

On the Role of Mentions on Tweet Virality

Soumajit Pramanik¹, Qinna Wang², Maximilien Danisch³, Sumanth Bandi¹, Anand Kumar¹,
Jean-Loup Guillaume⁴ and Bivas Mitra¹

¹Department of Computer Science & Engineering, IIT Kharagpur, India

²Sorbonne Universités, UPMC Université Paris 06, CNRS, LIP6 UMR 7606, 4 place Jussieu 75005 Paris, France

³Institut Mines Telecom, Telecom Paristech, CNRS, Paris

⁴L3I, University of La Rochelle, France

soumajit.pramanik@cse.iitkgp.ernet.in, qinna.wang@gmail.com, maximilien.danisch@telecom-paristech.fr,
sumanth232@gmail.com, anandsit043@gmail.com, jean-loup.guillaume@univ-lr.fr, bivas@cse.iitkgp.ernet.in

Abstract—In this paper, we investigate the role of mentions on tweet propagation. We propose a novel tweet propagation model SIR_{MF} based on a multiplex network framework, that allows to analyze the effects of mentioning on final retweet count. The basic bricks of this model are supported by a comprehensive study of multiple real datasets and simulations of the model show a nice agreement with the empirically observed tweet popularity. Studies and experiments also reveal that follower count, retweet rate & profile similarity are important factors in gaining tweet popularity and allow to better understand the impact of the mention strategies on the retweet count. Interestingly, we analytically identify a critical retweet rate regulating the role of mention on the tweet popularity. Finally, our data driven simulation demonstrates that the proposed mention recommendation heuristic *Easy-Mention* outperforms the benchmark *Whom-To-Mention* algorithm.

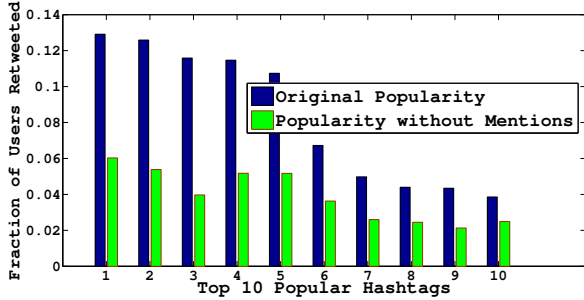
Keywords—Mention Recommendation; Multiplex Network; Information Diffusion.

I. INTRODUCTION

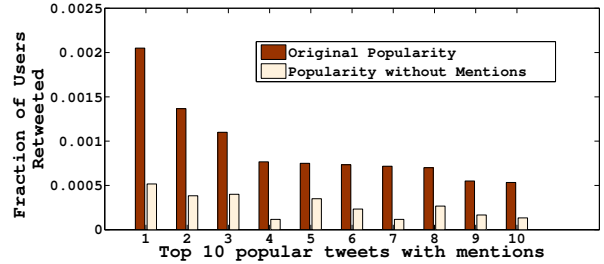
In recent times, Twitter has become one of the most influential micro-blogging systems for spreading and sharing breaking news, personal updates and spontaneous ideas [1]. However, it is observed that the popularity of tweets and hashtags follow a skewed distribution in any unbiased collection of tweets: only a small set of the tweets (or hashtags) are heavily popular [2]. In Twitter, propagation of a tweet or hashtag from one user to another occurs mainly via two activities: “retweeting” and “mentioning” [3]. In case of retweet, information is simply relayed to all the followers of the retweeting user. However, mention utility allows to spread an information far beyond the neighborhood and improves its visibility by making it available to the appropriate set of users. Furthermore, as mentions get listed in a separate tab, they gain higher attention than regular posts. Admittedly, mention utility plays a significant role in the cascading behavior of tweets and hashtags in Twitter. For instance, in our dataset, we observe that the probability that a mentioned user retweets a post is on average 32% higher than the one of a follower. Hence, investigating the role of mention utility behind popularizing a tweet is an interesting research question.

The problem of popularizing a tweet has two opposite facets. On the one hand, this is important to realize that artificially boosting popularity may immediately lead to spamming behavior [4]. Moreover, public mentions and direct message features have been exploited a lot for spamming hyperlinks and irrelevant content. Automatic mentioning through bots may further compound the problem and surely lead to annoyance. Hence, any attempt towards popularizing a tweet should be ready to deal with the possible mistreatment by the spammers. On the other hand, follower distribution exhibits the fact that most of the normal Twitter users only have a low to moderate number of followers [2]. Hence, any useful information, produced by a normal and trustworthy user, reaches only to a small population, even after a lot of deliberation.

Several studies have been carried out in understanding the dynamics behind the popularity of tweets. In [5] and [6], researchers investigated the role of content and contextual features of tweets and identified factors that are significantly associated with retweet rate and tweet popularity. In [7], Uysal and Croft proposed methods to recommend useful tweets that users are really interested in and more likely to retweet: given a tweet, they rank users based on retweet probability. Considering *mentioned user* as an influential information broker, several influence models have been explored [8], [9]. Importantly, in [10], Cha et al revealed that follower count is not necessarily the best metric to measure the influence. Following this line, another set of influence models [11], [12], [8] have been proposed to identify the influential nodes in a network. However, mentioning one influential user does not ensure that she retweets the post. This later part depends on several factors including information content of the post, profile of the tweeting user, etc. This motivates the community for the development of mention recommendation algorithms to identify the suitable users to mention. For instance, Wang et al. [13] proposed the *Whom-to-Mention* heuristic that uses features (such as user interest match, content-dependent user relationship and user influence) and uses machine learning to train a ranking function for extracting the best users to mention. Similar recommendation heuristics can be found in [14], [15], [16] and [17].



(a) Popularities (number of times posted) of top 10 popular hashtags in “Algeria” dataset with & without mentions.



(b) Popularities (retweet counts) of top 10 popular tweets (containing mentions) in “Egypt” dataset with & without mentions.

Fig. 1. Mention dependency for tweets and hashtags in “Algeria” & “Egypt” datasets.

Notably the aforementioned state of the art endeavors suffer from several limitations. In [13], the relevance function remains unchanged for different tweet messages, leading to same recommended ranked list for different tweets. Moreover, most of these heuristics rely on a large set of features to be calculated on a large population which is infeasible in real time; hence those approaches cannot be used to design an online mention recommendation system. More importantly, all these works fail to shed light on the interplay between the factors involved in the propagation of the tweets. For instance, it is not clear how exactly the mentioned user can make the tweet popular; does mentioning somebody in a tweet of her interest really helps in gaining popularity; how does the users’ activity (say retweet) rate influence the choice of the mention strategy? In order to address these questions, a simple model to mimic the tweet cascading process is necessary. This model can guide one to identify the role of individual factors on the tweet propagation and lead to the development of a simple recommendation system which may recommend users to mention. This paper takes an important step towards this direction.

In this paper, we dissect the impact of *mentioning* on tweet popularity. We start with a comprehensive data study to motivate the importance of mention utility on the popularity of a tweet. This study enables us to identify the important features of the mentioned user contributing to tweet popularity; her follower count, activity (retweet) rate, her profile similarity with the post, etc. (section II). We represent the tweet propagation process as a multiplex network [18] and propose an analytical framework SIR_{MF} to model the flow of tweets. Simulation of the model with suitable parameters show a nice agreement with the empirical tweet popularity observed in the dataset (section III). Moreover, the simulation model identifies a critical threshold on the retweet rate which demarcates the phase transition beyond which the retweet count increases exponentially. Our analytical framework nicely quantifies this observed critical threshold (section IV). Finally, taking cues from this model, we propose a simple mention recommendation heuristic which outperforms the *Whom-to-Mention* benchmark algorithm [13] (section V).

II. MOTIVATIONAL EXPERIMENTS

In this section we introduce the datasets and perform few motivational experiments to establish the importance of mention utility on the spread of tweets. This data study enables us to identify the key features of the mentioned users and works as a general guideline for identifying the right users to mention for maximizing the retweet count.

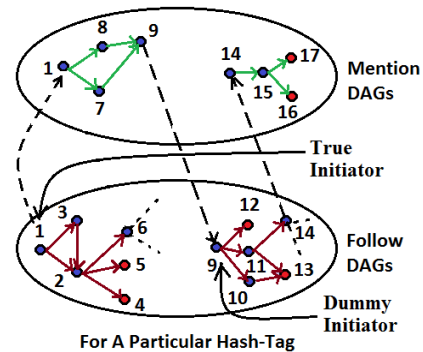


Fig. 2. Example of Mention-Follow Multiplex

A. Dataset

We collect the tweets posted during two particular real-life events - (a) Arab-Spring Movement-2011 and (b) World-Cup

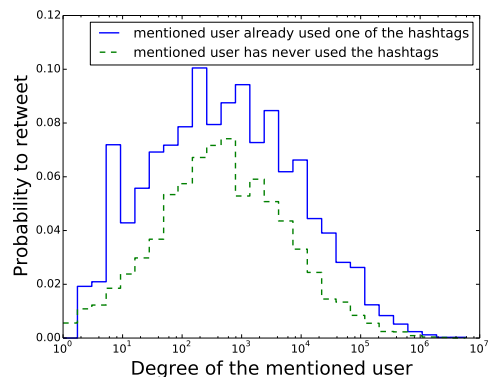
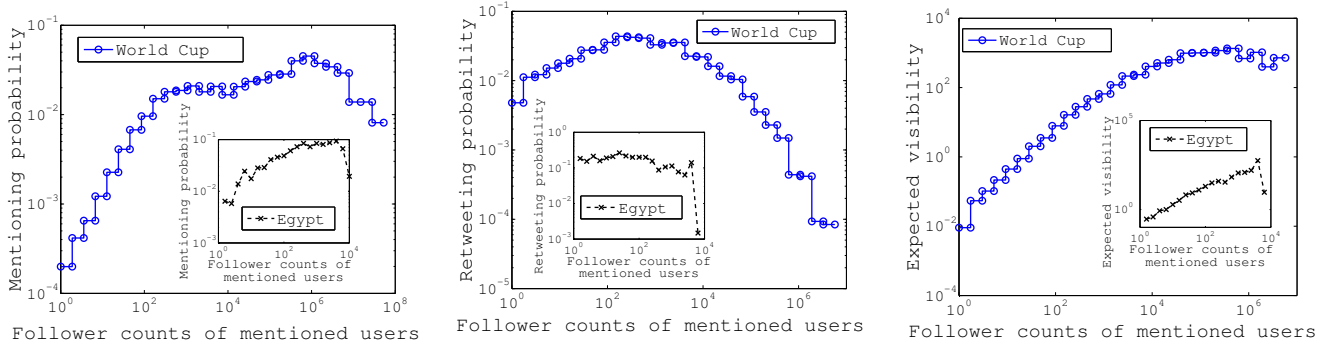


Fig. 4. Probability of retweeting for a mentioned user based on content-similarity in “World-Cup” dataset



(a) Probability of mentioning popular users (b) Retweeting probability of mentioned users (c) Expected visibility

Fig. 3. Users’ tendency & reasons to mention popular users in “World-Cup” & “Egypt” datasets

Football-2014. In both events, Twitter was used extensively to propagate news and opinions; however the domains, locations and time-spans of these two events are very different, hence tweet propagation in both events are completely independent. We may therefore assume that observed behaviors and results may hold more generally in Twitter.

(a) Arab-Spring Dataset: We collected two publicly available datasets [19] connected to these events - (i) “Algeria” Dataset is a collection of around 60K tweets (tweet-ids) and 20K users who posted them during the ‘Algeria movement’. We also crawled the tweet content, user profile and the corresponding follower network. (ii) “Egypt” dataset is a collection of around 2.6 million tweets (tweet-ids) posted during ‘Egypt uprising’. We crawled the tweet content of 0.2 million posts and the profile details and follower network of around 60K users who posted them.

(b) World-Cup Football Dataset: This dataset¹ consists of all tweets (2.8 million tweets) which are posted during the soccer World-Cup 2014 and contain official team-hashtags (#BRA, #CRO etc.) or match-hashtags (#BRACRO, #MEXCMR etc.).

B. Multiplex Representation

For a given hash-tag ‘#h’, the multiplex representation contains two layers: the bottom one represents tweet propagation via follow links, the top one via mention links (Fig. 2). More precisely, all users who tweet ‘#h’ appear as a node in the bottom (follow) layer. A directed link connects user ‘A’ to ‘B’ if ‘A’ (re)tweets ‘#h’ before ‘B’ further retweets and ‘B’ is a follower of ‘A’. In the top (mention) layer a directed link connects ‘C’ to ‘D’ if ‘C’ tweets ‘#h’ before ‘D’ further retweets and ‘C’ mentions ‘D’ in her post (‘D’ may or may not be a follower of ‘C’). One user is free to appear in both the layers.

A closer look reveals that both the layers are essentially collection of directed acyclic graphs (DAG). We denote the root of each DAG as an initiator since they are responsible for initiating the spreading process. We can identify two classes

of initiators, the ‘true initiators’ and the ‘dummy initiators’. A *true initiator* of ‘#h’ is a user who is a root in a Follow or a Mention DAG but never appears as non-root member of any DAG. These users have actually started the spreading process (for ‘#h’) as a result of some external influences. A *dummy initiator* is a user who is a root in a follow DAG but a non-root member of a mention DAG. Basically a dummy initiator gets the information from someone else via mention and subsequently initiates the spreading process to its followers.

C. Dependence on Mention

Given this multiplex representation, we measure the impact of the mentioned users on the popularity of hashtags. Let us define the popularity of a hashtag as the number of (re)tweets it receives. We select few popular hashtags for which we estimate the popularity reduction by dropping mentions. In this estimation, first we find the dummy initiators (set D) for a hashtag ‘#h’ and all the users (set S) who only belong to the DAGs rooted by dummy initiators. Obviously the retweet activity of the $S \cup D$ users is dependent on the mention layer. If hashtag ‘#h’ is tweeted by total n users, then mention dependency of ‘#h’ can be measured as $\frac{|S \cup D|}{n}$. Looking at the most popular hashtags (tweets) in Fig. 1(a), 1(b) we observe that such hashtags (tweets) are heavily mention dependent.

D. Properties of mentioned nodes

Next we turn our attention to the node level properties of the mentioned users. This may provide us some guideline to select proper users to mention.

Impact of popularity and retweet activity: In order to confirm whether people like to mention popular users, we plot the proportion of mentioned users with different follower counts (see Fig. 3(a)). The plot clearly depicts that a significant fraction of users mention popular people. On the other hand, Fig. 3(b) shows that the probability of getting a retweet from a mentioned user reduces sharply if her follower count is over 1000 (celebrities are choosy in retweeting). This clearly demonstrates that two opposite forces play a role in tweet

¹We received it from the “linkfluence” company (<http://linkfluence.com/en/>)

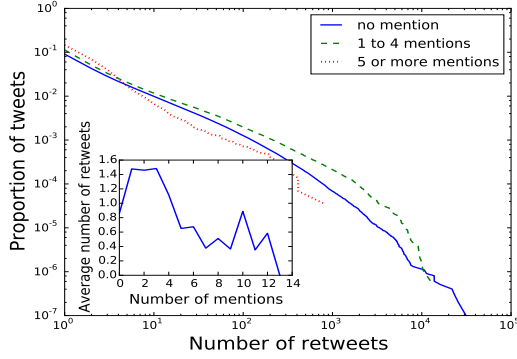


Fig. 5. The distribution of retweet counts of tweets containing different number of mentions in “World-Cup” Dataset. The Inset shows how the average retweet count changes with number of mentions in the tweets.

propagation through mentions; highly popular users are less likely to retweet but they provide high reachability when they retweet. In order to measure the combined effect of user popularity and retweet rate, we introduce *expected visibility*, which is the product of follower count and retweet probability of a mentioned user, and plot its distribution in Fig. 3(c). The peak of the curve demonstrates the existence of a balance between popularity and retweet rate, while mentioning some user.

Impact of content similarity: Content similarity between the profile of the mentioned user and the posted tweet is another factor which determines the propensity of retweeting. We compute the expectation that the mentioned user retweets in the “World-Cup” dataset (see Fig. 4), (a) if the tweet contains at least one hashtag that she has already posted (expected probability to retweet 0.029) and (b) if the tweet does not contain any hashtag which she has already posted (expected probability to retweet 0.017). Hence if the mentioned user has already posted the hashtag, her probability to retweet almost doubles. Moreover, Fig. 4 also reveals that this fact is independent to the follower count of the mentioned user.

Impact of the number of mentions: Mentioning the correct number of users is important to gain a high number of retweets. In the “Egypt” Dataset, we observe that 23.9% of all tweets in our dataset contain mentions. Out of them 80.5% of the tweets contain only one mention, 14.7% contain two, 3.2% contain three and the remaining 1.6% contain more than three. We also observe similar statistics in the “Algeria” & “World-Cup” datasets. Fig. 5 highlights the fact that mentioning few (say 2-3) intended users is always beneficial in gaining retweets; mentioning too many people makes the tweet content short and less interesting.

III. SIMULATION MODEL

In this section we model the tweet propagation dynamics in an epidemiological framework [20]. We validate the model with respect to the real retweet counts observed in the “Algeria” and “Egypt” datasets. Finally, we investigate the influence of the individual model parameters on the retweet count.

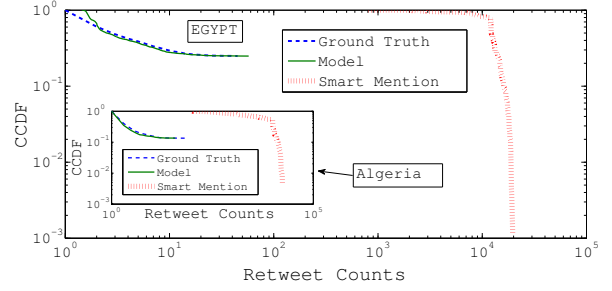


Fig. 6. Matching ground truth tweet popularities with the simulation model (with same α , β , λ & initiator) and comparing with ‘smart’ mention strategy for “Algeria” and “Egypt” datasets.

A. Model description

We propose a SIR based epidemic model SIR_{MF} to mimic the propagation of tweets on the mention-follow multiplex network (Fig. 2). Initially, all the nodes (representing users) are in the susceptible state. A node v gets infected by a tweet T if it retweets T in the next timestamp. A node once infected gets recovered instantaneously in the next timestamp. We assume that there is only one information (post) propagating in the system and any node can tweet / retweet it only once. The simulation stops when no more users can be infected.

In this framework, the infection of a node v for a tweet T is governed by three factors, (a) v has to be exposed to T , (b) v has to show interest in T and (c) v must have a certain retweet rate to retweet T ; even exposed to an interesting tweet, v may not retweet it. In more details, (a) A node v may get exposed to T by a node u in two different ways; (i) via follower links: if u posts T and v is a follower of u , (ii) via mention links: if v is not a follower of u but u mentions v while posting T . This forms the structure of the multiplex network (see Fig. 2). (b) The interest of v in tweet T depends on whether it has been exposed through mention or follow link. We model user interests (normalized between $[0,1]$) with two Poisson distributions with mean μ_1 and μ_2 respectively for the posts received through mention and follow links. Since mentions are more visible than normal posts, we keep $\mu_1 \geq \mu_2$. (c) The retweet rate κ_v (normalized between $[0,1]$) of node v is modeled by a power law distribution with exponent κ [21].

For a node v , we denote the infection rate (probability of retweeting) through mention links as α_v and through follow links as β_v . Notably, in Fig. 2 nodes get infected in the mention layer with average probability $\alpha_v = f(\kappa_v, \mu_{1v})$ and in the follow layer with average probability $\beta_v = f(\kappa_v, \mu_{2v})$. The general function f can simply be the product of both probabilities. One user is allowed to mention on average λ users in her tweet. The model parameters are summarized in Table I.

1) *Mention Strategies:* In SIR_{MF} model, we introduce three mention strategies, following which the user u can be chosen for mentioning in a tweet.

Random mention: The user u is chosen uniformly at random from the set of susceptible users.

Epidemic Propagation	Tweet Propagation
Susceptible	Users yet to post any tweet or retweet
Getting Infected	Tweeting/retweeting a post
Infected Individual	User who tweets/retweets a post
Model parameters	
κ	Exponent of Power-Law distribution representing user-activity
Infection Probability (via Mention) $\alpha_u = (\kappa_u \times \mu_{1_u})$	Probability that v has been mentioned in post T and v retweets T in the next timestep.
Infection Probability (via Follow) $\beta_u = (\kappa_u \times \mu_{2_u})$	Probability that v receives the post from followee and retweets the post in the next timestep
λ	Average Number of users mentioned in each tweet T

TABLE I
MAPPING THE TERMINOLOGIES AND PARAMETERS OF EPIDEMIC PROPAGATION & TWEET PROPAGATION.

Smart mention: The user u is chosen preferentially to her $f_u \times \alpha_u$ score where f_u is the follower count of u . Here the objective is to maximize the expected number users exposed to that tweet.

Parametric mention: The user u is chosen preferentially to her $(f_u \times \alpha_u)^\theta$ score where $(\theta \in [0, 1])$ is a tunable parameter. Admittedly, this is the generic mention strategy; with $\theta = 0$ & 1, we may emulate the ‘random’ and ‘smart’ mention strategies respectively.

B. Simulation setup and metrics

We simulate the SIR_{MF} model taking the follower networks from the “Algeria” and “Egypt” datasets. For all simulations, we fix λ (except for Fig. 6) to the average number of mentions in the dataset and $\kappa = -2.5$ considering that the average number of mentions per tweet and the retweet-rate of users do not change frequently over time [21]. We vary μ_1 and μ_2 to regulate the probabilities α (avg. of α_v) and β (avg. of β_v) respectively. Each result presented here is an average of 500 simulations.

1) *Evaluation Metrics:* We introduce the following four metrics to quantify the tweet propagation dynamics. These set of metrics will be further applied for evaluating the performance of different mention recommendation algorithms in section V:

(a) **Retweet count with Mentions (R_U)** is the average number of times tweets containing mentions are retweeted. In simulations we have a single tweet in the system and that tweet contains mentions (as $\lambda > 0$), therefore R_U is simply the infected population in the network.

(b) **Retweet count without Mentions (N_U)** is the average number of times tweets without mentions are retweeted.

(c) **Retweet Fraction by Mentioned Users (F_M)** is the average fraction of all the retweets (of the posts containing mentions) done by the mentioned users. In simulations this gives the fraction of retweeting users who has received the tweet via mention links and retweeted it.

(d) **Fraction of Mentioned Users Retweeted (F_C)** is the fraction of mentioned users who retweeted the post.

Note that N_U is not relevant for simulations since we only simulate tweets with mentions. Similarly F_C is not

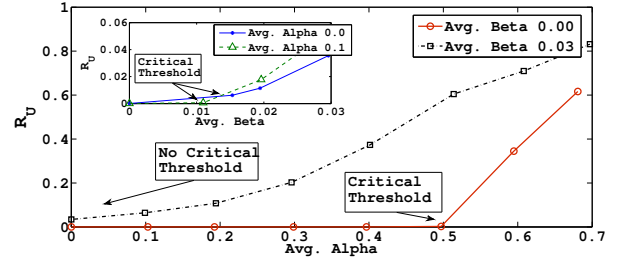


Fig. 7. Effect of varying Average α and Average β (see Inset) on R_U (Random Mentioning) in “Algeria” dataset.

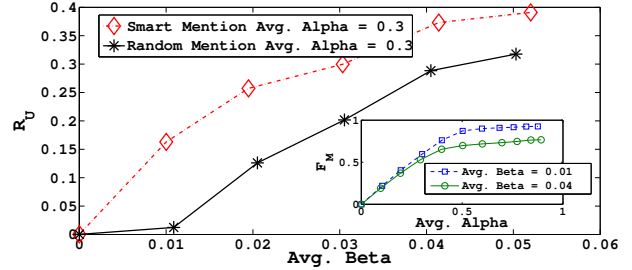


Fig. 8. Smart Mentioning vs. Random Mentioning w.r.t. R_U in “Algeria” dataset. **Inset:** Effect of varying Average α on F_M in “Algeria” dataset.

an observable metric in simulations; this simply depicts our model parameter α . However both metrics will play an important role to evaluate the performance of different mention recommendation algorithms in section V.

C. Model validation and insights

First we validate the SIR_{MF} model with respect to the retweet counts (R_U) of the tweets containing mentions in the “Algeria” & “Egypt” datasets. We implement the ‘Parametric’ mention strategy and simulate the model for each tweet (with a positive α) on the follower network obtained from the datasets. In order to run the model, we estimate suitable μ_1 and μ_2 to keep the average infection probabilities α and β close to the real data. Furthermore, we simulate each tweet diffusion with the same set of initiators and keeping the same number of mentions (λ) as in the real data. We adjust θ in the model to compute the total infected population, and fix θ which results to the best agreement with empirical R_U , estimated from the dataset.

In Fig. 6, we observe a nice agreement between the infected population of SIR_{MF} model and the real retweet count R_U estimated for both the “Algeria” and “Egypt” datasets. For most of the tweets, we fix $\theta \approx 0$, indicating that in reality, random mention strategy mostly gets followed. Nevertheless, the Fig. 6 demonstrates the fact that there is ample scope to boost the retweet count R_U by choosing the users to be mentioned, smartly.

Once we validate SIR_{MF} model with real data, next we investigate the role of individual parameters on the retweet count (R_U).

1) *Impact of infection rates α and β :* Inset of Fig. 7 shows that under a critical β , the tweet does not gain much retweet. Once it exceeds the threshold, the total retweet count increases almost linearly with β . However, the critical threshold value




Fig. 9. Transition probability of the states in the SIR_{MF} model.

of β decreases with increasing α . Similar effects can be seen if we keep β constant and vary α in X-axis (see Fig. 7). Also after a threshold value of α , R_U increases sharply and that critical threshold α value lowers if β is higher. Importantly, we analytically estimate these threshold points in the next section.

In the inset of Fig. 8, we observe that for same α , retweet fraction by the mentioned users (F_M) is lower for higher β values. This is intuitive because if β is high, more people retweet due to follow links which in turn lowers the fraction F_M . We note that F_M increases almost linearly with α up to a point and then converges.

2) *Impact of the mention strategies:* Fig. 8 shows that smart mention proves beneficial in low activity environment (low β). However, increase in β reduces the gap of R_U between the two mention strategies. This is because when β increases, mention-strategies become less important as most of the users start to get infected due to only follow links.

IV. ANALYTICAL REPRESENTATION OF SIR_{MF} MODEL

Inspired by the observation in Fig. 7, where the epidemic critical threshold exists for $\alpha = 0, 0.01, 0.5$, we now compute the epidemic threshold of our SIR_{MF} model by using the Microscopic Markov Chain Approach (MMCA). The MMCA equations are derived from the scheme shown in Fig. 9 which describes the transition probability of the states in the multiplex framework (Fig. 2). At time t , a susceptible user v on the top (mention) layer gets mentioned by an infected user and switch its state to mention-susceptible (MS); rest of the susceptible nodes on the follow (bottom) layer remain at the unmention-susceptible (US) state. In the next timestep ($t+1$), both kinds of nodes are allowed to retweet and switch to infected (I) state. However, there is a possibility that at time $t+1$ one user v does not get infected and hence switch back to susceptible (S) state. This event enables v to be mentioned or unmentioned in the next subsequent steps.

Let $A = (a_{vz})$ denote the adjacency matrix that describes the follow links between individuals. N is the total number of individuals. On the follow layer, each individual v has a certain probability of being in one of the three states at time t , given by $p_v^S(t)$, $p_v^I(t)$ and $p_v^R(t)$, respectively. Similarly, on the mention layer, for each (susceptible) individual v , the probabilities of being in one of the two states are denoted

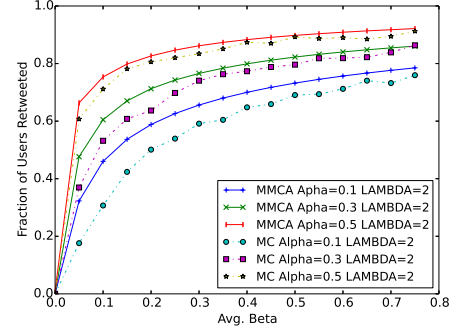


Fig. 10. Comparison of the MMCA model and the Monte Carlo approach w.r.t. R_U in "Algeria" dataset.

by $p_v^U(t)$ and $p_v^M(t)$, respectively for unmentioned-susceptible and mentioned-susceptible.

Let the total number of individuals infecting neighbors at time t be denoted by $I(t)$, that is, $I(t) = \sum_v p_v^I(t)$. Thereby, the probability that a (susceptible) individual v is in mentioned state at time t is

$$p_v^M(t) = \frac{\lambda I(t)}{N} \quad (1)$$

Since all susceptible users are either mentioned or unmentioned, given an individual v at time t , the probability to be in the mentioned-susceptible state $p_v^{MS}(t)$ is the product of the probabilities of being mentioned and susceptible. The same applies for the unmentioned-susceptible state $p_v^{US}(t)$. Hence

$$\begin{aligned} p_v^{MS}(t) &= p_v^M(t)p_v^S(t) = \frac{\lambda I(t)}{N} p_v^S(t) \\ p_v^{US}(t) &= (1 - p_v^M(t))p_v^S(t) = \frac{N - \lambda I(t)}{N} p_v^S(t) \end{aligned} \quad (2)$$

The transition probability for a susceptible individual v not to be infected by any neighbor through a follow link is $r_v(t) = \prod_z (1 - a_{zv} p_z^I(t) \beta)$ and the probability for v not to be infected through mention is $(1 - \alpha)$. It follows that the transition probabilities for an individual v unmentioned-susceptible ($q_v^{US}(t)$) or mentioned-susceptible ($q_v^{MS}(t)$) not to be infected are

$$\begin{aligned} q_v^{US}(t) &= r_v(t) \\ q_v^{MS}(t) &= r_v(t)(1 - \alpha) \end{aligned} \quad (3)$$

By using Eqs. 3, we can develop the Microscopic Markov Chains for the epidemic spreading process for each node v :

$$p_v^S(t+1) = p_v^{US}(t)q_v^{US}(t) + p_v^{MS}(t)q_v^{MS}(t) \quad (4)$$

$$p_v^I(t+1) = p_v^{US}(t)(1 - q_v^{US}(t)) + p_v^{MS}(t)(1 - q_v^{MS}(t)) - p_v^I(t) \quad (5)$$

$$p_v^R(t+1) = p_v^R(t) + p_v^I(t) \quad (6)$$

To validate our MMCA based analytical model, we compare them with Monte-Carlo simulation (MC). In MC, each simulation starts with a single infected node that is randomly chosen among the individuals, while in MMCA, each individual is initially infected with probability $p_v^I(0) = \frac{1}{N}$ ($p_v^I(0) \approx 0.00005$ in our simulations). Here, we represent the fraction of users retweeted as R_U . In MC, each simulation

stops when no susceptible node gets infected, while in MMCA, $R_U = \sum p_v^R(t)$ when $I(t) < 10^{-7}$. One can observe in Fig. 10 that the results based on the two approaches are in good agreement. We also note that MMCA upper-bounds MC simulations for a large range of β . This is caused by the mean field theory (MMCA) which assumes that events are independent [22].

Since MMCA provides the results that closely approximate MC simulation, we will derive the epidemic threshold from MMCA equations. The epidemic threshold determines whether the epidemic can outbreak or die out. Let us assume the existence of a critical point β_c for fixed parameters: α and λ , i.e. the epidemic will die out if $\beta < \beta_c$. The calculation of this critical point is performed by considering that when $\beta \rightarrow \beta_c$, the probability of nodes being infected $p_v^I \approx \epsilon_v \ll 1$. The smaller the probability of nodes being infected, the faster the epidemic dying out. Consequently, $q_v^{US} \approx 1 - \beta \sum_z a_{vz} \epsilon_z$, $q_v^{MS} \approx 1 - \alpha - \beta \sum_z a_{vz} \epsilon_z$ and $p_v^M = \frac{\lambda \sum_z \epsilon_z}{N} \approx \lambda \epsilon_v$. Inserting this and Eqs. 2 in Eq. 5, we obtain

$$\begin{aligned} \epsilon_v &= (p_v^{US} + p_v^{MS})\beta \sum_z a_{vz} \epsilon_z + \alpha p_v^{MS} \\ &= p_v^S [\beta \sum_z a_{vz} \epsilon_z + \alpha \lambda \epsilon_v] \end{aligned}$$

and therefore,

$$\sum_v \left[a_{vz} - \frac{\delta_{vz} (\frac{1}{p_v^S} - \alpha \lambda)}{\beta} \right] \epsilon_v = 0, \quad (7)$$

where δ_{vz} are the elements of the identity matrix such that $\delta_{vz} = 1$ if $v = z$; otherwise, $\delta_{vz} = 0$. We can rewrite the solution of Eq. 7 into the form: $A = \frac{1 - \alpha \lambda}{\beta} I_N$ by taking place for $t \rightarrow 0$ and $p_v^S(0) \rightarrow 1$. According to Frobenius theorem, the vector A is equal to the vector $I_N = (\delta_{vz})$ only if $\frac{1 - \alpha \lambda}{\beta}$ is equal to the maximum eigenvalue of A denoted by $\Lambda_{max}(A)$. Hence

$$\beta \times \Lambda_{max}(A) + \alpha \lambda = 1 \quad (8)$$

Intuitive Justification: Basically the terms in the left hand side of Eq. 8 represents the contribution of follow-links and mention-links on the total number of infected users. The infection via follow-links depends on two factors - (i) the density of the follow-network (represented by $\Lambda_{max}(A)$) and (ii) the probability of retweeting via follow-links i.e. β . Similarly, the infection via mention-links depends on the average number of mentions per tweet i.e. λ and the probability of retweeting via mention-links i.e. α . For both the terms if any of the factor is 0, the contribution of that part will be nullified regardless of how big the other factor is. For example, if α is 0 even if λ is very high, there will be no infection via mention-links and vice versa.

Eq. 8 is the key equation of our analysis which can be used to derive critical values of both α and β . For example, from Eq. 8 critical β can be derived as

$$\beta_c = \frac{1 - \alpha \lambda}{\Lambda_{max}(A)} \quad (9)$$

Note that β_c not only depends on the eigenvalue $\Lambda_{max}(A)$,

Content Attributes	Behavioral Attributes
#Words per tweet	#Followees
#Characters per tweet	#Followers
#URLs per tweet	#Followers/#followees
#Hashtags per tweet	#Tweets
#Users mentioned per tweet	Age of account
#Retweets per tweet	#times mentioned

TABLE II
EXAMPLES OF CONTENT AND BEHAVIORAL ATTRIBUTES USED FOR SPAMMER DETECTION

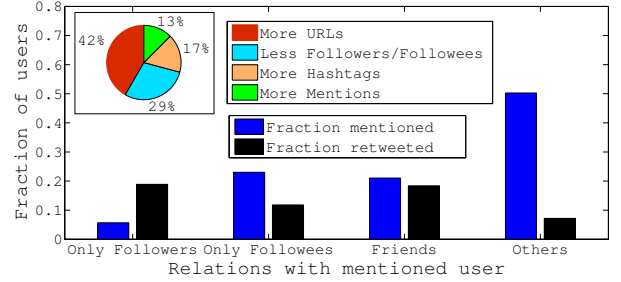


Fig. 11. Probabilities of mentioning users with different relations (reciprocal followers are denoted as 'Friends' here) and their probabilities of retweeting in "Egypt" dataset. Inset shows the annotators' major reasons of labeling users as spammers for "Egypt" dataset.

but also on α and λ . For $\alpha > \frac{1}{\lambda}$, β_c does not exist.

Next, we compare our computed critical epidemic threshold to the result obtained by using MC simulation, as shown in Fig. 7. Considering that MMCA upper-bounds MC simulations, the epidemic threshold obtained by MC will be larger than β_c given by Eq. 9. We extend Eq. 9 for MC simulations by following our observation and [22], which is

$$\beta_c^{MC} = \frac{1 - \alpha \lambda}{0.7 \Lambda_{max}(A)} \quad (10)$$

We obtain $\Lambda_{max}(A) \approx 117.5$ for the follow network of Algeria dataset. Consequently, for $\alpha = 0$, $\beta_c^{MC} = 0.012$; for $\alpha = 0.1$, $\beta_c^{MC} = 0.0097$; for $\alpha = 0.5$, $\beta_c^{MC} = 0$. Hence this is nice to observe that the results derived from Eq. 10 and our MC simulations in Fig. 7 are identical. Their good agreement verifies the effectiveness of our computed critical epidemic threshold. Similar agreement holds for critical α too.

V. EASY-MENTION: RECOMMENDATION HEURISTIC

In this section, we propose a Mention Recommendation heuristic named *Easy-Mention* which is easily deployable in online systems. The design of the *Easy-Mention* is mostly driven by the insights obtained from the model proposed in the previous sections. We show its effectiveness by comparing it with the benchmark algorithms.

A. Development of Easy-Mention

The objective of the *Easy-Mention* heuristic is to recommend the user, while she posts a tweet, the best set of users to mention in order to boost the retweet count of that tweet. Hence, the input of the heuristic is the submitted tweet and

the output is a ranked list of users to be mentioned. The three major stages of this recommendation are the following.

1) *Detect spammers*: The first stage of *Easy-Mention* is to protect the application from the malicious users. It is expected that any mention recommendation system has a high potential to be exploited by spammers for spreading their spam tweets. We implement a spammer detection algorithm (inspired from [23]) at the first stage, to refrain spammers from using our service². If this stage detects one user as a potential spammer, *Easy-Mention* terminates immediately. In this spammer detection algorithm, we crawl her recent tweets and focus on the following two class of features.

Content attributes: Content attributes are features of the tweet text posted by the users, which capture specific properties related to the way people write tweets. Studies show that, in general, spammers post tweets with higher number of hyperlinks, mentions and hashtags compared to non-spammers [23]. We analyze the tweet content characteristics based on the maximum, minimum, average, and median of the features shown in Table. II. In total, we consider 39 attributes related to the content of tweets for spammer classification.

Behavioral attributes: Behavioral attributes capture specific features connected to user behavior in terms of the posting frequency, social interactions and influence on the Twitter network. Admittedly, spammers have a lower followers to followees ratio than non-spammers and they generally possess recent accounts (less age) since Twitter continuously suspends potential spammers [23]). We consider 23 different features connected to user’s behavioral attributes as summarized in Table. II.

We evaluate the performance of the algorithm on the “Algeria” and “Egypt” datasets; the major challenge is the ground truth labeling of spammers and non-spammers. We train our model on a labeled spammer dataset available in [23]³. The model classifies 537 out of 20268 users in “Algeria” dataset and 27 out of 20092 users in “Egypt” dataset as spammers. During validation, this is comforting for us to notice that 10% of the detected accounts have already been suspended by Twitter. For the remaining 90% of the accounts detected as spammers, we perform a human survey with 3 volunteers and they labeled 89% of them as true spammers unanimously by manually going through their profiles. Their justification and rationale are summarized in the inset of Fig. 11. In summary, stage I efficiently performs the spammer detection in our datasets.

2) *Identify the candidate users*: In stage II, we narrow down the search space for ranking and recommending the users to be mentioned. This includes two steps. In the first step, we identify the keywords in the submitted tweet (hashtags and proper nouns) and search for the followers and followees, who recently posted them. This is a quick way to collect a reasonable set of active users who are interested in the post. In general, we find that if users are mentioned within one hop

neighborhood (happens in 50% of cases), they have higher probabilities of retweeting (see Fig. 11). Moreover, selecting the candidates from the one hop neighbors may significantly reduce the spamming threat for *Easy-Mention*.

In the second step, we exploit the critical thresholds obtained in *SIR_{MF}* model to further narrow down the candidate set. Specificity, we leverage on Eq. 10 and remove all the candidate users with α below the estimated threshold, from the candidate set. This step enhances the quality of the candidate set by keeping only the promising nodes (for virality) in the set. In case, we fail to find any suitable user with α above the critical threshold, we go ahead with the candidate set obtained in the first step and flash suitable warning message to the end user regarding the possibility of ‘non-cascading behavior’ of the post. This warning may help the user in deciding on the suitability and usefulness of the recommendation.

3) *Calculating a score for each candidate user*: In stage III, *Easy-Mention* assigns a quality score to each candidate user u . This score basically signifies the expected gain in popularity of tweet T , if u is mentioned in T . The data study and *SIR_{MF}* model show that the following factors may regulate the quality score (i) follower count (f_P): this is motivated from the smart mention strategy described in the previous section (ii) retweet rate (f_R): this factor captures the general retweet rate of an user. This is motivated from the retweet rate β of the *SIR_{MF}* model (iii) the content similarity (f_I) between the posted tweet T and the profile of the mentioned user u : a mentioned user with higher content similarity has high propensity to retweet. This essentially captures the notion of α in *SIR_{MF}* model.

Finally, the score $S(u, T)$ is computed for each candidate user u related to a submitted tweet T . In order to estimate the score $S(u, T)$, we simply use the regression models to suitably combine the key features $f_P(u)$, $f_R(u)$ and $f_I(u, T)$ ($f_P(u)$ is u ’s normalized follower-count, $f_R(u)$ is her normalized retweet rate and $f_I(u, T)$ is the similarity between the profile of u and the tweet T) to optimize ‘Relevance’ introduced in [13]. Relevance of a user-tweet pair (say, u & T) is calculated as the sum of the follower counts of the (re)tweeting users in the cascade subtree (of tweet T) rooted by u . In other terms, relevance for a user-tweet pair measures the visibility brought by the user u to the tweet T .

To represent the profile of u in real-time, we use the term-vector T_u^V created from the words (after stemming and stopwords-removal) in u ’s past retweets⁴. In the same way, we create another term-vector T_T^V for the submitted tweet T and finally calculate $f_I(u, T)$ as the cosine-similarity between these two term vectors (T_u^V & T_T^V). The score $S(u, T)$ assigned to each user u prepares the ranked list of candidate mention users who maximize the expected visibility.

B. Experimental setup

Now we aim to evaluate the performance of *Easy-Mention* with the competing benchmark algorithms. In order to accomplish the task, first (a) we implement a standard retweet model

²The details of the spammer detection methodology is out of the scope of this paper.

³<http://homepages.dcc.ufmg.br/~fabricio/spammerscollection.html>

⁴In the evaluation experiments, we compose the profile of a user from all her retweets in the dataset.

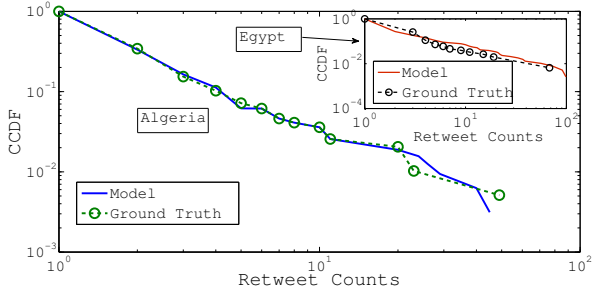


Fig. 12. Comparison of the tweet-popularity distribution from the “Algeria” dataset and the model. The inset shows the same for the “Egypt” dataset.

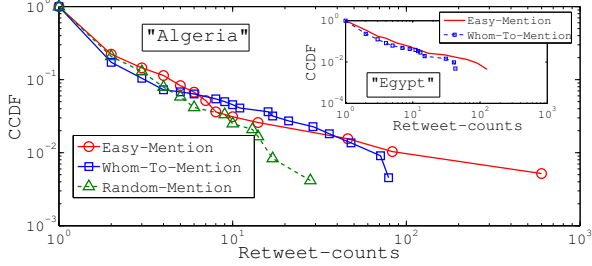


Fig. 13. CCDF of retweet-counts of tweets using different mention strategies for “Algeria” and “Egypt” datasets.

which simulates the propagation of the tweets via *retweet* activity. Next (b) on top of the retweet model, we implement the mention recommendation algorithms to evaluate the performance of *Easy-Mention*.

1) *Retweet Model*: We choose a well-accepted retweet model by Vespignani et al. [24]. It basically deals with competing memes in social networks and employs a parsimonious agent based model to study whether such a competition may affect the popularity of different memes. Since this is only a retweet model and does not handle the mention utility separately, we adapt it to include the mention utility in the following way. First we construct a tweet corpus D_T from each of the “Algeria” and “Egypt” datasets such that only 50% of tweets contain mentions. In order to post a new tweet or retweet, one user is chosen preferentially based on her retweet rate. If she chooses to post a new tweet, one tweet is selected randomly from D_T and she tweets the post with the same number of mentions (including zero) as in the original tweet. The specific users to be mentioned in that tweet is regulated by the specific “mention-recommendation” algorithm. The other possibility is that she opts to retweet an already received post. For each user u , we maintain a ‘screen window’ and a ‘mention window’ where tweets received via follow links (retweet from the followees of u) and tweets received via mention links (tweets where u has been mentioned) are stored respectively. If the selected user u chooses to retweet, one of these two windows is chosen based on its similarity with the profile of u (computed as cosine-similarity of term vectors) and then the most similar post (with respect to u ’s profile) in that window is retweeted. However, there is a fair possibility of not retweeting any post, if the context similarity is below a threshold. The value of the threshold is fixed externally

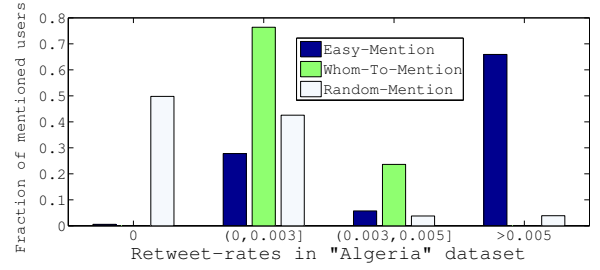


Fig. 14. Comparison of retweet-rates (in “Algeria” dataset) of users mentioned by competing recommendation algorithms.

Dataset	Algorithms	R_U	F_M	$R_U - N_U$	F_C
“Algeria”	<i>Easy-Mention</i>	2.52	0.136	0.69	0.087
	<i>Whom-To-Mention</i> [13]	1.77	0.012	-1.31	0.024
“Egypt”	<i>Easy-Mention</i>	2.32	0.588	1.22	0.195
	<i>Whom-To-Mention</i> [13]	1.38	0.307	-0.19	0.029

TABLE III
METRIC VALUES FOR DIFFERENT MENTIONING STRATEGIES APPLIED ON “ALGERIA” AND “EGYPT” DATASETS.

depending on the tweet environment.

In order to validate, we simulate this retweet model on “Algeria” & “Egypt” datasets with posts containing no mentions. It is comforting for us to observe that the result explains the heterogeneity in the tweet popularity distribution with reasonable accuracy (see Fig. 12). Now we are ready to use this retweet model to evaluate the performance of different mention recommendation heuristics.

2) *Competing algorithms*: On top of this retweet model, we apply the proposed mention recommendation heuristic *Easy-Mention* and compare its performance with the baseline algorithms *Whom-to-mention* [13] and *Random-mention*. The outline of the baseline algorithms are given below

1. Whom-to-Mention: To the best of our knowledge, *Whom-To-Mention* [13] is currently the most widely accepted state-of-the-art mention recommendation algorithm. In this algorithm, whenever a user u wishes to mention somebody in her tweet T , all the users in Twitter are considered as potential users to mention. In order to rank these potential users, three types of features are extracted - (a) Interest-match between the post and users’ recent tweets (b) Social-Tie and (c) User influence. Finally an SVR (Support Vector Regression) based system is used to rank these users, taking into account the average depth of the retweet cascades created by them.

2. Random Mention: This is a baseline algorithm where the recommended users to be mentioned are chosen randomly from the set of users in the dataset.

C. Performance evaluation

Finally, we perform the experiments on the “Algeria” and “Egypt” datasets (tweets and follower network); the evaluation metrics are already introduced in the section III. In this experiment, while posting a tweet T , we remove the original mentions from the tweet T and replace each mention by the

username selected by the specific mention recommendation algorithm. To ensure fairness, we keep the same number of mentioned users in each tweet as in the original tweet. Once the users to be mentioned are identified, we simulate the retweet model.

Fig. 13 clearly illustrates the fact that *Easy-Mention* outperforms the other competing algorithms in achieving tweets with higher retweet counts. Delving into the details, in Table. III we enumerate the observed evaluation metrics for different mentioning algorithms. Table. III uncovers the rationale behind the superiority of the *Easy-Mention*. It can be clearly observed that *Easy-Mention* is able to mention those users who not only frequently retweet that post (high F_C) but also are popular enough to give the tweet high visibility (the average follower count of users recommended by *Easy-Mention* is 158.4 whereas the same for *Whom-To-Mention* & *Random-Mention* are 93.1 & 28.2 respectively). This in turn helps *Easy-Mention* to achieve more retweets for the posts with mentions (R_U) than posts without mentions (N_U). Moreover, Fig. 14 points to the fact that, the mentioned users in case of *Easy-Mention* retweets more frequently compared to the competing algorithms; this directly contributes to the cascade size. In summary, all these properties help *Easy-Mention* to popularize tweets effectively by creating more cascades and larger ones. Importantly, Eq. 10 confirms virality of 86.5% of the tweets while using *Easy-mention* and the tweets satisfying this criteria get on average 30% more retweets than others.

VI. CONCLUSION

In this paper, we offer an in-depth study on explaining the role of mentions on tweet virality. We have identified that a significant fraction (sometimes even up to 50%-60%) of retweets might disappear if people stop using mentions (see Fig. 1(a), 1(b)). In order to have a detailed understanding, we have proposed a SIR based epidemic model SIR_{MF} to mimic the propagation of tweets on the mention-follow multiplex framework. Using a novel MMCA based analytical model, we have identified the critical threshold on retweet-rates (α & β), regulating the role of mention utility on tweet cascading effect. Exploiting the insights obtained from the motivational studies and modeling experiments, we have extracted the following three key-parameters controlling the effectiveness of mentioning: follower-count, retweet-rate & content-similarity and proposed our *Easy-Mention* recommendation heuristic. We have shown that the proposed approach outperforms the state of the art *Whom-To-Mention* algorithm [13] in the yardstick of performance. Nevertheless, the state of the art ‘Influence maximization’ algorithms ([12], [8]) may open up new possibilities for further improvement of the *Easy-Mention* heuristic.

VII. ACKNOWLEDGEMENTS

This work has been partially supported by the SAP Labs India Doctoral Fellowship program, DST - CNRS funded Indo - French collaborative project titled “Evolving Communities and Information Spreading” and the French National Research Agency contract CODDDE ANR-13-CORD-0017-01.

REFERENCES

- [1] S. González-Bailón, J. Borge-Holthoefer, A. Rivero, and Y. Moreno, “The dynamics of protest recruitment through an online network,” *Scientific reports*, vol. 1, 2011.
- [2] S. M. Kywe, T.-A. Hoang, E.-P. Lim, and F. Zhu, “On recommending hashtags in twitter networks,” in *Social Informatics*. Springer, 2012, pp. 337–350.
- [3] S. Kato, A. Koide, T. Fushimi, K. Saito, and H. Motoda, “Network analysis of three twitter functions: Favorite, follow and mention,” in *Knowledge Management and Acquisition for Intelligent Systems*. Springer Berlin Heidelberg, 2012, vol. 7457, pp. 298–312.
- [4] C. Freitas, F. Benevenuto, S. Ghosh, and A. Veloso, “Reverse engineering socialbot infiltration strategies in twitter,” in *ASONAM '15*. New York, NY, USA: ACM, 2015, pp. 25–32.
- [5] S. Petrovic, M. Osborne, and V. Lavrenko, “Rt to win! predicting message propagation in twitter,” in *ICWSM '11*. The AAAI Press, 2011.
- [6] A. Malhotra, C. K. Malhotra, and A. See, “How to get your messages retweeted,” *MIT Sloan Management Review*, vol. 53, no. 2, pp. 61–66, 2012.
- [7] I. Uysal and W. B. Croft, “User oriented tweet ranking: A filtering approach to microblogs,” in *CIKM '11*. New York, NY, USA: ACM, 2011, pp. 2261–2264.
- [8] D. Kempe, J. Kleinberg, and E. Tardos, “Maximizing the spread of influence through a social network,” in *SIGKDD '03*. New York, NY, USA: ACM, 2003, pp. 137–146.
- [9] K. Saito, R. Nakano, and M. Kimura, “Prediction of information diffusion probabilities for independent cascade model,” in *Knowledge-based intelligent information and engineering systems*. Springer, 2008, pp. 67–75.
- [10] M. Cha, H. Haddadi, F. Benevenuto, and P. K. Gummadi, “Measuring user influence in twitter: The million follower fallacy,” *ICWSM '10*, vol. 10, no. 10-17, p. 30, 2010.
- [11] J. Borge-Holthoefer, A. Rivero, and Y. Moreno, “Locating privileged spreaders on an online social network,” *Physical review E*, vol. 85, no. 6, p. 066123, 2012.
- [12] W. Chen, Y. Wang, and S. Yang, “Efficient influence maximization in social networks,” in *SIGKDD '09*. New York, NY, USA: ACM, 2009, pp. 199–208.
- [13] B. Wang, C. Wang, J. Bu, C. Chen, W. V. Zhang, D. Cai, and X. He, “Whom to mention: Expand the diffusion of tweets by @ recommendation on micro-blogging systems,” in *WWW '13*. New York, NY, USA: ACM, 2013, pp. 1331–1340.
- [14] Y. Gong, Q. Zhang, X. Sun, and X. Huang, “Who will you @?” in *CIKM '15*. New York, NY, USA: ACM, 2015, pp. 533–542.
- [15] G. Zhou, L. Yu, C. Zhang, C. Liu, Z. Zhang, and J. Zhang, “A novel approach for generating personalized mention list on micro-blogging system,” in *ICDM Workshop, 2015*, 2015, pp. 1368–1374.
- [16] K. Lee, J. Mahmud, J. Chen, M. Zhou, and J. Nichols, “Who will retweet this? detecting strangers from twitter to retweet information,” in *IUI '14*. ACM, 2014, pp. 247–256.
- [17] L. Tang, Z. Ni, H. Xiong, and H. Zhu, “Locating targets through mention in twitter,” *World Wide Web*, pp. 1–31, 2014.
- [18] C. Granell, S. Gómez, and A. Arenas, “Dynamical interplay between awareness and epidemic spreading in multiplex networks,” *Physical review letters*, vol. 111, no. 12, p. 128701, 2013.
- [19] P. N. Howard, A. Duffy, D. Freelon, M. M. Hussain, W. Mari, and M. Mazaid, “Opening closed regimes: what was the role of social media during the arab spring?” *Available at SSRN 2595096*, 2011.
- [20] M. E. Newman, “Spread of epidemic disease on networks,” *Physical review E*, vol. 66, no. 1, p. 016128, 2002.
- [21] K. Lerman and R. Ghosh, “Information contagion: An empirical study of the spread of news on digg and twitter social networks,” *ICWSM*, vol. 10, pp. 90–97, 2010.
- [22] M. Youssef and C. Scoglio, “An individual-based approach to sir epidemics in contact networks,” *Journal of theoretical biology*, vol. 283, no. 1, pp. 136–144, 2011.
- [23] F. Benevenuto, G. Magno, T. Rodrigues, and V. Almeida, “Detecting spammers on twitter,” in *Collaboration, Electronic messaging, Anti-Abuse and Spam Conference (CEAS)*, 2010.
- [24] L. Weng, A. Flammini, A. Vespignani, and F. Menczer, “Competition among memes in a world with limited attention,” *Sci. Rep.*, vol. 2, no. 335, 2012.