Front.Comput.Sci.
DOI10.1007/s11704 – 016 – 5198 – y

RESEARCH ARTICLE

The Power of Comments: Fostering Social Interactions in Microblog Networks

Tianyi Wang[§†], Yang Chen (✉)[⋆], Yi Wang[‡], Bolun Wang[†], Gang Wang[†], Xing Li[§], Haitao Zheng[†], Ben Y. Zhao[†]

[§]Depart. of Electronic Engineering, Tsinghua University, Beijing, China
[†]Depart. of Computer Science, UC Santa Barbara, CA, USA
[⋆]School of Computer Science, Fudan University, Shanghai, China
[‡]Depart. of Mathematics, Syracuse University, NY, USA

Abstract

Today's ubiquitous online social networks serve multiple purposes, including social communication (Facebook, Orkut), and news dissemination (Twitter). But what about a social network's design defines its functionality? Answering this would allow social network providers to take a proactive role in defining and guiding user behavior.

In this paper, we take a *first* step to answer this question with a data-driven approach, through measurement and analysis of the Sina Weibo microblogging service. Often compared to Twitter because of its format, Weibo is interesting for our analysis because it serves as *both* a social communication tool and a platform for news dissemination. While similar to Twitter in functionality, Weibo provides a distinguishing feature, *comments*, allowing users to form threaded conversations around a single tweet. Our study focuses on this feature, and how it contributes to interactions and encourages social engagement. We use analysis of comment interactions to uncover their role in social interactivity, and use *comment graphs* to demonstrate the structure of Weibo user interactions. Finally, we present a case study that shows the impact of comments in *malicious user detection*, a key application on microblogging systems. Using properties of comments significantly improves accuracy in both modeling and detection of malicious users.

Keywords Microblogs, comments, social and interaction graph, user behavior

Received May 20, 2015; accepted Oct. 29, 2015

E-mail: *chenyang@fudan.edu.cn*

1 Introduction

Today's online social networks (OSNs) pervade through all aspects of our daily lives, serving a variety of functions from social communication (Facebook, Renren), to information and news dissemination (Twitter), to professional development (LinkedIn).

While some networks are designed with specific usage scenarios in mind, *e.g.* LinkedIn and Pinterest, others are more general and support a variety of usage-agnostic features such as making friends, messaging, and content sharing. Even within these general frameworks, recent work has shown that networks can evolve and become more specialized along specific usage scenarios, *e.g.* Twitter is often considered more of a news media platform than a social network [1].

While a variety of factors clearly contribute to the formation of user behavior, can we determine what role specific user features played in this process? The answer to this question can reveal potential impacts of social networking features on user behavior, and whether OSN providers can proactively "guide" user behavior by introducing features or interface modifications.

In this paper, we present first efforts to answer this question using an empirical, data-driven approach. More specifically, we examine this question through detailed measurements and analysis of Sina Weibo, a social "microblogging network." Analysis of Weibo provides an interesting case study because it is very similar to Twitter in nearly all aspects of its basic functionality, but is often viewed by its users as a hybrid network for *both* news dissemination and social interactions [2]. This is somewhat surprising, since social media tools in China largely mirror the functionality and usage patterns of their international counterparts, *e.g.* Renren [3] provides near-identical functionality to Facebook.

Through our work, the first goal is to understand how Weibo and Twitter actually differ via quantitative metrics. We then dive deeper to determine if these differences can be *partially* attributed to any individual user feature (or design choice). More specifically, our analysis focuses on Weibo's *comment* feature [1)], which is considered as one of the most distinguishing features between Weibo and Twitter by many analysts[2)3)]. To demonstrate how comments are correlated with such differences, we show how it contributes to social interactions in Weibo at the level of individual users, and how it shapes user interaction patterns similar to those of Facebook at macroscopic network levels. Finally, we also evaluate the importance of the comment feature at the application level.

Our work makes four key contributions.

First, we use a large-scale dataset to quantify the difference in network structure between Weibo and Twitter. Weibo exhibits significant structural differences: not only are each Weibo user's incoming and outgoing links more balanced, but a much larger portion of each user's relationships are bidirectional, *i.e.* consistent with friendships in social networks such as Facebook.

Second, we analyze the correlation between user comments and users' social interactions at the *individual user level*. We find comment is a more prevalent type of interactions in Weibo than reposts and mentions. We also find that users who comment (*commentors*) are frequently *friends* (bi-directional social link) with the tweet author, and they usually form concentrated conversations. The observations demonstrate that comments contribute to social interactions among users, and are highly correlated with users' bidirectional friendships.

Third, we further explore this issue at the *macroscopic level*, by analyzing and comparing the structures of Weibo's comment graph, Weibo's repost graph and Facebook and Twitter interaction graphs. We find that Weibo's comment graph most closely resembles Facebook's interaction graph despite the different user populations on the two sites, which implies that comment interactions are strongly indicative of bidirectional friendship interactions. In addition, we find significant correlations (overlaps) between Weibo's comment graph and social graph, and comment graph acts as a good predictor of social links. These confirm our intuition that comments play an important role in enabling bidirectional friendships at the macroscopic level.

Finally, at the *application level*, we use a case study to quantify the impact of comments on analyzing and modeling user behavior. Specifically, we look at *malicious user detection*, a critical application for all online social networks. We apply machine-learning tools to build a detector of malicious user activity, and apply it to ground-truth data from Weibo. We find that compared to commonly used features, features based on comment activity provide much higher discriminatory power in modeling the difference between malicious and normal users. Not only are comment-based features sufficient to produce an accurate detector, but when added to common features, they significantly elevate the accuracy of the resulting detector.

To the best of our knowledge, our work provides a first effort to explore the impact of comments on user behavior in social networks. The results of our work show that single feature like comments can contribute to significantly higher levels of interaction between different users.

2 Background and Methodology

We briefly describe Sina Weibo and our dataset to provide background for our analysis. We first discuss common features that Weibo shares with Twitter, and highlight its unique features. We then explain our data collection methodology and present high-level statistics of our dataset. Finally we introduce the methodology of this paper.

2.1 Sina Weibo

Sina Weibo is the largest microblogging service in China, and the second largest microblogging service in the world. As of March 2013, it has more than 500 million registered users[4)] and generates more than 100 million tweets per day[5)]. Weibo is popular around the world with many international users like Brazilian football player Pele and organizations like the UN.

Weibo shares many features with Twitter. Users can post tweets with up to 140 Chinese characters or 280 English characters, and repost (retweet) others' tweets. Each tweet can tag specific topics, mention other users by using an '@', post short URLs, geographic information and even pictures. A user can subscribe to other users' tweets by following these users. If user A follows user B, we say that A is B's *follower*, and B is A's *followee*. Weibo tweets are public to its registered users, although the platform only places each user's followees' tweets on her timeline. By default, Weibo users can manually visit any user's home page, which contains the user's profile and published tweets.

The most notable feature that distinguishes Weibo from

[1)] comment allows users to create threads of comments centered around a single tweet or microblog. For simplicity, we refer to microblogs in both platforms as *tweets*.

[2)] Twitter vs. weibo: 8 things twitter can learn from the latter. http://www.hongkiat.com/blog/ things-twitter-can-learn-from-sina-weibo/

[3)] Twitter vs weibo: Differences. http://www.juanmarketing. com/twitter-vs-weibo-differences/2011/05/24/

[4)] Sina has more than 500 million registered users. http://news.xinhuanet.com/newmedia/2013-02/21/c-124369896. htm

[5)] Weibo has more than 100 milion tweets per day. http://www.washingtonpost.com/blogs/worldviews/wp/2013/03/08/how-china-censors-100-million-tweets-per-day/

Twitter is the *comment* feature[6]. By default, a Weibo user can comment directly on any published tweet, and reply to any comment. This is a broad type of interaction, since users do not need to follow the author and commentors of a tweet before interacting directly with them. A comment does not create new tweet, but is associated with the tweet in a comment list, which includes all comments and replies sorted by time. The threaded comments make it easier for users to have concentrated conversations with both tweet authors and other commentors within the same tweet. As described in next section, the threaded comment list also allows us to crawl comments efficiently.

The comment function in Weibo is significantly different from the reply function in Twitter. In Weibo, all comments on a tweet are associated with the original tweet. Thus users can easily view past comments and add follow-on comments. In contrast, the reply function in Twitter will generate an independent tweet, which makes it more difficult to trace back all replies from a new reply[7]. The result is a much more strongly threaded sequence of messages into conversations.

2.2 Datasets

To crawl Weibo, we used its open API to access user profiles, tweets and comments. The API provides a user's complete list of followees, the latest 2000 tweets, and up to 5000 followers. Our crawls created two Weibo datasets.

Crawling the Social Graph. Obtaining an unbiased sample of the Weibo social graph is nontrivial. An unbiased sample is desired because it would capture the graph properties (*e.g.,* degree distribution) while making the graph size small. In an unbiased dataset, each node in the graph is sampled with the same probability. Conventional algorithms like Breadth-First Sampling (BFS) and Random Walk are known to be biased towards high degree nodes. That is, the users with high degree are more likely to be sampled. Existing unbiased sampling methods [4, 5] require the complete follower/followee set for each user, which is limited by the fact that Weibo API returns at most 5000 followers.

Instead, we seed our crawl using a large number of randomized user IDs. We leverage the fact that Weibo's API provides fast access to public tweets, and each tweet contains IDs of its author and mentioned users. We performed an API call once every 3 seconds for one month and obtained roughly 60 million unique user IDs [8]. We then crawled the Weibo network using these IDs as seeds, and obtained 57.1 million user profiles, each profile containing the number of the user's followers, followees and friends[9].

[6] The comment feature has been in Weibo since its inception.

[7] Twitter supports the function of displaying the full conversations in Aug., 2013, https://blog.twitter.com/2013/keep-up-with-conversations-on-twitter.

[8] Each API request returns 200 public tweets.

[9] When users *A* and *B* follow each other in a bidirectional social link, we call them *friends*.

Crawling Reposts and Comments. We also crawled a smaller set of the tweets for detailed repost and comment data by accessing the comment and repost lists of a tweet. Since crawling reposts or comments requires multiple API requests per tweet, we had to reduce the size of the targeted crawl to limit load on Weibo servers. In addition, we need a connected subgraph of users in order to analyze interactions among them. Thus we used BFS to obtain a connected subgraph. While BFS can introduce bias in node selection, it is attractive under our scenario because it is efficient and provides a direct comparison to prior work on user interactions [1, 6] that also used BFS. In total, we obtained 61.5 million tweets from 723K users, with 118.1 million comments and 86.2 million reposts.

To the best of our knowledge, our dataset is the most comprehensive sample of the Weibo social network to date. More importantly, while prior efforts focused on the social graph and tweet content [2, 7–10], our dataset contains all reposts and comments for 61.5 million tweets.

2.3 Our methodology

We seek to understand potential factor(s) that drive Weibo, a microblogging network, into a hybrid platform for both news dissemination (like Twitter) and social friending and communication (like Facebook). We reveal this by demonstrating how *comments* contribute to social interactions and friending in Weibo.

We start our analysis by quantifying the differences between Weibo and Twitter (this section). Particularly, by comparing the follower-followee relationship between users, we find Weibo users not only have more balanced incoming and outgoing links, but maintain a much larger portion of *bidirectional* relationships, *i.e.* resembling the friendships in traditional social networks such as Facebook.

These results motivate us to understand what function (or design choice) is correlated with the difference. We draw our attention to the *comment* function in Weibo. We focus on comments for two reasons. First, comment is a key feature that distinguishes Weibo from Twitter, and it is also a common function supported by traditional social networks, *e.g.* Facebook. Second, comment is one of the most heavily used communication channels by Weibo users. Users generate an order of magnitude more comments than tweets [11].

We study Weibo comment and how it contributes to social interactions and friending from *three* levels: individual user level (Section 4), macroscopic network level (Section 5) and application level (Section 6). We aim to give a comprehensive view on how this single design choice impacts on user's way of using Weibo.

Individual User Level. We start from individual user's perspective to understand their usage patterns of Weibo comments. More specifically, we analyze the popularity of comment as a communication channel in comparing with

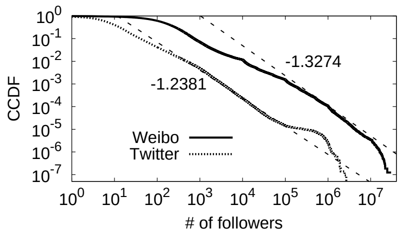

Fig. 1 Followers distribution.

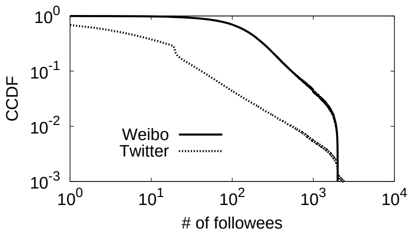

Fig. 2 Followees distribution.

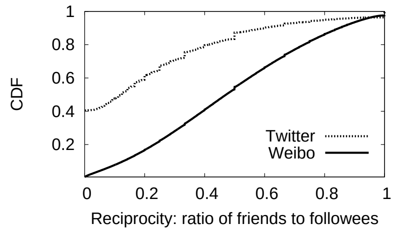

Fig. 3 Reciprocity distribution.

other channels such as tweet and repost. In addition, we analyze the relationship between users who participant in comment interactions, and explore whether users who comment are likely to be friends (bi-directional social link).

Macroscopic Network Level. We then move to a macroscopic view, by building network-wide interaction graphs, and comparing them to interaction graphs on Facebook and Twitter. We seek to understand whether it is the comment (or repost) function that defines users' interaction patterns (regardless of user background and culture). We compare Weibo's comment graph with that of Facebook, and Weibo's repost graph with that of Twitter, to explore their network-level similarities. In addition, we explore the correlation between comments and friendship at the network level by analyzing the overlaps between Weibo's comment graph and social graph. We also study the ability of comment graph to predict potential social links.

Application Level. Following observations that comments significant impact user behavior both at individual and network levels, we further study the importance of comment actions in user's overall behavioral profiles, and explore using comments to augment user-behavior based applications. In this paper we study two applications. First, we study machine-learning (ML) detectors of malicious users. Second, we study the influence maximization problem in social graph. For both applications, we do not focus on proposing any new methods or algorithms. Instead, we introduce comment, and evaluate how comment improves existing methods.

3 Weibo vs. Twitter

We start from a high-level comparison of Weibo and Twitter in terms of the social graph, focusing on the follower-followee relationship between users. We examine both user degree distributions and the ratio of relationship links that are bidirectional.

Our analysis uses our crawled and anonymized Weibo social graph, and an anonymized Twitter graph from [1] (41.7 million user profiles and 1.47 billion following relationships). Note that the Twitter graph is a complete crawl, and is unbiased and comparable to our Weibo data.

Degree Distribution In both Twitter and Weibo social graphs, a node's in-degree represents the number of followers of the user, and the out-degree represents the number of followees. In Figures 1 and 2 we plot the complementary cumulative distribution function (CCDF) of the follower and followee counts respectively. We make two key observations.

First, in terms of the follower count, the two networks display a similar pattern: the lines follow a commonly observed power-law distribution [1, 12], with similar scaling parameters [13] (α=1.3274 for Weibo and 1.2381 for Twitter). *Second* and more interestingly, the two networks differ significantly in terms of the followee count (out-degree). In Weibo, most users follow 10-100 users, while a significant portion of Twitter users follow no more than one user. One intuitive explanation is that Weibo users use it to communicate with friends, hence every user has some minimum number of users/friends they follow. To verify this intuition, we will check balance between followers and followees and reciprocity in later sections.

Balance between Followers and Followees To further study each individual user's follower/followee behavior, we define a *balance* metric, which is the ratio of a user's follower count to its followee count, *i.e.* a user's in-degree divided by her out-degree. Since a fully symmetric social network has a balance value of 1 for all users, for our analysis we define *well balanced users* to be those whose balance ratio is between 0.5 to 2.

We found that a higher fraction of users in Weibo belong to the balanced user category. Specifically, 57.4% of Weibo users are well balanced, while the ratio drops to 49% in Twitter. More broadly speaking, 80% of Weibo users have a balance ratio between 0.1 and 2 while only 60% of Twitter users do so. This again indicates that Weibo as a whole is more similar towards symmetric social networks.

Reciprocity We also consider *reciprocity*, another widely used metric for quantifying the symmetric relationship between users. It is defined as the ratio of a user's friend count (the number of bidirectional links) to her followee count (out-degree). It holds a value between 0 and 1: the higher the reciprocity, the higher the fraction of friends in a user's relationship.

Figure 3 plots the cumulative distribution function (CDF) of reciprocity for both Weibo and Twitter. Weibo has a

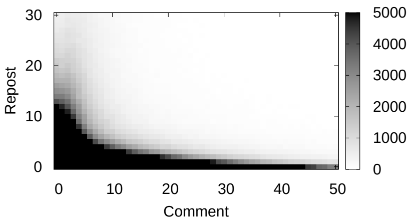

Fig. 4 Comment vs. repost.

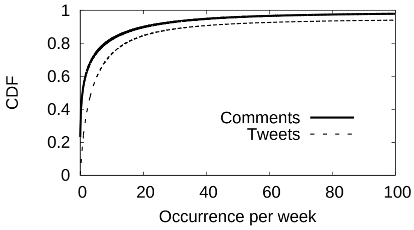

Fig. 5 User comment frequency.

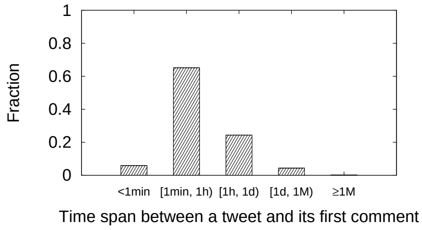

Fig. 6 First comment arrival time.

higher level of reciprocity than Twitter as a whole. For Weibo, 99.45% (40% for Twitter) users have a non-zero reciprocity, and more than half users (less than 20 for Twitter) have a reciprocity larger than 0.5.

From these results, we conclude that Weibo has a much higher level of reciprocity than Twitter, and shows more signs of supporting symmetric social relationships. The results motivate us to seek pontential factors which correlate with such differences.

4 Comment Analysis

With the differences between Twitter and Weibo, our next goal is to understand whether and how any design choice(s) in Weibo led to such differences. As discussed earlier, we will focus on "comment" feature, and look at the role of comments in the Weibo network. In this section, we focus on analysis at the *individual user level* and seek to answer the following questions.

- *Are comments a popular channel of user interactions*?
- *What are the temporal properties of comments, e.g. when do comments arrive following tweets*?
- *Who posts comments*? *What are the relationships between the author of a tweet and users who comment on them (commentors)*?
- *Do comments (and their replies) form intense social interactions, e.g. conversations among users*?

4.1 User Interactions in Weibo

There are three types of user interactions in Weibo: *Repost*, *Mention* and *Comment*. A user can *repost* an existing tweet, similar to *Retweet* in Twitter, *mention* another user in a tweet, or *comment* on an existing tweet (or other comments). Among the three, the comment feature is unique in Weibo. Our analysis in this subsection demonstrates that comment is the dominating form of user interactions in Weibo.

We present two key results that demonstrate the popularity of comments among Weibo users. The first result examines, for each user, the total number of comments, reposts and mentions for her latest 2000 tweets. Out of 723K users, the majority (65.8%) received more comments than reposts, and more than 55% users received at least 2 times more comments than reposts. A more detailed result is in Table 1, which lists the statistics of each feature in terms of the 80-, 50- and 20-percentile values across all the users. For all three metrics, the value corresponding to comments is significantly higher than that of reposts and mentions, often by more than 50%. These results show that users are more inclined to use comments.

	Comment	Repost	Mention
80 percentile of users	157	40	11
50 percentile of users	21	13	3
20 percentile of users	3	2	0

Table 1 Quantitative comparison among interactions.

Our second result examines interactions for each tweet. We plot the number of reposts versus that of comments in a heat map (Figure 4). For each point in the 2D plane, we count the number of the corresponding tweets and use the color to represent the tweet count. The darker the point is, the larger the number of tweets. For better visualization, we display only the significant part of the heat map by truncating it at 30 reposts and 50 comments[10]. We see that the black areas stretch widely along the x axis where there are considerably more comments than reposts.

These results confirm that comments are much more popular than reposts and mentions in Weibo. Next, we perform detailed analysis on this unique feature to gain a better understanding of users' comment behavior.

4.2 Frequency and Responsiveness

We begin by examining the temporal properties of comments. Specifically, we seek to quantify the *frequency*, *i.e.* how often a user receives comments, and *responsiveness*, *i.e.* how fast the comments arrive after a tweet.

To measure the comment frequency, we calculate for each user the total number of comments she received, normalized by the time span between her first and last tweets. We organize the results in terms of the average number of comments per week and show the CDF across all the users in Figure 5. Among all the users, 46.1% received at least 1

[10] This covers more than 99% of the tweets, because less than 1% of the tweets have more than 30 reposts or 50 comments.

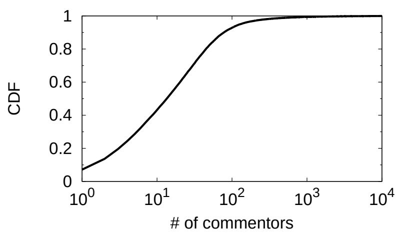

Fig. 7 Distribution of commentors for a user.

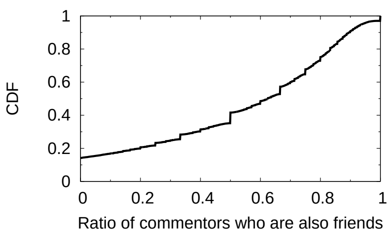

Fig. 8 Fraction of commentors who are also friends.

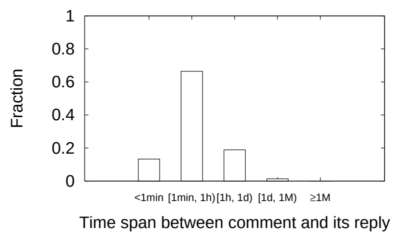

Fig. 9 Time span between comment and its reply.

comment per week and 17.5% received more than 10 comments per week. As a reference, we also plot the CDF of the average number of tweets per week for each user. The fact that the two CDF curves are close to each other indicates that Weibo users, by using comments, interact with each other almost as frequently as they tweet.

Next, we measure how quickly users interact via comments after a tweet. In Figure 6 we plot a histogram on the time span between each tweet and their first comment, excluding those without any comment. We observe that users act quickly to make comments. The large majority of tweets (95%) received their first comment within a day, and 71% got a comment within just an hour. Thus we can conclude that in Weibo, the comment feature enables fast (and concentrated) social interactions among users.

4.3 Composition of Commentors

For each user, we define a *commentor* as any other user who has posted at least one comment on her tweets. We now study the composition of commentors with four key questions: 1) *how many commentors does a user have*, 2) *what portion of these commentors are the user's friends* and 3) *what are the relationships among the commentors?*

Figure 7 plots the CDF of the number of commentors for each user. We see that 50% users have more than 10 commentors, and 10% received comments from more than 100 users. We also learn that the average number of commentors per user is 39.1, which is significantly larger than the average number of comments per tweet (1.92). We can infer that a commentor usually involves in commenting on multiple tweets.

To answer the second question, we plot in Figure 8 the CDF of the fraction of *friend commentors*, who are the commentors that are also the user's friends (bidirectional links). In our case, for 60% users, their friends contribute to more than half of their commentors. We see that the commentors are frequently friends with the tweet author.

For the third question, we define CC_u, a user u's commentor clustering coefficient, which measures the extent to which the commentors follow each other:

$$CC_u = \begin{cases} \frac{|F_{v,w}|}{c_u(c_u-1)} & \text{if } c_u > 1 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where c_u is the number of commentors and $F_{v,w}$ is the set of all following links between v and w (either direction) such that v, w are both u's commentors. CC_u helps us to evaluate how likely a user's commentors are friends (follow each other). A higher value of CC_u means user u's commentors are more likely friends with each other.

For our dataset, the average commentor clustering coefficient is 0.180. It is higher than the (general) clustering coefficient in Weibo (0.130, as we measure), Renren (0.063) [3], Facebook (0.164) [6] and Twitter (0.106) [14]. It indicates that a user's commentors are more likely to follow each other, reaffirming the fact that comments are indicative of strong social connections between users.

4.4 Conversations

The comment feature allows Weibo users to interact with each other at ease. In particular, users can reply to each other's comments. These replies (*i.e.* replying comments), if exist, arrive quickly, usually within an hour (Figure 9). Because comments and their replies reveal a unique type of (concentrated) interactions among users under the original tweet, we characterize them via *conversations* and study them in detail. Intuitively, a conversation contains a series of comments and the subsequent replies. To be more exact, a conversation is a chain of comments, where each comment replies to the former one.

In the following, we study the conversations in Weibo from three perspectives: 1) *how often do commentors form conversations*? 2) *how long does each conversation last*? and 3) *how often does a conversation involve the tweet author herself?*

We start by investigating, for each tweet with at least one comment, the portion of commentors involved in at least one conversation:

$$\text{ratio} = \frac{\text{\# of commentors in at least one conversation}}{\text{\# of commentors}} \quad (2)$$

Here we consider a tweet author as a commentor if she also participated in at least one conversation. Figure 10 shows the CDF of this metric across all the qualified tweets. We see that 60% tweets produced conversations (ratio>0). More specifically, in 50%+ tweets, more than 50% commentors were involved in conversations. This result is

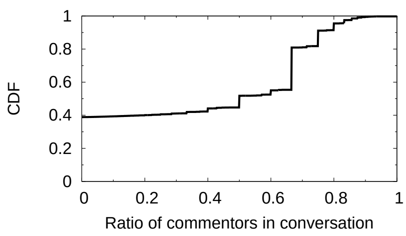

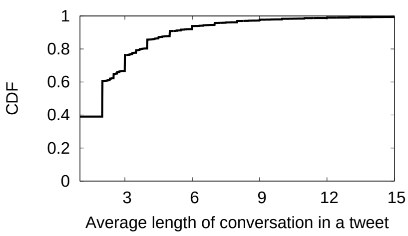

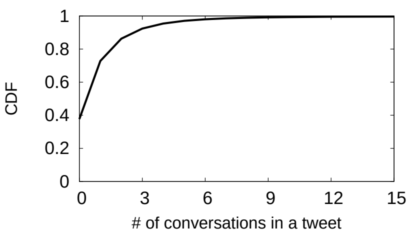

Fig. 10 Ratio of commentors who join in conversations.

Fig. 11 Average length of conversations in a tweet.

Fig. 12 Number of conversations per tweet.

Graphs	Nodes	Edges	Power-Law Fit α in/out	Exp. Cutoff λ in/out	RMSE in/out	Cluster. Coeff.	Assort.	Tie Strength
Comment	382,713	2,009,948	0.980/0.804	0.0796/0.0902	0.124/0.111	0.110	-0.011	7.74
Repost	352,557	1,529,967	2.09/1.13	**0.0015**/0.0769	0.193/0.145	0.070	-0.10	2.20
Facebook	273,497	1,627,253	0.927/0.956	0.0476/0.0596	0.112/0.167	0.081	0.23	6.54
Twitter	4,175,695	28,599,550	1.91/1.092	**0.0028**/0.0578	0.163/0.466	0.092	-0.027	3.14

Table 2 Basic properties of our interaction graphs, including the Comment and Repost Graphs from Weibo, plus the Facebook and Twitter Interaction Graph.

not surprising, considering the fact that users tend to respond to comments, effectively forming conversations.

Next, we quantify the length of a conversation by the number of comments and replies it contains. Since a conversation requires at least one comment and its reply, the minimum length is 2. Figure 11 plots the CDF of the average length of conversations in each tweet. We see that the majority of conversations are short, *e.g.* 80% of conversations contain 4 or fewer comments, and almost all the conversations have less than 15 threads. Note that a tweet may contain multiple conversations. Based on our analysis, 98% of the tweets have less than 6 conversations (Figure 12). These results show that the comment feature encourages concentrated interactions among users.

Finally, we look at the participants in each conversation, and examine how often the original tweet author gets involved. Interestingly, out of all the tweets that have any conversation, we observe 92% whose conversations all involve the tweet author, and 3% whose conversations never involve the tweet author. The latter case maps to tweets posted by celebrities and organizations, where the tweet authors just broadcast the news and the commentors initiate the subsequent conversations. From these results, we can conclude that the comment feature is highly effective for a tweet author to interact (bidirectionally) with other users.

4.5 Summary of Observations

Our detailed analysis provides four key observations:

- The comment feature, which is a key difference between user features in Weibo and Twitter, is also the most prevalent interaction mode in Weibo.
- Weibo users interact via comments nearly as frequent as they tweet, and usually provide comments very quickly within an hour of the original tweet.
- Each user receives comments mainly from friends.
- Comments on each tweet often form conversations, allowing users to interact intensively with each other in a short period. The mass majority of the conversations involve the original tweet authors.

Together, these key findings also confirm that the comment feature is a significant enabler (and contributor) to social interactions in Weibo. It makes Weibo essentially different from Twitter and much closer to classical social networks like Facebook.

5 Comment & Interaction Graphs

Next, we take our analysis on comments to a *macroscopic* level and examine the graph (or network) structure of comment activity in Weibo. We construct a *comment graph* and a *repost graph* from our Weibo data (Sec. 5.1), and compare their structures to different types of interaction graphs on Facebook and Twitter (Sec. 5.2). The comparison clearly reveal two patterns of interaction graphs: comment graph resembles that of Facebook and repost graph is close to that of Twitter. Then, we explore the correlations between user's interactions with their social relationships, by comparing Weibo's comment and repost graphs with Weibo's social graph from a statistical perspective (Sec. 5.3) and deploying the experiments of link prediction through an experimental perspective (Sec. 5.4).

5.1 Building Interaction Graphs

Prior works have examined social interactions in a number of social networks, focusing on both visible interactions such as wall posts and photo tags [6, 15, 16], and latent interactions such as social profile browsing [3, 17, 18]. For each form of interaction, we can build an interaction graph capturing the activity across the network [6, 19]. The resulting graph represents each user as a node, and each interaction taken by user A onto user B as a directed edge from A to B. If a user does not perform or receive any

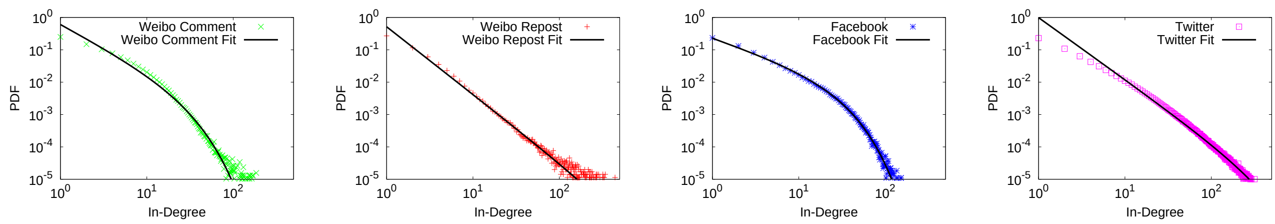

(a) Weibo Comment (b) Weibo Repost (c) Facebook (d) Twitter

Fig. 13 In-Degree Distribution and Fits using Power-Law with Exponential Cutoff.

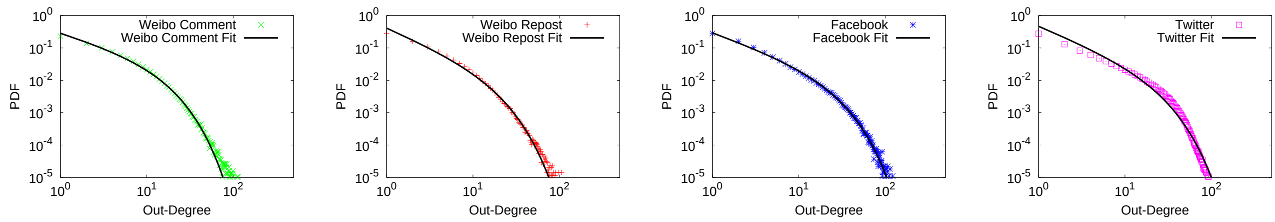

(a) Weibo Comment (b) Weibo Repost (c) Facebook (d) Twitter

Fig. 14 Out-Degree Distribution and Fits using Power-Law with Exponential Cutoff.

interactions, she becomes a singleton and is removed from the graph. For our analysis, the interaction graphs only need to capture the existence of interactions between users rather than the number of interactions. Thus the resulting graphs are directed and unweighted. In the following, we describe how we build interaction graphs in Weibo, Facebook and Twitter. Table 2 lists the basic properties of these graphs.

Weibo Interaction Graphs. To build Weibo's interaction graphs, we first obtain a connected subgraph (Weibo's *social graph*) of 723K users (see Section 2). We crawl the latest 2000 tweets from these users, along with all reposts and comments created by these users. These include comments or reposts made by users in our dataset on other users outside of our dataset. To build our graph, we only consider those interactions where both endpoints are within our user set. We denote each user as a node and a comment/repost from user A to B as a directed edge from A to B. In this way, we construct two interaction graphs from our dataset: Weibo *comment graph* and *repost graph*.

Facebook and Twitter Interaction Graphs. For our Facebook and Twitter graphs, we contacted the authors of prior papers on Facebook [6] and Twitter [19], and received permission to use their anonymized graphs as bases for comparison in our work. Wilson et al. built the visible interaction graph based on Facebook's wall posts [6] and compared it with the social graph of Facebook. We use the same dataset[11)], and built an anonymous interaction graph of Facebook. Unlike [6], our interaction graph is directed: when user A posts on user B's wall, we create a directed edge from A to B. For Twitter, we use the data set from previous work [19], which contains about 3 million users' profiles with social links and all of their tweets. We identify retweet interactions and built an interaction graph for Twitter from those events. More specifically, if user A publishes a tweet that includes "RT @" followed by user B's name, we create a directed edge from A to B in the interaction graph. We call this *Twitter interaction graph*.

5.2 Comparing Interaction Graphs

We compare and contrast the four interaction graphs in the context of graph metrics, including degree distribution, clustering coefficient, balance and reciprocity. Note that the data for all graphs were obtained through the same BFS algorithm, thus metrics of the three graphs are comparable.

Degree Distribution. Prior studies [13, 20] show that for most social networks, node degree follows a power-law distribution $P(k) \propto k^{-\alpha}$. We found, however, the power-law distribution with an exponential cutoff ($P(k) \propto k^{-\alpha}e^{-\lambda k}$), a generalized version of the power-law distribution, lowers fitting errors and is a better fit for our graphs.

With the probability distribution of power-law, we can estimate the expectation of the value of maximum degree. If the maximum degree is K, and the number of nodes in the graph is N, we have the following formula:

$$\int_K^{\infty} P(x)dx \approx \frac{1}{N}$$

It means: the expected number of nodes with degree $> K$ should be less than 1. If we set $P(x) = (\alpha - 1)x^{-\alpha}$, we can calculate an expected maximum degree:

$$K = N^{\frac{1}{\alpha-1}}$$

11) http://sandlab.cs.ucsb.edu/facebook/

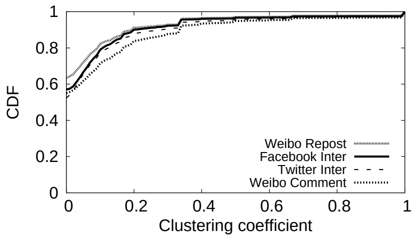

Fig. 15 Clustering coefficient distribution of interaction graphs.

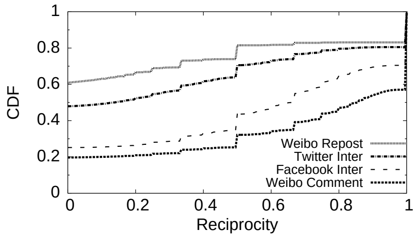

Fig. 16 Reciprocity comparison among interaction graphs.

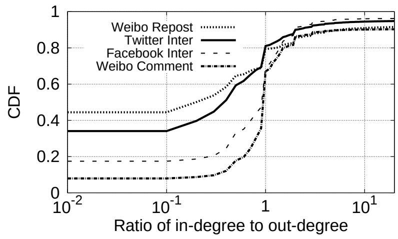

Fig. 17 Balance comparison among interaction graphs.

In practice, the value of α is usually between 2 and 3 [13]. If we have $\alpha = 2.5$, and we have 1 billion users, the maximum degree would be 1 million. This is normal for the number of Twitter followers, but would be too large for a user's Facebook friends.

Larger α would get smaller K. In the extreme case where $\alpha = 3$, we would get the smallest $K \approx 30000$. This would still be a little bit too large for a user's friends list. Prior studies on the Dunbar number [21, 22] shows that one can only maintain stable social relationships with a small set of people. As a result, a refinement with exponent cutoff ($P(k) \propto k^{-\alpha}e^{-\lambda k}$) needs to be introduced to the pure power law distribution.

Table 2 lists the fitting parameters (α and λ) and the root mean square error (RMSE) for these four interaction graphs. We use the Matlab function *cftool* to estimate the power-law fit parameter with the least mean squared error. We also plot the in- and out-degree distributions together with their fitted curves in Figure 13 and 14 respectively. From these results, we make two key observations.

First, the in-degree and out-degree distributions are similar for both Weibo comment graph and Facebook interaction graph. This means that like interactions in Facebook, comments in Weibo display a symmetric graph structure. On the other hand, for both Twitter interaction graph and Weibo repost graph, the in-degree and out-degree distributions are significantly different, indicating a strong asymmetry in user relationships.

Second, the in-degree distributions of Twitter interaction and Weibo repost graphs have very small λ values. A small λ means that the percentage of users with high node degrees drops slowly (or there are relatively higher number of users with high node degrees). This is because in both Twitter and Weibo, celebrities and popular organizations generate many tweets that are also highly retweeted, thus the resulting Twitter interaction and Weibo repost graphs have many nodes with high in-degrees. On the other hand, users are unlikely to maintain a large number of (symmetric) social interactions. In fact, this result aligns with prior studies on the Dunbar number [21, 22], which suggest that one can only maintain stable social relationships with a small set of people.

Clustering Coefficient. In Figure 15 we plot local clustering coefficients [23] for the four interaction graphs, and the average values are also shown in Table 2. The local clustering coefficient is defined as:

$$C_u = \begin{cases} \frac{|E_{v,w}|}{k_u(k_u-1)} & \text{if } k_u > 1 \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

where $E_{v,w}$ is the set of all edges between v and w (either direction) such that $v, w \in N_u$. Here N_u is the set of u's neighbors, *i.e.*, all nodes that are directly connected to or from u. Then $k_u = |N_u|$. The clustering coefficient is a fraction between 0 and 1 and characterizes the connectivity among one node's neighbors, where 0 represents a star shape around the local node, and 1 represents a full clique.

Interestingly, Weibo comment graph has a larger local clustering coefficient than Weibo repost graph. This means that the comment graph is more densely connected in its local structure, which confirms that comments are more prevalent than reposts.

Reciprocity and Balance. Figure 16 plots the reciprocity of our interaction graphs. The reciprocities for Weibo comment graph and Facebook interaction graph are close, and both are significantly higher than graphs based on information dissemination events, *i.e.*, Twitter interaction graph and Weibo repost graph. Bidirectional edges make up more than half of all edges for more than 70% of all nodes in both Weibo's comment graph and Facebook's interaction graph. This again demonstrates the impact of comments as a mechanism for bidirectional social interactions.

Figure 17 shows the balance of the interaction graphs, which represents the ratio of in-degree to out-degree in each graph. Again, we find that Weibo comment graph and Facebook interaction graph are similar, and are more balanced than the other two. Surprisingly, the ratio of balanced users is 72.8% in Weibo comment graph, which is slightly higher than that of Facebook (66%). This indicates that comments in Weibo are slightly more balanced and symmetrical than social interactions in Facebook.

Assortativity. Assortativity measures the homophily of users in a social network. A positive assortativity coefficient means nodes tend to connect with other nodes with similar degrees, which is usually considered a property of social network. Our result shows Weibo comment and Facebook have higher assortativity than Weibo repost and Twitter. Weibo comment does not produce very high assortativity, because comments partially serve as an interaction channel for users and celebrities, as we find in Section 4.

Tie strength. Tie strength measures how frequently users interact with each other. Our result shows users will interact more frequently via comments than reposts. This is consistent with the intuition. Friends will comment to each other frequently, while users repost only when there is a valuable tweet.

5.3 Interaction Graphs versus Social Graph

Next, we compare Weibo's interaction graphs (comment and repost) with Weibo's social graph. Our goal is to better understand, at the network level, whether user's interactions are correlated to whom users are making friends with.

We start by analyzing the overlaps between *comment* graph and social graph. Specifically, we examine how many edges in comment graph connect users with established social relationships (*e.g.* friend, follower or followee)[12]. The results are shown in the upside of Table 3. A quick observation is that users who have comment interactions (either one-way or bidirectional comment) have a high probability to be friends (76.8%), while the chance they are followers or followees is only 9.2%. This indicates that comment interaction has a strong correlation with the *bidirectional* social relationships in Weibo. Particularly, if two users comment to each other (bi-comment), 93.5% of the chance they are friends.

Next, we repeat the same analysis between *repost* graph and social graph. The results are shown in the bottom of Table 3. For *all* cases, users who have repost interactions are not necessarily to be socially connected: 70.1% of the chance they are not friends, and 50.2% of the chance they don't have any kind of social relationships between them. This indicates that repost is an interaction that also happens among strangers. Note that if two users repost each other's tweet (bi-repost), they still have a high probability to be friends (95.9%). However, bi-reposts are very rare in Weibo, with only 26K out of 1.3M repost edges bidirectional (1.9%).

Our analysis shows Weibo's comment graph has a bigger overlap with social graph than repost graph does, especially on bidirectional social relationships (76.8% to 29.9%). In another word, users who participate in comment interactions are typically friends, while repost is a interaction that often happens between non-friends or even strangers. This result helps to explain why Weibo (with comment function) has more balanced and symmetric social relationships than Twitter (with repost but no comment function).

Interac. Edges	Total #	Friend	Follower (e)	None
1-way comment	558,883	57.4%	16.6%	26.0%
Bi-comment	654,500	93.5%	2.8%	3.7%
All comment	1,213,383	**76.8%**	**9.2%**	14.0%
1-way repost	1,351,534	28.6%	20.3%	51.1%
Bi-repost	26,603	95.9%	3.1%	1.0%
All repost	1,378,137	29.9%	**19.9%**	**50.2%**

Table 3 Overlaps between Interaction and Social Graphs.

	Com. Neigh.(CN)	Jaccard(J)	Adamic(A)
Random	$8.33 * 10^{-6}$		
Comment graph	0.035	0.021	0.041
Social graph	0.026	0.014	0.031

Table 4 Accuracy of link prediction.

5.4 Link Prediction

In this section, we further confirm the strong correlation between comments and strength of social links via link prediction experiments. We follow experiments from [24] with widely-used metrics in prior work [25].

From the social graph we built, we randomly select 10% of edges as "missing" edges. We remove these edges from the social graph and comment graph. The link prediction problem is to find out deleted links based on the remaining graph structure, *i.e.* remaining social graph or comment graph after the edges are deleted. The idea is that two nodes with higher similarity are more likely to establish a link. We use three widely used metrics, *i.e.* Common Neighbors (CN), jaccard coefficient (J), and adamic (A) to measure node similarity [25]. If we use $\Gamma(x)$ to represent the neighbors of user x, then the three metrics can be calculated by the following three formulas: $CN = |\Gamma(x) \cap \Gamma(y)|$, $J = \frac{|\Gamma(x)\cap\Gamma(y)|}{|\Gamma(x)\cup\Gamma(y)|}$, and $A = \sum_{z\in\Gamma(x)\cap\Gamma(y)} \frac{1}{log|\Gamma(z)|}$. For all three metrics, a larger value means a higher similarity between the two nodes. To avoid unnecessary computational complexity, we only consider the nodes from deleted edges. We then rank node pairs by the similarity and get a list of top M edges. Here M is the number of the deleted edges. We define accuracy as the fraction of correctly predicted edges in the list out of all deleted edges.

We repeat the random selection of deleted edges 10 times, and show average accuracy in table 4. We have two observations. First, both graphs give a much better prediction than the baseline, *i.e.* random guessed edges from node pairs. Second, comment graph is more accurate at predicting social links than the social graph with an improvement of at least 30%. This implies that comments better capture strength of social ties, and those properties can be leveraged for more effective link prediction.

5.5 Summary of Observations

We have two key observations in our network-wised analysis.

- Weibo's comment graph resembles Facebook's interaction graph, with high bidirectionality; While Weibo's repost graph is similarity to Twitter's interaction graph, which represents one-way,

[12] Edges that indicate self-interaction, *i.e.* user comments to herself, are not considered in this analysis.

asymmetric interactions in information dissemination events.
- Weibo's comment graph has a much bigger overlap with social graph than repost graph. Users are more likely to be friends if they have comment interactions.

These results indicate the presence of comment function has played a significant role in enabling and encouraging bidirectional *interactions* and *friending* among users. It helps to transform Weibo into a platform not just for information dissemination (like Twitter), but also for social friending and interactions (like Facebook).

6 Comments in User Modeling

Thus far, our analysis has confirmed that Weibo's comment function has significantly changed how users interact. This in turn makes comment a potentially important dimension for modeling user behaviors. In this section, we take a deeper look at comment and user-behaviors based applications with two key questions: First, *how can we apply comment features in user-behavior based applications;* Second, *how significant the role do comments play in such applications?*

To answer these questions, we consider two practical application cases of comment feature. First, we use machine-learning detectors to identify malicious users. We aim to leverage comment features to build (or augment) user behavior models for detection. Second, we investigate the influence maximization problem with our comment and repost graph. We want to evaluate how information disseminates over diffrent types of interaction links.

6.1 Malicious User Detection

In the following, we first label ground-truth legitimate and malicious accounts. Then we describe our comment features to model user behavior. Finally, we use several machine learning techniques to build behavioral models for malicious user detection. As we will illustrate, our results show that the comment feature is a key factor in defining user behavior, which significantly boosts the accuracy across all of our machine learning based detectors.

6.1.1 Labeling Users

For our case study, we identify both malicious and legitimate accounts from crawled dataset as follows.

Malicious Accounts. Weibo relies in part on its users to report suspicious accounts, by submitting screenshots of their malicious behavior. Weibo administrators manually check reported accounts, and immediately block confirmed malicious accounts, making them inaccessible by Weibo's APIs. We leverage this to generate a ground-truth dataset of banned accounts. We performed a second round of crawls in September 2013 (seven months after the original crawl), and discovered 4639 accounts were blocked. We label these 4639 accounts as malicious accounts.

Legitimate Accounts. One way for users to verify their identities to Weibo is to bind or associate their accounts with either their Chinese national ID or cell phone number. These users can be identified by a special *Authenticated* label in their profiles. By providing their real-world identities, these users can be held legally for their actions, and are unlikely to behave maliciously. We identified 71,890 of these authenticated accounts in our crawls, and use them as our dataset of legitimate users.

6.1.2 Characterizing Comment Patterns

We first study the per-user comment activities of the malicious and legitimate user groups. In previous analysis (Section 4 and Section 5) we mentioned some general metrics to characterize user comment interactions, but here we focus on 4 metrics below that are more indicative of malicious users. The results are consistent, indicating that malicious accounts tend to make a significantly smaller number of comments than legitimate users.

Tweets with Comments. Our first metric is, for each user, the percentage of tweets that have user comments. As shown in Figure 18(a), legitimate users clearly are more likely to attract comments. For 71.9% of legitimate accounts, at least 20% of their tweets have comments. In contrast, only 24.4% of malicious accounts have comments on 20% or more of their tweets.

Incoming Comments per Tweet. We plot the average number of comments users received per tweet in Figure 18(b). Malicious accounts have fewer comments on their tweets. Only 16.1% of malicious accounts have one or more incoming comments for their tweets. The analogous number for legitimate accounts is 92.4%.

Ratio of Bi-commentors. Figure 18(c) plots the ratio of *bi-commentors* associated with each user. Two users are called bi-commentors if they have commented on each other's tweets or on the same tweet. The figure shows a significant difference between malicious and legitimate users. Nearly 80% of malicious accounts have zero bi-commentors, while bi-commentors are quite common for legitimate users.

Comment h-index. Our last metric is *comment h-index*. A user has a comment h-index of h if she has at least h tweets with no less than h comments. This metric is inspired by the "h-index" research publication impact metric [26]. Here we use it to measure user influence. As shown in Figure 18(d), malicious accounts have significantly smaller comment h-index values relative to legitimate accounts.

These results show that malicious accounts have much less comment interactions. There are two possible reasons: First, comments are a form of interaction that involves frequent exchanges with friends (as shown in Section 4).

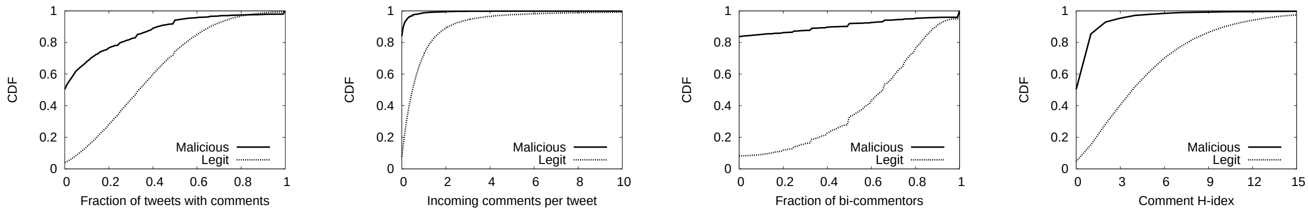

(a) Ratio of tweet with comments (b) In-comments per tweet (c) Ratio of bi-commentors (d) Comment h-index

Fig. 18 Comment patterns of malicious and legitimate accounts.

	EF Accuracy	CF Accuracy	ECF Accuracy
NaiveBayes	75.9%	87.1%	88.2%
SVM	92.2%	90.0%	93.1%
Random Forests	93.6%	90.2%	95.0%
Logistic Reg.	86.4%	89.3%	90.6%

Table 5 Detection accuracy using different algorithms.

These interactions come with a heavy overhead in time and energy, making them too costly for most malicious users. Second, comment does not really help malicious users to spam their followers, because unlike repost and tweet, comment does not generate new "events" on users' timeline, *i.e.* , zero impact on their followers. Thus comment is not an attractive activity for malicious users.

6.1.3 Detecting Malicious Accounts

To measure the extent to which comment interaction features can help with malicious user detection, we conduct three experiments, each applying a set of machine learning techniques but using different features. The first experiment uses 10 common user features already used in previous works [27, 28], including number of followers and followees, ratio of follower to followee, reciprocity, average, minimum and maximum number of tweets per day, number of mentions per tweet, ratio of tweets with mentions and ratio of tweets with URLs. We refer to these as the "existing features" (EF) set. For the second experiment, we use 9 features solely based on comment interactions, *i.e.* , the "comment features" (CF) set, including ratio of tweets with comments, ratio of outgoing to incoming comments, comment h-index, number of commentors, ratio of bi-commentors, ratio of friend commentors, ratio of conversations with tweet authors, number of incoming comments per tweet and number of outgoing commentors. Third, we experiment with the combination set of existing and comment-based features (ECF).

For detection, we use the banned 4639 malicious accounts and 4639 legitimate accounts randomly selected from the legitimate accounts set, and apply several widely used classification algorithms[13)], including NaiveBayes [30], Support Vector Machine (SVM) [31], Random Forests [32] and Logistic Regression [33] . We conduct the experiments with 10 fold cross-validation using EF, CF and ECF respectively. Accuracy is defined as the ratio of correctly predicted users to all users. Since we are using the same number of items for both categories, the accuracy is equal to the weighted average recall and F-measure [34].

[13)] Algorithm implementation by Weka toolkits [29]

Rank	Feature	χ^2	Type
1	# of out commentors	6144.9	CF
2	Ratio of friend commentors	5996.9	CF
3	Ratio of out to in comments	5905.2	CF
4	Avg # of in comments per tweet	5902.4	CF
5	Bi-commentors / all-commentors	5795.8	CF
6	# of followers	5215.5	EF
7	Comment h-index	4934.8	CF
8	Reciprocity	4932.8	EF
9	# of commentors	4675.7	CF
10	Ratio of tweets with mentions	4116.6	EF

Table 6 Top 10 features based on χ^2 statistic.

A summary of the experiment results are listed in Table 5. Our results show that using comment features (CF) alone can accurately distinguish malicious accounts from normal users with about 90% accuracy. In addition, adding comment features to existing features significantly boosts accuracy across all techniques, with an average improvement of more than 4%.

Feature Importance. Further, we want to examine the relative role comment-based features play in defining these detectors of malicious activities. Table 6 lists ECF's top 10 features ranked by χ^2 (Chi Square) statistic [35], which is a widely-used metric to measure feature's discriminative power. As shown, comment-based features account for all top five features and seven out of the top ten features. This indicates that comment-based features have stronger discriminative power than existing features.

However, despite comment feature's higher ranking, EF still outperforms CF on certain algorithms, *i.e.* Random Forests and SVM (Table 5). A possible explanation is that comment features are *individually* stronger but tend have overlapping effect when combined in a classifier. Figure 19 confirms this intuition: as we add more features to Random Forests classifier, EF provides higher incremental values than CF for boosting the overall accuracy. This indicates that

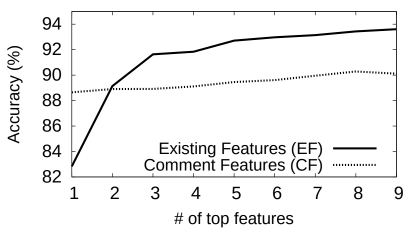

Fig. 19 Accuracy versus number of top features on Random Forests classifier.

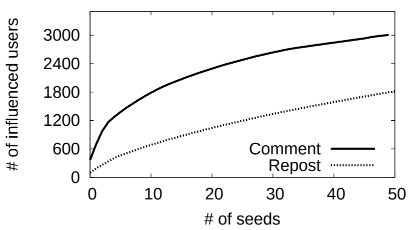

Fig. 20 Influence maximization in comment and repost graphs.

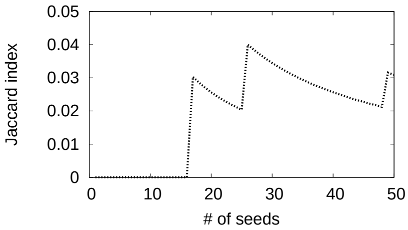

Fig. 21 Overlap of selected seeds.

the *diversity* of the feature set is still important for accurate user behavior modeling. Finally, Figure 19 also demonstrates the possibility to further shrink the feature set. In fact, if we only use the top-4 features from EF plus the top-1 feature from CF, Random Forests classifier can still produce accuracy as high as 94.1%.

Summary. The above analysis shows that comment interaction is a key factor in defining user behaviors. Particularly, it is a strong indicator for malicious behaviors of social network accounts. To this end, using comment-based features for malicious user detection is a big help for higher detection accuracy. In addition, comment-based features also make the detector more robust in practice – it is very difficult for malicious accounts to evade these features, because generating organic comment interactions take significant amount of efforts, time and even money (*e.g.* get comments via paid crowdsourcing [36]).

6.2 Influence Maximization

Understanding how information disseminates in OSNs is a critical problem for better marketing and user experience. The influence maximization problem is to decide a set of most influential people who can maximize the information diffusion in a social network. Formally, in a graph G and a random process which defines the information dissemination, the influence maximization problem is to optimize a set of seeds S, which maximizes the information propagation in G.

However, the influence maximization problem is proved to be NP-hard [37]. We can only use a greedy algorithm to get the best possible approximation. Varied algorithms are proposed to improve the efficiency of the greedy algorithm [38–40]. For example, CELF (Cost-Effective Lazy Forward) [38] optimizes the simple greedy algorithm based on the submodularity. In each round, CELF does not need to re-evaluate the incremental influence. As reported in [38], CELF is 700 times faster than the simple greedy algorithm empirically.

Goals and methodology. In this experiment, we aim to evaluate how information disseminates in two different graphs: comment graph and repost graph. We do not aim to propose any new algorithms for influence maximization. As a result, we leverage an existing and widely used algorithm CELF. This algorithm also serves as a baseline comparison in other works [39, 40]. For the information propagation, we use Independent Cascade Model (IC).

For our experiments, we run the CELF algorithm and consider the size of seeds set S varying from 1 to 100. For each S with n seeds, we evaluate the number of influenced nodes.

Results. We plot the evaluation results in Fig. 20. The x-axis is the number of seeds, and the y-axis is the number of users influenced by the seeds. From the figure, we find more users will be influenced in comment graph with the same size of seeds set. With 10 seeds, around 2000 users are influenced in comment graph. The corresponding value for repost graph is only 600. The result implies that comment graph is a more efficient representation for information diffusion in Weibo. This is not surprising, since comment is the dominating interaction type in Weibo. It reveals a closer relationship among users.

In Fig. 21 we evaluate the overlap of selected seeds between comment and repost graph. We find the selected nodes are quite different for different graphs. Out of 50 seeds, there are only 3 overlaps between the two graphs. That is, the representation of user interactions can lead to varied results in practical applications. The result suggests we should be careful when selecting the model of user interactions.

7 Related Work

Behaviors of Microblogging Networks. Recent studies have examined the behaviors of microblogging systems in detail, focusing primarily on Twitter. The work by Hwak et al. is the first to show that Twitter is a news media rather than a conventional social network [1]. A recent study [41] shows that Twitter has envolved into a hybrid of information network and social network, and it is based on only active users. Subsequent efforts have studied Twitter from different perspectives, ranging from information diffusion [42, 43], user influence [44, 45], to opinion mining [46] and user demographics [47]. Our work differs from these efforts by focusing on the unique feature of Weibo, a different (social) microblogging network. Using detailed data analysis, our study discovered viable structure differences between Weibo and Twitter, and identified the key user feature that causes such significant differences.

There have also been data-driven studies on Weibo, including video tweeting analysis [10], tweets deletion behavior [48, 49] and social influence [50]. Specifically, there are studies comparing Weibo and Twitter: One study compares Weibo and Twitter from several aspects of user behaviors (tweets in particular) [2], but did not consider social interactions, *i.e.* comments, among Weibo users. Another study examined information propagation in Weibo and Twitter [7], concluding that Twitter's information propagation is much faster and more frequent than Weibo. This conclusion is consistent with ours, which shows Weibo is not only a news media but also supports social interactions among users. We note that both studies were based on a very small fraction of the Weibo users (<1400) and tweets (1.5M). Our work differs from these by focusing on the comment feature. And our data in use is a much larger collection of both tweets and comments, *i.e.* 61.5M tweets from 723K users, and analysis of 57 million users.

Several works have examined Weibo from an application-centric perspective. Qu et al. studied the user behavior after a major earthquake, demonstrating the effectiveness of Weibo in providing quick responses to disasters [51]. Yu et al. tried to detect the sleeping times of users in Weibo according to their activities [52]. And Liao et al. studied the rumor propagation in Weibo and explored the information dynamics [9] while Yang et al. proposed automatic detection algorithms of rumors in Weibo [53]. These studies all conclude that an in-depth understanding of Weibo is critical in developing successful applications.

User Interactions in Social Networks. User interaction is a unique and critical feature of online social networks (OSNs), and has attracted attentions from the research community, including the visible interaction in Facebook [6] and mention-based interactions in Twitter [54]. Our work was motivated by these studies and their insights. Our work differs from these existing works by focusing on the impact of interactions and how these interactions are shaped. We found that Weibo users interact mostly via comments, not reposts, which form interesting social interactions among users similar to that of Facebook. We then built a series of Weibo interaction graphs to further understand such social interactions.

Malicious Account Detection. Researchers have devoted significant efforts to detect malicious accounts, (*e.g.*, spammers and sybils) in large OSNs and microblogging systems, including Facebook [55], Twitter [27, 28] and Renren [56–58]. One category of work [27, 28, 59] takes advantage of different classification algorithms of machine learning techniques. These works mostly focus on features of malicious accounts' (suspected) attacking behavior such as mentions, blacklist URLs and spam key words. Our work differs from these works by proposing effective new features of comments (user interactions).

8 Conclusion

This paper raises the question: *Can the design of a single feature affect user behavior in microblogging networks?* To answer this question, we perform a detailed measurement and analysis of the Sina Weibo microblogging network. Weibo is similar in all respects to Twitter, except for its support for *comments* on tweets. We believe that this key feature has led to significant amount of symmetric social interactions on Weibo, partially transforming it from primarily a news dissemination platform (like Twitter) into a hybrid social network that supports two primary use models: news dissemination and social interactions.

We test our hypothesis from a variety of perspectives using datasets from Weibo, Twitter and Facebook. Our results show that users display much more symmetric social connectivity than Twitter. We also show that Weibo's user comments are similar in many respects to Facebook's social interactions, and mark a clear departure from the asymmetric interactions typically found in Twitter. While limited to a single feature, our study demonstrates a strong correlation between a single design feature (comments) and a dramatic difference in user behavioral patterns.

References

1. Kwak H, Lee C, Park H, Moon S. What is twitter, a social network or a news media? In: Proceedings of the 19th International Conference on World Wide Web. 2010, 591–600
2. Gao Q, Abel F, Houben G J, Yu Y. A comparative study of user's microblogging behavior on sina weibo and twitter. In: Proceedings of the 20th Conference on User Modeling, Adaptation, and Personalization. 2012, 88–101
3. Jiang J, Wilson C, Wang X, Sha W, Huang P, Dai Y, Zhao B Y. Understanding latent interactions in online social networks. ACM Transactions on the Web (TWEB), 2013, 7(4): 18
4. Gjoka M, Kurant M, Butts C T, Markopoulou A. Practical recommendations on crawling online social networks. IEEE Journal on Selected Areas in Communications, 2011, 29(9): 1872–1892
5. Ribeiro B, Towsley D. Estimating and sampling graphs with multidimensional random walks. In: Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement. 2010, 390–403
6. Wilson C, Boe B, Sala A, Puttaswamy K P, Zhao B Y. User interactions in social networks and their implications. In: Proceedings of the 4th ACM European Conference on Computer Systems. 2009, 205–218
7. Gao Q, Abel F, Houben G J, Yu Y. Information propagation cultures on sina weibo and twitter. In: Proceedings of the 4th Annual ACM Web Science Conference. 2012, 157 – 162

8. Kw F, Chau M, Holme P. Reality check for the chinese microblog space: a random sampling approach. PLOS ONE, 2013, 8(3): e58356
9. Liao Q, Shi L. She gets a sports car from our donation: Rumor transmission in a chinese microblogging community. In: Proceedings of the 2013 Conference on Computer Supported Cooperative Work. 2013, 587–598
10. Guo Z, Huang J, He J, Hei X, Wu D. Unveiling the patterns of video tweeting: A sina weibo-based measurement study. In: Proceedings of Passive and Active Measurement. 2013, 166–175
11. Chen L, Zhang C, Wilson C. Tweeting under pressure: Analyzing trending topics and evolving word choice on sina weibo. In: Proceedings of the 1st ACM Conference on Online Social Networks (COSN). 2013, 89–100
12. Sala A, Zheng H, Zhao B Y, Gaito S, Rossi G P. Brief announcement: Revisiting the power-law degree distribution for social graph analysis. In: Proceedings of the 29th ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing. 2010, 400–401
13. Clauset A, Shalizi C R, Newman M E J. Power-law distributions in empirical data. SIAM Review, 2009, 51(4): 661–703
14. Java A, Song X, Finin T, Tseng B. Why we twitter: understanding microblogging usage and communities. In: Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 Workshop on Web Mining and Social Network Analysis. 2007, 56–65
15. Chun H, Kwak H, Eom Y H, Ahn Y Y, Moon S, Jeong H. Comparison of online social relations in volume vs interaction: a case study of cyworld. In: Proceedings of the 8th ACM SIGCOMM Conference on Internet Measurement. 2008, 57–70
16. Viswanath B, Mislove A, Cha M, Gummadi K P. On the evolution of user interaction in facebook. In: Proceedings of the 2nd ACM Workshop on Online Social Networks. 2009, 37–42
17. Benevenuto F, Rodrigues T, Cha M, Almeida V. Characterizing user behavior in online social networks. In: Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement Conference. 2009, 49–62
18. Schneider F, Feldmann A, Krishnamurthy B, Willinger W. Understanding online social network usage from a network perspective. In: Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement Conference. 2009, 35–48
19. Xu T, Chen Y, Jiao L, Zhao B Y, Hui P, Fu X. Scaling microblogging services with divergent traffic demands. In: Proceedings of the 12th International Middleware Conference. 2011, 20–39
20. Mislove A, Marcon M, Gummadi K P, Druschel P, Bhattacharjee B. Measurement and analysis of online social networks. In: Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement. 2007, 29–42
21. Dunbar R. Neocortex size as a constraint on group size in primates. Journal of Human Evolution, 1992, 22(6): 469–493
22. Dunbar R I. The social brain hypothesis. Foundations in Social Neuroscience, 2002, 5(71): 69
23. Watts D J, Strogatz S. Collective dynamics of'small-world'networks. Nature, 1998, 393(6684): 440–442
24. Zhao X, Chang A, Sarma A D, Zheng H, Zhao B Y. On the embeddability of random walk distances. VLDB Endowment, 2013, 6(14): 1690–1701
25. Liben-Nowell D, Kleinberg J. The link-prediction problem for social networks. Journal of the American Society for Information Science and Technology, 2007, 58(7): 1019–1031
26. Hirsch J E. An index to quantify an individual's scientific research output. National Academy of Sciences of the United States of America, 2005, 102(46): 16569–16572
27. Benevenuto F, Magno G, Rodrigues T, Almeida V. Detecting spammers on twitter. In: Proceedings of Collaboration, Electronic Messaging, Anti-Abuse and Spam Conference (CEAS). 2010, 12
28. Wang A H. Don't follow me: Spam detection in twitter. In: Proceedings of the 2010 International Conference on Security and Cryptography (SECRYPT). 2010, 1–10
29. Hall M, Frank E, Holmes G, Pfahringer B, Reutemann P, Witten I H. The weka data mining software: an update. ACM SIGKDD Explorations Newsletter, 2009, 11(1): 10–18
30. Lewis D D. Naive (bayes) at forty: The independence assumption in information retrieval. In: Proceedings of 10th European Conference on Machine Learning. 1998, 4–15
31. Chang C C, Lin C J. Libsvm: A library for support vector machines. ACM Transactions on Intelligent Systems and Technology (TIST), 2011, 2(3): 27
32. Breiman L. Random forests. Machine Learning, 2001, 45(1): 5–32
33. Le Cessie S, Van Houwelingen J C. Ridge Estimators in Logistic Regression. Applied Statistics, 1992, 191–201
34. Billsus D, Pazzani M J. Learning collaborative information filters. In: Proceedings of International Conference on Machine Learning (ICML). 1998, 46–54
35. Yang Y, Pedersen J O. A comparative study on feature selection in text categorization. In: Proceedings of International Conference on Machine Learning (ICML). 1997, 412–420
36. Wang G, Wilson C, Zhao X, Zhu Y, Mohanlal M, Zheng H, Zhao B Y. Serf and turf: Crowdturfing for fun and profit. In: Proceedings of the 21st International Conference on World Wide Web. 2012, 679–688
37. Kempe D, Kleinberg J, Tardos É. Maximizing the spread of influence through a social network. In: Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. 2003, 137–146
38. Leskovec J, Krause A, Guestrin C, Faloutsos C, VanBriesen J, Glance N. Cost-effective outbreak detection in networks. In: Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. 2007, 420–429
39. Chen W, Wang Y, Yang S. Efficient influence maximization in social networks. In: Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. 2009, 199–208
40. Goyal A, Lu W, Lakshmanan L V. Celf++: optimizing the greedy algorithm for influence maximization in social networks. In: Proceedings of the 20th International Conference Companion on World Wide Web. 2011, 47–48
41. Myers S A, Sharma A, Gupta P, Lin J. Information network or social

network?: the structure of the twitter follow graph. In: Proceedings of the Companion Publication of the 23rd International Conference on World Wide Web Companion. 2014, 493–498
42. Yang J, Counts S. Predicting the speed, scale, and range of information diffusion in twitter. In: Proceedings of the International AAAI Conference on Web and Social Media (ICWSM). 2010, 355–358
43. Lumezanu C, Feamster N, Klein H. #bias: Measuring the tweeting behavior of propagandists. In: Proc. of the International AAAI Conference on Web and Social Media (ICWSM). 2012
44. Cha M, Haddadi H, Benevenuto F, Gummadi P K. Measuring user influence in twitter: The million follower fallacy. In: Proceedings of the International AAAI Conference on Web and Social Media (ICWSM). 2010, 10–17
45. Bakshy E, Hofman J M, Mason W A, Watts D J. Everyone's an influencer: Quantifying influence on twitter. In: Proceedings of the 4th ACM International Conference on Web Search and Data Mining. 2011, 65–74
46. Boutet A, Kim H, Yoneki E. What's in your tweets? i know who you supported in the uk 2010 general election. In: Proc. of the International AAAI Conference on Web and Social Media (ICWSM). 2012
47. Mislove A, Lehmann S, Ahn Y Y, Onnela J P, Rosenquist J N. Understanding the demographics of twitter users. In: Proceedings of the International AAAI Conference on Web and Social Media (ICWSM). 2011
48. Bamman D, O'Connor B, Smith N. Censorship and deletion practices in chinese social media. First Monday, 2012, 17(3)
49. Zhu T, Phipps D, Pridgen A, Crandall J R, Wallach D S. The velocity of censorship: High-fidelity detection of microblog post deletions. In: Proceedings of the 22nd USENIX Conference on Security. 2013, 227–240
50. Zhang J, Liu B, Tang J, Chen T, Li J. Social influence locality for modeling retweeting behaviors. In: Proceedings of the 23rd International Joint Conference on Artificial Intelligence. 2013, 2761–2767
51. Qu Y, Huang C, Zhang P, Zhang J. Microblogging after a major disaster in china: a case study of the 2010 yushu earthquake. In: Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work. 2011, 25–34
52. Yu H, Sun G, Lv M. Users sleeping time analysis based on microblogging data. In: Proceedings of the 2012 ACM Conference on Ubiquitous Computing. 2012, 964–968
53. Yang F, Liu Y, Yu X, Yang M. Automatic detection of rumor on sina weibo. In: Proceedings of the ACM SIGKDD Workshop on Mining Data Semantics. 2012, 13
54. Macskassy S A. On the study of social interactions in twitter. In: Proceedings of the International AAAI Conference on Web and Social Media (ICWSM). 2012
55. Gao H, Hu J, Wilson C, Li Z, Chen Y, Zhao B Y. Detecting and characterizing social spam campaigns. In: Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement. 2010, 35–47
56. Wang G, Konolige T, Wilson C, Wang X, Zheng H, Zhao B Y. You are how you click: Clickstream analysis for sybil detection. In: Proceedings of the 22nd USENIX Conference on Security. 2013, 241–256
57. Wang G, Mohanlal M, Wilson C, Wang X, Metzger M, Zheng H, Zhao B Y. Social turing tests: Crowdsourcing sybil detection. In: Proceedings of the 20th Annual Network and Distributed System Security Symposium (NDSS). 2013
58. Yang Z, Wilson C, Wang X, Gao T, Zhao B Y, Dai Y. Uncovering social network sybils in the wild. ACM Transactions on Knowledge Discovery from Data (TKDD), 2014, 8(1): 2
59. Stringhini G, Kruegel C, Vigna G. Detecting spammers on social networks. In: Proceedings of the 26th Annual Computer Security Applications Conference. 2010, 1–9

Tianyi Wang is a PhD student in the Department of Electronic Engineering, Tsinghua University, China. He was a visiting student at the Department of Computer Science, UC Santa Barbara, USA. His research interests include data mining, the analysis and modeling of Online Social Networks.

Yang Chen is a Pre-tenure Associate Professor within the School of Computer Science at Fudan University. Before that he was a Postdoctoral Associate within the Department of Computer Science at Duke University from Apr. 2011 to Sep. 2014. He received his B.S. and Ph.D. degrees from the Department of Electronic Engineering, Tsinghua University, in 2004 and 2009, respectively. His research interests include online social networks, Internet architectures, and cloud computing. He is a senior member of the IEEE.

Yi (Grace) Wang received her Diploma in Mathematics from Huazhong University of Science and Technology, Wuhan, China, 2005. From 2006 to 2012, she was with the University of Minnesota, where she received her M.S. in Statistics and Ph.D. in Mathematics. Between 2012 and 2015, she was a postdoctoral researcher at the Statistical and Applied Mathematical Sciences Institute (SAMSI) and a visiting assistant professor at the Mathematics Department, Duke University, Durham, NC. Since 2015, she joined as an assistant professor the Department of Mathematics at Syracuse University, Syracuse, NY. Her research interests lie in applied harmonic analysis, machine learning, signal and image processing, as well as applications to real data.

Bolun Wang is a 1st year PhD student in the Computer Science Department at the University of California, Santa Barbara. He received his B.S. degree from Tsinghua University in 2009. His research interests are Online Social Networks, Data Security, and Privacy.

Gang Wang received his B.E. degree in electrical engineering from Tsinghua University, Beijing, China, in 2010. He is currently pursuing the Ph.D. degree in computer science in the University of California at Santa Barbara, CA, USA. His research interests are security and privacy, online social networks, mobile networks and crowdsourcing systems.

Xing Li received his PhD in Electrical Engineering from Drexel University. He is a Professor with the Electronic Engineering Department, Tsinghua University and the deputy director of China Education and Research Network (CERNET) Center.

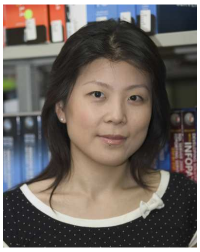

Haitao Zheng received her B.S.degree (with highest honor) from Xian Jiaotong University in July 1995, her M.S.EE and Ph.D degree in Electrical and Computer Engineering from University of Maryland, College Park, in May 1998 and July 1999, respectively. She joined wireless research lab, Bell-Labs, Lucent Technologies as a member of technical staff in August 1999, and moved to Microsoft Research Asia as a project lead and researcher, in March 2004. Since Sept. 2005, she has been a faculty member in Computer Science Department, University of California, Santa Barbara, where she is now a Professor. Dr. Zheng was named as the 2005 MIT Technology Review Top 35 Innovators under the age of 35 for her work on cognitive radios. Her work was selected by MIT Technology Review as one of the 10 Emerging Technologies in 2006. She also received 2002 Bell Laboratories President's Gold Award from Lucent Bell-Labs, and 1998-1999 George Harhalakis Outstanding Graduate Student Award from Institute of System Research, University of Maryland, College Park. She was elected IEEE fellow in 2014. Dr. Zheng's recent research interests include wireless systems and networking and social networks.

Ben Zhao is a Professor at the Computer Science department, U. C. Santa Barbara. He completed his M.S. and Ph.D. degrees in Computer Science at U.C. Berkeley (2000, 2004), and his B.S. from Yale (1997). He is a recipient of the National Science Foundation's CAREER award, MIT Technology Review's TR-35 Award (Young Innovators Under 35), and Computer-World Magazine's Top 40 Technology Innovators award. His work has been covered by media outlets such as New York Times, Boston Globe, MIT Tech Review, and Slashdot. He has published over 100 publications in areas of security and privacy, networked and distributed systems, wireless networks and data-intensive computing. Finally, he has served as program chair for top conferences (WOSN, WWW 2013 OSN track, IPTPS, IEEE P2P), and is a co-founder and steering committee member of the ACM Conference on Online Social Networks (COSN).