7.09s	898s
Size ≥ 500	$4.35 * 10^5 s$	4.33s	891s
Size ≥ 1000	$3.4 * 10^5 s$	3.30s	881s

demonstrate the significant advantage of NEWER in predicting large cascades in very early stage.

Comparatively, the Log-linear method does not achieve satisfactory results in this task. The main reason is that the coefficients in the Log-linear model are highly biased towards the dominant number of small-sized cascades, which is also argued by [3]. In our method, we successfully overcome this bias by shifting from macro cascade level features to micro behavioral dynamics. The substantial gain achieved by all behavioral dynamics based methods (including NEWER, Cox, Exponential and Rayleigh) exemplifies the importance of this micro mechanism for cascade prediction.

In order to demonstrate the efficiency of the proposed method, we also evaluate the computational cost of NEWER and Sampling-NEWER in the computational environment with 3.4GHZ Quad Core Intel i7-3770 and 16GB memory. We track the process of all cascades. The base cascade prediction model (Base) re-predicts the final size at every time points (in second), while the sampling-based cascade prediction model (Sampling) re-predicts the final size only when the observed cascade sizes increase. As shown in Table III, the Sampling model (with a 10 percent performance degradation tolerance) is much more efficient than Base model by almost 5 magnitudes. According to Section IV-D, it is guaranteed that the Sampling method can also improve with similar magnitudes than the Base model in cascading process prediction task. So we omit these results for brevity.

C. Outbreak Time Prediction

Another interesting problem is to predict when a cascading outbreak will happen. For example, in the early stage of a cascade, can we predict when the cascade reaches a specific size? Without loss of generality, we set the outbreak size threshold to be 1000. We evaluate the prediction performance with different number of observed nodes in the cascades. As shown in Figure 8, the NEWER model get the best performances in both *RMSLE* and $\Delta\sigma$ -Precision metrics. Although Exponential and Rayleigh models report better results than NEWER in very early stage (less than 50 observation nodes), the improvements of their performances with increasing number of observed nodes are not as significant as NEWER.

D. Cascading Process Prediction

The ultimate purpose of this paper is to predict the cascading process. For each cascade, we use Δt to represent the early stage window and \hat{t} to represent its ending time. Then we use the cascade information during $[0, \Delta t]$ to predict the cascading process during $[\Delta t, \hat{t}]$. At any time $t \in [\Delta t, \hat{t}]$, we

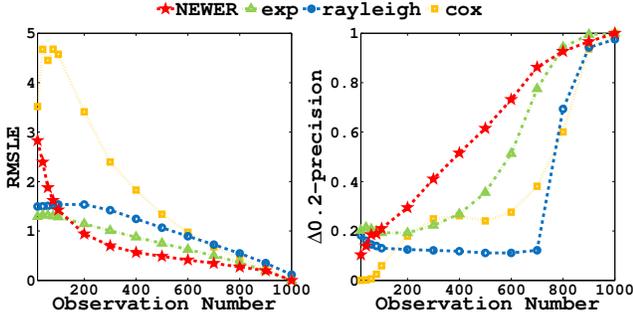


Fig. 8: Outbreak time prediction results of different methods with different number of observed nodes in cascades.

calculate whether the predicted cascade size at t is within the σ tolerance of the groundtruth size at t . Then we calculate the $\Delta\sigma$ -Precision by integrating t to describe the prediction accuracy for this cascading process. Finally, we average the $\Delta\sigma$ -Precision for all cascades and show the results in Figure 9. Here, we vary the early stage percentage (i.e. $\Delta t/\hat{t}$) from 0 to 50%, and discover that in all the settings of early stage percentage, NEWER always carries out the best performances in cascading process prediction. More over, the advantage of NEWER is more clear in smaller early stage percentage. When we set the early stage to be 15% of the whole cascade duration, we can get the $\Delta 0.2$ -Precision of 0.849. That means that we can correctly predict the cascade sizes at 84.9% time points, which indicates that the cascading process is predictable and the proposed method is adequate and superior in cascading process prediction. Furthermore, changing the precision tolerance value σ will not affect the relative results of all the methods in our experiments, and the precision value will be smaller when setting σ smaller. For abbreviation, we only report the results of $\sigma = 0.2$, which is a reasonable tolerance in most application scenarios.

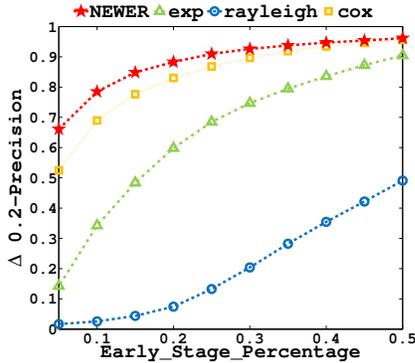


Fig. 9: Cascading process prediction accuracy of different methods under different early stage percentage settings.

E. Out-of-sample Prediction

In real applications, the interaction information between nodes is not always available, which makes some nodes' behavioral dynamics cannot be directly derived by maximum likelihood estimation from data. We call these nodes as *out-of-sample nodes*. This is the main reason why we propose NEWER to incorporate the covariates of behavioral dynamics. In order to evaluate the performance of NEWER in handling this case, we simulate the scenario by hiding the interaction information of randomly selected 10% users as out-of-sample

users, and then predict the final sizes of the cascades that these users involved in early stages.

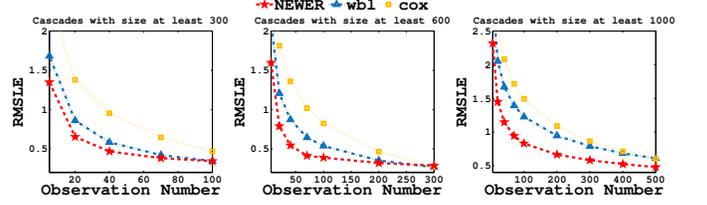


Fig. 10: Prediction result by unknown users.

In Cox model, the scale parameters in behavioral dynamics of out-of-sample users can be regressed by the covariates. For the shape parameter, we calculate the average value of shape parameters in observed users and apply this value to the shape parameters of out-of-sample users. In NEWER model, both of shape and scale parameters can be regressed by covariates with the learned β and γ . We also employ the standard Weibull Regression (Wbl) as a baseline, which can be derived by simply setting μ and η to be 0 in Equation 10. Then we use the averaged shape and scale parameters of observed users as the parameters of out-of-sample users.

As shown in Figure 10, the NEWER model can significantly and consistently outperform Cox and Wbl models in out-of-sample prediction, which demonstrates that the discovered covariates from behavioral features of a user's networked neighbors can effectively predict the user's behavioral dynamics. Also, we visualize the regression coefficients β and γ in Figure 11. It can be observed that the behavioral features of a user's followers plays more important roles in predicting both scale and shape parameters for the user, while the user's structural features are less important.

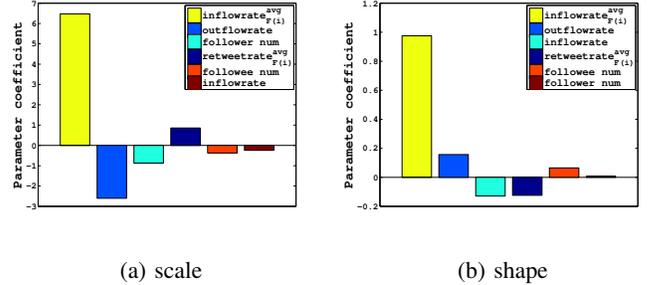


Fig. 11: Parameter coefficients for scale and shape parameter.

VI. CONCLUSIONS

In this paper, we raise an important and interesting question: beyond predicting the final size of a cascade, can we predict the whole cascading process if the early stage information of cascades is given? In order to address this problem, we propose to uncover and predict the macro cascading process with micro behavioral dynamics. Through data-driven analysis, we find out the common principles and important patterns laying in behavioral dynamics, and propose a novel NEWER model for behavioral dynamics modeling with good interpretability and generality. After that, we propose a scalable method to aggregate micro behavioral dynamics into macro cascading processes. Extensive experiments on a large scale

real data set demonstrate that the proposed method achieves the best results in various cascading prediction tasks, including cascade size prediction, outbreak time prediction and cascading process prediction.

ACKNOWLEDGMENT

This work is supported by the National Basic Research Program of China, No. 2015CB352300; National Natural Science Foundation of China, No. 61370022 and No. 61210008; International Science and Technology Cooperation Program of China, No. 2013DFG12870. Thanks for the support of NExT Research Center funded by MDA, Singapore, under the research grant, WBS:R-252-300-001-490 and the research fund of Tsinghua-Tencent Joint Laboratory for Internet Innovation Technology.

APPENDIX: PROOF OF THEOREM 1

Proof: It's evident that both $G_2(\beta, \lambda)$ and $G_3(\gamma, k)$ has global minimum value. Next we prove that $G_1(\lambda, k)$ also has global minimum value, or to prove $\log L(\lambda, k)$ has global maximum value.

Let $\lambda'_i = \lambda_i^{-k_i}$, $\log L'(\lambda', k) = \log L(\lambda, k) = \sum_{i=1}^N l'_i(\lambda'_i, k_i)$ where $l'_i(\lambda'_i, k_i) = m_i \log k_i + (k_i - 1) \sum_{j=1}^{m_i} \log T_{i,j} + m_i \log \lambda'_i - \lambda'_i \sum_{j=1}^{m_i} T_{i,j}^{k_i}$, the partial derivatives of the l'_i are given by:

$$\frac{\partial l'_i}{\partial \lambda'_i} = \frac{m_i}{\lambda'_i} - \sum_{j=1}^{m_i} T_{i,j}^{k_i}, \quad \frac{\partial^2 l'_i}{\partial \lambda'^2_i} = -\frac{m_i}{\lambda'^2_i} < 0 \quad (15)$$

$$\frac{\partial l'_i}{\partial k_i} = \frac{m_i}{k_i} + \sum_{j=1}^{m_i} \log T_{i,j} - \lambda'_i \sum_{j=1}^{m_i} T_{i,j}^{k_i} \log T_{i,j} \quad (16)$$

$$\frac{\partial^2 l'_i}{\partial k_i^2} = -\frac{m_i}{k_i^2} - \lambda'_i \sum_{j=1}^{m_i} T_{i,j}^{k_i} (\log T_{i,j})^2 < 0 \quad (17)$$

Since $\frac{\partial^2 l'_i}{\partial \lambda'^2_i} < 0$ and $\frac{\partial^2 l'_i}{\partial k_i^2} < 0$, the conditional marginal posterior densities of parameters λ'_i and k_i are log-concave. Moreover, when

$$0 < k_i < 1, 0 < \lambda'_i < \min \left(\frac{m_i}{\sum_{i=1}^{m_i} T_{i,j}^{k_i}}, \frac{1}{\sum_{j=1}^{m_i} T_{i,j} \log T_{i,j}} \right),$$

$$\frac{\partial l'_i}{\partial \lambda'_i} = \frac{m_i}{\lambda'_i} - \sum_{j=1}^{m_i} T_{i,j}^{k_i} \geq \frac{m_i}{\sum_{i=1}^{m_i} T_{i,j}^{k_i}} - \sum_{j=1}^{m_i} m_i T_{i,j}^{k_i} = 0 \quad (18)$$

$$\frac{\partial l'_i}{\partial k_i} = \frac{m_i}{k_i} + \sum_{j=1}^{m_i} \log T_{i,j} - \lambda'_i \sum_{j=1}^{m_i} T_{i,j}^{k_i} \log T_{i,j}$$

$$\geq \frac{m_i}{k_i} + \sum_{j=1}^{m_i} \log T_{i,j} - \frac{\sum_{j=1}^{m_i} T_{i,j}^{k_i} \log T_{i,j}}{\sum_{j=1}^{m_i} T_{i,j} \log T_{i,j}} \geq m_i + \sum_{j=1}^{m_i} \log T_{i,j} - 1 > 0$$

when $k_i \geq \max \left(1, \frac{m_i}{\lambda'_i \sum_{j=1}^{m_i} T_{i,j} \log T_{i,j} - \sum_{j=1}^{m_i} \log T_{i,j}} \right)$ and $\lambda'_i \geq$

$$\max \left(1, \frac{m_i}{\sum_{j=1}^{m_i} \frac{T_{i,j}^{k_i}}{m_i^{k_i}}} \right),$$

$$\frac{\partial l'_i}{\partial \lambda'_i} = \frac{m_i}{\lambda'_i} - \sum_{j=1}^{m_i} T_{i,j}^{k_i} \leq \frac{m_i}{\sum_{i=1}^{m_i} \frac{T_{i,j}^{k_i}}{m_i^{k_i}}} - \sum_{j=1}^{m_i} m_i T_{i,j}^{k_i} = 0 \quad (19)$$

$$\frac{\partial l'_i}{\partial k_i} = \frac{m_i}{k_i} + \sum_{j=1}^{m_i} \log T_{i,j} - \lambda'_i \sum_{j=1}^{m_i} T_{i,j}^{k_i} \log T_{i,j}$$

$$\leq \frac{m_i}{\lambda'_i \sum_{j=1}^{m_i} T_{i,j} \log T_{i,j} - \sum_{j=1}^{m_i} \log T_{i,j}} + \sum_{j=1}^{m_i} \log T_{i,j} - \lambda'_i \sum_{j=1}^{m_i} T_{i,j}^{k_i} \log T_{i,j}$$

$$= \lambda'_i \sum_{j=1}^{m_i} T_{i,j} \log T_{i,j} - \sum_{j=1}^{m_i} \log T_{i,j} + \sum_{j=1}^{m_i} \log T_{i,j} - \lambda'_i \sum_{j=1}^{m_i} T_{i,j}^{k_i} \log T_{i,j} \quad (20)$$

$$< 0$$

which means there should be a global maximum of l'_i , so does $\log L$. ■

REFERENCES

- [1] W. Chen, A. Collins, R. Cummings, T. Ke, Z. Liu, D. Rincon, X. Sun, Y. Wang, W. Wei, and Y. Yuan. Influence maximization in social networks when negative opinions may emerge and propagate. In *SDM*, volume 11, pages 379–390. SIAM, 2011.
- [2] W. Chen, Y. Wang, and S. Yang. Efficient influence maximization in social networks. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 199–208. ACM, 2009.
- [3] J. Cheng, L. Adamic, P. A. Dow, J. M. Kleinberg, and J. Leskovec. Can cascades be predicted? In *Proceedings of the 23rd international conference on World wide web*, pages 925–936. International World Wide Web Conferences Steering Committee, 2014.
- [4] R. Cohen, S. Havlin, and D. Ben-Avraham. Efficient immunization strategies for computer networks and populations. *Physical review letters*, 91(24):247901, 2003.
- [5] D. Cox. Regression models and life-tables. *Journal of the Royal Statistical Society. Series B (Methodological)*, 34(2):187–220, 1972.
- [6] P. Cui, S. Jin, L. Yu, F. Wang, W. Zhu, and S. Yang. Cascading outbreak prediction in networks: a data-driven approach. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 901–909. ACM, 2013.
- [7] P. Cui, F. Wang, S. Liu, M. Ou, S. Yang, and L. Sun. Who should share what?: item-level social influence prediction for users and posts ranking. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*, pages 185–194. ACM, 2011.
- [8] P. Domingos and M. Richardson. Mining the network value of customers. In *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 57–66. ACM, 2001.
- [9] N. Du, L. Song, H. Woo, and H. Zha. Uncover topic-sensitive information diffusion networks. In *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics*, pages 229–237, 2013.
- [10] A. Gionis, E. Terzi, and P. Tsaparas. Opinion maximization in social networks. *arXiv preprint arXiv:1301.7455*, 2013.
- [11] M. Gomez-Rodriguez, J. Leskovec, and B. Schölkopf. Modeling information propagation with survival theory. In *ICML*, 2013.
- [12] M. Gomez Rodriguez, J. Leskovec, and B. Schölkopf. Structure and dynamics of information pathways in online media. In *Proceedings of the sixth ACM international conference on Web search and data mining*, pages 23–32. ACM, 2013.
- [13] D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146. ACM, 2003.
- [14] R. G. Miller Jr. *Survival analysis*, volume 66. John Wiley & Sons, 2011.
- [15] S. Myers and J. Leskovec. On the convexity of latent social network inference. In *Advances in Neural Information Processing Systems*, pages 1741–1749, 2010.
- [16] J. E. Pinder III, J. G. Wiener, and M. H. Smith. The weibull distribution: a new method of summarizing survivorship data. *Ecology*, pages 175–179, 1978.
- [17] G. Rodriguez. Parametric survival models. *Lectures Notes, Princeton University*, 2005.
- [18] M. G. Rodriguez, D. Balduzzi, and B. Schölkopf. Uncovering the temporal dynamics of diffusion networks. *arXiv preprint arXiv:1105.0697*, 2011.
- [19] M. G. Rodriguez and B. Schölkopf. Influence maximization in continuous time diffusion networks. *arXiv preprint arXiv:1205.1682*, 2012.
- [20] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.