

# Clash of the Contagions: Cooperation and Competition in Information Diffusion

Seth A. Myers  
Stanford University  
samyers@stanford.edu

Jure Leskovec  
Stanford University  
jure@cs.stanford.edu

**Abstract**—In networks, contagions such as information, purchasing behaviors, and diseases, spread and diffuse from node to node over the edges of the network. Moreover, in real-world scenarios multiple contagions spread through the network simultaneously. These contagions not only propagate at the same time but they also interact and compete with each other as they spread over the network.

While traditional empirical studies and models of diffusion consider individual contagions as independent and thus spreading in isolation, we study how different contagions interact with *each other* as they spread through the network. We develop a statistical model that allows for competition as well as cooperation of different contagions in information diffusion. Competing contagions decrease each other’s probability of spreading, while cooperating contagions help each other in being adopted throughout the network.

We evaluate our model on 18,000 contagions simultaneously spreading through the Twitter network. Our model learns how different contagions interact with each other and then uses these interactions to more accurately predict the diffusion of a contagion through the network. Moreover, the model also provides a compelling hypothesis for the principles that govern content interaction in information diffusion. Most importantly, we find very strong effects of interactions between contagions. Interactions cause a relative change in the spreading probability of a contagion by 71% on the average.

## I. INTRODUCTION

As we progress further into the information age, we are constantly bombarded with information on an unparalleled scale. Due to this abundance of information, we cannot interact with or consume all the information to which we are exposed. Rather, as we read newspapers, watch TV, or browse news websites we are constantly making choices about which pieces of media content to consume and adopt, and which pieces to ignore. This is especially the case in social media where the users are exposed to long streams of posts from their friends and followers. Commonly, users sequentially examine their stream of posts (*i.e.*, the news feed) and for each post decide whether to interact with it or maybe even share the posts with their network friends. The process of a user examining the input stream of posts and sharing/forwarding some of them along the edges of their network results in information cascades as content travels from user to user across connections in these networks [20].

Much work has been done to understand how consumers of media react and interact with different types of content. This process has been studied both empirically [21], [28],

[29] as well as through mathematical models and simulations [8], [22], [24]. However, most works studying this phenomenon focus on a *single* piece of information, described as a *contagion*, as it propagates like an epidemic along edges of the network. While these models have been effective, they do not capture the effects of *multiple contagions* cascading through the network at the same time.

In the real-world there are multiple pieces of information spreading through the network simultaneously. Moreover, these pieces of information do not spread in *isolation*, independent of all other information currently diffusing in the network. For example, consider two different news stories on the same event that diffuse through the network at the same time. One would imagine that two such contagions would help each other diffusing in a sense that many users would see both stories. This would make them think the news is very important and thus they would be more likely to adopt and share it. In this case two contagions *cooperate* as they mutually help each other in spreading through the network.

On the other hand, imagine two very different but both very interesting stories. For example, the same company releasing a new exciting product, while also laying off a number of its employees. Even though both stories are very interesting by themselves, they would likely compete with each other for attention — some people would notice the release of a new product and completely overlook the layoffs, while others would discuss the layoffs while practically ignoring the exciting new product the company developed. Thus, even though many people would be exposed to both stories, they would likely only chose to adopt and share one of them. This is an example of *competition*, where one story suppresses the other.

Given these examples, there is a clear opportunity to extend the present understanding of information diffusion in networks by jointly studying multiple contagions as they simultaneously spread through the network. Thus, rather than considering each contagion in isolation from all the others, modeling the interactions between contagions could lead to more predictive diffusion models and better influence maximization strategies [13], [16].

**Cooperation and competition in information diffusion.** We study how different pieces of content interact with *each other*. Through empirical analysis of real-world social media

and the use of statistical modeling, we quantify relationships between pieces of content as they simultaneously cascade, or *diffuse*, across an actual large scale social network. The model we develop provides fundamental insight into the way one component of information can have a significant effect on the way users perceive another component as they both diffuse across the same social network.

We develop a statistical model that *learns* how different contagions interact with each other and then use these interactions to more accurately predict the diffusion of a contagion through the network. We consider a setting where a user examines pieces of content sequentially. A user examines a piece of content and then decides whether to interact with the content (*e.g.*, share it with her neighbors in the social network). The user then proceeds and examines the next piece of content and makes the same decision again. Under this setting our goal is to model the probability that the user will adopt/share a particular piece of content  $X$ . Generally, there are three types of signal that determine this probability. First is the inherent interestingness of content  $X$  (*i.e.*, *content virality*), second is the inherent likelihood of the user to share any kind of content  $X$  (*i.e.*, *user bias*), and third is the *content interaction term* that updates the probability of sharing  $X$  based on what other content the user has been exposed to in the past.

Our model considers all three signals. While the content virality and the user bias can simply be modeled as per user and per content parameters, modeling the interaction term is more challenging. Not only is there potentially significant interactions amongst any two contagions, but one might expect that these interactions could change the farther apart in time a user is exposed to either contagion. So the influence another contagion may have on a user’s perception may change considerably with each new contagion to which she is exposed [26]. Because of this, our model considers an entire *sequence* of contagions the user was exposed to leading up to the current contagion  $X$  of interest. To model the interactions between all contagions and how these interactions change with the timing of the contagions presents a major scalability challenge. For example, let  $K$  be the number of contagions back we go into a user’s exposure history. Thus, we need to model the probability of infection for each contagion in our dataset, conditioned on every possible sequence of contagions to which the user could be exposed. In our experiments, we consider  $W = 18,000$  different contagions, which means that there are  $W^K$  ( $\approx 10^{21}$  for  $K = 5$ ) possible exposure sequences.

We address this challenge and propose a model where we assume contagions probabilistically belong to multiple clusters, and then we only need to estimate the cluster-to-cluster interactions (quadratic in the number of clusters) rather than the full contagion-to-contagion interaction matrix. Rather than considering all possible exposure sequences, we consider them as conditionally independent, which significantly

reduces the number of parameters. This leads to a highly scalable model. We estimate its parameters through the use of a stochastic gradient descent-based parameter fitting method, and fit our model to large collections of contagions.

The applications of our model are quite compelling. Today many social media sites consider inserting advertisements in user’s social streams (*e.g.*, “promoted” and “trending” tweets and Facebook posts). Our model provides means to optimize click-through rates by placing promoted content in such positions that the preceding content enhances the click-through probability. Similarly, many times one would want to combat the spread of a particular negative piece of information. Here our model suggests ways to create a second contagion that suppresses the first. Experiments show that the content interactions are very significant – the adoption probability can be manipulated by up to 70% by controlling the context in which the content appears.

**Overview of results.** We frame our work in the context of Twitter. When a user logs on to Twitter, she sees the stream of “tweets” posted by the other users she follows, and from time to time she decides to interact with them by sharing a tweet with her followers by “re-tweeting” it. In our particular context, we examine how the re-tweet probability changes based on what other tweets the user has just seen.

We focus on tweets that contain URLs (whether it be a link to an online news article or to an amusing YouTube video), and so in this manor a user can share a link with all of their followers. If their followers enjoy the link then they can in turn share it with all of their followers by re-tweeting it. Through this mechanism, a single link (URL) can potentially spread across many users in the Twitter network. Here, we study these cascading URLs, treating each one as a different contagion. Our goal is to model how a user being exposed to one URL can either increase or decrease her probability of interacting with another URL.

We treat a URL posted in a tweet as a single contagion, we consider the Twitter users who posts these URLs as nodes, and we view the follower relationships as a network. When a user chooses to tweet a URL, we say she has become *infected* by the URL. In turn, all of the users who follow her and see the URL become *exposed* to the URL. With each exposure from their neighbors, users decide whether or not to tweet (to be become infected by) the URL.

We validate our model on the *complete* set of tweets from January of 2011, which amounts to more than 3 billion posts. From these tweets, we extracted and crawled highly tweeted 18,186 URLs that contain rich English text-based content. We fit our model to the interactions between these contagions, and we test how accurately it assigns probability of infection to each exposure event. In summary, our model marks an improvement in performance by 680% over the Independent Cascade model [9], [13] baseline and 400% over the Exposure Curve baseline [21].

Our model also provides a compelling hypothesis for the principles that govern content interaction in information diffusion. We find evidence that more infectious URLs have an adversely negative (suppressive) effect on less infectious URLs that are of unrelated content or subject matter, while at the same time they can dramatically increase the infection probability of URLs that are less infectious but are highly related in subject matter. Moreover, the interactions represent a very significant effect. Interactions with other cascading links can cause a relative increase or decrease in user adoption probability of a link on average by 71%!

## II. RELATED WORK

The sharing of content diffusing through social networks has been extensively studied in recent years [1], [15], [21], [27]. Several well studied models have been developed [6], [7], [24]. However, widely adopted models of information diffusion, like the Linear Threshold Model [8], [22], the Independent Cascade Model [9], [13], as well as research based on the concept of exposure curves [17], [21] all consider each contagion in *isolation*, independent of others.

The diffusion of several contagions has been the focus of several recent works [10], [12], [18], [19], [23]. In each of these works, however, it is assumed that being infected with one contagion is mutually exclusive to be infected by another. The common scenario discussed by these works is that there are two behaviors or technologies spreading through the network, and the user will ultimately choose to adopt only one of them. Such a system could be, for example, a set of users choosing between Skype and Microsoft Messenger. Often, a user gains some utility in choosing the same technology as her neighbors in the network [11]. The mutual exclusivity condition causes a competition between the contagions, which warrants a game theoretic approach for the analysis, and this line of study has applications in maximizing the spread of one contagion over another [3], [5] as well as limiting the spread of a competing contagion [4], [14]. The main difference to our research here is that our model does not assume contagions to be mutually exclusive. Moreover, rather than considering a game theoretic analysis with no real data, we empirically study the interactions between a large number of contagions.

Recently in [2], the mutually exclusive competition assumption between contagions is relaxed. Authors consider two viruses propagating through the network simultaneously where being infected by one virus gives a node partial immunity and only decreases (as opposed to eliminates) the chances of being infected by the other virus. Although mathematically clean, the model [2] is not meant to capture the rich dynamics that occur when many different contagions are spreading simultaneously, as is the case with most real-world information diffusion systems [15], [17], [27]. In our work we directly model the interactions between actual real-world contagions and gain insight into the nature of

the interactions themselves. We consider a large number of contagions and model the probability of adoption as a function of contagions a person has been exposed in the past. We estimate the interactions between tens of thousands of different contagions to learn what contagions help each other and which compete for attention.

The work most closely related to our own is [25], where authors examine the dynamics of many contagions diffusing across the Twitter network. They develop a simple agent-based user model in which each user has a limited capacity for contagions (*i.e.*, Twitter hashtags) and so they must choose only a few to tweet about over any given period of time. Authors show that the simulated data that this model generates matches the real data on several global statistics. This model, however, does not quantify the interactions between specific contagions — competition occurs implicitly as a result of finite user capacity. Additionally, as is the case in [2] the possibility of positive interaction, or *cooperation*, between contagions is never considered. Our model does all of this handedly, and we use our model to predict the propagation of each specific contagion as opposed to trying to match global aggregate statistics to the real data.

## III. THE MODEL

In order to study the interactions between pieces of content as they propagate from user to user across the network, we develop a mixing model that determines the probability of a user adopting a piece of content based on what other content the user was previously exposed to. This model will quantify how much exposing a user to one piece of content can increase or decrease their receptiveness to another piece of content. Not only this, but the model will capture how these interactions *change* with time and with the user’s exposure to additional pieces of content.

**Probability of infection.** We describe the act of a user tweeting a URL (contagion) after she has been exposed to it as an *infection*. When a user’s neighbor tweets a contagion, the contagion then shows up in the user’s news feed, and we refer to this as an *exposure*. The goal of this model is to estimate the probability of a user being infected by one contagion, given the sequence of contagions to which the user was previously exposed. We imagine that as a user reads through the contagions that her neighbor has tweeted, there is a sliding window going back  $K$  contagions that she keeps in mind. By considering this sequence of previous exposures, our model will not only be able capture the interaction between two contagions when a user sees one right after the other, but it will also quantify how this interaction changes as the user sees other intermittent contagions. See Fig. 1 for a visual representation of this sliding window process.

More formally, let the variable  $Y_k$  be the  $k^{th}$  most recent contagion to which the user was exposed, and let  $X$  be the contagion in which the user is currently examining (Fig. 1). Our model will assign a probability of the user

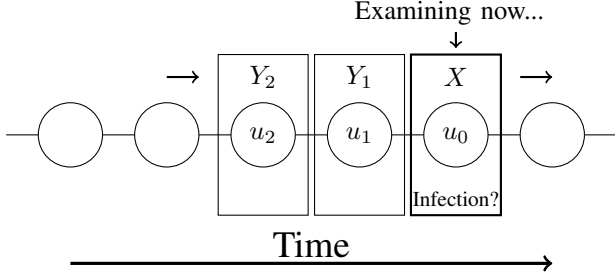


Figure 1. A visual representation of our model for  $K = 2$ . Here, a particular user has been exposed to the sequence of contagions. She is currently examining contagion  $u_0$ , but our model is assuming that she is still be affected by  $u_1$  and  $u_2$  that she previously saw. Our goal is to model the probability of the user adopting  $u_0$  as a function of which contagions she was exposed to in the past.

being infected by  $X$  upon being exposed to it, given her exposure history  $Y_1, Y_2, \dots, Y_k$ . In other words, we model

$$P(\text{infection by } X = u_0 \mid \text{exposed to the sequence } X = u_0, Y_1 = u_1, Y_2 = u_2, \dots, Y_K = u_K)$$

for any combination of contagions  $u_0, u_1, \dots, u_K$ . For the sake of brevity, we represent this probability as  $P(X|Y_1, \dots, Y_K)$  or

$$P\left(X \mid \{Y_k\}_{k=1}^K\right) \quad (1)$$

(i.e. exposure to  $X$  is always assumed).

Let  $W$  be the number of contagions that we are studying, so there are  $W^K$  different contagion combinations for which we need to calculate Eqn. (1). In our dataset, we have over  $W = 18,000$  real-world contagions, so for a fixed  $K$ , this is obviously infeasible. We make the assumption that  $Y_k$  is independent of  $Y_l$  i.e. the contagion  $k$  exposures ago is independent of the contagion  $l$  exposures ago, for any  $k, l$ . This assumption allows us to re-express Eq. 1:

$$\begin{aligned} P\left(X \mid \{Y_k\}_{k=1}^K\right) &= \frac{P(X) \cdot P\left(\{Y_k\}_{k=1}^K \mid X\right)}{P\left(\{Y_k\}_{k=1}^K\right)} \\ &= \frac{P(X) \cdot \prod_{k=1}^K P(Y_k \mid X)}{\prod_{k=1}^K P(Y_k)} \\ &= \frac{P(X) \cdot \prod_{k=1}^K \frac{P(X|Y_k) \cdot P(Y_k)}{P(X)}}{\prod_{k=1}^K P(Y_k)} \\ &= \frac{1}{P(X)^{K-1}} \prod_{k=1}^K P(X|Y_k). \end{aligned}$$

It should be noted here that

$$P(Y_k) \equiv P(k^{\text{th}} \text{ most recent exposure was } Y_k)$$

whereas

$$P(X) \equiv P(\text{infection by } X \text{ given just exposed to } X).$$

We refer to  $P(X)$  as the *prior infection probability*, and it can easily be computed empirically by counting the number

times a user was infected by  $X$  after being exposed to it and dividing by the number of times a user was exposed to  $X$ . Therefore, we only need to model  $P(X|Y_k)$  for each  $k = 1, \dots, K$ . This reduces the contagion combinations down to  $W \times W \times K$ , which is significantly less than before but still prohibitively large. The final step we make is that instead of modeling interactions between all pairs of contagions, we model the interactions between clusters (i.e., latent topics) of contagions. Specifically, we assume there exists a small number of latent clusters in which each contagion is a member with varying degree. Our approach is to parameterize each contagion's membership to these clusters while simultaneously parameterizing the interactions between these clusters. Here, our model currently assumes that the infection probability does not change from user to user. This is discussed later in the section "User Bias."

**Modeling interactions.** To begin, we assume that each contagion has some inherent infectiousness or virality (modeled by the prior infection probability  $P(X)$ ), and being exposed to other contagions either slightly increases or decreases the probability of infection. In other words, we model:

$$P(X = u_j \mid Y_k = u_i) \approx P(X = u_j) + \Delta_{cont.}^{(k)}(u_i, u_j) \quad (2)$$

where  $\Delta_{cont.}^{(k)}(u_i, u_j)$  is the interaction function that represents the effect contagion  $u_i$  has on contagion  $u_j$  from  $k$  exposures away, and  $P(X = u_j)$  is the empirically measured prior infection probability for contagion  $u_j$ . We can treat  $\Delta_{cont.}^{(k)}(u_i, u_j)$  as the  $i, j$  entry of the matrix  $\Delta_{cont.}^{(k)} \in \mathbb{R}^{W \times W}$ , which is far too large to model explicitly. Rather, we model the interactions between clusters (i.e., latent topics) of contagions.

Our strategy will be to identify clusters of the contagions that interact in similar ways, and then only model the interactions between these clusters. Let there be a small number of latent clusters (say there are  $T$  of them) to which each contagion is a member of in varying degree. If we know to which cluster each contagion belongs, then all we would have to do is model  $T \times T$  different interactions.

Given  $W$  contagions, we define the contagion to cluster membership matrix  $\mathbf{M} \in [0, 1]^{W \times T}$  such that

$$\mathbf{M}_{i,t} = P(\text{contagion } u_i \in \text{cluster } c_t)$$

and so  $\sum_t \mathbf{M}_{i,t} = 1 \forall i$ . To express the interactions between each latent cluster, for each  $k = 1, \dots, K$  we have a new interaction function  $\Delta_{clust}^{(k)}(c_t, c_s)$  to model the effect of cluster  $c_t$  on cluster  $c_s$ . Now, for  $u_i \neq u_j$

$$\Delta_{cont.}^{(k)}(u_i, u_j) = \sum_t \sum_s \mathbf{M}_{j,t} \cdot \Delta_{clust}^{(k)}(c_t, c_s) \cdot \mathbf{M}_{i,s}$$

and for  $u_i = u_j$

$$\Delta_{cont.}^{(k)}(u_j, u_j) = \sum_s \mathbf{M}_{j,s} \cdot \Delta_{clust}^{(k)}(c_s, c_s).$$

If we represent the cluster interactions as the matrix  $\Delta_{clust}^{(k)} \in \mathbb{R}^{T \times T}$  then the above expression can be represented as

$$\begin{bmatrix} \Delta_{cont.}^{(k)} \end{bmatrix} = \begin{bmatrix} \mathbf{M} \end{bmatrix} \times \begin{bmatrix} \Delta_{clust}^{(k)} \end{bmatrix} \times \begin{bmatrix} \mathbf{M}^T \end{bmatrix}$$

and it is apparent how this clustering method is equivalent to a low rank assumption. Technically speaking, this relationship holds only on the off-diagonal elements. As a shorthand, from now on we will use  $\Delta_{s,t}^{(k)} \equiv \Delta_{clust}^{(k)}(c_s, c_t)$ .

With this model, we can now express  $P(X|Y_k)$  as:

$$P(X = u_j | Y_k = u_i) = P(X = u_j) + \sum_t \sum_s \mathbf{M}_{i,t} \cdot \Delta_{t,s}^{(k)} \cdot \mathbf{M}_{j,s}$$

for when  $u_i \neq u_j$ , and when  $u_i = u_j$  then:

$$P(X = u_j | Y_k = u_j) = P(X = u_j) + \sum_s \mathbf{M}_{j,s} \cdot \Delta_{s,s}^{(k)}.$$

With this model, we specify the number of interaction clusters  $T$  and how many previous exposures  $K$  to include in the interaction model. Then, we learn the  $W \times T$  entries of the membership matrix  $\mathbf{M}$  and the  $T \times T$  entries of the cluster interaction matrix  $\Delta^{(k)}$  for  $k = 1, \dots, K$ . We enforce the constraints that  $0 \leq P(X = u_j | Y_k = u_i) \leq 1$  for all contagions during the fitting of the parameters. During the testing phase, it is possible (although extremely rare) for  $P(X = u_j | Y_k = u_i) \leq 0$ , in which we just set the probability to a minimum value of 1E-10;

Besides this additive model, we also performed extensive experimentation on several multiplicative models as well. Even though multiplicative models feel somewhat more natural in this setting, however, the additive model provided the best performance while also using far fewer parameters.

**Fitting the model.** For a given dataset, we can observe which users adopted (tweeted) which contagions (URLs) and when. With each adopted contagion, a user generates an exposure to that contagion for each of her neighbors in the network. Then, for each user we can observe the sequence of contagions she was exposed to, in between the contagions she adopts. For a given set of users, we can count the number of times  $X = u_i$  and  $Y_k = u_j$  for all  $u_i, u_j$  as well as whether or not this exposure pair led to the infection of  $X$ . Let all of these observations be contained in a set  $D$ . Then the log-likelihood function according to our model is

$$\begin{aligned} \mathcal{L}(D; \mathbf{M}, \{\Delta\}_{k=1}^K) = & \sum_{i,j,k} \alpha_{ij}^k \log \left[ P(X = u_j) + \sum_t \sum_s \mathbf{M}_{i,t} \cdot \Delta_{t,s}^{(k)} \cdot \mathbf{M}_{j,s} \right] \\ & + \beta_{ij}^k \log \left[ 1 - P(X = u_j) - \sum_t \sum_s \mathbf{M}_{i,t} \cdot \Delta_{t,s}^{(k)} \cdot \mathbf{M}_{j,s} \right] \end{aligned}$$

where  $\alpha_{ij}^k$  is the number of times in which  $Y_k = u_i$  and  $X = u_j$  and then the user adopted  $u_i$ , and  $\beta_{ij}^k$  is the number of times that exposure combination did not lead to infection. Our approach is to choose  $\mathbf{M}$  and  $\Delta_1, \dots, \Delta_K$  for some fixed  $T$  (the number of latent clusters) and  $K$  such that the log-likelihood is maximized.

Optimizing so many parameters over a such a large datasets is arduous. Specifically, for our dataset of  $W=18,000$  contagions, this likelihood function has just shy of 20 million observed pairs of  $\alpha_{ij}$  and  $\beta_{ij}$ . After employing numerous methods, the one that worked the best was a variation on stochastic gradient descent (of the *negative* log-likelihood function). Of all the pairs of contagions observed in users' exposure sequences, we randomly selected a subset of them to calculate the gradient and the line search in the direction of the negative gradient. We found that using a subset of around random 300,000 contagion pairs at each step produced good performance. We re-sampled these 300,000 pairs every 20 iterations. The line step length (how far into the direction of the negative gradient the solution moved at each iteration) was updated every 5-10 iterations using a simple bisection search in order to produce a sufficient decrease in the objective function.

In terms of optimization, it is important to observe that an entry in  $\mathbf{M}$  appears in many less objective function terms than an entry in  $\Delta^{(k)}$ . This means the membership matrix varies on a very different scale than the interaction matrices and thus has drastically different curvature. To account for this, we split each iteration such that first we take a step along the gradient of  $\mathbf{M}$ , and then we calculate and take a step along the gradient of  $\Delta^{(k)}$ . We calculated the line step length for either descent step independently.

**User bias.** It is reasonable to assume that different users will react to contagions in different ways. This would suggest it is prudent to consider a user bias in our model.

Now, we consider the probability  $P_n(X = u_j | Y_k = u_i)$  where  $n$  is a specific Twitter user:

$$P_n(X = u_j | Y_k = u_i) = P(X = u_j) + \gamma_n + \sum_t \sum_s \mathbf{M}_{i,t} \cdot \Delta_{t,s}^{(k)} \cdot \mathbf{M}_{j,s}.$$

Here, the user bias  $\gamma_n$  models how user  $n$  is more or less likely in general to adopt contagions. Including this term, however, makes fitting model difficult. Now, instead of the log-likelihood function having a term for each pair of exposed contagions and infected contagions, there are terms for every *user* and exposed contagion and infected contagion combination. Furthermore, we also found that the additional user bias parameter does not increase the overall accuracy.

#### IV. EXPERIMENTS

To validate our model, we run a set of experiments to determine how accurate it is at predicting contagion adoption by users. We use the model's success at this task

to demonstrate that it is expressing the correct structure in the way URLs interact with each other.

**The dataset.** Our dataset consists of *every* tweet sent in the month of January of 2011, which was more than 3 billion tweets in total. We parsed each tweet that contained a URL, and we aggregated every URL that was tweeted by at least 50 unique users, which was 191,650 URLs. Of these URLs we removed URLs by hand that were part of several different types of blatant spamming behavior. We then crawled all the URLs and discarded those that did not have large blocks of natural ASCII text (at least 50 tokens). This was a necessary step because we will later examine the relationship between interactions and URL page content similarity. This process left us with 39,771, of which we were able to classify  $W = 18,186$  URLs as English. From all of these  $W$  URLs combined, there are a total of 2,664,207 infection events and over 810 million exposure events. We then took the set of all users who tweet at least one of these  $W$  URLs, and then used the Twitter API to assemble the follower links amongst these users. In all, we were left with a subgraph with **1,087,033 nodes** and **103,112,438 edges**.

**The setup.** We treat each user as a series of ordered events. Each event can be: (1) the user’s exposure to a contagion posted by one of their neighbors, (2) the user tweeting a contagion that at least one of their neighbors has previously posted (infection), or (3) the user tweets a message that does not contain a contagion. Assuming that a user is not constantly logged on to Twitter to read messages posted by their neighbors the instant that they are posted, there is some lag between when a user sees the tweet (i.e. the next time she logs in). We assume that when a user posts a tweet, they have read all the tweets posted by her neighbors since the last time the user tweeted. Under this assumption, as the stream of exposures flow into a user, they are placed in a queue. The next time the user posts a tweet (whether or not it contains a contagion), the user reads through each of her exposures, and for each one she makes a decision of whether or not to tweet it. For each exposure in the queue, the user considers the  $K$  most recent exposures and the order in which they occurred.

This process creates a series of exposure intervals that are separated by user tweets, and our dataset contains 67,397,052 such intervals consisting 810,884,361 exposure events. Figure 2 shows a visual example of this process. We used 90% of these intervals to train the model as described in the previous section. We took another 5% to be a cross-validation set that we used to optimize a series of comparative baselines (discussed below). We reserved the final 5% as the test set with which we determine how accurate our model is in assigning probabilities of infection of contagions based solely on contagion interactions.

For each interval in the test set, we query our fitted model to assign a probability to each exposure in the interval that

a user will tweet that contagion. If the user does in fact tweet that contagion in the next set of tweets immediately following the interval, then an infection has occurred. For all exposures that occurred in the test set, we sort them by the probability assigned to them. The contagion exposure with the largest assigned probability becomes the model’s “first guess” for infection. The second largest probability is the next guess, and so on. We measure the accuracy of the model (as well as the baselines) through the precision and recall of this guessing process. The model is rewarded by assigning higher probabilities to exposures that resulted in infection as well as penalized for high probabilities that did not. The metrics we examine are the maximum F1-score across the entire precision-recall curve, as well as the area under the precision-recall curve. For both metrics, the closer the value is to 1, the better the performance. In addition, we also measured how well these assigned probabilities fit the actual infections in terms of log-likelihood.

**Baselines.** We compared our model to a series of commonly used diffusion models. If our model can out-perform these significantly, then we can assume the model is identifying true interactions in the information cascade process.

- **Infection Probability (IP).** This first baseline is the Independent Cascade Model [9], [13], which assigns a probability of infection of contagion  $X = u_i$  to be simply the prior infection probability (i.e., the virality):

$$P(X = u_i | Y_k = u_j) = P(X = u_i)$$

for all  $k, u_j$ . Considering the wide range of variation between the inherent infectiousness of different contagions, it might be effective to assign infection probabilities to each exposure event independent of all previous exposures.

- **Infection Probability + User Bias (UB).** For each user  $n$ , we fit a bias in infection probability  $\gamma_n$  using maximum likelihood. Thus,

$$P_n(X = u_i) = P(X = u_i) + \gamma_n.$$

This user bias accounts for the fact that certain users may be more prone to tweeting contagions in general, which could potentially have dramatic effects on contagion cascades. We induced sparsity across the  $\gamma_n$ ’s using an  $L_1$  regularization function. The  $L_1$  coefficient was chosen so as to optimize the log-likelihood on the validation dataset.

- **Exposure Curves (EC).** We use the concept of exposure curves [17], [21] as third baseline. An exposure curve is a function that assigns a probability of infection by a contagion based on the number of times the user was previously exposed to the contagion. This function represents how a user’s perception of a new piece of information changes each time she is exposed to it. For example, the probability of a user adopting a new piece

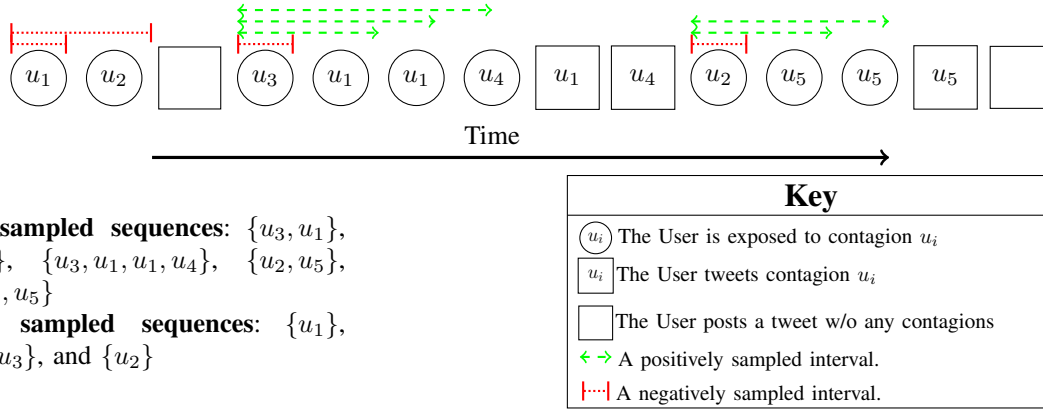


Figure 2. A representation of what one particular user sees in her Twitter news stream. Circles represent contagions that her neighbors have tweeted, and squares signify tweets by the user. We know every time she herself tweets, whether her tweet contains a contagion or not, she has seen and made a decision about the content to which she was previously exposed. Therefore, upon her tweeting a set of tweets (a set that is uninterrupted by exposures from her neighbors) we count all sequences in the previous exposure interval that ended in a contagion that she did tweet as positive examples. Similarly, we count all intervals that ended in a contagion she did not tweet as negative examples. We only considered samples of sequences of length at most  $K$ .

Model Name	Log-Like.	max $F_1$	Area under PR
IP	-335,550.39	0.0150	0.0157
UB	-338,821.54	0.0112	0.0123
EC	-338,367.86	0.0181	0.0250
<b>Our Model - With Prior</b>			
IMM K=1	-313,843.93	0.0412	0.0515
IMM K=2	-299,884.86	<b>0.0465</b>	<b>0.1238</b>
IMM K=3	<b>-299,352.32</b>	0.0380	0.0926
IMM K=4	-315,319.54	0.0321	0.0804
IMM K=5	-352,687.54	0.0386	0.0924

Table I  
RESULTS OF THE PROPOSED MODEL COMPARED TO THE BASELINES.

of technology may increase each time she is exposed to it as she learns more and more about it. On the other hand, the probability of a user retweeting a piece of celebrity gossip may decrease each time she sees it, as the gossip becomes stale and dated. For each contagion, the probability of infection after each number of observed infections is measured empirically in the training set. Since the sampling across all possible numbers of exposures can be sparse, we smoothed the exposure curve function using locally weighted linear regression. The smoothing coefficient was chosen so as to optimize the log-likelihood on the validation dataset.

**Results.** The results of the experiments are shown in Table I. The interaction models were fit using  $T = 20$  (i.e. the number of clusters is 20). This parameter was chosen as a balance between the performance on the validation set, and the runtime of fitting the model. Note that the number of parameters in the model is  $T \times T \times K + W \times T$ , so for  $T$  greater than 50 it takes a prohibitively long to fit the model.

The results of the performance experiments can be found in Table I. In all, our IMM model outperformed the baselines in all three metrics by a significant margin. E.g., our IMM model performed almost 400% better on the Area under PR Curve metric than the best baseline (the Exposure Curve)! Additionally, our model increased performance on

the Maximum  $F_1$  Score metric by 168% over best baseline.

**User bias.** Surprisingly, including a user bias term (UB) actually decreased performance compared to the Infection Probability baseline (IP). In fact, the UB baseline performed 25% worse than the IP baseline in max  $F_1$  score. Keep in mind that the user bias terms were fitted using  $L_1$  regularized maximum likelihood optimization and cross validation. Even still, the result of the user bias term was over-fitting. This suggests that there might not be significant enough fluctuations between each user’s probability of tweeting a contagion to justify the use of such of a parameter in the model. For example, the users that we included in the study were users that tweeted at least one contagion. In other words, all the users with which we trained and tested our model were chosen through a selection process that was biased towards users that tweeted more contagions, so this could produce less variation in infection probability from user to user. Regardless, this decrease in performance lead us to the decision not to include user bias in our interaction model.

The most important result of these experiments is how much better our model performed compared to pure infection probability model (IP). Recall that this baseline model is the result of a pure independent contagion assumption. The only additional information our model uses is interactions between contagions, and the result is a 210% improvement in max  $F_1$  score and a 680% improvement in the area under the PR curve. This implies that our model is effectively capturing interactions between contagions, and we now explore what our model implies about the interactions between contagions in general.

## V. DISCUSSION

The results of the experiments in the last section attribute significant gains in predictive performance (over the pure infection probability baseline) to modeling interactions between contagions. This begs the question of how much of

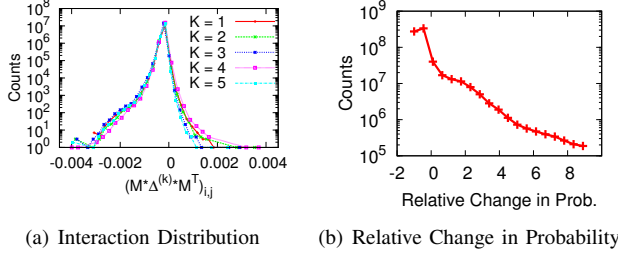


Figure 3. (a) The distribution in the expected interaction across all pairs of intersecting contagions for  $k = 1, \dots, 5$ . (b) The distribution of relative change in probability caused by including the interaction term of the model across all exposure events in the training dataset.

the probabilities of infection assigned to each exposure event comes from the contagion prior infection probability and how much comes from the contagion interactions.

Figure 3(a) shows the distribution of the expected interaction across all pairs of contagions in our dataset. Given this figure, it appears that on average, the expected interaction term of the model is very small. In fact, the average infection probability  $P(X)$  is 0.0029 with a standard deviation of 0.0055, whereas the average expected interaction is  $-6.637\text{E-}5$  with a standard deviation of  $7.630\text{E-}5$ . First, this validates the assumption presented in Eqn. 2;  $P(X|Y_k)$  is largely the prior infection probability, plus or minus the interaction term. This could, however, imply that the vast majority of the probability of infection that the interaction model assigns to each exposed contagion comes from the infection probability term and not from interactions between other contagions. In actuality, only a very few number of the possible pairs of contagions ever interact with each other (*i.e.*, are exposed to the same user at close to the same time). Furthermore, many of the pairs of contagions that do interact only do so a few number of times, whereas there are other contagion pairs that interact several thousand times. To account for this, we took every single exposure event in the training dataset and recorded the probability of infection given by our interaction model. We then recorded the relative change in probability of the interaction model compared to just the prior probability of infection  $P(X)$ :

$$\text{Relative Change} = \frac{(\text{Interact. Model Prob.}) - (\text{Prior Prob.})}{(\text{Prior Prob.})}.$$

With this measure, the contribution of each contagion pair to the distribution of interactions is proportional to how often they interact. Fig. 3(b) shows the distribution of this quantity across all exposure events in the training dataset. This plot tells a very different story in that the contribution of contagion interactions to the model vary widely, with the distribution showing a heavy tail that reaches 1,000% relative change in the infection probability. In fact, the average absolute value of relative change is 71%, indicating that on average more than half of the assigned probability comes from interactions between contagions. In short, the




Figure 4. (a) Relative Interaction versus the source contagion's prior infection probability minus the destination contagion's prior infection probability. (b) Relative Interaction versus content similarity between the contagions. The relative interaction in both figures is calculated using Eq 3.

aggregate of several small changes in each  $P(X|Y_k)$  creates a large change in  $P(X|Y_1, \dots, Y_k)$ .

We consider this strong evidence, in conjunction with the significant increase in infection prediction accuracy, that interactions between contagions should be a necessary component of any information cascade model.

**Why so negative?** Another interesting observation that comes from the distribution of interactions shown in Fig. 3(a) is that it is not centered at exactly 0. Specifically, there are more than a hundred times as many contagion pairs at the mode of this distribution (which is slightly negative) than there are at 0. In other words, there appears to be this default negative interaction between contagion pairs, and it is only some inherent interaction between specific pairs of contagions that changes this. This implies that if one studied the process of a contagion propagating across Twitter with no other contagions propagating at the same time, the infection probabilities would be higher and the final reach of the cascade would be larger compared to if another contagion was randomly chosen to propagate at the same time. This is intuitive. If a user has exactly one contagion in her news feed when she logs in, the chances she even sees the contagion is much higher, and the contagion will not have to share her focus with others.

**Why do contagions interact?** The most important question that can be asked of our interaction model is *what causes interactions between cascading contagions?*

To answer this question, we first look at how the prior infection probabilities of contagions affect their interactions. Fig. 4(a) shows for each pair of interacting contagions the expected interaction versus the prior infection probability of the source (exposing) contagion minus the infection probability of the destination (infecting) contagion. For the expected interaction, we normalize the the interaction across all pairs of URLs to a standard normal, *i.e.*, the normalized interaction between URLs  $i$  and  $j$  is

$$\text{Rel. Interaction}_k[i, j] = \frac{(M \cdot \Delta^{(k)} \cdot M^T)_{i,j} - \mu_{int}^{(k)}}{\sigma_{int}^{(k)}} \quad (3)$$

where  $\sigma_{int}^{(k)}$  is the standard deviation of interactions between



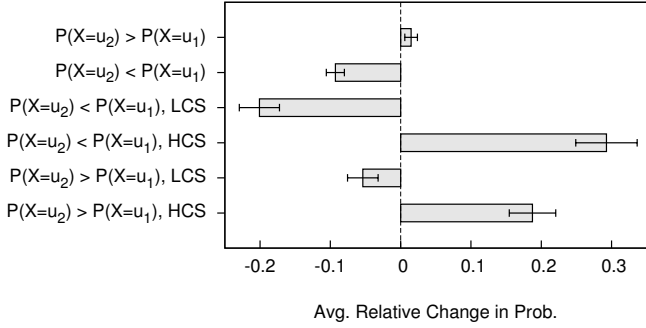


Figure 5. The empirically measured average relative change in infection probability for each URL  $u_2$  given certain conditions of the URL to which a user was previously exposed ( $u_1$ ). For each condition specified on the  $y$ -axis, we measure the average relative change in the infection probability  $P(X = u_2)$  when a user is first exposed to any URL  $u_1$  satisfying the condition before being exposed to  $u_2$ , and then we take the average across all  $u_2$ . “HCS” (high content similarity) implies that the  $u_1, u_2$  pair are in the 99<sup>th</sup> percentile in content similarity across all URL contagion pairs, and “LCS” (low content similarity) indicates they are in the 1<sup>st</sup> percentile. The error bars represent 95% confidence intervals.

URLs  $k$  exposures apart, and  $\mu_{int}^{(k)}$  is the mean value. This plot shows that for contagion pairs that have almost no interactions (0 on the  $x$ -axis), the destination contagion is much more infectious than the source contagion. If the interaction is strong, whether it is positive or negative, the source contagion is more infectious than the destination contagion. This implies that one contagion exposure interacts more strongly with contagions the user is later exposed to if this first contagion is *more infectious* than the later contagions. This makes intuitive sense, since when a user is reading through the tweets she has received, seeing new tweets will make her forget about older tweets, unless the older tweets are more influential/infectious. In other words, she will stay focused on a particular contagion unless something more infectious comes along. If it is only more infectious contagions that influence infection probabilities of less infectious contagions, what determines if this interaction is positive or negative?

The similarity in content also plays a crucial role in contagion interaction. Fig. 4(b) shows the expected interaction versus content similarity (the cosine similarity between the two URLs’ TF-IDF weighted word vectors), and it is clear that more positively interacting contagions are more closely related in content similarity and more negatively interacting contagions are unrelated. This, in combination with Fig. 4(a), is evidence of the following process:

- A user is exposed to a contagion, say  $u_1$ . Whether or not she retweets it, the contagion influences her.
- The user is then exposed to contagion  $u_2$ . If  $u_2$  is more infectious than  $u_1$ , the user shifts focus from  $u_1$  to  $u_2$ .
- On the other hand, if  $u_2$  is less infectious than  $u_1$ , then the user will still be influenced by  $u_1$  instead of switching to  $u_2$ . However, one of two things can happen:

- If  $u_2$  is highly related to  $u_1$  in content, then the user will be more receptive of  $u_2$  and this will increase the infection probability of  $u_2$ .
- If  $u_2$  is unrelated to  $u_1$ , then the user is less receptive of  $u_2$  and the infection probability is decreased.

For example, let’s say the user sees a compelling article about the popular uprising in Egypt. For the next period of time, the user is focused on this subject matter. If an uninteresting article about Justin Bieber comes along while the user is still thinking about Egypt, then she ignores it. However, while the user is in this state of mind, she will be more receptive to other articles about the Egyptian revolution than she otherwise would be.

To further validate this hypothesis, we can examine the empirically-measured conditional probability  $P(X = u_2 | Y_1 = u_1)$  *i.e.* the probability of being infected by each URL  $u_2$  given that the user was just exposed to first  $u_1$  and then to  $u_2$ , for  $u_1$  and  $u_2$  satisfying various conditions. For each  $u_2$ , we average the relative change in infection probability across all  $u_1$  satisfying a certain condition. Then, we average this relative change across each  $u_2$ , and the results are shown in Fig. 5. For example, the relative change in infectiousness of a URL when it follows the exposure of a less infectious URL is

$$\frac{1}{W} \sum_{u_2} \frac{1}{|S(u_2)|} \sum_{u_1 \in S(u_2)} \frac{P(X = u_2 | Y = u_1) - P(X = u_2)}{P(X = u_2)}$$

where  $P(X = u_2)$  is prior probability of infection of  $u_2$ , and  $S(u_2) \equiv \{u_1 | P(X = u_1) < P(X = u_2)\}$  (this value is the top row in Fig. 5).

In Fig. 5, we see that on average, when  $P(X = u_1) < P(X = u_2)$  then  $u_1$  has a very limited effect on  $u_2$ . For  $P(X = u_1) > P(X = u_2)$  on the other hand, the probability of infection of  $u_2$  decreases by almost 10%. This decrease in infectiousness more than doubles when  $u_1$  and  $u_2$  are completely dissimilar in content. In other words, the more infectious  $u_1$  suppresses the infection of  $u_2$ . When  $u_1$  and  $u_2$  are highly related in content, the effect is the opposite in that  $u_1$  increases the infection probability of  $u_2$  by almost 30%, and this is exactly in line with our hypothesis. Furthermore, we see that the effects of content between  $u_1$  and  $u_2$  decreases by half when  $P(X = u_1) < P(X = u_2)$  *i.e.*  $u_1$  has much less of an effect on  $u_2$  when  $u_2$  is more infectious.

To conclude, we look at the specific contagions and their interactions. For example, the pair of URLs that interact most positively at  $k = 1$  (according to the interaction model) have the following titles:

- $u_1$ : USA-Egypt Friendly CANCELED
- $u_2$ : U.S. Men’s National Team Match against Egypt Canceled

which are obviously highly related in content (they both are articles about a canceled soccer match due to the civil unrest

in Egypt). On the other hand, 3 of the 9 most negatively interacting URL pairs have

$u_1$ : *Is Yuri Milner A Threat To Silicon Valley?*

which is a TechCrunch.com article about a Silicon Valley investor’s controversial decision to blindly invest money into a large number of startup companies. The three URLs that it negatively interacts with are highly unrelated — they range from an article about the Egyptian army to travel deals available at Orbitz.com. This TechCrunch.com article effectively suppresses unrelated, less infectious URLs, just as our hypothesis would predict.

## VI. CONCLUSION

The majority of work on information diffusion has focused on a single contagion propagating through the social network, independent of any other contagion. We have presented a model that instead quantifies how different cascading contagions can interact with each other, either through competing or cooperating. In doing so, we can predict contagion infections with a 400% improvement over any of the standard baseline models.

Along with this dramatic boost in performance, our model predicts infection probabilities that are on average 71% more or less than the infection probability would be for a purely independent contagion. This is compelling evidence that looking at individual contagions by themselves does not offer a complete picture.

Our analysis has led us to the hypothesis that most interactions are governed by an underlying principle: contagions have an adversely negative (suppressive) effect on less infectious contagions that are of unrelated content or subject matter, while at the same time they can dramatically increase the infection probability of contagions that are less infectious but are highly related in subject matter.

In future work it would be very interesting see if this hypothesis holds for other types of social networks and contagion interactions. Twitter provides a very structured mechanism for how a user is exposed to the contagions of her neighbors — there is a guarantee that the contagion will appear in her news feed, and we know the order in which the exposures arrive. This is not the case for say the Blogosphere, in which it is ambiguous if and when one blogger is exposed to the contagions of another blogger. Adjustments would be required to model such a system.

**Acknowledgements.** This research has been supported in part by NSF CNS-1010921, IIS-1016909, CAREER IIS-1149837, IIS-1159679, DARPA SMISC, Brown Institute for Media Innovation, Albert Yu & Mary Bechmann Foundation, Boeing, Allies, Samsung, Intel, Alfred P. Sloan Fellowship and the Microsoft Faculty Fellowship.

## REFERENCES

- [1] E. Bakshy, J. Hofman, W. Mason, and D. Watts Everyone’s an influencer: quantifying influence on Twitter. *WSDM 2011*.
- [2] A. Beutel, B.A. Prakash, R. Rosenfeld, C. Faloutsos. Interacting Viruses in Networks: Can Both Survive? *KDD’12*.
- [3] S. Bharathi, D. Kempe, and M. Salek. Competitive influence maximization in social networks. *WINE 2007*.
- [4] C. Budak, D. Agrawal, and A. El Abbadi. Limiting the spread of misinformation in social networks. *WWW 2011*.
- [5] T. Carnes, C. Nagarajan, S. Wild, and A. van Zuylen. Maximizing Influence in a Competitive Social Network: A Follower’s Perspective. *ICEC 2011*
- [6] D. Centola and M. Macy. Complex contagions and the weakness of long ties. *American Journal of Sociology 2007*.
- [7] D. Cosley, D. Huttenlocher, J. Kleinberg, X. Lan, S. Suri. Sequential influence models in social networks. *ICWSM’10*.
- [8] M. Granovetter. Threshold models of collective behavior. *American Journal of Sociology 1978*.
- [9] J. Goldenberg, B. Libai, and E. Muller. Talk of the Network: A Complex Systems Look at the Underlying Process of Word-of-Mouth. *Marketing Letters 2001*.
- [10] S. Goyal, and M. Kearns. Competitive Contagion in Networks. *STOC 2012*.
- [11] S. Judd, M. Kearns, Y. Vorobeychik. Behavioral dynamics and influence in networked coloring and consensus. *PNAS’10*.
- [12] B. Karrer, and M. Newman. Competing epidemics on complex networks. *Phys Rev E 2011*.
- [13] D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the Spread of Influence in a Social Network. *KDD 2003*.
- [14] C. Kuhlman, V. Kumar, M. Marathe, S. Swarup, G. Tuli, S. Ravi, and D. Rosenkrantz. Inhibiting the Diffusion of Contagions in Bi-Threshold Systems: Analytical and Experimental Results. *AAAI 2011*.
- [15] J. Leskovec, M. McGlohon, C. Faloutsos, N. Glance, and M. Hurst. Cascading behavior in large blog graphs. *SDM’07*.
- [16] J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, J. VanBriesen, and N. Glance. Cost-effective Outbreak Detection in Networks *KDD 2007*.
- [17] S. Myers, C. Zhu, and J. Leskovec Information Diffusion and External Influence in Networks. *KDD 2012*.
- [18] N. Pathak, A. Banerjee, and J. Srivastava. A Generalized Linear Threshold Model for Multiple Cascades. *ICDM 2010*.
- [19] B.A. Prakash, A. Beutel, R. Rosenfeld, and C. Faloutsos. Winner Takes All: Competing Viruses or Ideas on fair-play Networks. *WWW 2012*.
- [20] E. Rogers. *Diffusion of Innovations*. Free Press, 4th Ed,’95.
- [21] D. M. Romero, B. Meeder, and J. M. Kleinberg. Differences in the Mechanics of Information Diffusion Across Topics: Idioms, Political Hashtags, and Complex Contagion on Twitter. *WWW 2011*.
- [22] T. Schelling. *Micromotives and Macromotives*. Norton, 1978.
- [23] K. Sneppen, A. Trusina, M. Jensen, and S. Bornholdt. A Minimal Model for Multiple Epidemics and Immunity Spreading. *PLoS ONE 2010*.
- [24] D. J. Watts. A simple model of global cascades on random networks. *PNAS 2002*.
- [25] L. Weng, A. Flammini, A. Vespignani, and F. Menczer. Competition among memes in a world with limited attention. *Nature 2012*.
- [26] F. Wu and B. Huberman. Novelty and collective attention. *PNAS 2007*.
- [27] S. Wu, J. Hofman, W. Mason, and D. Watts Who says what to whom on Twitter. *WWW 2011*.
- [28] J. Yang, and J. Leskovec. Patterns of temporal variation in online media. *WSDM 2011*.
- [29] J. Yang, and J. Leskovec. Modeling Information Diffusion in Implicit Networks. *ICDM 2010*.