

# Who is Retweeting the Tweeters? Modeling, Originating, and Promoting Behaviors in the Twitter Network

PALAKORN ACHANANUPARP, EE-PENG LIM, JING JIANG, and TUAN-ANH HOANG,  
Singapore Management University

13

Real-time microblogging systems such as Twitter offer users an easy and lightweight means to exchange information. Instead of writing formal and lengthy messages, microbloggers prefer to frequently broadcast several short messages to be read by other users. Only when messages are interesting, are they propagated further by the readers. In this article, we examine user behavior relevant to information propagation through microblogging. We specifically use retweeting activities among Twitter users to define and model *originating* and *promoting* behavior. We propose a basic model for measuring the two behaviors, a *mutual dependency* model, which considers the mutual relationships between the two behaviors, and a *range-based* model, which considers the *depth* and *reach* of users' original tweets. Next, we compare the three behavior models and contrast them with the existing work on modeling influential Twitter users. Last, to demonstrate their applicability, we further employ the behavior models to detect interesting events from sudden changes in aggregated information propagation behavior of Twitter users. The results will show that the proposed behavior models can be effectively applied to detect interesting events in the Twitter stream, compared to the baseline tweet-based approaches.

Categories and Subject Descriptors: J.4 [Social and Behavioral Sciences]:

General Terms: Human Factors, Measurement, Experimentation

Additional Key Words and Phrases: Twitter, retweet, weak retweet, behavior modeling, originating behavior, promoting behavior, information propagation, event detection

## ACM Reference Format:

Achananuparp, P., Lim, E.-P., Jiang, J., and Hoang, T.-A. 2012. Who is retweeting the tweeters? Modeling, originating, and promoting behaviors in the Twitter network. *ACM Trans. Manage. Inf. Syst.* 3, 3, Article 13 (October 2012), 30 pages.

DOI = 10.1145/2361256.2361258 <http://doi.acm.org/10.1145/2361256.2361258>

## 1. INTRODUCTION

### 1.1 Motivation

Twitter, one of the most popular microblogging services, has attracted many users to broadcast very short text messages (up to 140 characters), also known as *tweets*, to one another. Other than publishing tweets themselves, users can subscribe (or *follow* in Twitter terminology) other users so as to receive the tweet feeds from the latter. Users who have follow links to other users are known as *followers* while users who have follow links from other users are known as *friends* or *followees* of the latter. The act of tweeting can be motivated by many different reasons. Users tweet about their

---

This research is supported by the Singapore National Research Foundation under its International Research Centre @ Singapore Funding Initiative and administered by the IDM Programme Office.

Author's address: P. Achananuparp; email: palakorna@smu.edu.sg

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permission may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701, USA, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

© 2012 ACM 2158-656X/2012/10-ART13 \$15.00

DOI 10.1145/2361256.2361258 <http://doi.acm.org/10.1145/2361256.2361258>

interests and hobbies, e.g., sharing information about their idols [Kwak et al. 2010]. Psychology research has also suggested that some users tweet to relieve existential anxiety [Qiu et al. 2010].

Despite the fact that a multitude of tweets are published every minute, a large portion of tweets do not carry informative content [Naaman et al. 2010]. Nevertheless, due to its very lightweight and real-time nature, Twitter has been extremely popular for fast, sometimes even serious, exchange of information. For example, Mendoza et al. [2010] analyzed the usage of Twitter in reporting the earthquake and rescue activities in Chile. While these are good examples of tweeting about serious events, there are many challenging issues that prevent the useful and interesting tweets from being uncovered. First, there are simply too many tweets being generated by a massive number of users. The large-scale data coupled with the real-time data characteristics prevent a complete analysis of tweet content. Second, each tweet has little content, and the writing style can be highly informal. Low quality and informal content contributes to noise in the Twitter data. Due to its noisy nature, traditional natural language (NLP) techniques do not work effectively when applied to tweet content. Finally, tweet data are difficult to collect. Using the APIs provided by Twitter can allow researchers to collect a snapshot of a tweets collection, instead of a complete set.

## 1.2 Objectives and Contributions

In this article, we view tweeting as a means of propagating information from one user to another. A piece of information is first mentioned in some tweet written by a user and gets propagated when another user sees the first tweet and mentions the information to her followers through retweeting (or passing on the content of the *original* or *source tweet*). Writing tweets interesting to other users coupled with retweeting others' tweets are therefore the two basic mechanisms that directly contribute to information propagation. We call the user behaviors associated with the two mechanisms the *originating* and *promoting* behaviors, respectively. A strong originating user is one who is likely to attract others to retweet her tweets. A strong promoting user is one who is likely to retweet others.

Studies have shown that users share content with one another for both altruistic and self-enhancement reasons [Berger and Milkman 2009]. By sharing content effectively, users strengthen their relationships with others and invite new ones. These are important motivations to Twitter users and users in social media generally. The way originating and promoting behaviors are defined is to characterize the interaction pattern a user demonstrates with his followers when sharing content with the latter.

Both originating and promoting behaviors are novel concepts not studied in the literature. The originating behavior is not equivalent to productive tweeting. Writing many tweets that do not interest others does not make an originating user. Similarly, prolific tweeters are not promoters if they do not tweet based on others' content.

We believe that the originating and promoting behaviors directly contribute to interesting tweeting and retweeting activities. We propose a behavior modeling framework that begins by finding connections among tweets through retweeting. As retweeting is relatively rare [Romero et al. 2011], we introduce the notion of *weak retweet* to find propagation connections among tweets using comentions of keywords, URLs, or hashtags. To model the originating and promoting behaviors, we propose to measure the two behaviors quantitatively using the basic and mutual dependency models. In addition, by tracking the behavior of users who mention some topics over time, one may be able to detect interesting events.

Behavior modeling and mining is a new topic in social media research. To the best of our knowledge, there has not been any work on modeling information propagation

behavior in the context of Twitter. Our work has specifically made the following research contributions.

- We propose a behavior modeling framework that can be used to study user behavior in Twitter. The framework divides the modeling process into, (a) interesting content item extraction, (b) retweet relationship construction or retweet linkage, and (c) behavior modeling.
- We introduce a way to extract interesting items from the tweet contents. This allows us to define weak retweet, which is different from the general retweet, also known as strict retweet. As shown in our experiments, weak tweets have much wider coverage of information propagation; approximately 14 times the number of strict retweets, in our experiments.
- We identify the originating and promoting behaviors and develop three models for measuring them. All behaviors are defined based on users' retweeting activities. The basic model measures the behavior directly using the propagation of retweets. The mutual dependency model exploits the mutual relationship between originating and promoting behaviors. Last, the range-based model incorporates information about the entire diffusion paths of the users' original tweets.
- We evaluate the models using a set of tweets generated by Singapore-based Twitter users and contrast the models with other related measures, including the recent Influential-Passivity (IP) algorithm of Romero et al. [2011]. We propose a method to detect events in Twitter by measuring changes in aggregated behavior. Our results will show that the originating and promoting behaviors can be effectively applied to find interesting events.

This article is organized as follows. First, we review related work in Section 2. Section 3 briefly characterizes a set of tweets used in this study. Next, we outline the proposed behavior modeling framework in Section 4. Section 5 introduces the originating and promoting behaviors in detail and presents the analysis of the behavior models in Section 6. Then, we discuss the application of behavior models in detecting interesting events from Twitter data in Section 7 and conclude the article in Section 8.

## 2. RELATED WORK

Modeling and analyzing user behavior has been one of the major topics in information retrieval research for many years [Agichtein et al. 2006; Morita and Shinoda 1994]. This work, though studying the behavior models, focuses on an entirely different problem domain. To the best of our knowledge, there has not been much work on modeling information propagation behavior in the context of Twitter. To some extent, this research conceptually builds upon the framework in previous work on modeling *engagement* and *responsiveness* behaviors in email exchange [On et al. 2010]. We describe the related work in two main research areas.

*Measuring User Influence.* Previous work has been done in measuring user influence in various social networks [Agarwal et al. 2008; Cha et al. 2010; Ghosh and Lerman 2010; Goyal et al. 2010; Kempe et al. 2003; Kwak et al. 2010; Romero et al. 2011; Subbian and Melville 2011; Yang and Leskovec 2010]. For example, Cha et al. [2010] and Kwak et al. [2010] measure the influence of Twitter users based on the sheer number of retweets spawned from the users' tweets. Weng et al. [2010] apply the PageRank algorithm [Page et al. 1999] to quantify the topic-sensitive influence of Twitter users. Recently, Wu et al. [2011] have studied the elite users who control a significant portion of the production, flow, and consumption of information in the Twitter network. Although one of the goals in this work is similar to that of Wu et al. [2011], we use a bottom-up approach to identify the top users based on their tweeting and

retweeting behaviors while Wu et al. [2011] use a top-down approach by identifying the top users based on how frequently the users appear in multiple user-generated lists. Both our work and that of Romero et al. [2011] employ retweeting interactions among Twitter users to compute their scores. However, our proposed models focus exclusively on the interactions between the users and their followers/followees. Furthermore, our mutual dependency model can be thought of as an inverse variant of the Romero et al. Influential-Passivity model. Conceptually, the mutual dependency model is closely related to Kleinberg's hubs and authorities model [Kleinberg 1999].

Our work is related but not identical to finding influential users. Influential users are those whose actions can influence or impact the actions of other users. On the other hand, we focus on the strength of their behavior. Users who exhibit strong originating and/or promoting behaviors are not necessarily influential and vice versa.

*Information Diffusion and Event Detection in Social Media.* Previous work has been done to analyze various aspects of information diffusion in social media [De Choudhury 2010; Gruhl et al. 2004; Leskovec et al. 2009; Yang and Leskovec 2010]. For instance, Leskovec et al. [2009] proposed a framework to track variants of short phrases being used by the mainstream media over the Web. Suh et al. [2010] studied factors contributing to the likelihood that a tweet will be retweeted. While some past research has focused on the dissemination of information at a message level, our work considers the propagation of information at a user level with respect to how the user's propagation behavior can help identify interesting events. Generally, the goal of an event-detection task is to detect sudden surges in some time series data and flag them as potential events. Many burst-like event detection methods have been proposed in the literature [Kleinberg 2003; Lin et al. 2010; Platakis et al. 2009; Wang et al. 2007; Zhao et al. 2007]. They have been broadly classified into, (a) offline and online, and (b) threshold-based and automaton-based burst detection methods [Kleinberg 2006]. Our work only focuses on offline burst detection using threshold-based methods. Recently, Sakaki et al. [2010] proposed a method utilizing the tweet contents as well as the associated geographical information as social sensors to detect earthquakes. Our work does not either rely on analyzing contents or location information to detect events. Furthermore, in contrast to the previous work in event-detection in text and social media streams [Becker et al. 2010; Mathioudakis and Koudas 2010; Popescu and Pennacchiotti 2010; Wang et al. 2007; Weng and Lee 2011; Zhao et al. 2007], which identifies events by employing message-level approaches, e.g. bursty keywords, classifiers trained on linguistic features, and so on, our work focuses on detecting events from changes in aggregated user behavior.

### 3. OVERVIEW OF SINGAPORE TWITTER DATASET

We begin by briefly introducing the Twitter dataset used in this study. We obtained the Singapore Twitter dataset by crawling tweets and follow-links of some selected Singapore-based Twitter users. A user is considered Singapore-based if she specifies "Singapore" as the time zone in her Twitter profile. The following procedures are employed.

- (1) Obtain the top 1000 Singapore-based Twitter users with the highest number of followers from twitterholic.com. Denote this set as  $V^*$ .
- (2) For each of the users in  $V^*$ , obtain her Singapore-based followers and friends within two hops. Aggregate the set of users obtained in this step and set  $V^*$ . Denote this aggregated set as  $V$ .
- (3) For each user in  $V$ , obtain all her published tweets whose content is in English. Denote the set of all the tweets as  $T$ .




Fig. 1. Number of followees and followers.

Twitter REST API<sup>1</sup> is used to facilitate the data collection. It is noted that the tweets collected in Step (3) are constrained by the Twitter API limitation, which sets the upper limit of each user's tweets available for collection at 3200 even though a user may have published more than 3200 tweets—resulting in a partial snapshot of  $V$ . The majority of the tweets collected are published for a 20-week period from the week of December 1, 2009 through the week of April 18, 2010. There are a total of 2,660,803 tweets collected. Those tweets are published by 18,703 unique users. All the results reported in the rest of this article are based on this range of data.

According to the log-log scatter plot of the numbers of incoming and outgoing follow-links of all users in Figure 1, a very small fraction of users follow more than 2000 users. Due to an upper limit currently imposed by Twitter, once users have followed 2000 users or followees, the number of additional users they can follow is constrained by their follower-to-followee ratio.<sup>2</sup> Only those who have more than 2000 followers are allowed to follow more than 2000 users.

In terms of follower-to-followee ratio, we observed that 42.6% of the users have more followers than followees while 53.8% have more followees than followers. In terms of tweet production and user popularity, we found that there is a low correlation between number of tweets and number of followees/followers, with Pearson's  $r < 0.05$ . It is reasonable to believe that users' information propagation behaviors do not need to correlate closely with their follow-link popularity. This observation holds true for larger datasets used in other studies as well [Cha et al. 2010; Kwak et al. 2010]. The most active user, who published 2943 tweets, has followed 648 users and has been followed by 1445 users. On the other hand, popular users who have more than 1000 followers have published tweets in the range of [10,100].

#### 4. BEHAVIOR MODELING FRAMEWORK

Figure 2 displays a proposed framework for modeling information propagation behavior in Twitter. We derive and measure behavior of Twitter users utilizing both follow-links and tweet content. We assume that the follow-links are static while tweet messages are timestamped and their contents are available. The proposed framework comprises three major steps as follows.

—*Step 1. Interesting Item Extraction.* We first extract interesting items from tweet messages that are discussed or shared among Twitter users. These interesting items allow us to select relevant tweets for the subsequent processing steps, thus

<sup>1</sup>Twitter API: <http://dev.twitter.com/doc>.

<sup>2</sup><http://support.twitter.com/articles/68916-following-rules-and-best-practices>



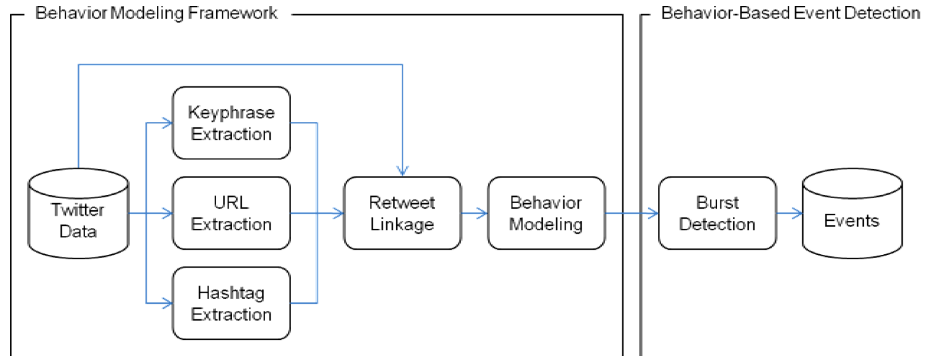


Fig. 2. The proposed behavior modeling framework for Twitter.

reducing the amount of noise in the raw tweet data. The interesting items include keyphrases, URLs, and hashtags. The details of this step are described in Section 4.1.

- *Step 2. Retweet Linkage.* In this step, we determine retweet relationships among tweet messages. Such relationships are indicative of meaningful interaction among Twitter users and we use them for modeling propagation behavior. Other than the *strict retweet linkage*, we also introduce a *weak retweet linkage*, which involves a user writing a tweet that mentions some interesting item in an earlier tweet written by another user. Section 4.2 describes this step in greater detail.
- *Step 3. Behavior Modeling.* Based on tweeting and retweeting activity, we formulate two information propagation behaviors for Twitter users, namely *originating* and *promoting*. They are formally defined in Section 5.

#### 4.1 Extraction of Interesting Items

The purpose of extracting interesting items from tweet messages is to establish the interactions among Twitter users. As suggested by Suh et al. [2010], tweets that contain interesting items, such as URLs and hashtags, are more likely to be retweeted by others. In this work, we consider three types of interesting items, namely (1) URL, (2) hashtag, and (3) keyphrase.

Twitter users often share interesting URLs with their followers. Due to the 140 character constraint, the original URLs are usually shortened using services such as bit.ly before being put into a tweet message. Next, hashtag, in the form of *#term*, is embedded in a tweet message to relate the message to some specific topic. Both URLs and hashtags can be easily extracted from the tweet messages. To fully distinguish unique URLs, we expand the shortened URLs to their full versions.

In addition, we also consider keyphrases in tweet messages as interesting items for propagation. The steps of extracting keyphrases are as follows. First, we preprocess tweet messages by removing tweet-specific terms, such as *@*, *RT*, and *via*. In addition, common stop words and internet slang words,<sup>3</sup> such as *lol*, *rofl*, *omg*, etc., are also removed from the tweet messages. Then, we extract a set of single-word tokens  $W$  from the preprocessed tweets and construct a cooccurrence matrix  $C$  from the set of tokens. Each row and column in  $C$  represents a single word while each element  $c_{ij}$  in  $C$  represents the number of times  $w_i$  cooccurs with  $w_j$  within a window size of  $N$  words. We set the value of  $N$  to 2 based on the optimal result obtained from the literature

<sup>3</sup>[http://simple.wikipedia.org/wiki/List\\_of\\_Internet\\_slang\\_words](http://simple.wikipedia.org/wiki/List_of_Internet_slang_words)




Fig. 3. Elapsed time when strict retweets are published with respect to the published time of the original tweets.

[Mihalcea and Tarau 2004]. Next, we rank each single-word token according to its score derived from the TextRank algorithm [Mihalcea and Tarau 2004] as described in the following.

Using the cooccurrence matrix  $C$  as an input, the algorithm iteratively performs random walks over an undirected lexical graph  $G(V, E)$ , where a set of vertices  $V$  represents the set of single words  $W$  and a set of edges  $E$  represents a cooccurrence relation between the vertices. The algorithm is run until convergence, after which we select the top- $k$  scoring words. Note that  $k$  is dynamically determined as a fraction of the number of vertices in the graph. In this case, a third of the number of vertices is selected. Then, we obtain candidate keyphrases by searching matrix  $C$  for the cooccurring top-ranked words. If two or more highly-ranked words occur within the predefined window size, we concatenate them to form a multiword keyphrase. If the top-ranked words do not cooccur with other top-ranked words, we include them in a set of single-word keyphrases. Ultimately, 6885 candidate keyphrases have been identified from all tweets in the dataset.

#### 4.2 Constructing Strict and Weak Retweet Linkages

In this step, we are constructing two kinds of retweet relationships: (a) strict retweet linkage, and (b) weak retweet linkage. The strict retweet linkage refers to finding tweets containing *RT @[username]* or *via @[username]* string. When a tweet  $tw_i$  contains any of these strings, we will locate an earlier tweet  $tw_j$  such that the  $tw_i$  is reposted based on  $tw_j$ . A strong retweet relationship  $tw_j \rightarrow tw_i$  is then constructed. In some cases,  $tw_j$  cannot be found due to unavailable data or  $tw_j$ 's user is not covered by the dataset. From our Singapore Twitter dataset, the strict retweet linkage step generates 7099 strict retweet relationships.

We examine both the elapsed time between tweets and their retweets and the depth of retweets. A temporal analysis of strict retweets in Figure 3 has shown that the time interval distribution for strict retweet relationships is skewed toward a small number of hours. Over 95% of retweets are posted within 24 hours after the original tweets (the source tweets of the retweet linkages) have been published. Next, a retweet linkage has a depth-defined based on the number of reposts from the original tweet to its retweets. The depth distribution of strict retweets is shown in Figure 4. As we can see, over 98% of the strict retweet linkages are direct reposts of other original tweets (depth of 1). Interestingly, the highly skewed distribution observed in our dataset is similar to that of a larger Twitter dataset [Kwak et al. 2010]. Nonetheless, there are a few factors affecting the distribution of our strict retweets. First, it is challenging to accurately infer a complete retweet linkage from the content of a strict retweet due to




Fig. 4. Depth of strict and weak retweet linkages.

the noisy nature of the tweet content. Second, since we restricted the Twitter users in this study to Singapore-based users only, we had to ignore many retweets referencing non Singapore-based users.

Given that strict retweets are not very common and may be too strict for capturing user interactions, we now introduce retweet linkages based on the common interesting items found among tweets. An original tweet is linked to its weak retweets by some interesting item mentioned in their contents. Suppose  $d$  represents an interesting item, and  $tw_i$  and  $tw_j$  represent two tweet messages, we say that  $tw_i$  is a weak retweet of  $tw_j$  ( $tw_j \mapsto tw_i$ ) if:

- $tw_i$  and  $tw_j$  contain  $d$ ;
- $tw_i$  and  $tw_j$  are published by  $u_i$  and  $u_j$ , respectively, and  $u_i$  follows  $u_j$ ;
- $tw_i$  and  $tw_j$  have timestamps  $t_i$  and  $t_j$ , respectively,  $0 < t_i - t_j < \delta$ ; and
- there are no other tweet messages  $m_k$  from other followers of  $u_i$  containing  $d$  and having timestamp  $t_k > t_j$  and  $0 < t_i - t_k < \delta$ .

According to the definition, a tweet can generate many weak retweets but a weak retweet can be associated with only one tweet.  $\delta$  is a delay threshold parameter value. Based on the temporal analysis of strict retweets, we have empirically set  $\delta = 24$  hours given that most strict retweets happen less than 24 hours after the original tweet message (see Figure 3). Depending on the types of interesting item  $d$ , a weak retweet can be (a) URL-based, (b) hashtag-based, or (c) keyphrase-based. We only focus on modeling user behavior from the keyphrase-based weak retweets ( $kp$ ) in this study, due to the higher interpretability of keyphrases, compared to that of URLs or hashtags, upon which we rely in the subsequent evaluation of the proposed behavior models.

Figure 5 graphically illustrates how weak retweet linkages are constructed given the publishing timelines. As depicted in Figure 5, each tweet is represented by a node. Newer tweets are placed progressively to the right along the timeline. In the top panel,  $tw_{i,1}$  and  $tw_{i,2}$  are linked to  $tw_j$  ( $tw_j \mapsto tw_{i,1}$  and  $tw_j \mapsto tw_{i,2}$ ) since they were published within  $\delta$  delay of  $tw_j$ . In the bottom panel,  $tw_{i,1}$ ,  $tw_{i,2}$ , and  $tw_{i,3}$  are linked to  $tw_k$  instead of  $tw_j$  ( $tw_k \mapsto tw_{i,1}$ ,  $tw_k \mapsto tw_{i,2}$ , and  $tw_k \mapsto tw_{i,3}$ ) since they were published within  $\delta$  delay of  $tw_k$  and  $tw_k$  was published after  $tw_j$ . As shown in Figure 4, the weak retweets tend to have more linkages with greater depth than the strict retweets. Nevertheless, it is extremely rare to observe the weak retweet linkages with a depth of 4 or greater.

Moreover, we compare the numbers of strict and weak retweets mentioning the popular keyphrases extracted by the method described in the previous section in Table I. In this case, “Justin Bieber” is the most widely mentioned multiword keyphrase while “CNY” (an abbreviated form of “Chinese New Year”) is the most






Fig. 5. An illustration of weak retweet linkages.

Table I. Examples of the Top Multiword and Single-Word Keyphrases (kp) According to the Number of Weak Retweets in which they are Mentioned

Numbers in parentheses represent strict retweets mentioning the same keyphrases.

Multi-word	# weak retweets	Single-word	# weak retweets
Justin Bieber	194 (16)	CNY	635 (10)
Jack Neo	189 (19)	Jonghyun	487 (6)
Apple iPhone	114 (4)	David	239 (26)
Universal Studios	82 (13)	Glee	180 (11)
Star Awards	62 (7)	Idol	178 (14)




Fig. 6. Interaction networks generated from strict retweets (left) and weak retweets data (right).

frequently used single-word keyphrase. We can see that the numbers of weak retweets containing these keyphrases are much larger, compared to those of strict retweets.

Next, we compare the characteristics of interaction networks based on strict retweet and weak retweet linkages. Figure 6 displays two directed graphs,  $G^R$  (left) and  $G^W$  (right), constructed from the strict retweet and weak retweet linkages, respectively.  $G^R = (V^R, E^R)$  is the strict retweet graph, where  $V^R$  is a set of vertices

Table II. Network Statistics of the Strict Retweet ( $G^R$ ) and Weak Retweet ( $G^W$ ) Graphs

Network Statistics	$G^R$	$G^W$
Number of vertices ( $ V $ )	2,870	7,139
Number of edges ( $ E $ )	4,092	29,925
Diameter	18	23
Average degree	1.426	4.192
Average weighted degree	2.467	12.006
Number of connected components	454	445

Table III. Singapore Twitter Dataset Statistic

# users ( $ U $ )	18,703
# follow-links	2,346,285
# tweets	2,660,803
# retweets	7,099
# unique hashtags	15.7K
# unique full URLs	321.2K
# keyphrase-based weak retweets ( $kp$ )	91.6K
# hashtag-based weak retweets ( $ht$ )	10.8K
# URL-based weak retweets ( $ur$ )	2.5K
# weak retweets ( $kp+ht+ur$ )	101K
# users of weak retweets ( $kp+ht+ur$ )	9K

representing Twitter users and  $E^R$  is a set of directed edges where an edge  $v_i \rightarrow v_j$  representing a retweet relation between a retweeting user  $v_i$  and a source user  $v_j$ . An edge weight represents the number of times  $v_i$  retweets  $v_j$ 's tweets. Similarly,  $G^W = (V^W, E^W)$  is derived from the weak retweet linkages. A directed edge in  $E^W$  represents a weak retweet relation  $v_i \mapsto v_j$ . As we can see, the weak retweet network is much denser than the strict retweet network.

Table II displays the network statistics of  $G^R$  and  $G^W$ . Overall,  $G^W$  is more connected than  $G^R$ . Moreover, users in  $G^W$  generally interact with one another more than those in  $G^R$ , as indicated by a higher average degree. Each user in  $G^W$  interacts (retweets from/is retweeted by) with 4.192 users while that in  $G^R$  interacts with 1.426 users on average. Next, we further examine the retweeting patterns among the notable users. Using the retweet relations in  $G^R$  and  $G^W$ , we compute for each user a *hub* score and an *authority* score [Kleinberg 1999] and calculate density [Wasserman and Faust 1994]  $d(G^R)$  and  $d(G^W)$  of the subgraphs induced from retweet interactions between the top 500 hub and authority users in each network,  $G^R$  and  $G^W$ . The finding suggests that there is a much higher mutual interaction between good hub and authority users in  $G^W$  than  $G^R$ , as the density of  $G^W$  is much higher than that of  $G^R$ ,  $d(G^W) = 0.033$  and  $d(G^R) = 0.002$ .

Last, since the weak retweets, constructed according to our definitions, are more likely to form a relatively longer chain with greater interactions among the users, compared to the strict retweets', they serve as a good alternative to the strict retweets for studying the users' information propagation behavior, especially given the technical challenges of parsing the strict retweets' content to inferring a complete linkage. Table III summarizes the overall statistics of the constructed Singapore-based weak retweets dataset. The number of weak retweets is about 14 times greater than that of strict retweets. Henceforth, we only use weak retweets as the underlying data for studying the behavior.

Table IV. Symbols

$t_i$	Number of unique tweets published by $u_i$
$t_{ij}$	Number of unique tweets published by $u_i$ retweeted by $u_j$
$T_i$	Set of tweets published by $u_i$
$T_i$	Set of tweets published by $u_i$ retweeted by her followers
$T_i$	Set of tweets published by $u_i$ 's followees retweeted by $u_i$
$u_i^{in}$	Number of unique followers retweeting $u_i$ 's tweets
$u_i^{out}$	Number of unique followees whose tweets are retweeted by $u_i$
$F_i^{in}$	Set of followers of $u_i$
$F_i^{out}$	Set of followees of $u_i$
$F_i^{out}$	Set of followees of $u_i$ whose tweets are retweeted by $u_i$
$F_{ik}^{in}$	Set of followers of $u_i$ retweeting $u_k$ 's tweets
$F_{ik}^{out}$	Set of followees of $u_i$ whose tweets are retweeted by $u_k$
$U_{ik}^{in}$	Set of $k$ th-hop unique followers retweeting $u_i$ 's tweets

## 5. BEHAVIOR MODELING

To measure the strength of information propagation behavior of Twitter users, we introduce originating and promoting behavior. Originating users are expected to seed new topics used in tweets published by other users while promoting users are eager to disseminate others' interesting tweets. We use the weak retweet relationships, described in the previous section, to model the originating and promoting behavior. In particular, the retweet interactions are confined to the Twitter users and their followees/followers only. Preliminarily, a few symbols to be used to formally describe the models are shown in Table IV.

*Robustness Assumption.* As the behavior models to be introduced are new and not yet published, we assume that the users have no incentives to abuse the mechanisms for quantifying behavior scores in order to achieve a higher ranking position.

Particularly, we assume that the retweeting activities are carried out by Twitter users with legitimate accounts. That is, for a set of Twitter users  $U$ , each user  $u_i$  in  $U$  is uniquely identified by  $uid_i$  given by the system. A legitimate user is one whose  $uid_i$  has a one-to-one mapping with a real-world unique identifier  $uid_i^R$ , which can be represented by any form of government-issued identification number, e.g., a social security number, a driver's license number, and so on. Simply put, one person is only allowed to have one user account. Given that the assumption holds, we can ensure that the behavior scores are contributed by valid propagation of information among legitimate users only.

Should the behavior models be made publicly known to Twitter users, we may need a more robust model to prevent abuse, especially from a potential case of retweet spamming. We discuss this and other related issues in greater detail in Section 5.4.

### 5.1 Basic Model

A basic model for the originating and promoting behaviors is to assign originator and promoter measures to a user  $u_i$ . Conceptually, we adopt two assumptions underlying the originating and promoting behaviors as follows.

- A user exhibits a strong originating behavior if a large fraction of her tweets have been retweeted by a large fraction of her followers. That is, a strong originator frequently publishes original tweets that are pertinent to many of her followers' interests.

- A user exhibits a strong promoting behavior if she retweets a large fraction of interesting tweets, i.e., those that have been retweeted by others, which she has seen from a large fraction of her followees. In other words, a strong promoter frequently propagates many interesting tweets from many unique sources that she follows.

Formally, we define the originator score of user  $u_i(O_i)$  and the promoter score of user  $u_i(P_i)$  as follows.

$$O_i = \sum_{j \in F_i^{in}} \frac{t_{ij}}{t_i} \times \frac{u_i^{in}}{|\bigcup_k F_{ik}^{in}|} \quad (1)$$

$$P_i = \sum_{j \in F_i^{out}} \frac{t_{ji}}{\sum_{k \in F_i^{out}} \sum_l t_{kl}} \times \frac{u_i^{out}}{|\bigcup_k F_{ik}^{out}|}. \quad (2)$$

The first component of  $O_i$  measures a proportion of tweets of user  $u_i$  being retweeted while the second component measures a proportion of followers of  $u_i$  who ever retweeted. Together,  $O_i$  is high when a significant portion of tweets by  $u_i$  has been retweeted by a significant portion of her unique followers. Similarly,  $P_i$  is high when  $u_i$  retweets a significant portion of unique and interesting tweets that  $u_i$  has seen in her tweet stream from a significant portion of her followees who have ever been retweeted.

## 5.2 Mutual Dependency Model

The mutual dependency model introduces a mutual reinforcement relationship between the originating and promoting behaviors by assuming the following.

- A user is a strong originator when most tweets by her are retweeted by many strong promoters.
- A user is a strong promoter when most retweets she has published are based on tweets published by strong originators.

As we can see, the underlying motivation of the mutual dependency model follows the classical hub and authority definitions [Kleinberg 1999] whereby a good authority (originator) is linked to many good hubs (promoters), and vice versa. With these assumptions, we formally define the new  $O_i^m$  and  $P_i^m$  measures as follows.

$$O_i^m = \sum_{j \in F_i^m} \frac{t_{ij} \times P_j^m}{t_i} \times \frac{u_i^{in}}{|\bigcup_k F_{ik}^{in}|} \quad (3)$$

$$P_i^m = \sum_{j \in F_i^{out}} \frac{t_{ji} \times O_j^m}{\sum_{k \in F_i^{out}} \sum_l t_{kl}} \times \frac{u_i^{out}}{|\bigcup_k F_{ik}^{out}|}. \quad (4)$$

The power iteration method is used to compute these measures. We first initialize  $O_i^m$  with the same value for all users. Then we compute new values for  $P_i^m$ 's, which are in turn used to compute new values for  $O_i^m$ 's. The process repeats until all values converge. On average, it takes less than 10 iterations for the values to reach convergence.

## 5.3 Range-Based Model

In addition, we extend the basic models by incorporating the indirect effects of users' originating and promoting behavior along the diffusion chains of tweets. Similar to the diffusion-based measures of Yang and Leskovec [2010], we first define the range of user's propagation behaviors as a function of average depth and reach of her tweets




Fig. 7. Diffusion chain of a tweet with depth of 2 and reach of 6. A colored node represents a user who published original tweets while other nodes represent retweeting users. Edges indicate the direction in which the user's tweet diffused.

propagated through the extended followers, or the  $K$ -hop followers of followers, where depth is the number of hops in a diffusion chain of a tweet and reach is the total number of unique users in her extended follower networks who help spread her original tweets. An example of depth and reach of a diffusion chain is shown in Figure 7. Based on these measures, the range-based behavior model assumes the following.

- A user is a strong originator if a large fraction of her tweets have been retweeted by a large fraction of her followers and her original tweets tend to be propagated in great depth and reach by her extended followers.
- A user is a strong promoter if she retweets a large fraction of interesting tweets she has seen from a large fraction of her followees and those interesting tweets tend to be propagated in greater depth and reach by the extended followers of her followees.

We formally define range-based originating ( $O_i^R$ ) and promoting ( $P_i^R$ ) behaviors of user  $u_i$  as follows.

$$O_i^R = \left( \frac{\sum_{j \in T_i} d_j}{|T_i|} \times \log \left| \bigcup_{k=1}^K U_{in}^{ik} \right| \right) \times O_i \quad (5)$$

$$P_i^R = \left( \frac{\sum_{j \in T_i} d_j}{|T_i|} \times \log \left| \bigcup_{j \in F_i^{out}} \bigcup_{k=1}^K U_{jk}^{in} \right| \right) \times P_i, \quad (6)$$

where  $K$  is the maximum depth of the diffusion chains of a given user,  $d_j$  is the depth of a tweet  $j$ , and  $U_{ik}^{in}$  is the set of  $k$ th-hop unique followers retweeting  $u_i$ 's tweets. The first component in the parenthesis is the average depth of the retweet chains generated by  $u_i$ 's original tweets (in the case of originating behavior) or  $u_i$ 's retweets (in the case of promoting behavior) while the second component is the total reach or magnitude of the impact. As we can see, users who tend to originate or promote long-range diffusion chains will be scored highly according to this model.

#### 5.4 Discussion

To conclude this section, we highlight some of the key challenges pertaining to the robustness and practicality of our proposed models in order to suggest some future directions of the work. First, as previously mentioned at the beginning of this Section, our proposed behavior models rely on the assumption that the retweeting activities are carried out by legitimate Twitter users only. However, in most real world scenarios, it is almost impossible to assume the ways in which user accounts are created and used. Most Web-based systems, including Twitter, only require a unique email address with which to associate a unique user account during the registration process. Other systems have tried to link a user account with a real-world identifier and had to abandon



the plan due to strong resistance from their users.<sup>4</sup> More importantly, many deviant users have exploited the fact that new user accounts can be easily registered in Twitter given an email address and have proceeded to populate the systems with spam, irrelevant or unsolicited bulk tweets, retweets, or follow links. This simple but persistent strategy can potentially inflate the originator and promoter scores of the users, given that such deceptive behavior cannot be differentiated from that which is expected.

Spamming is a well-known problem shared by all user behavior models and user importance models. Even for link analysis algorithms, such as PageRank [Page et al. 1999], HITS [Kleinberg 1999], and our mutual dependency models, a node can increase its importance or behavior score by injecting artificial links into the graph, effectively deteriorating the quality of the ranking results. The problem of link spam detection has been extensively studied, especially in a Web search domain [Castillo et al. 2007; Guha et al. 2004; Gyöngyi et al. 2004; Ntoulas et al. 2006]. As spamming detection itself is a specialized study, we leave the investigation of link spam detection methods on behavior models in Twitter for the future. Next, assuming that the majority of the retweeting activities are from legitimate users, the behavior models can be extended further to consider the various ways in which the retweeting activities occur. For example, it may be more interesting in some contexts to distinguish the rising originators and promoters whose current behavior scores have increased sharply from the top originators and promoters whose behavior scores are quite consistent over time. Next, among the high-ranking promoters, one may want to recognize early adopters or viral spotters, who tend to join the retweet cascades early on, from the late adopters or bandwagon jumpers, those who tend to join the cascades later after a large portion of promoters have already retweeted about the popular tweets. That is, by incorporating time dynamicity into the behavior models, more interesting types of users may potentially be uncovered.

Last, robustness in the computation of the mutual dependency scores can be further improved by introducing the normalization step in the same fashion as the HITS algorithm at the end of each iteration. This will help ensure that the behavior scores will converge at a faster rate, while the overall rankings remain largely unchanged. For instance, each  $O_m^i$  and  $P_m^i$  will be normalized by  $\sqrt{\sum_{i=1}^{|M|} O_i^{m^2}}$  and  $\sqrt{\sum_{i=1}^{|N|} P_i^{m^2}}$ ;  $M$  and  $N$  are the numbers of originators and promoters, at the end of the iteration, respectively.

## 6. ANALYSIS OF ORIGINATING AND PROMOTING BEHAVIORS

Throughout this section, we focus the analysis on active users only. To select the active users, we restrict the set of users to those who have published at least 10 tweets over a 20-week period. To ensure active user interaction, we only consider the users who have retweeted/been retweeted by an average of 7 unique users per week. In addition, we distinguish power users from non-power users. Specifically, given a set of users  $U^*$ , a user  $u_i$  is said to be a power user if the number of followers of  $u_i$  is greater than the number of followees by at least two orders of magnitude;  $u_i \in \{u_k \in U^* \mid f_k^{in} / f_k^{out} \geq 10^2\}$ . All power users are marked by asterisk.

### 6.1 Comparisons of Behavior in the Strict Retweet Linkages Data

First, we compare the originating and promoting users extracted from the weak retweet linkages to those from the strict retweet linkages. Table V displays the top-10 originators and promoters of the strict and weak retweet networks. Overall, the top users are relatively similar. Power users, news media users, and religious users

<sup>4</sup>[http://news.cnet.com/8301-17938\\_105-20010198-1.html](http://news.cnet.com/8301-17938_105-20010198-1.html)

Table V. Top-10 Basic Originators and Promoters Computed Using the Strict Retweet Linkages (top) and Weak Retweet Linkages (bottom)

Asterisk denotes a power user.

Rank	Originator		Promoter	
1	<i>stcom*</i>	Media	<i>trevidzi</i>	Student
2	<i>mrbrown</i>	Blogger	<i>victortan</i>	Blogger
3	<i>McDanielOng</i>	Blogger	<i>FeeKeeD</i>	Singer
4	<i>Seowhow</i>	Pastor	<i>CalvinTimo</i>	Blogger
5	<i>TweetSG</i>	Service	<i>STUMPBO</i>	Blogger
6	<i>hardwarezone*</i>	Portal	<i>shenheng</i>	Designer
7	<i>chcsg*</i>	Church	<i>mediumshawn</i>	Student
8	<i>cnalatest*</i>	Media	<i>winstongoh</i>	Blogger
9	<i>TaufikBatisah</i>	Singer	<i>thelensmen</i>	Entrepreneur
10	<i>TODAYonline*</i>	Media	<i>Ckaeteo</i>	Student

Rank	Originator		Promoter	
1	<i>stcom*</i>	Media	<i>rochorbeancurd</i>	Entrepreneur
2	<i>ChannelNewsAsia*</i>	Media	<i>alkanphel</i>	Blogger
3	<i>MediaAsia</i>	Media	<i>charlesyeo</i>	Entrepreneur
4	<i>chcsg*</i>	Church	<i>MediaAsia</i>	Media
5	<i>arieszulkarnain</i>	Pastor	<i>CalvinTimo</i>	Blogger
6	<i>mrbrown</i>	Blogger	<i>minzmint</i>	Student
7	<i>hardwarezone*</i>	Portal	<i>callsg</i>	N/A
8	<i>TaufikBatisah</i>	Singer	<i>LisahK</i>	Student
9	<i>Seowhow</i>	Pastor	<i>congyuan</i>	Student
10	<i>lynettegoh</i>	Student	<i>trevidzi</i>	Student

mostly dominate the top originator lists while students and bloggers comprise the majority of the top promoters. There are more common top originators than common top promoters between the two networks. Kendall's tau rank order correlation coefficient [Kendall 1938] between the top-100 originators from the strict retweet linkages and the weak retweet linkages is 0.54 while that of the top-100 promoter lists is 0.31.

The composition of the top promoters is more affected by the relaxed definition of propagation used in the weak retweet linkages than those of the top originators. As we can see, the common top originators in both networks are either news media or popular bloggers who tend to publish many tweets related to well-known events, such as breaking news, and/or interesting information sources. Therefore, their tweets and embedded topical keywords are likely to be propagated by similar sets of users through both strict and weak retweets. That is, the followers of the top originators will explicitly retweet the original tweets and/or subsequently publish their own tweets mentioning the same topics used in the popular tweets. As a result, these same originators will be scored highly in their originating behavior in the strict and weak retweet linkages.

On the other hand, most top promoters in both networks are non-media users who tweet about varieties of topics. Those tweets may include mentioning well-known events as well as personal conversations between the promoters and their followees. Therefore, the propagation behavior of the top promoters in the two networks is more divergent than those of the originators. Especially, the top promoters in the weak retweet network are more likely to propagate topical keywords from both formal (tweets about popular events seen within a 24-hour window) and informal tweets (conversational tweets carrying topical keywords from their followees). Because the weak retweet linkages cover both explicit and implicit propagation behavior, the subsequent analyses will be done on the propagation behavior of users in the weak retweet networks only.

Table VI. Kendall's tau Rank Order Correlation of the Top-100 Users  
The correlation values  $\geq 0.5$  are highlighted in boldface.

	$O$	$P$	$O^m$	$P^m$	$O^R$	$P^R$	$TCount$	$RTCount$
$O$	1	0.21	<b>0.56</b>	0.16	<b>0.65</b>	0.25	0.17	-0.11
$P$	-	1	0.21	<b>0.54</b>	0.16	<b>0.68</b>	0.11	0.11
$O^m$	-	-	1	0.17	0.46	0.24	0.15	0.05
$P^m$	-	-	-	1	0.16	0.36	0.14	0.34
$O^R$	-	-	-	-	1	0.10	0.01	-0.16
$P^R$	-	-	-	-	-	1	0.24	0.36
$TCount$	-	-	-	-	-	-	1	0.12
$RTCount$	-	-	-	-	-	-	-	1

Table VII. Top-10 Users Ranked by the Basic Originator and Promoter Scores, Respectively

Asterisk denotes a power user.

Rank	Originator		Promoter	
1	<i>stcom</i> *	Media	<i>rochorbeancurd</i>	Entrepreneur
2	<i>ChannelNewsAsia</i> *	Media	<i>alkanphel</i>	Blogger
3	<i>MediaAsia</i>	Media	<i>charlesyeo</i>	Entrepreneur
4	<i>chcsg</i> *	Church	<i>MediaAsia</i>	Media
5	<i>arieszulkarnain</i>	Pastor	<i>CalvinTimo</i>	Writer
6	<i>mrbrown</i>	Blogger	<i>minzmint</i>	Student
7	<i>hardwarezone</i> *	Portal	<i>callsg</i>	N/A
8	<i>TaufikBatisah</i>	Singer	<i>LisahK</i>	Student
9	<i>Seowhow</i>	Pastor	<i>congyuan</i>	Student
10	<i>lynettegoh</i>	Student	<i>trevidzi</i>	Student

## 6.2 Correlations of Different User Rankings

We examine the characteristics of the ranked lists of users produced by different measures in the weak retweet networks. For each user, we compute aggregated weekly scores for six behavior measures, including the originator ( $O$ ), promoter ( $P$ ), mutual originator ( $O^m$ ), mutual promoter ( $P^m$ ), ranged-based originator ( $O^R$ ), and ranged-based promoter ( $P^R$ ). For the baseline comparison, we simply rank the users by the number of tweets ( $TCount$ ) and retweets ( $RTCount$ ) they published.

According to Table VI, Kendall's tau rank order correlation coefficients between different ranked lists are not high, indicating that the behavior of the top originators is quite distinct from the top promoters'. This suggests that most users either behave as an originator or a promoter of interesting tweets.  $O$  produces more similar rankings to  $O^m$  and  $O^R$  and so do  $P$ ,  $P^m$ , and  $P^R$ . Last, the users with high behavior scores are not the same as those who actively tweet or retweet according to low correlation coefficients between the top behavior users and the baselines. This implies that highly active users are not necessarily the ones who produce interesting tweets.

## 6.3 Comparison of the Top-10 Users

Tables VII, VIII, and IX list the top-10 users ranked by different measures from the highest to the lowest scores. First, we examine the composition of power users and non-power users among the top-10 users. According to Tables VII, VIII, and IX, there are six unique power users on the top-10 users lists based, on  $O$ ,  $O^m$ , and  $O^R$ . On the other hand, no power user appears on the top-10 users lists based on either  $P$ ,  $P^m$ , or  $P^R$ . As we can see, the top users whose tweets have been retweeted many times tend to be power users while the top users who frequently retweet tend to be regular users.

Table VIII. Top-10 Users Ranked by the Mutual Originator and Promoter Scores, Respectively

Asterisk denotes a power user.

Rank	Mutual Originator		Mutual Promoter	
1	<i>stcom</i> *	Media	<i>CalvinTimo</i>	Writer
2	<i>SingaporeClub</i>	Aggregator	<i>sgbreakingnews</i>	Aggregator
3	<i>mrbrown</i>	Blogger	<i>LisahK</i>	Student
4	<i>lynettegoh</i>	Student	<i>ZairaOng</i>	Blogger
5	<i>chcsg</i> *	Church	<i>SingaporeClub</i>	Aggregator
6	<i>angjiehui</i>	Student	<i>rochorbeancurd</i>	Entrepreneur
7	<i>sginfomap</i>	Media	<i>alkanphel</i>	Blogger
8	<i>sglatestnews</i> *	Aggregator	<i>cynthiachul</i>	Student
9	<i>patlaw</i>	Entrepreneur	<i>Seowhow</i>	Pastor
10	<i>cnalatest</i> *	Media	<i>preius</i>	Student

Table IX. Top-10 Users Ranked by the Range-Based Originator and Promoter Scores, Respectively

Asterisk denotes a power user.

Rank	Range-Based Originator		Range-Based Promoter	
1	<i>stcom</i> *	Media	<i>alkanphel</i>	Blogger
2	<i>MediaAsia</i>	Media	<i>charlesyeo</i>	Entrepreneur
3	<i>chcsg</i> *	Church	<i>LisahK</i>	Student
4	<i>ChannelNewsAsia</i> *	Media	<i>rochorbeancurd</i>	Entrepreneur
5	<i>hardwarezone</i> *	Portal	<i>AngMoGirl</i>	Designer
6	<i>angjiehui</i>	Student	<i>CalvinTimo</i>	Writer
7	<i>cnalatest</i> *	Media	<i>moemasri</i>	Singer
8	<i>TaufikBatisah</i>	Singer	<i>deckstor</i>	Student
9	<i>SGnews</i> *	Media	<i>JadeChen85144</i>	Student
10	<i>patlaw</i>	Entrepreneur	<i>trevizdi</i>	Student

The top-ranked power users generally consist of news media, news aggregators, and celebrities, for example, *stcom* (The Straits Times, Singapore's best-selling newspaper), *ChannelNewsAsia* (a major local news network), and *HadyMirzaAmir* (Hady Mirza, a Singaporean singer). Similar findings of retweet-based influential users are also reported in Cha et al. [2010], Kwak et al. [2010], and Romero et al. [2011]. Furthermore, there are more news-media types of users in the top-10 range-based originators than the top-ten originators. This is not surprising since tweets published by news media tend to be diffused in a chain with longer range than those published by non-media users. Interestingly, *angjiehui*, whose interesting tweets are mostly related to religious topics, is ranked highly as a range-based originator only since the average depth of her tweets is more than 4.7, the highest among the top-10 ranged-based originators.

#### 6.4 Basic Model vs. Mutual Dependency Model

As the top mutual originators/promoters are relatively more distinct, compared to the top basic originators/promoters, we further discuss their differences in this section. Figure 8 displays six ego networks of the top-3 promoters (top row) and the top-3 mutual promoters (bottom row). A top promoter is located at the center of each network (*focal* node) while its neighbors are the users whose tweets are retweeted by the promoting user. To distinguish the ego networks of the basic promoters and the mutual promoters, we scale the neighboring node size proportionately to an originator score of the corresponding user. In addition, we mask out edges between neighboring nodes to declutter the graphs. Visually, the top mutual promoters frequently retweeted from many top originators judging by thicker edges between the focal nodes and the large




Fig. 8. Ego networks of the top-3 promoters (top row) and the top-3 mutual promoters (bottom row); the bigger the node size, the higher the originator score.




Fig. 9. Pairwise retweet distributions of the mutual originators (*SingaporeClub*, *sginfomap*, and *sglatestnews*) and the basic originators (*ChannelNewsAsia*, *MediaAsia*, and *arieszulkarnain*) given their top-5 followers, ranked by the promoter scores.

neighboring nodes. In contrast, the basic promoters retweeted from more users, most of which have low originator scores.

We further investigate the difference in the interaction patterns of the basic model users and the mutual dependency model users by analyzing their pairwise retweet distributions.

Figure 9 displays the pairwise retweet distributions between the mutual originators/originators and their top-5 followers who rank the highest by their promoter scores. According to the figure, tweets published by the mutual originators, *SingaporeClub*, *sginfomap*, and *sglatestnews*, have been frequently retweeted by a few top promoting followers. For example, about 90% of interesting tweets by *SingaporeClub* have been exclusively retweeted by its 2nd promoting follower. Cumulatively, virtually all *SingaporeClub*'s interesting tweets have been retweeted by its top-5 promoting followers. Similarly, over 80% and 90% of interesting tweets published by *sginfomap*






Fig. 10. Pairwise retweet distributions of the mutual promoters (*sgbreakingnews*, *ZairaOng*, and *Seowhow*) and the basic promoters (*charlesyeo*, *minzmint*, and *callsg*), given their top-5 followees ranked by the originator scores.

and *sghlatestnews* have been retweeted by their top-5 promoting followers, respectively. This explains why the three users have been ranked higher by  $O^m$  than  $O$ .

A similar relationship can be observed in the top mutual promoters. Figure 10 displays the pairwise retweet distributions between the mutual promoters (*sgbreakingnews*, *ZairaOng*, and *Seowhow*), promoters (*charlesyeo*, *minzmint*, and *callsg*), and their top-5 followees who rank the highest by their originator scores. As we can see, *sgbreakingnews*, *ZairaOng*, and *Seowhow* are ranked higher on the mutual promoter list than the basic promoter list since they mostly retweeted from a few top-originating followees. In contrast, *charlesyeo*, *minzmint*, and *callsg*, who are among the top-10 promoters on the basic list, tend to retweet from many different users. Fewer than 50% of their retweets are from the top-5 originating followees. As a result, these users are ranked lower on the mutual promoter list.

### 6.5 Comparison with the Influence-Passivity Model

Next, we compare the user rankings generated by the proposed models to the ones generated by another retweet-based user influence model called the *Influence-Passivity* (IP) algorithm [Romero et al. 2011]. The goal is to examine whether the top originators/mutual originators and promoters/mutual promoters are the same as the influential users, as defined by the IP algorithm, even when the influential users are determined by the retweet information. A user's influence score ( $IP^I$ ) is determined by the number of people she influences, considering their passivity while a user's passivity score ( $IP^P$ ) is determined by the number of people, she is not influenced by considering their influence. To make it more comparable to our proposed models, we slightly modify the original IP formulation by restricting that influence only propagates among the follow-links as opposed to any arbitrary links. They are formally described in the following equations.

$$IP_i^I \leftarrow \sum_{uj \in F_i^{in}} u_{ij} IP_j^P \quad (7)$$

$$IP_i^P \leftarrow \sum_{uj \in F_i^{out}} v_{ji} IP_j^I, \quad (8)$$

where the acceptance rate is defined as  $u_{ij} = \frac{w_{ij}}{\sum_{u_k \in F_j^{out}} w_{kj}}$  and the rejection rate is defined as  $v_{ji} = \frac{1-w_{ji}}{\sum_{u_k \in F_j^{in}} (1-w_{jk})}$  and  $w_{ij} = \frac{t_{ij}}{t_i}$ . According to the definitions, the IP algorithm can be

Table X. Kendall's tau of the Top-100 Users Ranked by the IP Algorithm and the Behavior Models

	$O$	$P$	$O^m$	$P^m$	$O^R$	$P^R$
$IP^I$	0.37	0.08	0.19	-0.04	0.48	0.12
$IP^P$	-0.04	-0.24	-0.07	0.00	-0.1	-0.22

conceptually thought of as an inverse dependency variant of our mutual dependency model in which strong originating users are retweeted by many weak promoting users and vice versa. We have experimented with this notion of dependency. Interestingly, we found that the top-100 users scored by the inverse dependency model are quite similar, sharing over 90% of common users, to that of the basic model. This can be explained using the examples of the top basic promoters' ego networks in Figure 8. According to the three ego networks in the top row, the majority of users retweeted by the top basic promoters are weak originators. We believe the inverse dependency model presents an interesting information propagation behavior which is fundamentally different from the mutual dependency model proposed in this work. Therefore, we plan to further investigate such a model as well as other behavior definitions in the future work.

According to Table X, the top-ranked influence users are more similar to the top-ranked originators ( $O$ ) and the top-ranked range-based originators ( $O^R$ ) than the top-ranked mutual originator ( $O^m$ ) based on the correlation coefficients. The top-ranked promoters, mutual promoters, and ranged-based promoters ( $P$ ,  $P^m$ , and  $P^R$ ) are not the same as the top-ranked influence users. This is to be expected judging from their formal definitions.

Next, Table XI shows that among the top-10 users, the  $IP^I$ 's list contains more celebrities and entertainers (5 out of 10), compared to that of  $O$ ,  $O^m$ , and  $O^R$ . This suggests that highly passive users, who are mostly students, tend to be influenced by stars and celebrities. In contrast, the top promoters come from varied backgrounds and generally tend to retweet interesting tweets from various groups of originators. Among the top originators who are also highly influential, *stcom's* and *mrbrown's* tweets tend to attract interest from a large number of highly passive users who do not usually retweet other users' tweets. These two Twitter users are in fact one of the most well-known mainstream media and bloggers in Singapore, respectively.

Interestingly, through the mutual dependency model, we uncover strong mutual interactions among specific groups of Twitter users who do not appear on the top-10 list ranked by  $IP^I$ . For example, members of two local churches, City Harvest Church (CHC) and Heart of God Church (HoGC), actively use Twitter to share information with one another. The in-group interaction behaviors mutually reinforce their behavior scores. These top-ranked Twitter users are *chcsg* (CHC's official account), *arieszulka-rnain* (CHC pastor), *SeowHow* (Tan Seow How, pastor and cofounder of HoGC), *lynnetegoh*, *angjiehui*, *LisahK*, and *ZairaOng* (HoGC members). These users consistently play a significant role either as strong originators or promoters in their networks.

As the underlying assumptions of the IP algorithm and the proposed behavior models are different, one method can be more useful than the other when applied to different contexts, such as types of networks, definitions of top/interesting users, the process by which information is propagated, and so on. The goal of comparing the top users identified by the originating/promoting measures and the IP algorithm, is not to find a simple answer as to which method is the best to rank Twitter users, but rather to explore whether or not the top users ranked by one method tend to be ranked highly by the other. Overall, we observe that some of the top originators, who are generally retweeted many times by many users, are likely to be the top influential users as well,

Table XI. Top-10 Users Ranked by the IP Algorithm, Compared to Those Ranked by the Originator and Promoter Scores

Asterisk denotes a power user.

Rank	Influential User		Passive User	
1	<i>TaufikBatisah</i>	Singer	<i>nualanoowool</i>	N/A
2	<i>stcom*</i>	Media	<i>sylviaaa.tan</i>	N/A
3	<i>TweetSG</i>	Service	<i>Fhnsignature</i>	N/A
4	<i>FauzieLaily</i>	Singer	<i>skhadijahkhalid</i>	Student
5	<i>HadyMirzaAmir*</i>	Singer	<i>09yana</i>	N/A
6	<i>effs</i>	Singer	<i>Bonchjela</i>	Student
7	<i>mrbrown</i>	Blogger	<i>sasyame</i>	Student
8	<i>McDanielOng</i>	N/A	<i>nannzz</i>	Student
9	<i>987MisterYoung</i>	Media	<i>thelioncityboy</i>	Artist
10	<i>JoannePeh*</i>	Actress	<i>naddy91</i>	Student

Rank	Originator		Promoter	
1	<i>stcom*</i>	Media	<i>rochorbeancurd</i>	Entrepreneur
2	<i>ChannelNewsAsia*</i>	Media	<i>alkanphel</i>	Blogger
3	<i>MediaAsia</i>	Media	<i>charlesyeo</i>	Entrepreneur
4	<i>chcsg*</i>	Church	<i>MediaAsia</i>	Media
5	<i>arieszulkarnain</i>	Pastor	<i>CalvinTimo</i>	Writer
6	<i>mrbrown</i>	Blogger	<i>minzmint</i>	Student
7	<i>hardwarezone*</i>	Portal	<i>callsg</i>	N/A
8	<i>TaufikBatisah</i>	Singer	<i>LisahK</i>	Student
9	<i>Seowhow</i>	Pastor	<i>congyuan</i>	Student
10	<i>lynettegoh</i>	Student	<i>trevidzi</i>	Student

Rank	Mutual Originator		Mutual Promoter	
1	<i>stcom*</i>	Media	<i>CalvinTimo</i>	Writer
2	<i>SingaporeClub</i>	Aggregator	<i>sgbreakingnews</i>	Aggregator
3	<i>mrbrown</i>	Blogger	<i>LisahK</i>	Student
4	<i>lynettegoh</i>	Student	<i>ZairaOng</i>	Blogger
5	<i>chcsg*</i>	Church	<i>SingaporeClub</i>	Aggregator
6	<i>angjehui</i>	Student	<i>rochorbeancurd</i>	Entrepreneur
7	<i>sginfomap</i>	Media	<i>alkanphel</i>	Blogger
8	<i>sglatestnews*</i>	Aggregator	<i>cynthiachul</i>	Student
9	<i>patlaw</i>	Entrepreneur	<i>Seowhow</i>	Pastor
10	<i>cnalatest*</i>	Media	<i>preius</i>	Student

Rank	Range-Based Originator		Range-Based Promoter	
1	<i>stcom*</i>	Media	<i>alkanphel</i>	Blogger
2	<i>MediaAsia</i>	Media	<i>charlesyeo</i>	Entrepreneur
3	<i>chcsg*</i>	Church	<i>LisahK</i>	Student
4	<i>ChannelNewsAsia*</i>	Media	<i>rochorbeancurd</i>	Entrepreneur
5	<i>hardwarezone*</i>	Portal	<i>AngMoGirl</i>	Designer
6	<i>angjehui</i>	Student	<i>CalvinTimo</i>	Writer
7	<i>cnalatest*</i>	Media	<i>moemasri</i>	Singer
8	<i>TaufikBatisah</i>	Singer	<i>deckstor</i>	Student
9	<i>SGnews*</i>	Media	<i>JadeChen85144</i>	Student
10	<i>patlaw</i>	Entrepreneur	<i>trevidzi</i>	Student

being able to attract highly passive users to retweet their content. These top users include mainstream media and celebrities. On the other hand, some other top originators, such as religious organizations and business entrepreneurs, are not as globally influential, as they tend to attract retweets from specific sets of users. Next, promoting users are those who are consistent in their retweeting activities. Their tweets and

retweets do not tend to attract additional retweets from highly passive users at all, judging by their near zero correlation with the influence measure. This suggests that they are distinctively a power consumer, who helps spread certain types of tweets, rather than a producer of interesting tweets.

## 7. DETECTING INTERESTING EVENTS VIA INFORMATION PROPAGATION BEHAVIOR

Does a sudden surge of interest among Twitter users, measured in terms of their aggregated information propagation behavior, possibly point to an interesting event? How effective is the performance of the proposed behavior models in the event-detection task, compared to the message-based approaches, such as surges of tweet and retweet volumes? We investigate these questions by applying the proposed behavior models to identify events from Singapore users' tweet streams. An event is described as an occurrence  $e_i(s, t)$ , specific to a topic of interest  $s$  and a time period  $t$ , that significantly differs from other occurrences  $e_j(s, t')$  determined by some baseline  $\theta$ . In this work, we derive a set of topics from highly frequent keyphrases extracted in step 1 of our framework. The challenge of the event detection task is to distinguish between significant events and those that may have occurred by chance. Detecting bursty activities in time series data has been extensively investigated [Kleinberg 2003; Platakis et al. 2009]. As the primary goal of this work is to model information propagation behavior, we simply employ a threshold-based approach to detect significant events.

### 7.1 Burst Detection

We first derive a time series using weekly tweeted messages containing a specific topic, and compute time series values from the tweet messages. Specifically, we have six different baseline and six aggregated-behavior time series values to track, as shown in Table XII.

The aggregated behavior measures are derived by adding the behavior scores of the users mentioning a specific topic in their tweet content in order to obtain a single value for the week. For example, the following equation derives the aggregated originator score value  $v_s(t)$  or a topic  $s$  during the time period  $t$ .

$$v_s(t) = \sum_{tw \in TW_{s(t)}} O_{u(tw)}, \quad (9)$$

where  $TW_s(t)$  denotes the set of tweets containing  $s$  at time period  $t$ , and  $u(tw)$  denotes the user who publishes tweet  $tw$ . The more users who are originators mentioning the topic  $s$  at time  $t$ , the higher the aggregated originator value for  $s$  at  $t$ . The same formula applies to other behavior-based time series. In this work, we use an offline approach to estimate the individual user's behavior scores. That is, we assume that we have complete knowledge of each user's tweeting/retweeting activities and derive her behaviour scores from all her data. In other words,  $O_{u(tw)}$  of a user  $u_i$  will be the same for all  $t$  in the originator-based time series. Alternatively, an online approach can also be employed in order to estimate the user's behavior scores from her activities from the first time period up to the current time  $t$ .

To detect bursts, we look out for the significant changes to time series values between consecutive time periods  $t$  and  $t-1$ . Two criteria of changes are used in our work, namely *absolute change* and *relative change* denoted in short by  $\Delta_A$  and  $\Delta_R$ , respectively. Absolute change is computed by  $\Delta_A(s, t) = v_s(t) - v_s(t-1)$ , where  $v_s(t)$  refers to the time series value at time period  $t$  for topic  $s$ . A burst is detected if the following conditions are satisfied.

$$(\Delta_A(s, t) > \theta_{\Delta_A(s)}) \wedge (\Delta_A(s, t) > 2v_s(t-1)), \quad (10)$$

Table XII. Different Time Series Used in Event Detection

Abbreviation	Time series' values ( $v_s(\mathbf{t})$ )
$O$	Weekly sum of originator scores
$P$	Weekly sum of promoter scores
$O^m$	Weekly sum of mutual originator scores
$P^m$	Weekly sum of mutual promoter scores
$O^R$	Weekly sum of range-based originator scores
$P^R$	Weekly sum of range-based promoter scores
$TCount$	Weekly tweet count
$RTCount$	Weekly retweet count
$Hub$	Weekly sum of hub scores
$Authority$	Weekly sum of authority scores
$IP-I$	Weekly sum of IP's influence scores
$IP-P$	Weekly sum of IP's passivity scores

Where  $\theta_{\Delta_A(s)} = \mu(\Delta_A(s, t')) + \sigma(\Delta_A(s, t'))$ ,  $t' \in TSet$ ,  $TSet$  denotes the set of all time periods,  $\mu$  is an arithmetic mean of absolute change values and  $\sigma$  is a standard deviation of absolute change values;  $\Delta_A(s, t') > 0$ .

Relative change is computed by  $\Delta_R(s, t) = \frac{v_s(t) - v_s(t-1)}{v_s(t-1) - v_s(t-2)}$ . A burst is detected if the following conditions are satisfied.

$$(\Delta_R(s, t) > \theta_{\Delta_R(s)}) \wedge (\Delta_R(s, t) > 2v_s(t-1)), \quad (11)$$

where  $\theta_{\Delta_R(s)} = \mu(\Delta_R(s, t')) + \sigma(\Delta_R(s, t'))$ ,  $t' \in TSet$ ,  $\mu$  is an arithmetic mean of relative change values and  $\sigma$  is a standard deviation of relative change values;  $\Delta_R(s, t') > 1$ .

## 7.2 Empirical Analysis of Detected Events

To show the difference among the candidate events detected by various methods, we examine the levels of user interest these events might have generated using the numbers of tweets and retweets mentioning the events as proxies. The greater the numbers of tweets/retweets, the higher the levels of interest. For each set of events detected by a given method, we find the corresponding sets of tweets and retweets mentioning each event that was published within 7 days after its occurrence. Then, we compute the fractions of those events having generated numbers of tweets and retweets greater than the specific threshold values.

Tables XIII and XIV show the fractions of events detected by  $\Delta_A$  and  $\Delta_R$ , respectively, given the number of tweets and retweets mentioning the events as thresholds. For instance, 46% of events detected by the originating behavior ( $O$ ) using  $\Delta_A$  criteria generated more than 100 tweets, while 29% of those generated more than 200 tweets after 7 days. Since  $\Delta_R$  uses a more stringent condition, the fractions of events detected under  $\Delta_R$  are mostly fewer than those detected under  $\Delta_A$ .

As we can see, the promoting behaviors tend to detect the larger fractions of candidate events that generally attract many tweets and retweets rather than other methods. Next, the basic models consistently outperform the mutual dependency models and the range-based models across most threshold values as shown in Table XIII. The range-based models appear to detect the smallest fractions of interesting events out of the three models. This suggests that signals from some long-range users may not be helpful in identifying highly tweeted/retweeted events. When  $\Delta_R$  is used in Table XIV, the ranged-based and mutual promoting behaviors perform better than their basic counterparts. Especially, the ranged-based promoting behavior tends to perform well on the low-threshold events, e.g. those attracting 100–200 tweets and



Table XIII. Fractions of Events Detected by Different Methods Using  $\Delta_A$  Given the Numbers of Tweets and Retweets about the Events Published within 7 days after their Occurrence

The best results are in boldface.

Method	Tweet Threshold					Retweet Threshold				
	100	200	300	400	500	10	20	30	40	50
<i>O</i>	0.46	0.29	0.21	0.21	0.18	0.29	0.21	0.14	0.07	0.04
<i>P</i>	<b>0.57</b>	<b>0.43</b>	<b>0.38</b>	<b>0.33</b>	<b>0.29</b>	0.29	<b>0.29</b>	<b>0.19</b>	<b>0.10</b>	<b>0.10</b>
<i>O<sup>m</sup></i>	0.42	0.24	0.21	0.21	0.18	0.30	0.21	0.15	0.09	0.06
<i>P<sup>m</sup></i>	0.48	0.30	0.26	0.22	0.22	0.22	0.22	0.15	0.07	0.07
<i>O<sup>R</sup></i>	0.40	0.24	0.20	0.14	0.14	0.28	0.14	0.06	0.04	0.02
<i>P<sup>R</sup></i>	0.41	0.24	0.22	0.17	0.15	0.30	0.22	0.09	0.07	0.07
<i>Hub</i>	0.42	0.29	0.24	0.20	0.16	<b>0.35</b>	0.16	0.09	0.05	0.04
<i>Authority</i>	0.42	0.28	0.25	0.21	0.17	0.34	0.17	0.09	0.06	0.04
<i>IP-I</i>	0.34	0.19	0.15	0.13	0.11	0.27	0.15	0.08	0.02	0.00
<i>IP-P</i>	0.47	0.27	0.22	0.18	0.15	0.31	0.15	0.09	0.05	0.04

Table XIV. Fractions of Events Detected by Different Methods Using  $\Delta_R$  Given the Numbers of Tweets and Retweets about the Events Published within 7 days after their Occurrence

The best results are in boldface.

Method	Tweet Threshold					Retweet Threshold				
	100	200	300	400	500	10	20	30	40	50
<i>O</i>	0.24	0.19	0.14	0.14	0.14	0.14	0.10	0.05	0.05	0.05
<i>P</i>	0.33	0.22	0.22	<b>0.22</b>	<b>0.22</b>	0.17	0.17	0.06	0.06	0.06
<i>O<sup>m</sup></i>	0.29	0.19	0.14	0.14	0.14	0.14	0.14	0.05	0.05	0.05
<i>P<sup>m</sup></i>	0.37	0.26	<b>0.26</b>	0.21	0.21	0.21	<b>0.21</b>	<b>0.16</b>	<b>0.11</b>	<b>0.11</b>
<i>O<sup>R</sup></i>	0.37	0.23	0.17	0.13	0.13	0.17	0.07	0.03	0.03	0.03
<i>P<sup>R</sup></i>	<b>0.45</b>	<b>0.26</b>	0.23	0.16	0.16	<b>0.39</b>	0.19	0.10	0.06	0.06
<i>Hub</i>	0.33	0.24	0.21	0.21	0.15	0.21	0.09	0.06	0.06	0.06
<i>Authority</i>	0.29	0.24	0.21	0.21	0.15	0.21	0.09	0.06	0.06	0.06
<i>IP-I</i>	0.28	0.22	0.08	0.08	0.06	0.14	0.06	0.06	0.03	0.03
<i>IP-P</i>	0.38	0.16	0.13	0.09	0.09	0.16	0.09	0.03	0.03	0.03

10 retweets while the mutual promoting behavior tends to perform well on the high-threshold events, e.g., those attracting 300 tweets or higher and 20 retweets or higher.

Nevertheless, interesting events may not necessarily be found via the sheer numbers of tweets/retweets alone; we further investigate the accuracy of the proposed event detection methods with respect to the ground truth events in the next section.

### 7.3 Evaluation of Interesting Event Detection

*The Ground Truth Events.* A set of ground truth events is needed in order to evaluate the effectiveness of event detection methods. Given a set of frequent keyphrases  $S$  ( $|S| = 47$ ), those that have been mentioned in at least 10 retweets, and a set of tweets over a 20-week period, we construct a set of candidate events  $C$  ( $|C| = 940$ ) based on a combination of keyphrase and weekly time period. Next, we ask three postdoctoral researchers to annotate the sets of tweets associated with the candidate events. The three annotators are fairly knowledgeable in global and local news and quite familiar with Twitter. For a candidate event, each annotator examines a set of tweets related to the event, i.e., those containing a specific keyphrase, and provides a

Table XV. Examples of Positive Ground Truth Events

Keyphrase	Weekly Period	Description
American Idol	Jan 10, 2010	The premiere of the 9th season of American Idol show
Google China	Jan 10, 2010	Announcement about Google's operations in China
Apple iPad	Jan 24, 2010	Announcement of Apple iPad tablet device
Alexander McQueen	Feb 7, 2010	Breaking news about Alexander McQueen's suicide
Jack Neo	Mar 7, 2010	Press conference by a local film director Jack Neo
Universal studio	Mar 14, 2010	Official opening of Universal Studio Singapore
Changi Airport	Mar 21, 2010	Top award given to Changi Airport in Singapore
Apple MacBook	Apr 11, 2010	Launch of a new Apple MacBook laptop model
Jay Chou Concert	Apr 11, 2010	Sales of Jay Chou's concert tickets in Singapore
Star Awards	Apr 18, 2010	Broadcast of a local entertainment industry awards

Table XVI. Methods to Combine Events

Abbreviation	Event Combinations
$O \cap P$	Intersection of the sets of events detected by $O$ and $P$
$O \cup P$	Union of the sets of events detected by $O$ and $P$
$O^m \cap P^m$	Intersection of the sets of events detected by $O^m$ and $P^m$
$O^m \cup P^m$	Union of the sets of events detected by $O^m$ and $P^m$
$O^R \cap P^R$	Intersection of the sets of events detected by $O^R$ and $P^R$
$O^R \cup P^R$	Union of the sets of events detected by $O^R$ and $P^R$

binary judgment as to whether the tweets express a common user interest or not. We compute the Fleiss' Kappa ( $K$ ) [Fleiss 1981] to measure interrater agreement between the annotators. In this case,  $K = 0.34$ , reflecting a fair agreement. The low  $K$  value suggests that event judgment is highly subjective.

In summary, 21 events are selected as the positive ground truth events according to the unanimous agreements while the other 919 events are assigned as the negative ground truth examples. We denote this set of ground truth events by  $G_1$ . Moreover, 31 events are selected as a set of relaxed positive ground truths based on the partial agreements of at least two annotators, while the other 909 events are selected as the negative ground truths. We denote this set of ground truth events by  $G_2$ ;  $G_1 \subset G_2$ . These events represent both local as well as global interests of the Singapore-based Twitter users. Table XV displays examples of the positive ground truth events. Similar to other studies [Asur et al. 2011; Kwak et al. 2010; Weng and Lee 2011], we observe that the majority of the ground truth events in the Singapore Twitter dataset are related to real-world news.

*Experiment Setup.* The experiment is designed to evaluate the effectiveness of different methods in terms of ranking of interesting events. Given a set of 940 candidate events in the ground truth data  $C$ , we evaluate the sets of events detected from six baseline time-series data and six single behavior-based time-series data as summarized in Table XII. In the experiment,  $TCount$  and  $RTCount$  are defined as the message-level baseline methods, while  $Hub$ ,  $Authority$ ,  $IP-I$ , and  $IP-P$  are the user-level baseline methods. Additionally, we propose six behavior-based methods, shown in Table XVI, that utilize the combined signals from both the originating and promoting behaviors to detect events. For instance,  $O \cap P$  method creates a joint set of events that are detected by both the  $O$  and  $P$  methods. To derive a ranking of a joint event, we simply use an average ranking.

Table XVII. Average Precision, Precision at Top-5, and Top-10 of  $\Delta_A$  and  $\Delta_R$ -Based Methods on  $G_1$  set ( $|G_1| = 21$ )

The best results are in boldface.

Method	$\Delta_A$			$\Delta_R$		
	AP	Prec@5	Prec@10	AP	Prec@5	Prec@10
<i>Single</i>	0.4608	<b>0.8</b>	<b>0.8</b>	0.111	0.4	0.3
<i>Multi</i>	<b>0.6057</b>	<b>0.8</b>	<b>0.8</b>	<b>0.2012</b>	<b>0.6</b>	<b>0.5</b>
<i>TCount</i>	0.4672	<b>0.8</b>	0.7	0.0532	0.0	0.1
<i>RTCount</i>	0.2792	0.4	0.3	0.1769	0.4	0.3
<i>Hub</i>	0.444	<b>0.8</b>	0.5	0.0638	0.4	0.2
<i>Authority</i>	0.1558	0.4	0.3	0.1102	0.4	0.3
<i>IP-I</i>	0.0692	0.0	0.2	0.0797	0.2	0.1
<i>IP-P</i>	0.141	0.6	0.3	0.024	0.2	0.1

To measure the effectiveness of each method in identifying interesting events, we consider the evaluation of event detection methods as event ranking tasks. Therefore, we employ two information-retrieval-based metrics, average precision ( $AP$ ) and precision at the top- $k$  ranked events ( $Prec@k$ ), to evaluate the overall quality of ranked lists of events produced by each method. Since the threshold levels  $\theta_{\Delta_A(s)}$  and  $\theta_{\Delta_R(s)}$  differ depending on the topic  $s$ , we consider ranking each detected event according to its relative importance  $\Delta_A(s, \hat{t})/\theta_{\Delta_A(s)}$  and  $\Delta_R(s, \hat{t})/\theta_{\Delta_R(s)}$ , where  $\hat{t}$  is the time period in which the event is found, for the sets of events detected by  $\Delta_A$  and  $\Delta_R$ , respectively. This is to further distinguish between events with comparable change values by considering the proportion of their change values relative to the respective thresholds for different topics  $s$ .

Given the binary judgment for the sets of events from the ground truth and the ranked lists of detected events, we compute  $AP$  and  $Prec@k$  as follows. Let  $tp_k$  be the number of correctly predicted events at the top- $k$  events,  $fp_k$  be the number of incorrectly predicted events at the top- $k$  events,  $AP$  and  $Prec@k$  are formally defined in the following equations.

$$AP = \frac{1}{|G|} \sum_{i=1}^{|G|} Prec(G_i) \quad (12)$$

$$Prec@k = \frac{tp_k}{tp_k + fp_k}, \quad (13)$$

where  $G$  is a set of ground truth events and  $Prec(G_i)$  is the precision computed for a ranked list up to the position where the  $i$ th ground truth event is found. A good method should rank more interesting ground truth events at a higher position than non-events on the ranked list. Both  $AP$  and  $Prec@k$  output a value within the  $[0, 1]$  range. Due to space limitations, we only report the performance of the best methods. These metrics are computed with respect to two sets of ground truth events,  $G_1$  and  $G_2$ . Last, the performance difference between the proposed methods and the baselines is statistically tested using a paired t-test.

#### 7.4 Experiment Results

Tables XVII and XVIII display the values of  $AP$ ,  $Prec@5$ , and  $Prec@10$  of the ranked lists of events detected by the best aggregated single behavior ( $P$  for  $\Delta_A$  and  $O^R$  for  $\Delta_R$ ), denoted *Single*, the best combined events ( $O \cup P$  for  $\Delta_A$  and  $O^R \cup P^R$  for  $\Delta_R$ ), denoted *Multi*, and the baseline methods on the ground truth events  $G_1$  and

Table XVIII. Average Precision, Precision at Top-5, and Top-10 of  $\Delta_A$  and  $\Delta_R$ -Based Methods on  $G_2$  set ( $|G_2| = 31$ )

The best results are in boldface.

Method	$\Delta_A$			$\Delta_A$		
	AP	Prec@5	Prec@10	AP	Prec@5	Prec@10
<i>Single</i>	0.53	<b>1</b>	0.9	0.1096	0.4	0.3
<i>Multi</i>	<b>0.6952</b>	<b>1</b>	<b>1</b>	<b>0.1676</b>	<b>0.6</b>	<b>0.6</b>
<i>TCount</i>	0.5418	0.8	0.8	0.0791	0.2	0.2
<i>RTCount</i>	0.308	0.4	0.4	0.1327	0.4	0.3
<i>Hub</i>	0.3744	0.8	0.5	0.0432	0.4	0.2
<i>Authority</i>	0.1427	0.4	0.3	0.0747	0.4	0.3
<i>IP-I</i>	0.0805	0.0	0.2	0.0608	0.2	0.1
<i>IP-P</i>	0.1293	0.6	0.3	0.0163	0.2	0.1

$G_2$ , respectively. Overall, the best combined events methods significantly outperform the baselines, p-value  $< 0.05$ , according to the AP metrics. They also perform equally well and, in most cases, significantly better than, the baselines, p-value  $< 0.05$ , on the *Prec@5* and *Prec@10* metrics. This suggests that both the originator and promoter measures are helpful. For example, for the lists of events detected using  $\Delta_A$  criteria and evaluating against  $G_1$ , the average precision of  $O \cup P$  is 22.87% and 54.9% higher than that of *TCount* and *RTCount*, respectively. The mutual dependency model, though useful in distinguishing different user behaviors, does not perform as well as the basic model in detecting interesting events. Next, the range-based originators and promoters performed better than their basic counterparts when  $\Delta_R$  is used as the criterion. Nevertheless, the best overall performance is still achieved by the basic models using  $\Delta_A$  criteria. Last, the user-level aggregation baselines, i.e. *Hub*, *Authority*, *IP-I*, and *IP-P*, also perform poorly, compared to the simpler message-level baselines. In summary, the results indicate that using multiple behavior measures to detect events in Twitter streams is a more robust approach against noise than the message-count-based approaches.

One major advantage of the behavior-based methods over the message-based methods is that they can uncover some ground-truth events that did not generate enough surges in the number of tweets/retweets to be considered as events by the message-based methods. Some of these events involve the official announcements of a local sport event, a new theme park, and a local entertainment industry award show.

We further examine the participation rates of active users, active originators, and active promoters in the events. To select active users, we employ the same filtering rules described in Section 6. According to Table XIX, there are 319 unique active users in total, 230 and 231 of which are active originators and promoters, respectively. It can be seen that the active originators and promoters have much higher participation rates in the events than normally expected considering the participation rate from all users.

Given that 556 out of 7910 users (7.03%) have tweeted about the events and there are 230 active originators and 231 active promoters in total, the expected number of active originators participating in the events is 16.17 while that of active promoters is 16.24. In contrast, 91 active originators and 83 active promoters actually tweeted about the events. The same observation holds true for the top-10 originators and promoters as well. Furthermore, the participation rates of the active originators and promoters are higher than that of active users in general. Interestingly, 80% of the top-10 originators had participated in 16 out of 21 events (76.2%) while 30% of the top-10

Table XIX. Numbers of Users and Tweets Participating in the Ground Truth Events  $G$ 

Statistics	Total	Total $G_1$	Total $G_1$ /Total
# users	7910	556	7.03%
# active users	319	96	30.1%
# active originators	230	91	39.57%
# active promoters	231	83	35.93%
# top-10 tweeting users	10	3	30%
# top-10 originators	10	8	80%
# top-10 promoters	10	3	30%
# tweets	91638	3644	3.98%

promoters had participated in 7 events (33.3%). Notice that only a small fraction (3.98%) of total tweets are related to the events. This finding suggests that a great portion of users with strong originating and promoting behavior are likely to tweet about some interesting events. Therefore, keeping track of their behavior provides another way to discover interesting events from social media streams.

## 8. CONCLUSION AND FUTURE WORK

In this article, we propose a novel framework to model information propagation behavior of Twitter users. Specifically, we introduce two propagation behaviors, namely originating and promoting behaviors, derived from the users' tweeting and retweeting activities. To overcome the scarcity of retweets, we introduce the notion of weak retweet linkages which can be extracted from the mentioning of interesting items among tweets published by Twitter users and their follow-links. Three behavior models are proposed. The basic model measures the strength of user behaviors according to the propagation of retweets. The mutual dependency model additionally incorporates mutual dependency between the originating and promoting behaviors. The range-based model extends the basic model by incorporating indirect effects of the basic propagation behavior in the diffusion chains of tweets. Next, we compare the user rankings produced by the proposed models with those of the existing measures, namely the Influence-Passivity model, and demonstrate how the behavior models can be applied to find interesting events from Twitter streams. Future work includes improving the robustness of the models' formulation to prevent abuse such as retweet spamming, and so on.

## ACKNOWLEDGMENTS

The authors thank the reviewers for their helpful comments and suggestions.

## REFERENCES

- AGARWAL, N., LIU, H., TANG, L., AND YU, P. S. 2008. Identifying the influential bloggers in a community. In *Proceedings of the ACM International Conference on Web Search and Web Data Mining (WSDM)*. 207–218.
- AGICHTEN, E., BRILL, E., AND DUMAIS, S. 2006. Improving web search ranking by incorporating user behavior information. In *Proceedings of the Annual ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*. 19–26.
- ASUR, S., HUBERMAN, B. A., SZABO, G., AND WANG, C. 2011. Trends in social media: Persistence and decay. In *Proceedings of the National AAAI Conference on Weblogs and Social Media (ICWSM)*. 434–437.



- BECKER, H., NAAMAN, M., AND GRAVANO, L. 2010. Learning similarity metrics for event identification in social media. In *Proceedings of the ACM International Conference on Web Search and Web Data Mining (WSDM)*.
- BERGER, J. A. AND MILKMAN, K. L. 2009. *Social Transmission, Emotion, and the Virality of Online Content*. Social Science Research Library (SSRN) eLibrary.
- CASTILLO, C., DONATO, D., GIONIS, A., MURDOCK, V., AND SILVESTRI, F. 2007. Know your neighbors: Web spam detection using the web topology. In *Proceedings of the Annual ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*.
- CHA, M., HADDADI, H., BENEVENUTO, F., AND GUMMADI, K. P. 2010. Measuring user influence in Twitter: The million follower fallacy. In *Proceedings of the National AAAI Conference on Weblogs and Social Media (ICWSM)*.
- DE CHOUDHURY, M. 2010. Discovery of information disseminators and receptors on online social media. In *Proceedings of the ACM Conference on Hypertext and Hypermedia (HT)*.
- FLEISS, J. L. 1981. *Statistical Methods For Rates and Proportions* 2nd Ed. John Wiley, New York.
- GHOSH, R. AND LERMAN, K. 2010. Predicting influential users in online social networks. In *Proceedings of the ACM Workshop on Social Network Mining and Analysis and Knowledge Discovery and Data Mining (SNAKDD)*.
- GOYAL, A., BONCHI, F., AND LAKSHMANAN, L. V. S. 2010. Learning influence probabilities in social networks. In *Proceedings of the ACM International Conference on Web Search and Web Data Mining (WSDM)*. 241–250.
- GRUHL, D., GUHA, R., LIBEN-NOWELL, D., AND TOMKINS, A. 2004. Information diffusion through blogspace. In *Proceedings of the International Conference on World Wide Web (WWW)*. ACM Press, New York, NY, 491.
- GUHA, R., KUMAR, R., RAGHAVAN, P., AND TOMKINS, A. 2004. Propagation of trust and distrust. In *Proceedings of the International Conference on World Wide Web (WWW)*. ACM Press, New York, NY, 403.
- GYÖNGYI, Z., GARCIA-MOLINA, H., AND PEDERSEN, J. 2004. Combating web spam with trustrank. In *Proceedings of the International Conference on Very Large Databases (VLDB)*. 576–587.
- KEMPE, D., KLEINBERG, J., AND TARDOS, É. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the International SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*. ACM Press, New York, NY, 137.
- KENDALL, M. 1938. A new measure of rank correlation. *Biometrika* 30, 81–89.
- KLEINBERG, J. 1999. Authoritative sources in a hyperlinked environment. *J. ACM* 46, 604–632.
- KLEINBERG, J. 2003. Bursty and hierarchical structure in streams. In *Proceedings of the International SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*. 373–397.
- KLEINBERG, J. 2006. *Data Stream Management: Processing HighSpeed Data*. Springer.
- KWAK, H., LEE, C., PARK, H., AND MOON, S. 2010. What is Twitter, a social network or a news media. In *Proceedings of the International Conference on World Wide Web (WWW)*.
- LESKOVEC, J., BACKSTROM, L., AND KLEINBERG, J. 2009. Memetracking and the dynamics of the news cycle. In *Proceedings of the International SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*. 497.
- LIN, C. X., ZHAO, B., MEI, Q., AND HAN, J. 2010. PET: A statistical model for popular events tracking in social communities. In *Proceedings of the International SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*.
- MATHIOUDAKIS, M. AND KOUDAS, N. 2010. TwitterMonitor. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*. ACM Press, New York, NY, 1155.
- MENDOZA, M., POBLETE, B., AND CASTILLO, C. 2010. Twitter under crisis: Can we trust what we RT? In *Proceedings of the Workshop on Social Media Analytics (SOMA)*.
- MIHALCEA, R. AND TARAU, P. 2004. TextRank: Bringing order into texts. In *Proceedings of the Meeting of Association of Computational Linguistics (ACL)*.
- MORITA, M. AND SHINODA, Y. 1994. Information filtering based on user behavior analysis and best match text retrieval. In *Proceedings of the Annual ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*. 272–281.
- NAAMAN, M., BOASE, J., AND LAI, C.-H. 2010. Is it really about me?: Message content in social awareness streams. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW)*.
- NTOULAS, A., NAJORK, M., MANASSE, M., AND FETTERLY, D. 2006. Detecting spam web pages through content analysis. In *Proceedings of the International Conference on World Wide Web (WWW)*. ACM Press, New York, NY, 83.

- ON, B.-W., LIM, E.-P., JIANG, J., PURANDARE, A., AND TEOW, L.-N. 2010. Mining interaction behaviors for email reply order prediction. In *Proceedings of the Conference on Advances on Social Networks Analysis and Mining (ASONAM)*. 306–310.
- PAGE, L., BRIN, S., MOTWANI, R., AND WINOGRAD, T. 1999. *The PageRank Citation Ranking: Bringing Order to the Web*. Stanford InfoLab.
- PLATAKIS, M., KOTSAKOS, D., AND GUNOPULOS, D. 2009. Searching for events in the blogosphere. In *Proceedings of the International Conference on World Wide Web (WWW)*. 1225–1226.
- POPESCU, A.-M. AND PENNACCHIOTTI, M. 2010. Detecting controversial events from twitter. In *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM)*. ACM Press, New York, NY, 1873.
- QIU, L., LEUNG, A. K.-Y., HO, J. H., YEUNG, Q. M., FRANCIS, K. J., AND CHUA, P. F. 2010. Understanding the psychological motives behind microblogging. In *Proceedings of the 15th Annual CyberTherapy & CyberPsychology Conference*.
- ROMERO, D. M., GALUBA, W., ASUR, S., AND HUBERMAN, B. A. 2011. Influence and passivity in social media. In *Proceedings of the International Conference on World Wide Web (WWW)*.
- SAKAKI, T., OKAZAKI, M., AND MATSUO, Y. 2010. Earthquake shakes Twitter users: Realtime event detection by social sensors. In *Proceedings of the International Conference on World Wide Web (WWW)*.
- SUBBIAN, K. AND MELVILLE, P. 2011. Supervised rank aggregation for predicting influence in networks. In *Proceedings of the IEEE Conference on Social Computing (SocialCom)*.
- SUH, B., HONG, L., PIROLI, P., AND CHI, E. H. 2010. Want to be retweeted? large scale analytics on factors impacting retweet in Twitter network. In *Proceedings of the IEEE Conference on Social Computing (SocialCom)*. 177–184.
- WANG, X., ZHAI, C., HU, X., AND SPROAT, R. 2007. Mining correlated bursty topic patterns from coordinated text streams. In *Proceedings of the International SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*. 784–793.
- WASSERMAN, S. AND FAUST, K. 1994. *Social Network Analysis: Methods and Applications*. Cambridge University Press.
- WENG, J. AND LEE, B.-S. 2011. Event detection in Twitter. In *Proceedings of the National AAAI Conference on Weblogs and Social Media (ICWSM)*.
- WENG, J., LIM, E.-P., JIANG, J., AND HE, Q. 2010. TwitterRank: Finding topicsensitive influential twitterers. In *Proceedings of the ACM International Conference on Web Search and Web Data Mining (WSDM)*. 261–270.
- WU, S., HOFMAN, J. M., MASON, W. A., AND WATTS, D. J. 2011. Who says what to whom on Twitter. In *Proceedings of the International Conference on World Wide Web (WWW)*. ACM Press, New York, NY, 705.
- YANG, J. AND COUNTS, S. 2010. Predicting the speed, scale, and range of information diffusion in Twitter. In *Proceedings of the National AAAI Conference on Weblogs and Social Media (ICWSM)*.
- YANG, J. AND LESKOVEC, J. 2010. Modeling information diffusion in implicit networks. In *Proceedings of the IEEE International Conference on Data Mining (ICDM)*. 599–608.
- ZHAO, Q., MITRA, P., AND CHEN, B. 2007. Temporal and information flow based event detection from social text streams. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.

Received July 2011; revised December 2011, March 2012, July 2012; accepted July 2012