
COEVOLVE: A Joint Point Process Model for Information Diffusion and Network Co-evolution

Mehrdad Farajtabar* Yichen Wang* Manuel Gomez-Rodriguez†
Shuang Li* Hongyuan Zha* Le Song*
Georgia Institute of Technology* MPI for Software Systems†
{mehrdad, yichen.wang, sli370}@gatech.edu manuelgr@mpi-sws.org
{zha, lsong}@cc.gatech.edu

Abstract

Information diffusion in online social networks is affected by the underlying network topology, but it also has the power to change it. Online users are constantly creating new links when exposed to new information sources, and in turn these links are alternating the way information spreads. However, these two highly intertwined stochastic processes, information diffusion and network evolution, have been predominantly studied *separately*, ignoring their co-evolutionary dynamics.

We propose a temporal point process model, COEVOLVE, for such joint dynamics, allowing the intensity of one process to be modulated by that of the other. This model allows us to efficiently simulate interleaved diffusion and network events, and generate traces obeying common diffusion and network patterns observed in real-world networks. Furthermore, we also develop a convex optimization framework to learn the parameters of the model from historical diffusion and network evolution traces. We experimented with both synthetic data and data gathered from Twitter, and show that our model provides a good fit to the data as well as more accurate predictions than alternatives.

1 Introduction

Online social networks, such as Twitter or Weibo, have become large information networks where people share, discuss and search for information of personal interest as well as breaking news [1]. In this context, users often forward to their *followers* information they are exposed to via their *followees*, triggering the emergence of information *cascades* that travel through the network [2], and constantly create new links to information sources, triggering changes in the network itself over time. Importantly, recent empirical studies with Twitter data have shown that both information diffusion and network evolution are coupled and network changes are often triggered by information diffusion [3, 4, 5].

While there have been many recent works on modeling information diffusion [2, 6, 7, 8] and network evolution [9, 10, 11], most of them treat these two stochastic processes independently and separately, ignoring the influence one may have on the other over time. Thus, to better understand information diffusion and network evolution, there is an urgent need for joint probabilistic models of the two processes, which are largely inexistent to date.

In this paper, we propose a probabilistic generative model, COEVOLVE, for the joint dynamics of information diffusion and network evolution. Our model is based on the framework of temporal point processes, which explicitly characterize the continuous time interval between events, and it consists of two interwoven and interdependent components (refer to Appendix B for an illustration):

- I. Information diffusion process.** We design an “identity revealing” multivariate Hawkes process [12] to capture the mutual excitation behavior of retweeting events, where the intensity of such events in a user is boosted by previous events from her time-varying set of followees. AI-

though Hawkes processes have been used for information diffusion before [13, 14, 15, 16, 17, 18, 19], the key innovation of our approach is to explicitly model the excitation due to a particular source node, hence revealing the identity of the source. Such design reflects the reality that information sources are explicitly acknowledged, and it also allows a particular information source to acquire new links in a rate according to her “informativeness”.

- II. **Network evolution process.** We model link creation as an “information driven” survival process, and couple the intensity of this process with retweeting events. Although survival processes have been used for link creation before [20, 21], the key innovation in our model is to incorporate retweeting events as the driving force for such processes. Since our model has captured the source identity of each retweeting event, new links will be targeted toward the information sources, with an intensity proportional to their degree of excitation and each source’s influence.

Our model is designed in such a way that it allows the two processes, information diffusion and network evolution, unfold simultaneously in the same time scale and excise bidirectional influence on each other, allowing sophisticated coevolutionary dynamics to be generated (*e.g.*, see Figure 5). Importantly, the flexibility of our model does not prevent us from efficiently simulating diffusion and link events from the model and learning its parameters from real world data:

- **Efficient simulation.** We design a scalable sampling procedure that exploits the sparsity of the generated networks. Its complexity is $O(nd \log m)$, where n is the number of samples, m is the number of nodes and d is the maximum number of followees per user.
- **Convex parameters learning.** We show that the model parameters that maximize the joint likelihood of observed diffusion and link creation events can be found via convex optimization.

Finally, we experimentally verify that our model can produce coevolutionary dynamics of information diffusion and network evolution, and generate retweet and link events that obey common information diffusion patterns (*e.g.*, cascade structure, size and depth), static network patterns (*e.g.*, node degree) and temporal network patterns (*e.g.*, shrinking diameter) described in related literature [22, 10, 23]. Furthermore, we show that, by modeling the coevolutionary dynamics, our model provide significantly more accurate link and diffusion event predictions than alternatives in large scale Twitter dataset [3].

2 Backgrounds on Temporal Point Processes

A temporal point process is a random process whose realization consists of a list of discrete events localized in time, $\{t_i\}$ with $t_i \in \mathbb{R}^+$ and $i \in \mathbb{Z}^+$. Many different types of data produced in online social networks can be represented as temporal point processes, such as the times of retweets and link creations. A temporal point process can be equivalently represented as a counting process, $N(t)$, which records the number of events before time t . Let the history $\mathcal{H}(t)$ be the list of times of events $\{t_1, t_2, \dots, t_n\}$ up to but not including time t . Then, the number of observed events in a small time window dt between $[t, t+dt)$ is $dN(t) = \sum_{t_i \in \mathcal{H}(t)} \delta(t-t_i) dt$, and hence $N(t) = \int_0^t dN(s)$, where $\delta(t)$ is a Dirac delta function. More generally, given a function $f(t)$, we can define the convolution with respect to $dN(t)$ as

$$f(t) \star dN(t) := \int_0^t f(t-\tau) dN(\tau) = \sum_{t_i \in \mathcal{H}(t)} f(t-t_i). \quad (1)$$

The point process representation of temporal data is fundamentally different from the discrete time representation typically used in social network analysis. It directly models the time interval between events as random variables, and avoid the need to pick a time window to aggregate events. It allows temporal events to be modeled in a more fine grained fashion, and has a remarkably rich theoretical support [24].

An important way to characterize temporal point processes is via the conditional intensity function — a stochastic model for the time of the next event given all the times of previous events. Formally, the conditional intensity function $\lambda^*(t)$ (intensity, for short) is the conditional probability of observing an event in a small window $[t, t+dt)$ given the history $\mathcal{H}(t)$, *i.e.*,

$$\lambda^*(t)dt := \mathbb{P} \{ \text{event in } [t, t+dt) | \mathcal{H}(t) \} = \mathbb{E}[dN(t) | \mathcal{H}(t)], \quad (2)$$

where one typically assumes that only one event can happen in a small window of size dt , *i.e.*, $dN(t) \in \{0, 1\}$. Then, given a time $t' \geq t$, we can also characterize the conditional probability that no event happens during $[t, t')$ and the conditional density that an event occurs at time t'

as $S^*(t') = \exp(-\int_t^{t'} \lambda^*(\tau) d\tau)$ and $f^*(t') = \lambda^*(t') S^*(t')$ respectively [24]. Furthermore, we can express the log-likelihood of a list of events $\{t_1, t_2, \dots, t_n\}$ in an observation window $[0, T]$ as

$$\mathcal{L} = \sum_{i=1}^n \log \lambda^*(t_i) - \int_0^T \lambda^*(\tau) d\tau, \quad T \geq t_n. \quad (3)$$

This simple log-likelihood will later enable us to learn the parameters of our model from observed data.

Finally, the functional form of the intensity $\lambda^*(t)$ is often designed to capture the phenomena of interests. Some useful functional forms we will use later are [24]:

(i) Poisson process. The intensity is assumed to be independent of the history $\mathcal{H}(t)$, but it can be a time-varying function, *i.e.*, $\lambda^*(t) = g(t) \geq 0$;

(ii) Hawkes Process. The intensity models a mutual excitation between events, *i.e.*,

$$\lambda^*(t) = \mu + \alpha \kappa_\omega(t) \star dN(t) = \mu + \alpha \sum_{t_i \in \mathcal{H}(t)} \kappa_\omega(t - t_i), \quad (4)$$

where $\kappa_\omega(t) := \exp(-\omega t) \mathbb{I}[t \geq 0]$ is an exponential triggering kernel, $\mu \geq 0$ is a baseline intensity independent of the history. Here, the occurrence of each historical event increases the intensity by a certain amount determined by the kernel and the weight $\alpha \geq 0$, making the intensity history dependent and a stochastic process by itself. We will focus on the exponential kernel in this paper. However, other functional forms for the triggering kernel, such as log-logistic function, are possible, and our model does not depend on this particular choice; and,

(iii) Survival process. There is only one event for an instantiation of the process, *i.e.*,

$$\lambda^*(t) = g^*(t)(1 - N(t)), \quad (5)$$

where $\lambda^*(t)$ becomes 0 if an event already happened before t .

3 Generative Model of Information Diffusion and Network Co-evolution

In this section, we use the above background on temporal point processes to formulate our probabilistic generative model for the joint dynamics of information diffusion and network evolution.

3.1 Event Representation

We model the generation of two types of events: tweet/retweet events, e^r , and link creation events, e^l . Instead of just the time t , we record each event as a triplet

$$e^r \text{ or } e^l := \begin{matrix} & \text{source} \\ & \downarrow \\ (u, & s, & t) \\ \uparrow & & \uparrow \\ \text{destination} & & \text{time} \end{matrix}. \quad (6)$$

For retweet event, the triplet means that the destination node u retweets at time t a tweet originally posted by source node s . Recording the source node s reflects the real world scenario that information sources are explicitly acknowledged. Note that the occurrence of event e^r does *not* mean that u is directly retweeting from or is connected to s . This event can happen when u is retweeting a message by another node u' where the original information source s is acknowledged. Node u will pass on the same source acknowledgement to its followers (*e.g.*, “I agree @a @b @c @s”). Original tweets posted by node u are allowed in this notation. In this case, the event will simply be $e^r = (u, u, t)$. Given a list of retweet events up to but not including time t , the history $\mathcal{H}_{us}^r(t)$ of retweets by u due to source s is $\mathcal{H}_{us}^r(t) = \{e_i^r = (u_i, s_i, t_i) | u_i = u \text{ and } s_i = s\}$. The entire history of retweet events is denoted as $\mathcal{H}^r(t) := \cup_{u,s \in [m]} \mathcal{H}_{us}^r(t)$.

For link creation event, the triplet means that destination node u creates at time t a link to source node s , *i.e.*, from time t on, node u starts following node s . To ease the exposition, we restrict ourselves to the case where links cannot be deleted and thus each (directed) link is created only once. However, our model can be easily augmented to consider multiple link creations and deletions per node pair, as discussed in Section 8. We denote the link creation history as $\mathcal{H}^l(t)$.

3.2 Joint Model with Two Interwoven Components

Given m users, we use two sets of counting processes to record the generated events, one for information diffusion and the other for network evolution. More specifically,

- I. Retweet events are recorded using a matrix $\mathbf{N}(t)$ of size $m \times m$ for each fixed time point t . The (u, s) -th entry in the matrix, $N_{us}(t) \in \{0\} \cup \mathbb{Z}^+$, counts the number of retweets of u due to source s up to time t . These counting processes are “identity revealing”, since they keep track of the source node that triggers each retweet. This matrix $\mathbf{N}(t)$ can be dense, since $N_{us}(t)$ can be nonzero even when node u does not directly follow s . We also let $d\mathbf{N}(t) := (dN_{us}(t))_{u,s \in [m]}$.
- II. Link events are recorded using an adjacency matrix $\mathbf{A}(t)$ of size $m \times m$ for each fixed time point t . The (u, s) -th entry in the matrix, $A_{us}(t) \in \{0, 1\}$, indicates whether u is directly following s . That is $A_{us}(t) = 1$ means the directed link has been created before t . For simplicity of exposition, we do not allow self-links. The matrix $\mathbf{A}(t)$ is typically sparse, but the number of nonzero entries can change over time. We also define $d\mathbf{A}(t) := (dA_{us}(t))_{u,s \in [m]}$.

Then the interwoven information diffusion and network evolution processes can be characterized using their respective intensities $\mathbb{E}[d\mathbf{N}(t) | \mathcal{H}^r(t) \cup \mathcal{H}^l(t)] = \mathbf{\Gamma}^*(t) dt$ and $\mathbb{E}[d\mathbf{A}(t) | \mathcal{H}^r(t) \cup \mathcal{H}^l(t)] = \mathbf{\Lambda}^*(t) dt$, where $\mathbf{\Gamma}^*(t) = (\gamma_{us}^*(t))_{u,s \in [m]}$ and $\mathbf{\Lambda}^*(t) = (\lambda_{us}^*(t))_{u,s \in [m]}$. The sign $*$ means that the intensity matrices will depend on the joint history, $\mathcal{H}^r(t) \cup \mathcal{H}^l(t)$, and hence their evolution will be coupled. By this coupling, we make: (i) the counting processes for link creation to be “information driven” and (ii) the evolution of the linking structure to change the information diffusion process. Refer to Appendix B for an illustration of our joint model. In the next two sections, we will specify the details of these two intensity matrices.

3.3 Information Diffusion Process

We model the intensity, $\mathbf{\Gamma}^*(t)$, for retweeting events using multivariate Hawkes process [12]:

$$\gamma_{us}^*(t) = \mathbb{I}[u = s] \eta_u + \mathbb{I}[u \neq s] \beta_s \sum_{v \in \mathcal{F}_u(t)} \kappa_{\omega_1}(t) \star (A_{uv}(t) dN_{vs}(t)), \quad (7)$$

where $\mathbb{I}[\cdot]$ is the indicator function and $\mathcal{F}_u(t) := \{v \in [m] : A_{uv}(t) = 1\}$ is the current set of followers of u . The term $\eta_u \geq 0$ is the intensity of original tweets by a user u on his own initiative, becoming the source of a cascade and the term $\beta_s \sum_{v \in \mathcal{F}_u(t)} \kappa_{\omega}(t) \star (A_{uv}(t) dN_{vs}(t))$ models the propagation of peer influence over the network, where the triggering kernel $\kappa_{\omega_1}(t)$ models the decay of peer influence over time.

Note that the retweet intensity matrix $\mathbf{\Gamma}^*(t)$ is by itself a stochastic process that depends on the time-varying network topology, the non-zero entries in $\mathbf{A}(t)$, whose growth is controlled by the network evolution process in Section 3.4. Hence the model design captures the influence of the network topology and each source’s influence, β_s , on the information diffusion process. More specifically, to compute $\gamma_{us}^*(t)$, one first finds the current set $\mathcal{F}_u(t)$ of followers of u , and then aggregates the retweets of these followers that are due to source s . Note that these followers may or may not *directly* follow source s . Then, the more frequently node u is exposed to retweets of tweets originated from source s via her followers, the more likely she will also retweet a tweet originated from source s . Once node u retweets due to source s , the corresponding $N_{us}(t)$ will be incremented, and this in turn will increase the likelihood of triggering retweets due to source s among the followers of u . Thus, the source does *not* simply broadcast the message to nodes directly following her but her influence propagates through the network even to those nodes that do not directly follow her. Finally, this information diffusion model allows a node to repeatedly generate events in a cascade, and is very different from the independent cascade or linear threshold models [25] which allow at most one event per node per cascade.

3.4 Network Evolution Process

We model the intensity, $\mathbf{\Lambda}^*(t)$, for link creation using a combination of survival and Hawkes process:

$$\lambda_{us}^*(t) = (1 - A_{us}(t))(\mu_u + \alpha_u \kappa_{\omega_2}(t) \star dN_{us}(t)) \quad (8)$$

where the term $1 - A_{us}(t)$ effectively ensures a link is created only once, and after that, the corresponding intensity is set to zero. The term $\mu_u \geq 0$ denotes a baseline intensity, which models when a node u decides to follow a source s spontaneously at her own initiative. The term $\alpha_u \kappa_{\omega_2}(t) \star dN_{us}(t)$ corresponds to the retweets of node u due to tweets originally published by source s , where the triggering kernel $\kappa_{\omega_2}(t)$ models the decay of interests over time. Here, the higher the corresponding retweet intensity, the more likely u will find information by source s useful and will create a *direct* link to s .

The link creation intensity $\Lambda^*(t)$ is also a stochastic process by itself, which depends on the retweet events, and is driven by the retweet count increments $dN_{us}(t)$. It captures the influence of retweets on the link creation, and closes the loop of mutual influence between information diffusion and network topology.

Note that creating a link is more than just adding a path or allowing information sources to take shortcuts during diffusion. The network evolution makes fundamental changes to the diffusion dynamics and stationary distribution of the diffusion process in Section 3.3. As shown in [14], given a fixed network structure \mathbf{A} , the expected retweet intensity $\mu_s(t)$ at time t due to source s will depend of the network structure in a highly nonlinear fashion, *i.e.*, $\mu_s(t) := \mathbb{E}[\Gamma_s^*(t)] = (e^{(\mathbf{A}-\omega_1\mathbf{I})t} + \omega_1(\mathbf{A} - \omega_1\mathbf{I})^{-1}(e^{(\mathbf{A}-\omega_1\mathbf{I})t} - \mathbf{I}))\boldsymbol{\eta}_s$, where $\boldsymbol{\eta}_s \in \mathbb{R}^m$ has a single nonzero entry with value η_s and $e^{(\mathbf{A}-\omega_1\mathbf{I})t}$ is the matrix exponential. When $t \rightarrow \infty$, the stationary intensity $\bar{\mu}_s = (\mathbf{I} - \mathbf{A}/\omega)^{-1}\boldsymbol{\eta}_s$ is also nonlinearly related to the network structure. Thus given two network structures $\mathbf{A}(t)$ and $\mathbf{A}(t')$ at two points in time, which are different by a few edges, the effect of these edges on the information diffusion is not just simply an additive relation. Depending on how these newly created edges modify the eigen-structure of the sparse matrix $\mathbf{A}(t)$, their effect can be drastic to the information diffusion.

Remark 1. In our model, each user is exposed to information through a time-varying set of neighbors. By doing so, we couple information diffusion with the network evolution, increasing the practical application of our model to real-network datasets. The particular definition of exposure (*e.g.*, a retweet’s neighbor) will depend on the type of historical information that is available. Remarkably, the flexibility of our model allows for different types of diffusion events, which we can broadly classify into two categories. In a first category, events corresponds to the times when an information cascade hits a person, for example, through a retweet from one of her neighbors, but she does not explicitly like or forward the associated post. In a second category, the person decides to explicitly like or forward the associated post and events corresponds to the times when she does so. Intuitively, events in the latter category are more prone to trigger new connections but are also less frequent. Therefore, it is mostly suitable to large event dataset for examples those ones generated synthetically. In contrast, the events in the former category are less likely to inspire new links but found in abundance. Therefore, it is very suitable for real-world sparse data. Consequently, in synthetic experiments we used the latter and in the real one we used the former. It’s noteworthy that Eq. (8) is written based on the latter category, but, Fig. 7 in appendix is drawn based on the former.

4 Efficient Simulation of Coevolutionary Dynamics

We can simulate samples (link creations, tweets and retweets) from our model by adapting Ogata’s thinning algorithm [26], originally designed for multidimensional Hawkes processes. However, a naive implementation of Ogata’s algorithm would scale poorly, *i.e.*, for each sample, we would need to re-evaluate $\Gamma^*(t)$ and $\Lambda^*(t)$, thus, to draw n samples, we would need to perform $O(m^2n^2)$ operations, where m is the number of nodes.

We designed a sampling procedure that is especially well-fitted for the structure of our model. The algorithm is based on the following key idea: if we consider each intensity function in $\Gamma^*(t)$ and $\Lambda^*(t)$ as a separate Hawkes process and draw a sample from each, it is easy to show that the minimum among all these samples is a valid sample from the model [12]. However, by drawing samples from all intensities, the computational complexity would not improve. However, when the network is sparse, whenever we sample a new node (or link) event from the model, only a small number of intensity functions, in the local neighborhood of the node (or the link), will change. As a consequence, we can reuse most of the samples from the intensity functions for the next new sample and find which intensity functions we need to change in $O(\log m)$ operations, using a heap. Finally, we exploit the properties of the exponential function to update individual intensities for each new sample in $O(1)$: let t_i and t_{i+1} be two consecutive events, then, we can compute $\lambda^*(t_{i+1})$ as $(\lambda^*(t_i) - \mu) \exp(-\omega(t_{i+1} - t_i)) + \mu$ without the need to compare all previous events.

The complete simulation algorithm is summarized in Algorithm 2 in Appendix C. By using Algorithm 2, we reduce the complexity from $O(n^2m^2)$ to $O(nd \log m)$, where d is the maximum number of followees per node. That means, our algorithm scales logarithmically with the number of nodes and linearly with the number of edges at any point in time during the simulation. We also note that the events for link creations, tweets and retweets are generated in a temporally intertwined and inter-

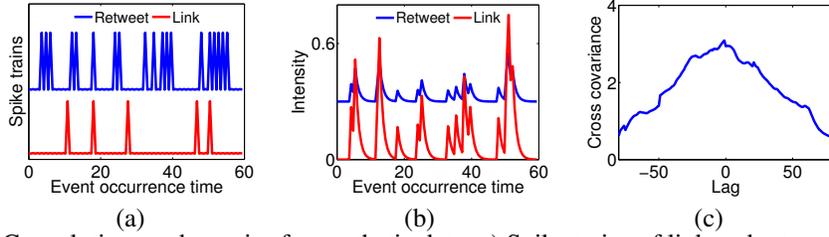


Figure 1: Coevolutionary dynamics for synthetic data. a) Spike trains of link and retweet events. b) Link and retweet intensities. c) Cross covariance of link and retweet intensities.

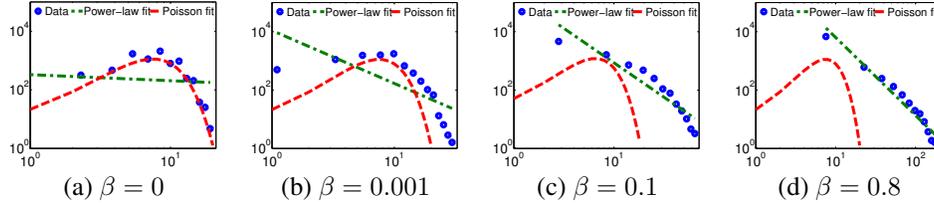


Figure 2: Degree distributions when network sparsity level reaches 0.001 for fixed $\alpha = 0.1$.

leaving fashion by Algorithm 2. This is because every new retweet event will modify the intensity for link creation, and after each link creation we also need to update the retweet intensities.

5 Efficient Parameter Estimation from Coevolutionary Events

Given a collection of retweet events $\mathcal{E} = \{e_i^r\}$ and link creation events $\mathcal{A} = \{e_i^l\}$ recorded within a time window $[0, T]$, we can easily estimate the parameters needed in our model using maximum likelihood estimation. Here, we compute the joint log-likelihood $\mathcal{L}(\{\mu_u\}, \{\alpha_u\}, \{\eta_u\}, \{\beta_s\})$ of these events using Eq. (3), *i.e.*,

$$\underbrace{\sum_{e_i^r \in \mathcal{E}} \log(\gamma_{u_i s_i}^*(t_i)) - \sum_{u, s \in [m]} \int_0^T \gamma_{us}^*(\tau) d\tau}_{\text{tweet / retweet}} + \underbrace{\sum_{e_i^l \in \mathcal{A}} \log(\lambda_{u_i s_i}^*(t_i)) - \sum_{u, s \in [m]} \int_0^T \lambda_{us}^*(\tau) d\tau}_{\text{links}}. \quad (9)$$

For the terms corresponding to retweets, the log term only sums over the actual observed events, but the integral term actually sums over all possible combination of destination and source pairs, even if there is no event between a particular pair of destination and source. For such pairs with no observed events, the corresponding counting processes have essentially survived the observation window $[0, T]$, and the term $-\int_0^T \gamma_{us}^*(\tau) d\tau$ simply corresponds to the log survival probability. Terms corresponding to links have a similar structure to those for retweet.

Since $\gamma_{us}^*(t)$ and λ_{us}^* are linear in the parameters (η_u, β_s) and (μ_u, α_u) respectively, then $\log(\gamma_{us}^*(t))$ and $\log(\lambda_{us}^*)$ are concave functions in these parameters. Integration of $\gamma_{us}^*(t)$ and λ_{us}^* still results in linear functions of the parameters. Thus the overall objective in Eq. (9) is concave, and the global optimum can be found by many algorithms. In our experiments, we adapt the efficient algorithm developed in previous work [18, 19]. Furthermore, the optimization problem decomposes in m independent problems, one per node u , and can be readily parallelized.

6 Properties of Simulated Co-evolution, Networks and Cascades*

In this section, we perform an empirical investigation of the properties of the networks and information cascades generated by our model. In particular, we show that our model can generate co-evolutionary retweet and link dynamics and a wide spectrum of static and temporal network patterns and information cascades. Appendix D contains additional simulation results and visualizations. Appendix E contains an evaluation of our model estimation method in synthetic data.

Retweet and link coevolution. Figures 1(a,b) visualize the retweet and link events, aggregated across different sources, and the corresponding intensities for one node and one realization, picked at random. Here, it is already apparent that retweets and link creations are clustered in time and often follow each other. Further, Figure 1(c) shows the cross-covariance of the retweet and link creation intensity, computed across multiple realizations, for the same node, *i.e.*, if $f(t)$ and $g(t)$ are two intensities, the cross-covariance is a function of the time lag τ defined as $h(\tau) = \int f(t+\tau)g(t) dt$. It can be seen that the cross-covariance has its peak around 0, *i.e.*, retweets and link creations are

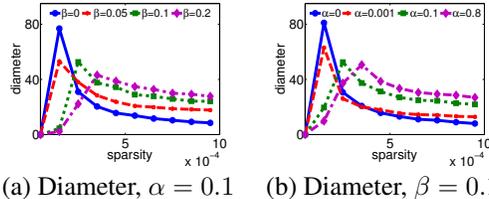


Figure 3: Diameter for network sparsity 0.001. Panels (a) and (b) show the diameter against sparsity over time for fixed $\alpha = 0.1$, and for fixed $\beta = 0.1$ respectively.

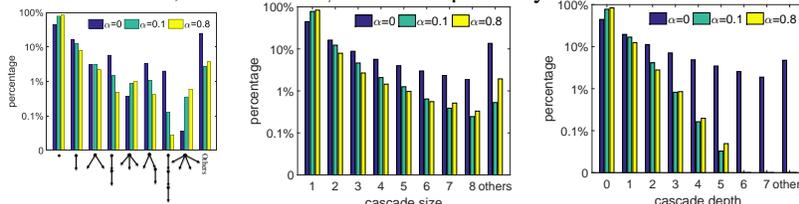


Figure 4: Distribution of cascade structure, size and depth for different α values and fixed $\beta = 0.2$.

highly correlated and co-evolve over time. For ease of exposition, we illustrated co-evolution using one node, however, we found consistent results across nodes.

Degree distribution. Empirical studies have shown that the degree distribution of online social networks and microblogging sites follow a power law [9, 1], and argued that it is a consequence of the rich get richer phenomena. The degree distribution of a network is a power law if the expected number of nodes m_d with degree d is given by $m_d \propto d^{-\gamma}$, where $\gamma > 0$. Intuitively, the higher the values of the parameters α and β , the closer the resulting degree distribution follows a power-law; the lower their values, the closer the distribution to an Erdos-Renyi random graph [27]. Figure 2 confirms this intuition by showing the degree distribution for different values of β .

Small (shrinking) diameter. There is empirical evidence that the diameter of online social networks and microblogging sites exhibit relatively small diameter and shrinks (or flattens) as the network grows [28, 9, 22]. Figures 3(a-b) show the diameter on the largest connected component (LCC) against the sparsity of the network over time for different values of α and β . Although at the beginning, there is a short increase in the diameter due to the merge of small connected components, the diameter decreases as the network evolves. Here, nodes *arrive* to the network when they follow (or are followed by) a node in the largest connected component.

Cascade patterns. Our model can produce the most commonly occurring cascades structures as well as heavy-tailed cascade size and depth distributions, as observed in historical Twitter data [23]. Figure 4 summarizes the results. The higher the α value, the shallower and wider the cascades.

7 Experiments on Real Dataset

In this section, we validate our model using a large Twitter dataset containing nearly 550,000 tweet, retweet and link events from more than 280,000 users [3]. We will show that our model can capture the co-evolutionary dynamics and, by doing so, it predicts retweet and link creation events more accurately than several alternatives. Appendix F contains detailed information about the dataset and additional experiments.

Retweet and link coevolution. Figures 5(a, b) visualize the retweet and link events, aggregated across different sources, and the corresponding intensities given by our trained model for one node, picked at random. Here, it is already apparent that retweets and link creations are clustered in time and often follow each other, and our fitted model intensities successfully track such behavior. Further, Figure 5(c) compares the cross-covariance between the empirical retweet and link creation intensities and between the retweet and link creation intensities given by our trained model, computed across multiple realizations, for the same node. The similarity between both cross-covariances is striking and both has its peak around 0, *i.e.*, retweets and link creations are highly correlated and co-evolve over time. For ease of exposition, as in Section 6, we illustrated co-evolution using one node, however, we found consistent results across nodes (see Appendix F).

Link prediction. We use our model to predict the identity of the source for each test link event, given the historical (link and retweet) events before the time of the prediction, and compare its performance with two state of the art methods, denoted as TRF [3] and WENG [5]. TRF measures

* Implementation codes are available at <https://github.com/farajtabar/Coevolution>

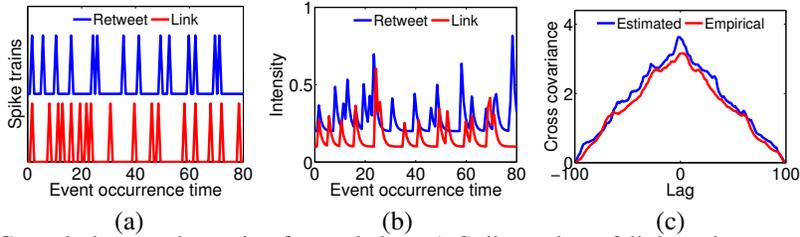


Figure 5: Coevolutionary dynamics for real data a) Spike trains of link and retweet events. b) Estimated link and retweet intensities. c) Empirical and estimated cross covariance of link and retweet intensities

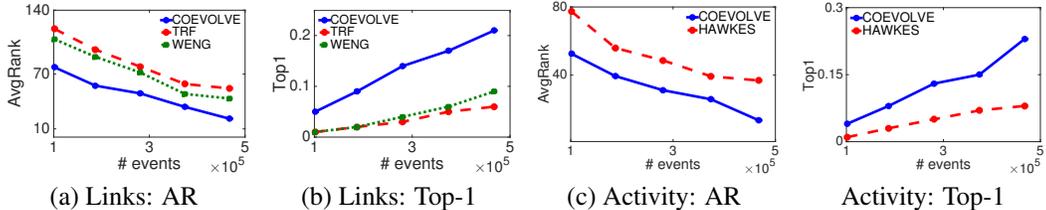


Figure 6: Prediction performance in the Twitter dataset by means of average rank (AR) and success probability that the true (test) events rank among the top-1 events (Top-1).

the probability of creating a link from a source at a given time by simply computing the proportion of new links created from the source with respect to the total number of links created up to the given time. WENG considers different link creation strategies and makes a prediction by combining them.

We evaluate the performance by computing the probability of all potential links using different methods, and then compute (i) the average rank of all true (test) events (AvgRank) and, (ii) the success probability (SP) that the true (test) events rank among the top-1 potential events at each test time (Top-1). We summarize the results in Fig. 6(a-b), where we consider an increasing number of training retweet/tweet events. Our model outperforms TRF and WENG consistently. For example, for $8 \cdot 10^4$ training events, our model achieves a SP 2.5x times larger than TRF and WENG.

Activity prediction. We use our model to predict the identity of the node that is going to generate each test diffusion event, given the historical events before the time of the prediction, and compare its performance with a baseline consisting of a Hawkes process without network evolution. For the Hawkes baseline, we take a snapshot of the network right before the prediction time, and use all historical retweeting events to fit the model. Here, we evaluate the performance the via the same two measures as in the link prediction task and summarize the results in Figure 6(c-d) against an increasing number of training events. The results show that, by modeling the co-evolutionary dynamics, our model performs significantly better than the baseline.

8 Discussion

We proposed a joint continuous-time model of information diffusion and network evolution, which can capture the coevolutionary dynamics, mimics the most common static and temporal network patterns observed in real-world networks and information diffusion data, and predicts the network evolution and information diffusion more accurately than previous state-of-the-arts. Using point processes to model intertwined events in information networks opens up many interesting future modeling work. Our current model is just a show-case of a rich set of possibilities offered by a point process framework, which have been rarely explored before in large scale social network modeling. For example, we can generalize our model to support link deletion by introducing an intensity matrix $\Xi^*(t)$ modeling link deletions as survival processes, *i.e.*, $\Xi^*(t) = (g_{us}^*(t)A_{us}(t))_{u,s \in [m]}$, and then consider the counting process $\mathbf{A}(t)$ associated with the adjacency matrix to evolve as $\mathbb{E}[d\mathbf{A}(t)|\mathcal{H}^r(t) \cup \mathcal{H}^l(t)] = \mathbf{\Lambda}^*(t) dt - \Xi^*(t) dt$. We also can consider the number of nodes varying over time. Furthermore, a large and diverse range of point processes can also be used in the framework without changing the efficiency of the simulation and the convexity of the parameter estimation, *e.g.*, condition the intensity on additional external features, such as node attributes.

Acknowledge

The authors would like to thank Demetris Antoniadis and Constantine Dovrolis for providing them with the dataset. The research was supported in part by NSF/NIH BIGDATA 1R01GM108341, ONR N00014-15-1-2340, NSF IIS-1218749, NSF CAREER IIS-1350983.

References

- [1] H. Kwak, C. Lee, H. Park, and others. What is Twitter, a social network or a news media? *WWW*, 2010.
- [2] J. Cheng, L. Adamic, P. A. Dow, and others. Can cascades be predicted? *WWW*, 2014.
- [3] D. Antoniadou and C. Drovolis. Co-evolutionary dynamics in social networks: A case study of twitter. *arXiv:1309.6001*, 2013.
- [4] S. Myers and J. Leskovec. The bursty dynamics of the twitter information network. *WWW*, 2014.
- [5] L. Weng, J. Ratkiewicz, N. Perra, B. Goncalves, C. Castillo, F. Bonchi, R. Schifanella, F. Menczer, and A. Flammini. The role of information diffusion in the evolution of social networks. *KDD*, 2013.
- [6] N. Du, L. Song, M. Gomez-Rodriguez, and H. Zha. Scalable influence estimation in continuous-time diffusion networks. *NIPS*, 2013.
- [7] M. Gomez-Rodriguez, D. Balduzzi, and B. Schölkopf. Uncovering the temporal dynamics of diffusion networks. *ICML*, 2011.
- [8] M. Gomez-Rodriguez, J. Leskovec, A. Krause. Inferring networks of diffusion and influence. *KDD*, 2010.
- [9] D. Chakrabarti, Y. Zhan, and C. Faloutsos. R-mat: A recursive model for graph mining. *Computer Science Department*, page 541, 2004.
- [10] J. Leskovec, D. Chakrabarti, J. Kleinberg, C. Faloutsos, and J. Leskovec. Kronecker graphs: An approach to modeling networks. *JMLR*, 2010.
- [11] J. Leskovec, L. Backstrom, R. Kumar, and others. Microscopic evolution of social networks. *KDD*, 2008.
- [12] T.J. Liniger. *Multivariate Hawkes Processes*. PhD thesis, ETHZ, 2009.
- [13] C. Blundell, J. Beck, K. Heller. Modelling reciprocating relationships with hawkes processes. *NIPS*, 2012.
- [14] M. Farajtabar, N. Du, M. Gomez-Rodriguez, I. Valera, H. Zha, and L. Song. Shaping social activity by incentivizing users. *NIPS*, 2014.
- [15] T. Iwata, A. Shah, and Z. Ghahramani. Discovering latent influence in online social activities via shared cascade poisson processes. *KDD*, 2013.
- [16] S. Linderman and R. Adams. Discovering latent network structure in point process data. *ICML*, 2014.
- [17] I. Valera, M. Gomez-Rodriguez, Modeling adoption of competing products and conventions in social media. *ICDM*, 2015.
- [18] K. Zhou, H. Zha, and L. Song. Learning social infectivity in sparse low-rank networks using multi-dimensional hawkes processes. *AISTATS*, 2013.
- [19] K. Zhou, H. Zha, and L. Song. Learning triggering kernels for multi-dimensional hawkes processes. *ICML*, 2013.
- [20] D. Hunter, P. Smyth, D. Q. Vu, and others. Dynamic egocentric models for citation networks. *ICML*, 2011.
- [21] D. Q. Vu, D. Hunter, P. Smyth, and A. Asuncion. Continuous-time regression models for longitudinal networks. *NIPS*, 2011.
- [22] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graphs over time: densification laws, shrinking diameters and possible explanations. *KDD*, 2005.
- [23] S. Goel, D. J. Watts, and D. G. Goldstein. The structure of online diffusion networks. *EC*, 2012.
- [24] O. Aalen, O. Borgan, and H. Gjessing. *Survival and event history analysis: a process point of view*, 2008.
- [25] D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. *KDD*, 2003.
- [26] Y. Ogata. On lewis' simulation method for point processes. *IEEE TIT*, 27(1):23–31, 1981.
- [27] P. Erdos and A Rényi. On the evolution of random graphs. *Hungar. Acad. Sci.*, 5:17–61, 1960.
- [28] L. Backstrom, P. Boldi, M. Rosa, J. Ugander, and S. Vigna. Four degrees of separation. *WebSci*, 2012.
- [29] M. Granovetter. The strength of weak ties. *American journal of sociology*, pages 1360–1380, 1973.
- [30] D. Romero and J. Kleinberg. The directed closure process in hybrid social-information networks, with an analysis of link formation on twitter. *ICWSM*, 2010.
- [31] J. Ugander, L. Backstrom, and J. Kleinberg. Subgraph frequencies: Mapping the empirical and extremal geography of large graph collections. *WWW*, 2013.
- [32] D.J. Watts and S.H. Strogatz. Collective dynamics of small-world networks. *Nature*, 1998.
- [33] T. Gross and B. Blasius. Adaptive coevolutionary networks: a review. *Royal Society Interface*, 2008.
- [34] P. Singer, C. Wagner, and M. Strohmaier. Factors influencing the co-evolution of social and content networks in online social media. *Modeling and Mining Ubiquitous Social Media*, pages 40–59. Springer, 2012.

A Further related work

The works most closely related to ours are empirical studies of information diffusion and network evolution [33, 34, 5, 3, 4]. Among them, [5] was the first to show experimental evidence the likelihood that a user u starts following a user s increases with the number of messages from s seen by u . [3] investigated the temporal and statistical characteristics of retweet-driven connections within the Twitter network and then identified the number of retweets as a key factor to infer such connections. Finally, [4] showed that the Twitter network can be characterized by steady rates of change, interrupted by sudden bursts of new connections, triggered by retweet cascades. They also developed a method to predict which retweets are more likely to trigger these bursts.

However, there are fundamental differences between the above mentioned studies and our work. First, they only characterize the effect that information diffusion has on the network dynamics, but not the bidirectional influence. In contrast, our probabilistic generative model takes into account the bidirectional influence between information diffusion and network dynamics. Second, previous studies are mostly empirical and only make binary predictions on link creation events. For example, [5, 3] predict whether a new link will be created based on the number of retweets; [4] predict whether a burst of new links will occur based on the number of retweets and users' similarity. In contrast, our model can learn parameters from real world data, and predict the precise timing of both diffusion and new link events.

B Illustration of the joint dynamics of information diffusion and network evolution

Figure 7 demonstrates the coevolution schematically in a toy social network. Consider node A as a source of information and node G as a potential follower. Node G is first exposed to a piece of information originated from A via a retweet from B at time t_1 (Figure 7(a)). The intensity $\lambda_{GA}^*(t)$ of G establishing a link to A increases. Then at time t_2 , when B retweets another tweet posted by A , node G again is exposed to information from A and the intensity $\lambda_{GA}^*(t)$ increases again (Figure 7(b)). After G observes sufficient information from A , she finds her a valuable source and thus G decides to follow A directly. This is how information diffusion affects network formation.

On the other side, network evolution will affect information diffusion. After G decides to follow A directly, she opens a new path of information for her downstream nodes. Now, when G retweets tweets originally posted by A , node I will be exposed and may retweet them too. Here, the exposure at time t_3 and t_4 is sufficient for I to retweet a tweet originated by A (Figures 7(c) and 7(d)).

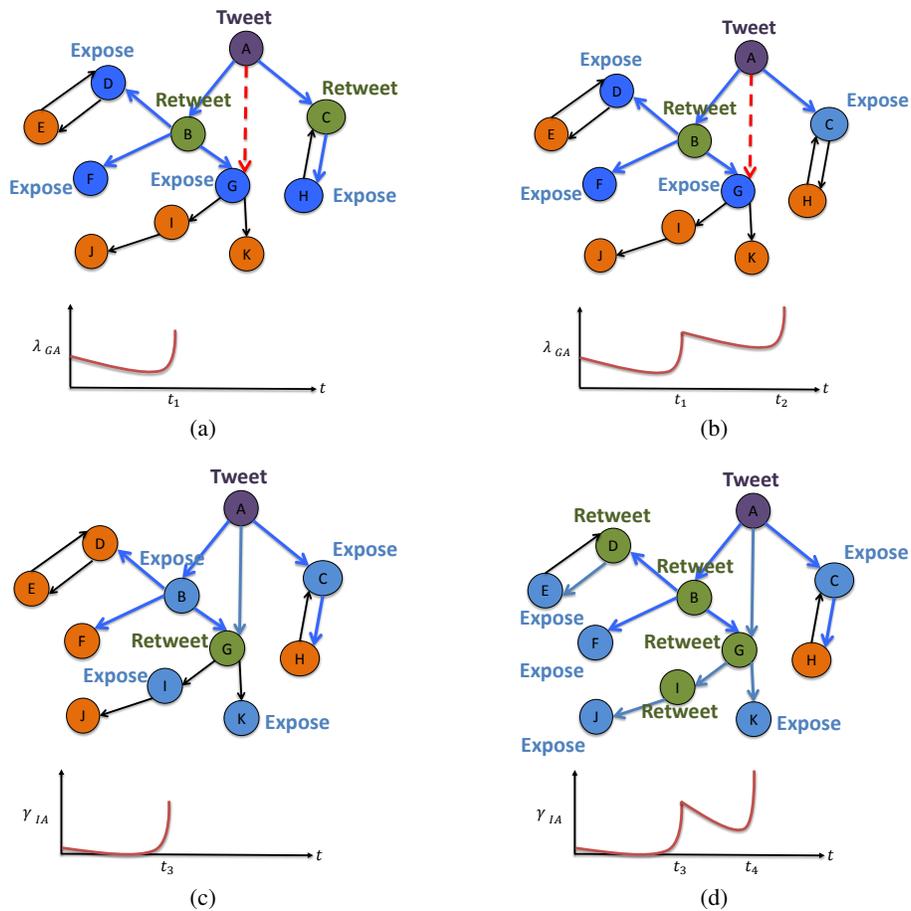


Figure 7: Joint dynamics of information diffusion and network evolution. Blue links form the diffusion paths. Green and blue circles are nodes which see the information. Green circles are those who re-share and propagate the information. Orange circles are nodes unaware of A 's post.

C Model Simulation

Algorithm 1 Efficient Intensity Computation

Global Variabels:
 Last time of intensity computation: t
 Last value of intensity computation: I
Initialization:
 $t = 0, I = \mu$
function *get_intensity*(t')
 $I' = (I - \mu) \exp(-\omega(t' - t)) + \mu$
 $t = t', I = I'$
return I
end function

Algorithm 2 Efficient Simulation

Initialization:
 Initialize the priority queue Q
for $\forall u, s \in [m]$ **do**
 Sample next link event e_{us}^l from A_{us} (Algorithm 3)
 $Q.insert(e_{us}^l)$
 Sample next retweet event e_{us}^r from N_{us} (Algorithm 3)
 $Q.insert(e_{us}^r)$
end for
General Subroutine:
 $t \leftarrow 0$
while $t < T$ **do**
 $e \leftarrow Q.extract_min()$
if $e = (u, s, t')$ is a retweet event **then**
 Update the history $\mathcal{H}_{us}^r(t') = \mathcal{H}_{us}^r(t) \cup \{e\}$
for $\forall v$ s.t. $u \rightsquigarrow v$ **do**
 Update event intensity: $\gamma_{vs}(t') = \gamma_{vs}(t'^-) + \beta$
 Sample retweet event e_{vs}^r from N_{vs} (Algorithm 3)
 $Q.update_key(e_{vs}^r)$
if NOT $s \rightsquigarrow v$ **then**
 Update link intensity: $\lambda_{vs}(t') = \lambda_{vs}(t'^-) + \alpha$
 Sample link event e_{vs}^l from A_{vs} (Algorithm 3)
 $Q.update_key(e_{vs}^l)$
end if
end for
else
 Update the history $\mathcal{H}_{us}^l(t') = \mathcal{H}_{us}^l(t) \cup \{e\}$
 $\lambda_{us}(t) \leftarrow 0 \forall t > t'$
end if
 $t \leftarrow t'$
end while

Algorithm 3 Sampling

Input: Current time: t
Output: Next event time: s
 Set $s \leftarrow t$ and $\hat{\lambda} \leftarrow \lambda(s)$ (Algorithm 1)
while $s < T$ **do**
 Generate $g \sim \exp(\hat{\lambda})$ set $s \leftarrow s + g$
 Set $\bar{\lambda} \leftarrow \lambda(s)$ (Algorithm 1)
 Rejection test: Generate $d \sim \mathcal{U}(0, 1)$
if $d \times \hat{\lambda} < \bar{\lambda}$ **then return** s **else** $\hat{\lambda} = \bar{\lambda}$
end while
return s

D Additional Experiments on Co-Evolution, Network Properties and Cascade Patterns

Simulation settings. Throughout this section, we simulate the evolution of a 8,000-node network as well as the propagation of information over the network by sampling from our model using Algorithm 2. We set the exogenous intensities of the link event and diffusion event intensities to $\mu_u = \mu = 4 \times 10^{-6}$ and $\eta_u = \eta = 1.5$ respectively, and the triggering kernel parameter to $\omega_1 = \omega_2 = 1$. The parameter μ determines the independent growth of the network, roughly speaking, the expected number of links each user establishes spontaneously before time T is μT . Whenever we investigate a static property, we choose the same sparsity level of 0.001.

Degree Distribution. Figure 8 shows the degree distribution against the sparsity of the network over time for different values of α . For sufficiently high values of α , the degree distribution follows a power-law. The lower the α values, the closer the network is to an Erdos-Renyi random graph [27], and as a consequence the closer is the degree distribution to a Poisson. In fact, it is easy to show that our model outputs Erdos-Renyi graphs for $\alpha = 0$.

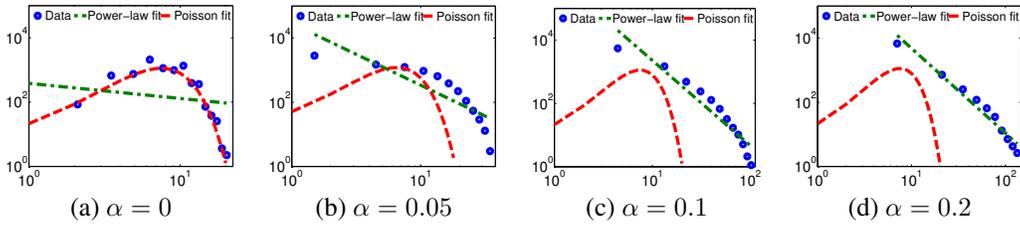


Figure 8: Degree distributions when network sparsity level reaches 0.001 for different α values and fixed $\beta = 0.1$.

Cascade Patterns. Figure 9 shows the distribution of the seven most common cascade structure, the cascade size and the cascade depth for different values of β . The model create the most commonly occurring cascades structures as well as heavy-tailed cascade size and depth distributions, as observed in real diffusion data [23]. The larger the value of β , the more shallow and wider cascades are.

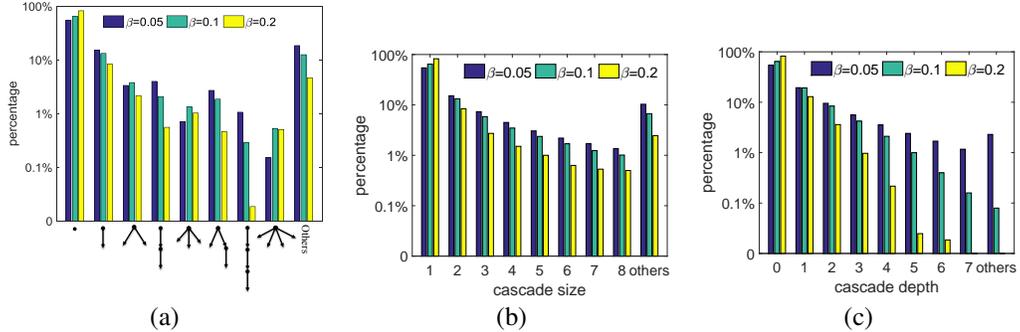


Figure 9: Distribution, size and depth of cascade structures for different β values and fixed $\alpha = 0.8$.

Clustering coefficient. Triadic closure [29, 11, 30] has been often presented as a plausible link creation mechanism. However, different social networks and microblogging sites present different levels of triadic closure [31]. Importantly, our method is able to generate networks with different levels of triadic closure, as shown by Figure 10(a-b), where we plot the clustering coefficient [32], which is proportional to the frequency of triadic closure, for different values of α and β .

Network Visualization. Figure 11 visualizes several snapshots of the largest connected component (LCC) of two 300-node networks for two particular realizations of our model, under two different values of β . In both cases, we used $\mu = 2 \times 10^{-4}$, $\alpha = 1$, and $\eta = 1.5$. The top two rows correspond to $\beta = 0$ and represent one end of the spectrum, *i.e.*, Erdos-Renyi random network. Here, the network evolves uniformly. The bottom two rows correspond to $\beta = 0.8$ and represent the

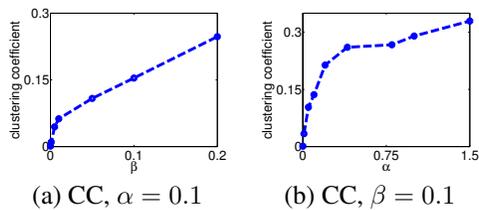


Figure 10: Clustering coefficient for network sparsity 0.001. Panels (a) and (b) show the clustering coefficient (CC) against β and α , respectively.

other end, *i.e.*, scale-free networks. Here, the network evolves locally, and clusters emerge naturally as a consequence of the local growth. They are depicted using a combination of forced directed and Fruchterman Reingold layout with Gephi¹. This figure clearly illustrates that by careful choice of parameters we can generate networks with a very different structure.

Furthermore, the retweet events (from others as source) of two nodes, A and B, are demonstrated at the bottom row. These two nodes arrive almost at the same time and establish links to two other nodes. However, node A's followers are more central, therefore, A is being exposed to more retweets. As an expected consequence node A will perform more retweets than B does.

¹<http://gephi.github.io/>

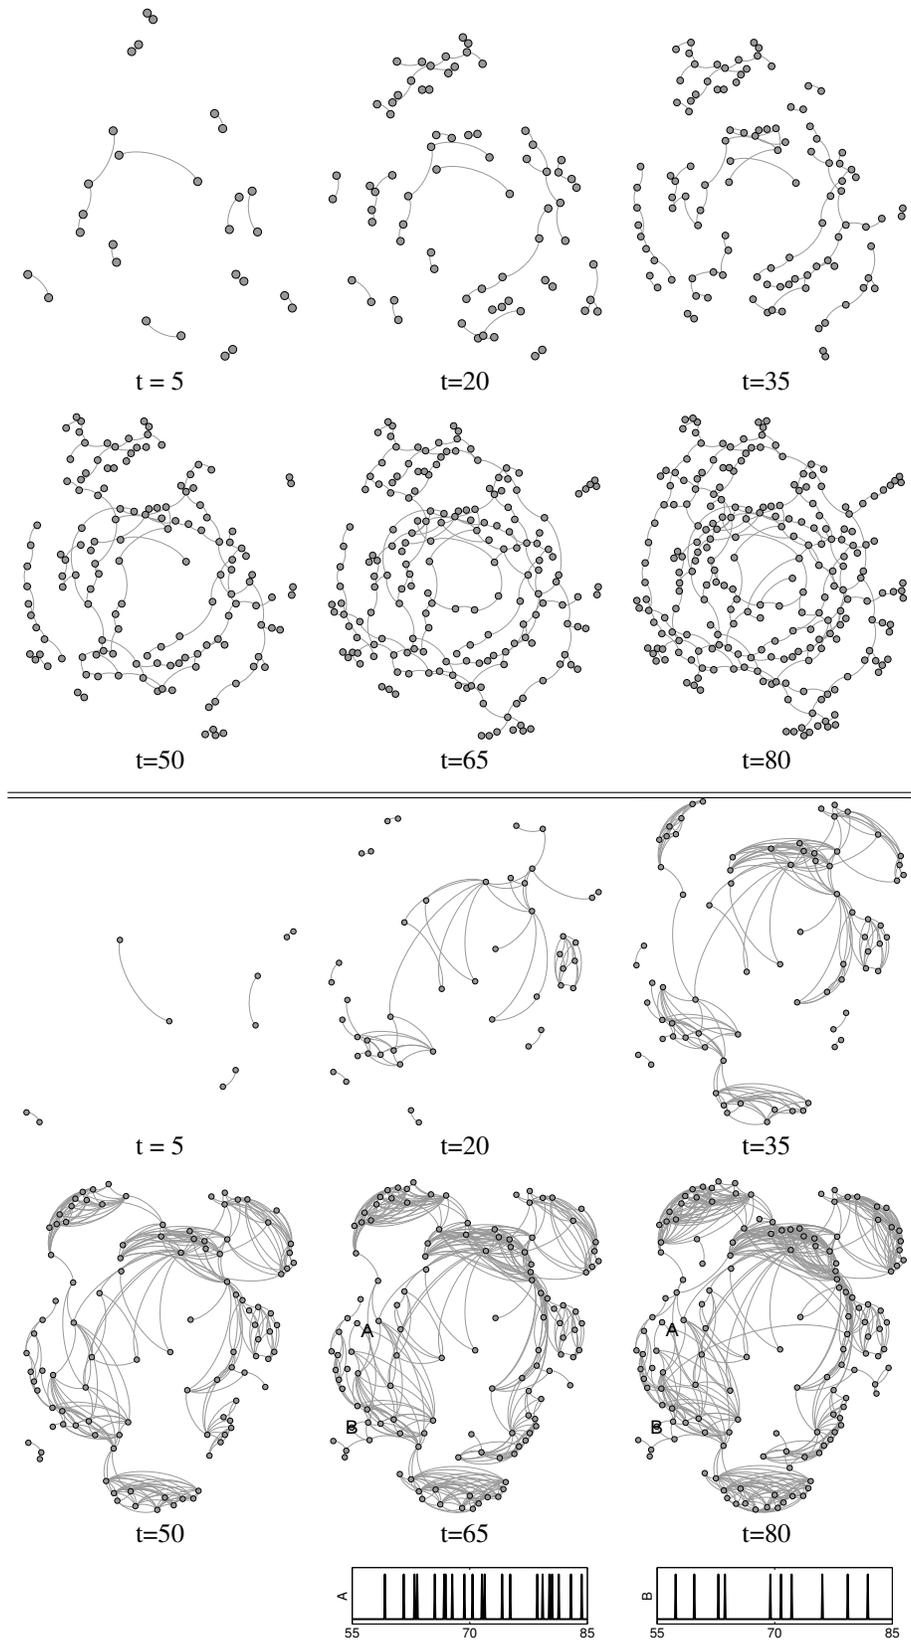


Figure 11: Evolution of a network with $\beta = 0$ (1st and 2nd rows) and $\beta = 0.8$ (3rd and 4th rows) and spike trains of nodes A and B (5th row). Node A is connected to more central nodes, therefore, will perform more retweets (originated by others) than B does.

E Experiments on Parameter Estimation

Experimental Setup. Throughout this section, we experiment with our model considering $m=400$ nodes. We set the model parameters for each node in the network by drawing samples from $\mu \sim U(0, 0.0004)$, $\alpha \sim U(0, 0.1)$, $\eta \sim U(0, 1.5)$ and $\beta \sim U(0, 0.1)$. We then sample up to 60,000 link and information diffusion events from our model using Algorithm 2 and average over 8 different simulation runs.

Model Estimation. We evaluate the accuracy of our model estimation procedure via two measures: (i) the relative mean absolute error (*i.e.*, $\mathbb{E}[|x - \hat{x}|/x]$, MAE) between the estimated parameters (x) and the true parameters (\hat{x}), (ii) the Kendall's rank correlation coefficient between each estimated parameter and its true value, and (iii) test log-likelihood. Figure 12 shows that as we feed more events into the estimation procedure, the estimation becomes more accurate.

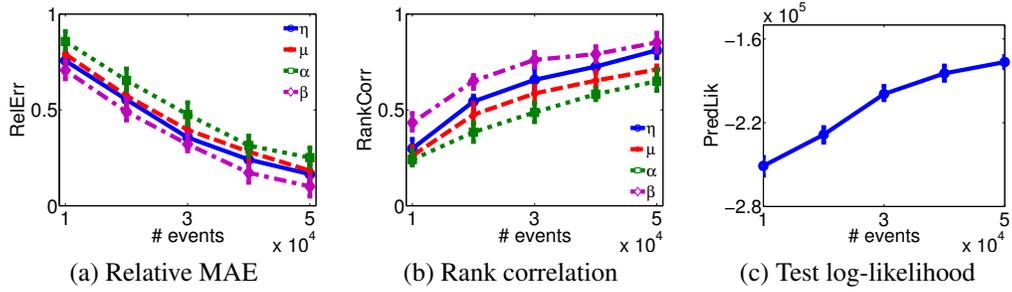


Figure 12: Performance of model estimation for a 400-node synthetic network.

F Additional Experiments on Real Data

Dataset Description.

We use a dataset that contains both link events as well as tweets/retweets from millions of Twitter users [3]. In particular, the dataset contains data from three set of users in 20 days; nearly 8 million tweet, retweet, and link events by more than 8 million users. The first set of users (8,779 users) are source nodes s , for whom all their tweet times were collected. The second set of users (77,200 users) are the followers of the first set of users, for whom all their retweet times (and source identities) were collected. The third set of users (6,546,650 users) are the users that start following at least one user in the first set during the recording period, for whom all the link times were collected.

In our experiments, we focus on all events (and users) during a 10-day period (Sep. 21 2012 - Sep. 30 2012) and used the information before Sep 21 to construct the initial social network (original links between users). We model the co-evolution in the second 10-day period using our framework. More specifically, in the coevolution modeling, we have 5,567 users in the first layer who post 221,201 tweets. In the second layer 101,465 retweets are generated by the whole 77,200 users in that interval. And in the third layer we have 198,518 users who create 219,134 links to 1978 users (out of 5567) in the first layer.

Experimental setup. We split events into a training set (covering 85% of the retweet and link events) and a test set (covering the remaining 15%) according to time, *i.e.*, all events in the training set occur earlier than those in the test set. We then use our model estimation procedure to fit the parameters from an increasing proportion of events from the training data.

Coevolution Behavior. Figure 13 demonstrates the link and retweet behavior of four typical users chosen randomly in the real-world dataset. For each user, the left panel (real activity) contains the links to a user and retweets her posts receive in 50 hours, and the right panel (real intensity) shows the corresponding intensity learned in the model. By visual inspection, it is already apparent a links or a retweets usually follow each other, but there are some link or retweet events triggered by exogenous causes.

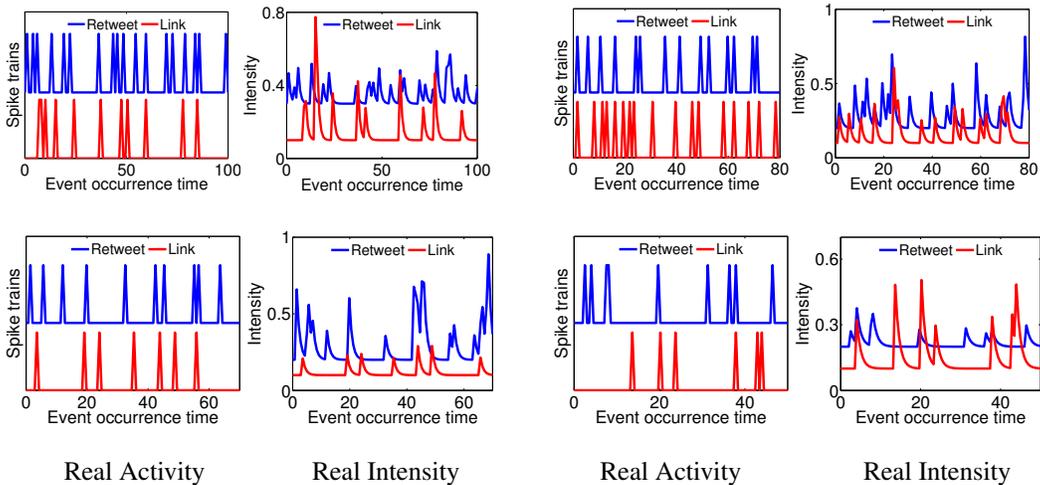


Figure 13: Link and retweet behavior of 4 typical users in the real-world dataset

Retweet and Link Event Co-evolution. Figure 14 shows the cross correlation of retweet and link intensity for 4 typical users chosen randomly from the Twitter dataset. If $f(t)$ and $g(t)$ are the intensity of the retweets of a user's post and links created to her then the cross-correlation is a function of the time lag τ defined as $h(\tau) = \int_t f(t + \tau)g(t) dt$. The empirical cross-correlation is drawn by interpolating the real events. The estimated cross-correlation is computed by simulating from the model using the parameters learned from the real dataset. As expected, the cross-correlation has its peak around 0, *i.e.*, links and activity are highly correlated, and the empirical cross-correlation coincide with the cross-correlation derived from the intensities in the learned model.

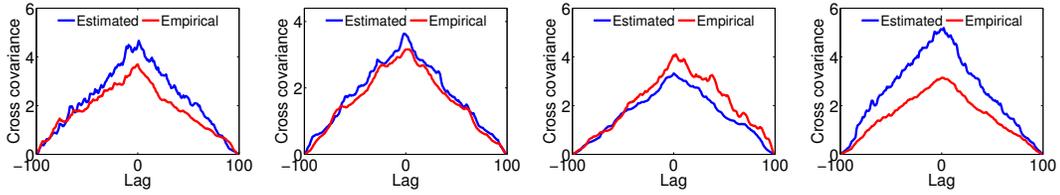


Figure 14: Empirical and simulated crosscorrelation for 4 typical users

To further verify that our model can capture the coevolution, we compute the average value of the empirical cross covariance function, denoted by m_{cc} , per user. Intuitively, one could expect that our model estimation method should assign higher α and/or β values to users with high m_{cc} . Figure 15 confirms this intuition on 1,000 users, picked at random users. Whenever a user has high α and/or β value, she exhibits a high cross covariance between her created links and retweets.

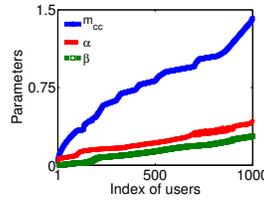


Figure 15: Empirical cross correlation versus the learned parameters

Link and Activity Prediction. Figure 16 shows the success probability that the true (test) events rank among the top-10 potential events at each test time for both link and activity prediction in the Twitter dataset.

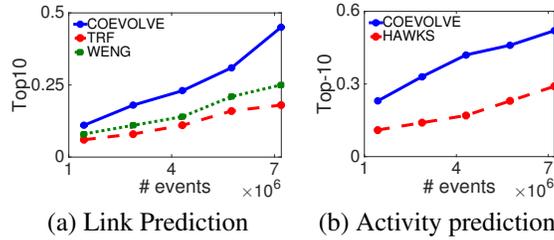


Figure 16: Prediction performance in the Twitter dataset by means of the success probability that the true (test) events rank among the top-10 events.

Test log-likelihood. We compute the link test log-likelihood using our trained models on the third set of users. Figure 17 summarizes the results, where we consider an increasing number of training retweet/tweet events from the first and second set.

Model Checking. Given all t_i and t_{i+1} subsequent event times generated using a Hawkes process, then, by the time changing theorem [12], the intensity integrals $\int_{t_i}^{t_{i+1}} \lambda(t) dt$ should conform to the unit-rate exponential distribution. Figure 18 presents the quantiles of the intensity integrals computed using the intensities with the parameters estimated from the real Twitter data against the quantiles of the unit-rate exponential distribution. It clearly shows that the points approximately lie on the same line, giving empirical evidence that a Hawkes process is the right model to capture the real dynamics.

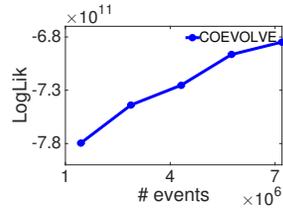


Figure 17: Test link log-likelihood in the Twitter dataset.

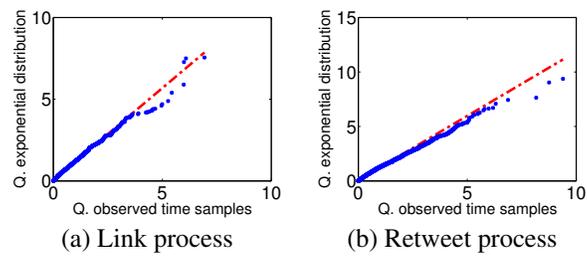


Figure 18: Quantile plots of the intensity integrals from the real link and retweet event time