

Artificial Inflation: The Real Story of Trends and Trend-setters in Sina Weibo

Louis Lei Yu, Sitaram Asur and Bernardo A. Huberman
Social Computing Lab, Social Computing Lab, Hewlett-Packard Labs
Email: {louis.yu, sitaram.asur, bernardo.huberman}@hp.com

Abstract—There has been a tremendous rise in the growth of online social networks all over the world in recent years. This has facilitated users to generate a large amount of real-time content at an incessant rate, all competing with each other to attract enough attention and become trends. While Western online social networks such as Twitter have been well studied, characteristics of the popular Chinese microblogging network Sina Weibo have not been. In this paper, we analyze in detail the temporal aspect of trends and trend-setters in Sina Weibo, contrasting it with earlier observations on Twitter. First, we look at the formation, persistence and decay of trends and examine the key topics that trend in Sina Weibo. One of our key findings is that retweets are much more common in Sina Weibo and contribute a lot to creating trends. When we look closer, we observe that a large percentage of trends in Sina Weibo are due to the continuous retweets of a small percentage of fraudulent accounts. These fake accounts are set up to artificially inflate certain posts causing them to shoot up into Sina Weibo’s trending list, which are in turn displayed as the most popular topics to users.

I. INTRODUCTION

In the past few years, social media services as well as the users who subscribe to them, have grown at a phenomenal rate. This immense growth has been witnessed all over the world with millions of people of different backgrounds using these services on a daily basis to communicate, create and share content on an enormous scale. This widespread generation and consumption of content has created an extremely complex and competitive online environment where different types of content compete with each other for the attention of users. Thus it is very interesting to study how certain types of content such as a viral video, a news article, or an illustrative picture, manage to attract more attention than others, thus bubbling to the top in terms of popularity. Through their visibility, these popular items and topics contribute to the collective awareness reflecting what is considered important. This can also be powerful enough to affect the public agenda of the community.

There have been prior studies on trends and trend-setters in Western online social media [1] [2]. In this paper, we examine in detail a significantly less-studied but equally fascinating online environment: Chinese social media, in particular, Sina Weibo: China’s biggest microblogging network.

Over the years there have been news reports on various Internet phenomena in China, from the surfacing of certain viral videos to the spreading of rumors [3] to the so called “human

flesh search engines” [4] ¹ in Chinese online social networks. These stories seem to suggest that many events happening in Chinese online social networks are unique products of China’s culture and social environment.

Due to the vast global connectivity provided by social media, netizens ² all over the world are now connected to each other like never before. This means that they can now share and exchange ideas with ease. It could be argued that this means the manner in which the sharing occurs is similar across countries. However, China’s unique cultural and social environment suggests that the way individuals share ideas might be different than that in Western societies. For example, the age of Internet users in China is a lot younger, they may respond to different types of content than Internet users in Western societies. The number of Internet users in China is larger than that in the U.S, and the majority of users lives in large urban cities. One would expect that the way these users share information can be even more chaotic. An important question to ask is to what extent would topics have to compete with each other in order to capture users’ attention in this dynamic environment. Furthermore, it is known that the information shared between individuals in Chinese social media is monitored [6]. Hence another interesting question to ask is what types of content would netizens respond to and what kind of popular topics would emerge under such circumstances.

Given the above questions, we present an analysis on the evolution of trends in Sina Weibo. We have monitored the evolution of the top trending keywords in Sina Weibo for 30 days. First, we analyze the model of growth of these trends and examine the persistence of these topics over time. We investigate if topics initially ranked higher tend to stay in the top 50 trending list longer. Subsequently, by analyzing the timestamps of posts containing the keywords, we look at the propagation and decaying process of the trends in Sina Weibo and compare it to earlier observations on Twitter [1].

Our findings are as follows:

- We establish that retweets play a greater role in Sina Weibo than on Twitter, contributing more to the generation and persistence of trends.
- When we examine the retweets in detail, we make an

¹is a primarily Chinese Internet phenomenon of massive search using online media such as blogs and forums [4].

²A netizen is a person actively involved in online communities [5].

important discovery. We observe that many of the trending keywords in Sina Weibo are heavily manipulated and controlled by certain fraudulent accounts that are setup for this purpose.

- We found significant evidence suggesting that a large percentage of trends in Sina Weibo are actually due to artificial inflation by these fake users, thus making certain posts trend and be more visible to other users.
- The users we identified as fraudulent were 1.08% of the total users sampled, but they were responsible for 49% of the total retweets (32% of the total tweets).
- We evaluate some methods to identify these fake users and demonstrate that by removing the tweets associated with the fake users, the evolution of the tweets containing the trending keywords follow the same persistent and decaying process as the one in Twitter.

The rest of the paper is organized as follows. In Section 2 we detail some related work. In Section 3 we provide background information on the Sina Weibo network. We perform a detailed analysis of trending topics, growth and persistence of trends and the identities of trend-setters on Sina Weibo in Section 4. In Section 5, we conclude and discuss our findings.

II. RELATED WORK

The Study of Chinese Online Social Networks: Jin [3] has studied the Chinese online Bulletin Board Systems (BBS), and provided observations on the structure and interface of Chinese BBS and the behavioral patterns of its users. Xin [7] has conducted a survey on BBS's influence on the University students in China and their behavior on Chinese BBS. Yu and King, [8] has looked at the adaptation of interests such as books, movies, music, events and discussion groups on Douban, an online social network and media recommendation system frequently used by the youth in China.

Trends and Trend-setters in Twitter: There are various studies on trends on Twitter [2] [9] [10]. Recently, Asur et al.[1] have examined the growth and persistence of trending topics on Twitter. They discovered that traditional media sources are important in causing trends on twitter. Many of the top retweeted articles that formed trends on Twitter were found to arise from news sources such as the New York Times. In this work, we evaluate how the trending topics in China relate to the news media.

Trend-setters in Sina Weibo: Yu et al. gave the first known study of trending topics in a Chinese online microblogging social network (Sina Weibo) [11]. They discovered that there are vast differences between the content that is shared on Sina Weibo and that on Twitter. In China, people tend to use microblogs to share jokes, pictures and videos and a significantly large percentage of posts are retweets. The trends that are formed are almost entirely due to the retweeting of such media content. This is contrary to what was observed on Twitter, where the trending topics have more to do with current events and the effect of retweets is not as high. Another

relevant work is by Chen et al. [12] on the so called "Internet water army" in China: paid users hired by companies to post comments and articles in online communities with the agenda of influencing other users' opinions toward events, people or products. The authors point out that these paid users are typically college students and unemployed workers from third-tier cities. A large quantity of said users can be hired with cheap wages. They also share some news stories as examples to illustrate how online paid posters hired by companies can be an affective marketing strategy in either a positive or damaging manner. Paid posters represent a new type of online opportunities in China. Finally, they collect posts from Sina Weibo of an reported event where paid users were hired to affect the outcome. They used the posts as training data to analyze the behavior of paid posters. We compare our findings with the findings from Chen et al [12] in Section 4.8.

III. BACKGROUND ON SINA WEIBO

Sina Weibo was launched by the Sina corporation, China's biggest web portal, in August 2009. It has been reported³, that Sina Weibo now has 250 million registered accounts and generates 90 million posts per day. Similar to Twitter, a user profile on Sina Weibo displays the user's name, a brief description of the user, the number of followers and followees the user has. There are three types of user accounts on Sina Weibo, regular user accounts, verified user accounts, and the expert (star) user account. A verified user account typically represents a well known public figure or organization in China.

Although both Twitter and Sina Weibo enable users to post messages of up to 140 characters, there are some differences in terms of the functionalities offered. While Twitter users can post tweets consisting of text and links, Sina Weibo users can post messages containing text, pictures, videos and links.

Twitter users can address tweets to other users and can mention others in their tweets [13]. A common practice on Twitter is "retweeting", or rebroadcasting someone else's messages to one's followers. The equivalent of a retweet on Sina Weibo is instead shown as two amalgamated entries: the original entry and the current user's actual entry which is a commentary on the original entry.

Sina Weibo has another functionality absent from Twitter: the comment. When a Weibo user makes a comment, it is not rebroadcasted to the user's followers. Instead, it can only be accessed under the original message.

IV. ANALYSIS OF TRENDS AND TREND-SETTERS ON SINA WEIBO

A. The Trending Keywords

Sina Weibo offers a list of 50 keywords that appear most frequently in users' tweets. They are ranked according to the frequency of appearances in the last hour. This is similar to Twitter, which also presents a constantly updated list of trending topics: keywords that are most frequently used in

³"Microblog Revolutionizing China's Social Business Development" by the Sina corporation and CIC is available at <http://www.ciccorporate.com>

tweets over a period of time. We extracted these keywords over a period of 30 days and retrieved all the corresponding tweets containing these keywords from Sina Weibo.

We first monitored the hourly evolution of the top 50 keywords in the trending list for 30 days. We observed that the average time spent by each keyword in the hourly trending list is 6 hours. And the distribution for the number of hours each topic remains on the top 50 trending list follows the power law (as shown in Figure 1 a). The distribution suggests that only a few topics exhibit long-term popularity.

Another interesting observation is that a lot of the key words tend to disappear from the top 50 trending list after a certain hour and then later reappear. We examined the distribution for the number of time keywords have reappeared in the top 50 trending list (Figure 1 b). We observe that this distribution follows the power law as well.

Both the above observations are very similar to our earlier study of trending topics on twitter [1]. However, one important difference with Twitter is that the average trending time is significantly higher on Sina Weibo (on Twitter it was 20-40 minutes). This suggests that Weibo may not have as many topics competing for attention as Twitter.

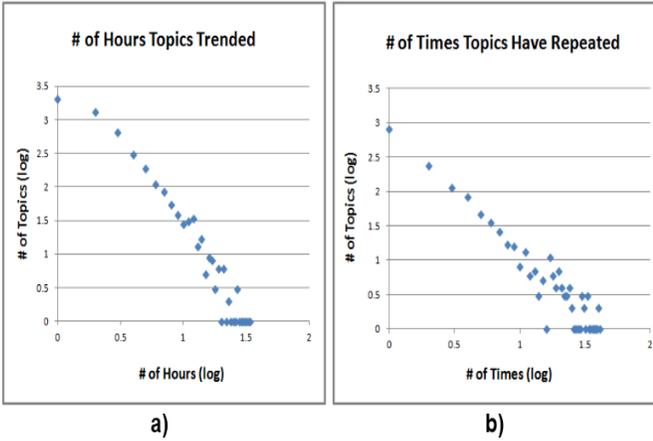


Fig. 1. Distributions for the trending time and the number of times topics reappeared

Following our observation that some keywords stay in the top 50 trending list longer than others, we wanted to investigate if topics that ranked higher initially tend to stay in the top 50 trending list longer. We separated the top 50 trending keywords into two ranked sets of 25 each: the top 25 and the bottom 25. Figure 2 illustrates the plot for the percentage of topics that placed in the bottom 25 relating to the number of hours these topics stayed in the top 50 trending list. We can observe that topics that do not last are usually the ones that are in the bottom 25. On the other hand, the long-trending topics spend most of their time in the top 25, which suggests that items that become very popular are likelier to stay longer in the top 50. This intuitively means that items that attract phenomenal attention initially are not likely to dissipate quickly from people’s interests.

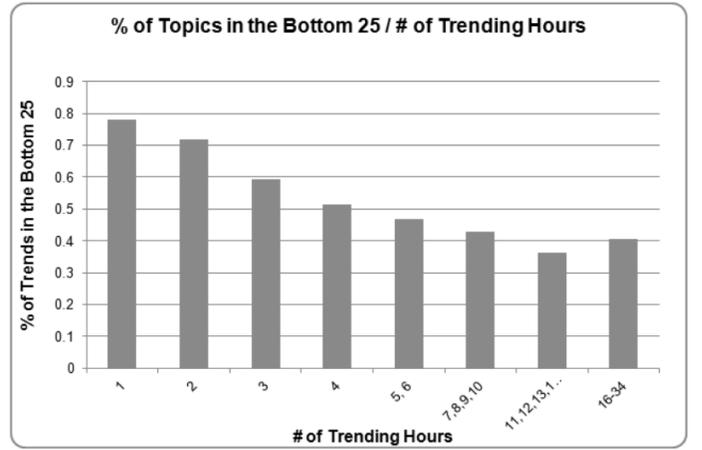


Fig. 2. Distribution of trending times for topics in the bottom 25 of the top 50 trend list

B. The Evolution of Tweets

Next, we want to investigate the process of persistence and decay for the trending topics on Sina Weibo. In particular, we want to measure the distribution for the time intervals between tweets containing the trending keywords. We continuously monitored the keywords in the top 50 trending list. For each trending topic we retrieved all the tweets containing the keyword from the time the topic first appeared in the top 50 trending list until the time it disappeared. Accordingly, we monitored 811 topics over the course of 30 days. In total we collected 574,382 tweets from 463,231 users. Among the 574,382 Tweets, 35% of the tweets (202,267 tweets) are original tweets, and 65% of the tweets (372,115 tweets) are retweets. 40.3% of the total users (187130 users) retweeted at least once in our sample.

We measured the number of tweets that each topic gets in 10 minutes intervals, from the time the topic starts trending until the time it stops. From this we can sum up the tweet counts over time to obtain the cumulative number of tweets $N_q(t_i)$ of topic q for any time frame t_i . This is given as :

$$N_q(t_i) = \sum_{\tau=1}^i n_q(t_\tau), \quad (1)$$

where $n_q(t)$ is the number of tweets on topic q in time interval t . We then calculate the ratios $C_q(t_i, t_j) = N_q(t_i)/N_q(t_j)$ for topic q for time frames t_i and t_j .

Figure 3 shows the distribution of $C_q(t_i, t_j)$'s over all topics for two arbitrarily chosen pairs of time frames: (10, 2) and (8, 3) (nevertheless such that $t_i > t_j$, and t_i is relatively large, and t_j is small).

These figures suggest that the ratios $C_q(t_i, t_j)$ are distributed according to the log-normal distributions. We tested and confirmed that the distributions indeed follow the log-normal distributions.

This finding agrees with the result from a similar experiment on Twitter trends. Asur et al [1] argued that the log-normal distribution occurs due to the multiplicative process involved

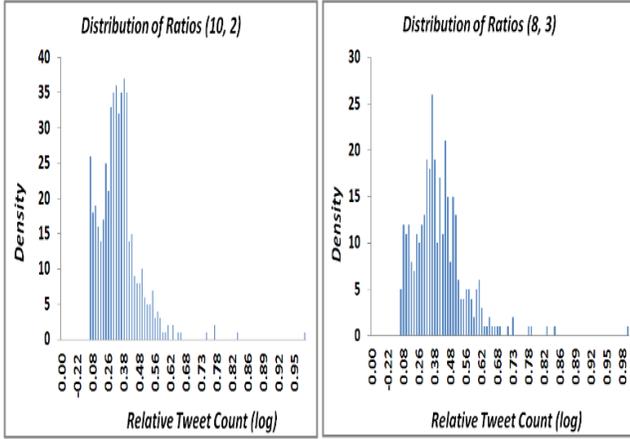


Fig. 3. The distribution of $C_q(t_i, t_j)$'s over all topics for two arbitrarily chosen pairs of time frames: (10, 2) and (8, 3)

in the growth of trends which incorporates the decay of novelty as well as the rate of propagation. The intuitive explanation is that at each time step the number of new tweets (original tweets or retweets) on a topic is multiplied over the tweets that we already have. The number of past tweets, in turn, is a proxy for the number of users that are aware of the topic up to that point. These users discuss the topic on different forums, including Twitter, essentially creating an effective network through which the topic spreads. As more users talk about a particular topic, many others are likely to learn about it, thus giving the multiplicative nature of the spreading. On the other hand, the monotonically decreasing decaying process characterizes the decay in timeliness and novelty of the topic as it slowly becomes obsolete.

However, while only 35% of the tweets in Twitter are retweets [1], there is a much larger percentage of tweets that are retweets in Sina Weibo. From our sample we observed that a high 65% of the tweets are retweets. This implies that the topics are trending mainly because of some content that has been retweeted many times. Thus, Sina Weibo users are more likely to learn about a particular topic through retweets.

C. Trend-setters on Sina Weibo

For every new trending keyword we retrieved the most retweeted tweets in the past hour and compiled a list of most retweeted users. Table I illustrates the top 20 most retweeted authors appearing in at least 10 trending topics each. We define an author's retweet ratio as the number of times the authors' tweets are retweeted divided by the number of trending topics these tweets appeared in. The influential authors are ranked according to their retweet ratios.

From Table I we observed that only 4 out of the top 20 influential authors were verified accounts. The 4 verified accounts represent an urban fashion magazine, a fashion brand, an online travel magazine, and a Chinese celebrity. The other 16 influential authors are unverified accounts. They all seem to have a strong focus on collecting user-contributed jokes, movie

trivia, quizzes, stories and so on. This is in sharp contrast to the topics that are popular on Twitter [1].

The *Tweets* column in Table I gives the unique tweets that have been retweeted. We can observe that the rate at which they have been retweeted is phenomenal. For example, the top retweeted user posted 37 tweets which were totally retweeted 1194999 times.

D. The Evolution of Retweets and Original Tweets

We separate the tweets in Sina Weibo into original tweets and retweets and calculated the densities of ratios between cumulative retweets/original tweets counts measured in different time frames. Figure 4 show the distribution of original tweets/retweets ratio over all topics for two arbitrarily chosen pairs of time frames: (10, 2) and (8, 3).

We find (as the examples in Figure ?? show) that the distribution of original tweets ratios follows the log-normal distribution. However, from Figure ?? we observe that for retweets, this does not satisfy all the properties of the log-normal distribution. There seems to be a large amount of low retweet ratios in the distribution. Furthermore, there are high spikes in the lower ratios area of the distribution.

E. Identifying Spam Activity in Sina Weibo

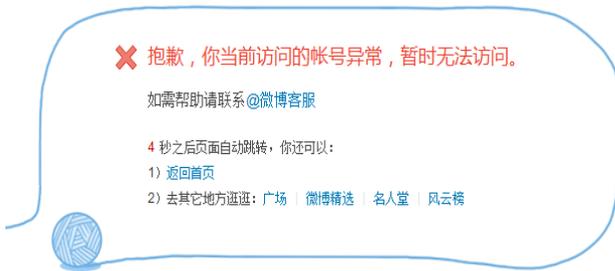
From Figure 4 in the previous Section we observed that there is a high percentage of low ratios in the distribution of retweet ratios. This suggests that for a lot of the topics, there is an initial flurry of retweets. We hypothesize that this is due to the activities of certain users in Sina Weibo. As these accounts post a tweet, they tend to set up many other fake accounts to continuously retweet this tweet, expecting that the high retweet numbers would propel the tweet to place in the Sina Weibo hourly trending list. This would then cause other users to notice the tweet more after it has emerged as the top hourly, daily, or weekly trend setter. We attempt to verify the above hypothesis empirically. We define a **spamming account** as one that is set up for the purpose of repeatedly retweeting certain messages, thus giving these messages artificially inflated popularity. According to our hypothesis, the users who retweet abnormally high amounts are more likely to be spam accounts.

Figure 6 a) illustrates the distribution for the number of users and their corresponding number of retweets (over all topics). Figure 6 b) illustrates the distribution for the number of users and the numbers of topics that they caused to trend by their retweets. We observe that both distributions in Figure 6 follow the power law. This implies that there are certain users who retweet a lot, and a small number of users responsible for a large number of topics. Next, we investigate who these users are. We manually checked the top 40 accounts who retweeted the most. To our surprise, 37 of these 40 accounts could no longer be accessed. That is, when we queried the accounts' IDs, we retrieved a message from Sina Weibo stating that the account has been removed and can no longer be accessed (see Figure 5). According to Sina Weibo's frequently asked question page, if a user sends a tweet containing illegal or

TABLE I
TOP 20 RETWEETED USERS IN AT LEAST 10 TRENDING TOPICS

	Author Description	Verified Account	Retweets	Tweets	Topics	Retweet-Ratio
1	Urban Fashion Magazine	Yes	1194999	37	12	99583.25
2	Fashion Brand VANCL	Yes	849404	21	13	65338.77
3	Online Travel Magazine	Yes	127737	123	21	57987.48
4	Gourmet Factory	No	553586	86	12	46132.17
5	Horoscopes	No	1545955	101	38	40683.13
6	Silly Jokes	No	3210130	258	81	39631.23
7	Good Movies	No	1497968	140	38	39420.21
8	Wonderful Quotes	No	602528	39	17	35442.82
9	Global Music	No	697308	116	22	31695.81
10	Funny Jokes Countdown	No	3667566	438	121	30310.46
11	Creative Ideas	No	742178	111	25	29687.12
12	Famous Chinese singer	Yes	284600	25	10	28460
13	Good Music	No	323022	52	12	26918.5
14	Movie Factory	No	1509003	230	59	25576.32
15	Strange Stories	No	1668910	250	66	25286.52
16	Beautiful Pictures	No	435312	33	18	24184
17	Global Music	No	432444	65	18	24024.67
18	Female Fashion	No	809440	87	34	23807.06
19	Useful Tips	No	735070	153	31	23711.94
20	Funny Quizzes	No	589477	77	25	23579.08

sensitive information, such tweet will be immediately deleted by Sina Weibo’s administrators, however, the users’ accounts will still be active. For the above reason we assume that if an account was active one month ago and can no longer be reached, This indicates that this account has very likely performed malicious activities such as spamming and has hence been deleted.



Sorry, the account can not be visited

Fig. 5. An error message

Next, we inspect the user accounts with the most retweets in our sample and the number of accounts they retweeted. We see that although the number of times these accounts retweeted was very high, they mostly only retweet messages from a few users. We re-organize the users who retweeted by the ratio between the number of times he/she retweeted and the number of users he/she retweeted. We refer to this as the **user-retweet ratio**. Table II illustrates the top 10 users with the highest user-retweet ratios. We note that for all these users, they each retweet posts from only one account. We observe that this is true for the top 30 accounts with the highest user-retweet

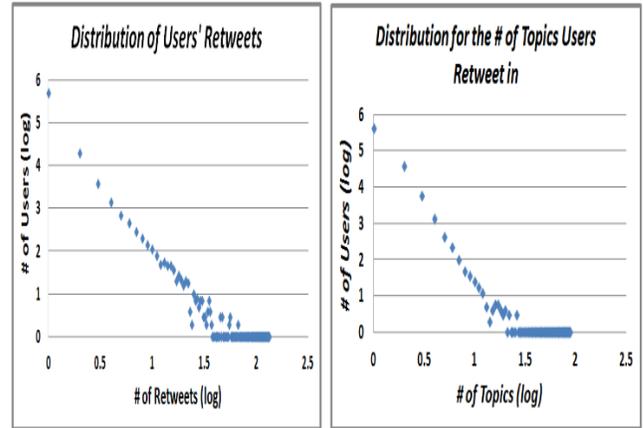


Fig. 6. The distribution for the number of users’ retweets and the number of topics users’ retweets trend in

ratios.

User ID	# Retweets	# Retweeted	U-R Ratio
1840241580	134	1	134
2241506824	125	1	125
1840263604	68	1	68
1840237192	64	1	64
1840251632	64	1	64
2208320854	55	1	55
2208320990	51	1	51
2208329370	48	1	48
2218142513	47	1	47
1843422117	44	1	44

TABLE II
THE TOP 10 ACCOUNTS WITH THE HIGHEST USER-RETWEET RATIOS (U-R RATIO)

Next, we conduct the following experiment: starting from the users with the highest user-retweet ratios, we used a crawler to automatically visit and retrieve each user’s Sina

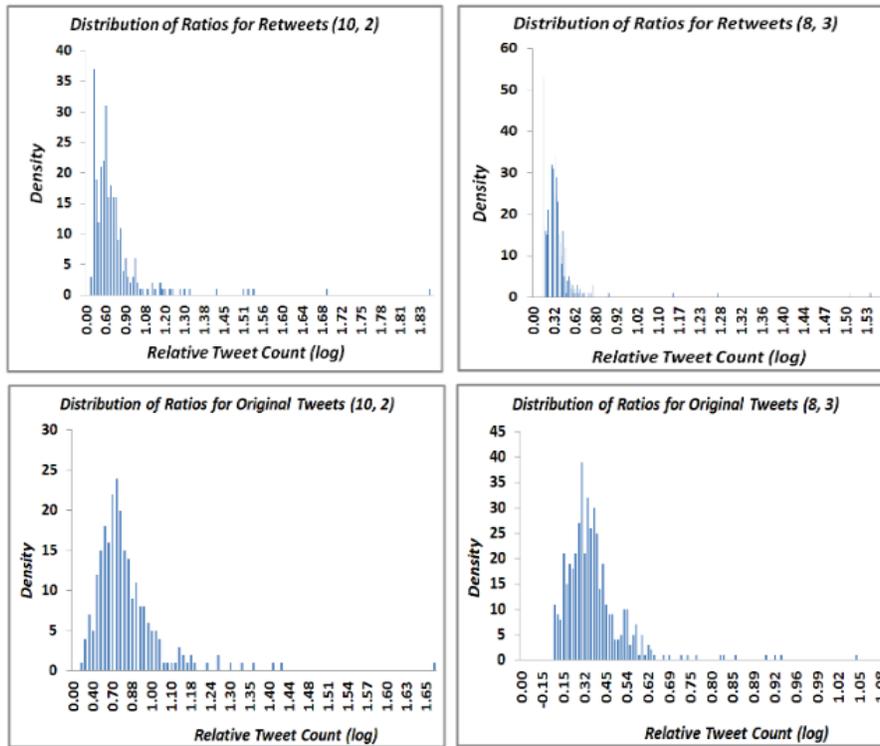


Fig. 4. The densities of ratios between cumulative original tweets/retweets counts measured in two arbitrary time frames: (10, 2) and (8, 3)

Weibo account. Thus we measured the percentage of user accounts that can still be accessed (as opposed to be directed to the error page) organized by user-retweet ratios (Table III). We observe that only 12% of the accounts with user-retweet ratios of above 30 are active. And, as user-retweet ratios decrease, the percentages of active accounts slowly increase. We consider this to be strong evidence for the hypothesis that user accounts with high user-retweet ratios are likely to be spam accounts.

Ratio	% Active Accounts	% Inactive Accounts
≥30	12%	88%
20 – 29	38%	63%
11 – 19	16%	84%
10	22%	78%
9	12%	88%
8	16%	84%
7	15%	85%
6	21%	79%
5	30%	70%
4	58%	42%
3	80%	20%
2	96%	4%
1	92%	8%

TABLE III
THE PERCENTAGE OF ACCOUNTS WHOSE PROFILES CAN STILL BE ACCESSED, ORGANIZED BY USER-RETWEET RATIO

We observe that in some cases, accounts with lower user-retweet ratios can still be a spam account. For example, an account could retweet a number of posts from other spam accounts, thus minimizing the suspicion of being detected as a spam account itself.

F. Removing Spammers in Sina Weibo

From our sample, after automatically checking each account, we identified 4985 accounts that were deleted by the Sina Weibo administrator. We called these 4985 accounts “suspected spam accounts”. There were 463,231 users in our sample, and 187,130 of them retweeted at least once. Thus we identified 1.08% of the total users (2.66% of users that retweeted) as suspected spam accounts.

Next, in order to measure the effect of spam on the Weibo network, we removed all retweets from our sample disseminated by suspected spam accounts as well as posts published by them (and then later retweeted by others). We hypothesize that by removing these retweets, we can eliminate the influences caused by the suspected spam accounts. We observed that after these posts were removed, we were left with only 189,686 retweets in our sample (51% of the original total retweets). In other words, by removing retweets associated with suspected spam accounts, we **successfully removed 182,429 retweets, which is 49% of the total retweets and 32% of total tweets (both retweets and original tweets) from our sample**. This result is very interesting because it shows that a large amount of retweets in our sample are associated with suspected spam accounts. The spam accounts are therefore artificially inflating the popularity of topics, causing them to trend.

To see the difference after the posts associated with suspected spam accounts were removed, we re-calculated the distribution of user-retweet ratios again for arbitrarily chosen

pairs of time frames. Figure 7 illustrates the distribution for time frames (10, 2). We observed that the distribution is now much smoother and seem to follow the log-normal distribution. We performed the log-normal test and verified that this is indeed the case.

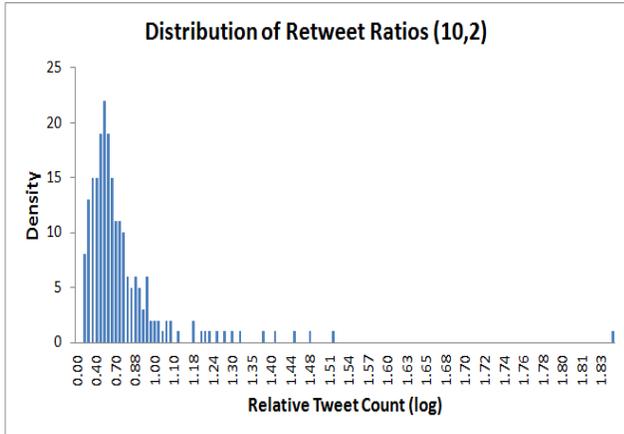


Fig. 7. The distribution of retweet ratios for time frame (10, 2) after the removal of tweets associated with suspected spam accounts

G. Spammers and Trend-setters

We found 6824 users in our sample whose tweets were retweeted. However, the total number of users who retweeted at least one person’s tweet was 187130, which is very skewed. Figure 8 illustrates the distribution for the number of times users were retweeted. This distribution follows the power law.

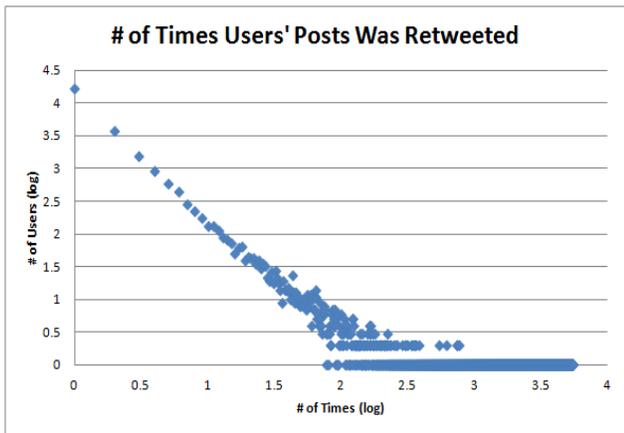


Fig. 8. The distribution for the frequency of retweets of user posts

We discovered that the number of users whose tweets were retweeted by the suspected spam accounts was 4665, which is a surprising **68%** of the users who were retweeted in our sample. This shows that the suspected spam accounts affect a majority of the trend-setters in our sample, helping them raise the retweet number of their posts and thereby making their posts appear on the trending list. The overall effect of the spammers is very significant. We also observed that a high

98% of the total trending keywords can be found in posts retweeted by suspected spam accounts. Thus it can also be argued that **many of the trends themselves are artificially generated**, which is a very important result.

H. Examples of Spam Accounts

Next, we investigate the activities of typical spam accounts Sina Weibo. We have shown that accounts with high retweet ratios are likely to be spam accounts. Although the majority of the accounts had already been deleted by the administrator, we manually inspected 100 currently existing accounts with high retweet ratios and found that 95 clearly participate in spamming activities. The other 5 were regular users supporting their favorite singers and celebrities by repeatedly retweeting their posts, which can also be construed as spam⁴. Figure 9 illustrates two examples of the activities from suspected spam accounts.

First, we observe that the suspected spam accounts we inspected tend to repeatedly retweet the same post with the goal of increasing the retweet number of said post. Next, the interval time of these repeated retweets tend to be very close to each other with long breaks between each set. Finally, we observe that the replies left from spam accounts often do not make any sense (see the comments circled in Figure 9). Chen et al. [12] had similar findings, and explained that this was because the paid posters are mainly interested in finishing the job as quickly as possible, thus they tend to retweet multiple times in short bursts and leave gibberish as replies. We observe that the replies in 9 a) and b) are not proper sentences.

For the 4665 users whose tweets were retweeted by at least one suspected spam account, we calculate the percentage of retweets from spam accounts and the percentage of suspected spam accounts involved. We selected only accounts whose tweets were retweeted by at least 50% of the accounts that are suspected spam accounts. From our manual inspection we found mainly three types of accounts:

- 1) Verified accounts from celebrities and reality show contestants : We hypothesize that they employ spam accounts to boast the popularity of their posts, making it seem like the posts were retweeted by many fans;
- 2) Verified accounts from companies : We hypothesize that they employ spam accounts to boast the perceived popularity of their products;
- 3) Unverified accounts with posts consist of ads for products : We hypothesize that these accounts employ spam accounts to distribute the ads and to boast the perceived popularity of their products, hoping other users will notice and distribute (see Figure 10 for an example).

V. DISCUSSION AND FUTURE WORK

We have examined the tweets relating to the trending topics in Sina Weibo. First we analyzed the growth and persistence of trends. When we looked at the distribution of tweets over time, we observed that there was a difference when contrasted with

⁴We exclude these from our list of suspected spam accounts



Fig. 9. Example of a spam account



Fig. 10. Example of an account using spam

Twitter. The main reason for the difference was that the effect of retweets on Sina Weibo was significantly higher than on Twitter. We also found (as our previous work [11] suggests) that many of the accounts that contribute to trends tend to operate as user contributed online magazines, sharing amusing pictures, stories and antidotes. Such posts tend to receive a large amount of responses from users and thus retweets.

When we examined the retweets in more detail, we made an important discovery. We found that 49% of the retweets in Sina Weibo containing trending keywords were actually associated with fraudulent accounts. We observed that these accounts comprised of a small amount (1.08% of the total users) of users but were responsible for a large percentage of the total retweets for the trending keywords. These fake accounts are responsible for artificially inflating certain posts, thus creating fake trends in Sina Weibo.

We relate our finding to the questions we raised in the introduction. There is a strong competition among content in online social media to become popular and trend and this gives motivation to users to artificially inflate topics to gain a competitive edge. We hypothesize that certain accounts in Sina Weibo employ fake accounts to repeatedly repeat their tweets in order to propel them to the top trending list, thus gaining prominence as top trend setters (and more visible to other users). We found evidence suggesting that the accounts that do so tend to be verified accounts with commercial purposes.

We have found that we can effectively detect suspected spam accounts using retweet ratios. This can lead to future work such as using machine learning to identify spamming

techniques. In the future, we would like to examine the behavior of these fake accounts that contribute to artificial inflation in Sina Weibo to learn how successful they are in influencing trends.

REFERENCES

- [1] S. Asur, B. A. Huberman, G. Szabo, and C. Wang, "Trends in social media - persistence and decay," in *5th International AAAI Conference on Weblogs and Social Media*, 2011.
- [2] B. A. Huberman, D. M. Romero, and F. Wu, "Social networks that matter: Twitter under the microscope," *Computing Research Repository*, 2008.
- [3] L. Jin, "Chinese outline BBS sphere: what BBS has brought to China," Master's thesis, Massachusetts Institute of Technology, April 2009.
- [4] F.-Y. Wang, D. Zeng, J. Hendler, Q. Zhang, Z. Feng, Y. Gao, H. Wang, and G. Lai, "A Study of the Human Flesh Search Engine: Crowd-Powered Expansion of Online Knowledge," *Computer*, vol. 43, 2010.
- [5] F. Y. Wang, "Beyond x 2.0: where should we go?" *IEEE Intelligent Systems*, vol. 24, no. 3, pp. 2-4, 2009.
- [6] T. Z. Xue, *The Internet in China : Cyberspace and Civil Society*. Routledge, 2006.
- [7] M. Xin, "Chinese bulletin board system's influence upon university students and ways to cope with it (in chinese)," *Journal of Nanjing University of Technology (Social Science Edition)*, vol. 4, pp. 100-104, 2003.
- [8] L. Yu and V. King, "The evolution of friendships in chinese online social networks," in *Proceedings of the 2010 IEEE Second International Conference on Social Computing*, ser. SOCIALCOM '10, 2010, pp. 81-87.
- [9] H. Kwak, C. Lee, H. Park, and S. Moon, "What is twitter, a social network or a news media?" in *Proceedings of the 19th international conference on World wide web*, ser. WWW '10, 2010, pp. 591-600.
- [10] M. Mathioudakis and N. Koudas, "Twittermonitor: trend detection over the twitter stream," in *Proceedings of the 2010 international conference on Management of data*, ser. SIGMOD '10, 2010, pp. 1155-1158.

- [11] L. Yu, S. Asur, and B. A. Huberman, "What trends in chinese social media," in *Proceedings of The fifth ACM workshop on Social Network Mining and Analysis*. ACM, 2011.
- [12] C. Chen, K. Wu, V. Srinivasan, and X. Zhang, "Battling the internet water army: Detection of hidden paid posters," *CoRR*, vol. abs/1111.4297, 2011.