

She Gets a Sports Car from Our Donation: Rumor Transmission in a Chinese Microblogging Community

Qinying Liao

IBM Research - China
Bldg. 19 Zhongguancun Software Park
Beijing, 100193
liaoqy@cn.ibm.com

Lei Shi*

State Key Laboratory of Computer Science
Institute of Software, Chinese Academy of Sciences
Beijing, 100190
shil@ios.ac.cn

ABSTRACT

In this paper we report on a case study of rumor transmission during a nationwide scandal via China's most popular microblogging service, weibo.com. Specifically, we explore dynamics of the rumor discourse by characterizing different statement types and their evolution over time. We examine the roles that different user groups play in the rumor discussions. Through qualitative and statistical analyses, our results identify seven reaction patterns to rumors and their different development trends. We reveal a three-stage pattern of the change of leadership during the rumor discussions. By connecting social theories on rumor transmission to the large scale social platform, this paper offers insight into understanding rumor development in social media, as well as utilizing microblogging data for effectively detecting, analyzing and controlling public rumors.

Author Keywords

Microblogging, social media, rumor transmission, China

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

General Terms

Human Factors.

INTRODUCTION

Rumors are unconfirmed messages, but prevalent in everyday communications. With the flourishing of social networking services (SNS, e.g. Facebook, Twitter), one has good reasons to believe that rumors could travel much faster and wider today than ever before. Researchers have been interested in examining the credibility of information on SNS (e.g., [11][20]) and developing technologies to identify rumors (e.g., [20][28]), under a common view that rumors are "misinformation" and "disinformation".

In many ways, however, rumors are not merely destructive.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CSCW '13, February 23–27, 2013, San Antonio, Texas, USA.
Copyright 2013 ACM 978-1-4503-1331-5/13/02...\$15.00.

Social scientists have believed rumors to be an attempt to relieve anxieties [25]. Recently others suggest that rumors are a communication strategy [33]. CSCW research has however overlooked this perspective of understanding rumoring behavior, which is common on SNS. Furthermore, little is known so far about how a rumor is transmitted or discussed on SNS. Practically, we believe understanding various rumor discussions could help to design and develop technologies for identifying, tracking and responding to rumors that could have a massive impact in the society.

In this paper, we report findings from a case study on rumor transmission in China's largest microblogging system, Sina Weibo. In 2011, rumors around a young lady's wealth and her possible connections with a government-owned charity caused a nationwide credibility crisis. We are interested in understanding how people participated in the rumor discussions and how the rumor developed over time. Specifically, we examine two aspects of the rumor transmission: evolution of content patterns and different user roles in the rumor-related discussions. We categorize different rumor statement types and analyze their development trends. Key leaders are identified that drive major discussions during the course. By bringing together social theories and empirical studies on microblogging, we propose a model of rumor discourse evolution for understanding large-scale rumor transmission processes, and examine the potential of using microblogging data to respond to rumors that are of public concern.

The rest of this paper is structured as follows. We first briefly review the existing research on rumor transmission and microblogging. We then describe the background of our case, as well as our methods of data collection and analyses. After that we report and discuss our results.

RELATED WORK

We review literatures on two main topics related to our study: First, we summarize known characteristics of rumor transmission from researches in psychology and sociology. Second, we examine recent work on the role of microblogs in generating and disseminating breaking news. We then introduce our research questions.

* Corresponding author

Rumor Theory

Rumor discourse has been regarded as a form of collective behavior in social sciences (e.g., [7][8][12][25]). While the authenticity of the unconfirmed message could be a matter of debate, whether a rumor is truthful or not is unimportant in studying its transmission [10]. It is the building-up of collectively sanctioned explanations of ambiguity that drives social scientists' research interest. Despite the existence of different definitions of rumor, they share common ideas of two essential features: *importance* and *uncertainty*. Rumors tend to arise in situations that bear topical relevance to people involved but are cognitively unclear and largely outside of their own control [4][5][10][23][24]. Such conditions are associated with people's stress of "collective ambiguity and anxiety" [25]. When people fail to obtain necessary and assured information from credible sources, they circulate unofficial news in the community to relieve the tensions of such uncertainty and anxiety [25]. Rumor transmission, as a social process for people to collect, evaluate and interpret information [15], provides meanings to the ambiguity and offers its participants a sense of control over the situations [23][25].

Despite the development of the theoretic frameworks, previous rumor research has been suffered from the lack of empirical evidence due to the difficulties in conducting natural observations. Allport and Postman [4] have explored the social factors in rumor transmission using the serial reproduction paradigm in experimental situations. In their theory, rumors would lose most details (*leveled*) during the discourse, while only certain details are selected and emphasized (*sharpened*). Moreover, according to them, rumors are often distorted in the telling to fit one's unconscious affection and intellectual motivations (*assimilation*). This framework, however, has been criticized for not resembling rumor transmission in real life which tends to be an exchange [10][23] and involves feedback from listeners actively.

More recent work on rumor spread utilizes small scale data collected in computer-mediated communication (CMC) settings (e.g., [6][7][8][9]) and indicates rumor interactions to be a rich conversational process. In these studies, different transmission patterns (*aka.* the Rumor Interaction Analysis System) are identified [7], including apprehensive, interrogatory, authenticating, sensemaking, belief and disbelief, etc. Four stages of interactions are suggested. In the first stage, the rumor is tentatively introduced; in the second and third stage, sensemaking increases and then dominates; in the fourth stage, the rumor fades out and other topics take over [7]. Besides, temporal analysis on group dynamics reveals that rumor discussions via CMC largely follow the face-to-face group development patterns [8]. These findings are consistent with the sociological view of rumoring as group problem-solving.

Researchers have also been interested in understanding individuals' role in rumor discourse. Shibutani discussed the communication styles of *messenger* and *interpreter* [25]

although without any empirical evidence. Bordia and Rosnow [9] analyzed eight clusters of communication profiles on rumor messages in CMC groups, for example, skeptical disbelievers, investigators, curious responses, etc. More recently, Bordia and DiFonzo [7] created a topology of what they called *communicative postures* with eleven clusters, such as evaluating, verifying, information-reporting, accepting, etc.

In retrospect, the rumor phenomenon is an old problem that seemed intractable due to the difficulty of study in naturalistic settings until the use of data in CMC groups. However, existing CMC studies only examined rumors in small (e.g. 30 participants [8][9]) and relatively closed groups (e.g. discussion groups [7]), and the rumors discussed were very light (e.g. number of most postings <60 [7][9]). With the latest development of SNS (e.g. Facebook, Twitter), substantial content especially breaking news nowadays is distributed widely in a speed that one can never imagine before. This not only offers an opportunity for observing the new patterns of rumor transmission, but also poses novel challenges of analyzing large scale social data of rumoring, which have not been addressed in literature. Next we review the work on news spread and crisis response with microblogging, mainly on Twitter.

Microblogging

Since its release in 2006, Twitter has been the mainstream microblogging service worldwide. People share their everyday activities, report news, and seek for information on Twitter [14]. Research has suggested that Twitter is more an information spreading medium than a social networking tool. The study by Kwak *et al.* [16] with a large-scale data collection showed a much lower reciprocity in the follower-followee links on Twitter than Facebook-like SNS. Recently more research has been conducted to understand and model the information flows on Twitter. Wu *et al.* [31] discovered that 50% URLs received on Twitter were created by the top 20,000 elite users and the information flows on Twitter possessed high homophily in that one category of people (e.g. celebrities) tended to follow people of the same kind.

One typical social impact of Twitter is to break grand news [13][16]. For example, the work by Hu *et al.* reveals that Twitter has convinced the mass audience of bin Laden's death even earlier than the traditional media [13]. In order to identify such emerging news, several methods have been developed to track the short-term topic trends on Twitter [17][32].

More recent work has paid intensive attention to the role that Twitter and other microblogging systems play during large public events, such as disasters and crises. In [27][29], the Twitter usages after the Red River flood and Oklahoma Grassfires are studied. Researchers have found and discussed an information production life-cycle of generation, derivation, synthesis and innovation on Twitter

[27]. Information vital to situational awareness, such as vulnerable geographical area, is shown to be traceable on Twitter for a better emergency management [29]. As a whole, microblogging has been considered as a wide-scale interaction platform [21] for hosting situation updates, communicating emotions [22], distributed problem-solving [21], as well as self-organizing rescuing activities [26] during large public events.

More relevant to our work, Mendoza *et al.* investigated the dissemination of false rumors and valid news on Twitter [19] under the similar context of disasters. By manually tagging user interactions over a few typical rumors and news, they obtained preliminary results that the Twitter users tend to question the baseless rumors while few ever deny the confirmed news. Another piece of work by Morris *et al.* has systematically studied the credibility issue of microblogging [20]. Through carefully designed surveys and online experiments, they verified a list of critical features (e.g. message topic, user name and user image) that could heavily impact the credibility of Twitter's messages. These findings suggested feasible methods for Twitter users and social search engines to enhance their credibility.

To the best of our knowledge, none of the existing work has studied the crowd's various reactions to rumors and how rumor discussions develop over time via microblogging. Taking the social perspective of rumor transmission, we are interested in exploring the following research questions:

Q1: How are public rumors generated and how do they develop in the microblogging community? What kinds of statements do people post on microblogs and how do they change over time during rumor discussions?

Q2: Do people play different roles in transmitting rumors with microblogging? Who leads the discussions? Do their roles change in the different phases of the rumor discourse?

STUDY OVERVIEW

In this section we first introduce the background of Sina Weibo and the Guo Meimei Incident, and then describe our method of data collection and analyses.

Sina Weibo (Sina Microblogging)

Launched in 2009, just about three years after Twitter, Sina Weibo [1] is by far the most popular microblogging platform in China. Till March 2012, Sina Weibo has more than 300 million registered users with over 100 million messages (i.e., "tweets" or "*weibos*") sent every day [2].

Sina Weibo is similar to Twitter, but different in many senses. One of the most striking differences is the volume of information communicated via each microblog in Chinese versus English. Both Sina Weibo and Twitter restrain the word count to 140 characters, but 140 characters in Chinese could convey much more information than that in English [30]. Such a language efficiency turns microblogging in China into a blog-like platform. Moreover,

besides re-posting ("re-tweeting" or "*zhuanfa*") a message, Sina Weibo allows users to comment on others' posts or add a comment when re-posting a message (with an extra 140 words). These features could promote more communications on Weibo by invoking new discussions and spreading messages.

Similar to [22], in this paper we use the word "post" for the publishing of a message, and "re-post" for the forwarding of an existing message. A "message" (*weibo*) could be a post or a re-post.

Guo Meimei and the Red Cross Incident

Guo Meimei (@ 郭美美 Baby), a 20-year-old woman, stirred China's Internet in summer 2011. In a number of messages on her microblog, she posted pictures of her piles of luxury handbags and sports cars (e.g. Figure 1). At the same time, she appeared in her verified Weibo profile to hold the position of "commercial general manager" at China's Red Cross Society, a government-backed charity.



Figure 1. Two sample posts of Guo Meimei about her luxurious lifestyle on Sina Weibo. Left: "Just got my 20-year-old birthday gift. The Year of 2011 is a lucky year, right?" Right: "These were my mommy's loving bags and now they are mine!!! If one day I left home with my lover and took away some of them, we can live off them, hahah"

On the night of June 21st when some anonymous microbloggers began to notice and pointed out Guo Meimei's scandalous posts, the young lady ignited a firestorm. On June 20th, Guo Meimei had only 2,126 people following her. In just three days, her followers increased to 190,000. By June 27th, the number had grown to 700,000.

Suspecting that Guo might have financed her lifestyle out of the money from charity, furious people started a *human flesh search*¹ hoping to find out who Guo Meimei was and how she became connected with the Red Cross. Rumors flushed out all over the Internet, with more and more stories told and details added. Guo afterwards denied her having any ties with the Society and posted apologetic messages declaring that her title was fabricated for fun out of her

¹ *Human flesh search* is a Chinese Internet phenomenon of collectively identifying and exposing individual's profile to the public using social media, e.g., forums and microblogs.

ignorance. The Red Cross also held numerous press conferences to deny the rumor. But the suspicions were not relieved. The public did not even buy in the police's investigation which said Guo and the Society had no direct ties with each other. For months, disclosures of inside stories and disputes continued to flood. Ms. Guo and the Red Cross had been the most talked-about subject among China's five hundred million netizens through the year. The Guo Meimei Incident, as it is called, ranked as the largest public affair generated from Sina Weibo in 2011.

It is not our intent to examine the credibility of information or the trust issue associated with this incident on Sina Weibo. As prior research [10] pointed out, the authenticity of a rumor is not important for studying its transmission. In our paper, we focus on understanding different discourse behaviors and user roles in the rumor spreading. Next, we describe our method of data collection and analyses.

Data Collection

A keyword search of posts with "Guo Meimei" (郭美美) on Sina Weibo now returns 9.7 million messages. Due to the huge volume of data, we decided to employ a sampling method for data collection. We only collected the relevant data from the website for 41 days, starting from June 21 when Guo Meimei's posts first came into the public notice, until July 31. According to Baidu Index² [3], the number of online posts about Guo Meimei dramatically went down in late July. Although there were discussions related to Guo Meimei in August as well, most posts during this period were around Guo Meimei's appearance on several TV shows instead of her rumored connection with the Red Cross. We thus decided to exclude the messages posted after July and focus on the major rumor discourse.

For practical reasons, we constrained our sample to the original posts that contain "Guo Meimei" (郭美美) in their messages. We are aware that we miss a portion of messages which do not contain "Guo Meimei". But that would require intensive technical work to identify text meanings and we believe the number of such posts is small.

During these 41 days, a total of 474,000 original posts had been published on Sina Weibo according to the search result. However, only a list of 50 pages with 20 posts per page (1,000 posts in total) was available for each search on Sina Weibo. We then made advanced search queries on a per-hour and per-province basis. The peak hour for original posts from Beijing was around 10AM on June 23 and the number of posts was 878, which is smaller than 1,000. Theoretically we have made a full coverage of the relevant data by making such advanced searches. Based on the search results, we periodically collected posts for 10

consecutive seconds in every minute, and sampled roughly 16% of all the posts (76,128) during those 41 days.

To exclude the spamming messages (e.g., advertisement) that also contain "Guo Meimei" but are irrelevant to the rumor discussions, two researchers examined 17,000 posts in the sample and identified about a hundred spam-patterns. These patterns were then applied to all the collected posts with the help of SQL queries to remove the spams. The final cleaned sample contains 49,725 posts, which we call the *large sample* throughout this paper.

For each post in the large sample, we collected data on author's name, time of publication, message content, number of comments, number of reposts, and full lists of associated comments, reposts and their authors, all through Sina Weibo API. For each user appearing in these 49,725 posts, we collected their profiles including geographical location, background, total number of posts, followers, followings and the status of certification on Sina Weibo. As a result, we collected profiles for 38,221 unique users in total.

Data Analysis

Statement Coding

Our first goal of the study was to understand different content of rumor discussions and the evolution over time. To identify the types of statements people posted during the Guo Meimei Incident, a smaller sample with 1,000 messages was created for manual coding and the in-depth qualitative analysis. A uniform sampling method similar to the large sample is adopted so that the temporal development trend of the *small sample* strictly follows that of the large sample.

We mixed a top-down and a bottom-up strategy to develop our coding framework. First, we built an initial coding scheme based on the Rumor Interaction Analysis System (RIAS) [7]. Two researchers applied the scheme to a random collection of 150 messages independently. Then they discussed their results and revised the coding scheme together. After three rounds of iteration, the RIAS scheme was revised and a final list of seven categories of statement was determined (see Table 2). Specifically, we combined three RIAS categories (*prudent*, *authenticating*, *personal involvement*) into *providing information*, replaced *apprehensive* with *emotional statements* to cover both the positive and negative, and added a new category of *opinions* which included the former *belief*, *disbelief* and *sarcastic statements*. The reason for such revisions was that some of the original categories in RIAS appeared too rare in our sample. The basic unit of our coding scheme was a 'statement' instead of a message as a message could contain several complete thoughts and belong to multiple categories.

² *Baidu index* is a keyword research and trend analysis tool developed by Baidu, China's largest search company.

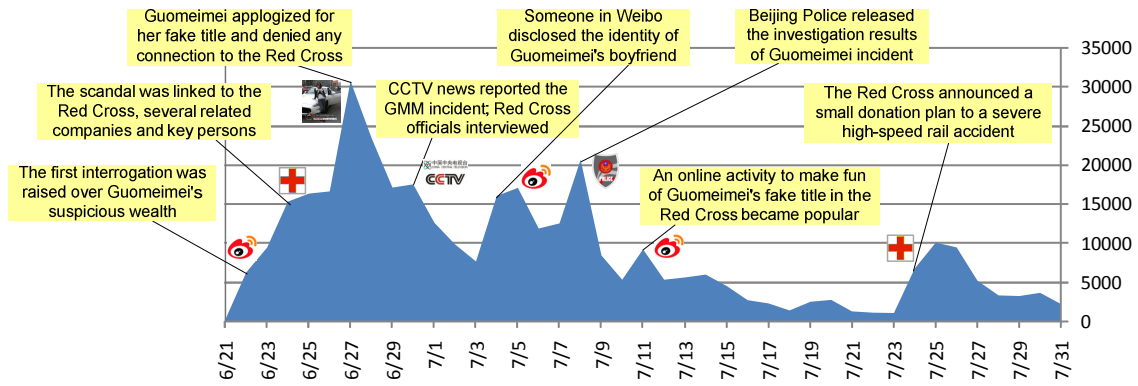


Figure 2. The trend of the original messages related to the rumor (49,725 posts).

The two researchers then applied the final coding scheme independently to all the 1,000 messages in the small sample. After that, they discussed together and found about 15% classifications different. They examined each of the disagreements until reaching consensus. As a result, a total of 1,186 statements were coded from the 1,000 messages.

Category	Verified Type	# Follower	Example
Celebrity	Personal	≥ 10000	Entertainment/Sports Star, President of large corporations
Certified	Personal	< 10000	Manager of small companies
Mass Media	Media	(Xinhua news, Times
Organization	Enterprise/Academic/Non-profit	-	Shopping mall, car company
Website	Website	-	Guangzhou portal
Internet Star	N/A	≥ 10000	Bloggers, personalized content providers, online opinion leaders
Ordinary	N/A	< 10000	Other Microblog users

Table 1. Sina Weibo user categorization.

User Categorization

Our second research goal was to understand the different roles microblogging users play in the rumor discourse. We first determined a method to categorize Sina Weibo users according to their profiles and online behaviors. Then we applied this method to our large sample of 49,725 messages from 38,221 unique users.

Sina Weibo provides publicly visible “VIP-titles” to verified users. There are multiple types of certifications, including individuals, media agencies, organizations (enterprise, academic and non-profit) and websites. To build our user categories, we incorporated users’ microblogging behavior with their certifications provided by Sina Weibo. We applied a cut-value of 10,000 in the

number of followers to identify the top 1.96% users, who attracted 79.9% followers in the entire user base. The final user categories are shown in Table 1.

RESULTS

What Did People Talk About?

In this section, we first present the overall trend of rumor discussions on the Guo Meimei Incident. Then we examine different types of statements people posted during the rumor discussions and their evolutions over time.

Overall Trends

We analyzed the overall trends of the 49,725 posts collected during the 41 days (June 21~July 31). 28.5% of these posts have at least one comment, 27.5% are reposted at least once. In total, these posts have 133,318 comments and 343,596 reposts by a total of 245,847 unique users.

We then examined the development curve of these posts upon some key dates and the associated events during the rumor discussions, as illustrated in Figure 2. The rumor was generated on the late night of June 21 and immediately exploded. People gathered rapidly and engaged heavily in the discussion. Multiple peaks appeared during the 41 days. The first peak emerged around June 27 after Guo Meimei posted her apologetic messages. Around June 30, a second peak was achieved, following the mass media’s interview with the Secretary General of the Red Cross Society. The post number dropped a bit in the coming few days, and then surged again on July 4 when another rumor was leaked out indicating Guo’s boyfriend to be a high-level official at the Red Cross. Intensive discussions continued and then reached another peak with the police investigation results released on July 8. Later on July 11, a small peak was caused by a popular online competition activity which made fun of Guo’s fabricated title at the Red Cross (messages were automatically posted to the microblogs of people who joined this activity, e.g., “Now I am seriously showing off my wealth: I am pleased to accept the offer of the position of the Commercial General Manager at the Red Cross Society. I have beaten 1087 people!”). After that, discussions about Guo Meimei became quieter until the Society announced a small donation plan for a severe high-

speed rail accident which touched upon the public's sensitive nerve again on July 24.

Statement Types and Trends

Table 2 shows the overall distribution of the 1,186 statements we manually coded. Most statements fell into three major categories: providing information (410, 34%), expressing emotions (226, 19%) and opinions (315, 27%). Notably, about 10% (122) of the statements were those attempting to analyze and interpret the situation, which was

Categories	Description	%
Providing information	Definition: Statements adding material to the rumor, e.g. 'factual' details, evidences, situation updates, inside stories, personal involvements, etc.	34
	Example: "Insiders leaked the identity of Guo Meimei's boyfriend [URL]"	
Opinion related statements	Definition: Statements indicating people's attitude towards the rumor, e.g. belief/disbelief, criticizing/agreement, etc.	27
	Example: "You suspect when they don't deny. But when they deny, they are actually telling you the rumors are true."	
Emotional statements	Definition: Rumor-related emotions, both positive and negative, including fear, anxiety, apprehension, surprise, stress, anger, wish, excitement, etc.	19
	Example: "Guo Meimei P.K. Red Cross! I am totally astonished!!!"	
Sense-making statements	Definition: Statements attempting to analyze and interpret the situation, e.g. disputing, justifying one's views, etc.	10
	Example: "During the whole course, the biggest problem lies in Sina Weibo's verification system."	
Interrogatory statements	Definition: Question seeking information	3
	Example: "#Guo Meimei incident# How did all this happen? Can anybody tell me?"	
Directive statements	Definition: Statements suggesting action.	3
	Example: "Why don't we put Guo Meimei into a human flesh search?"	
Digressive statements	Definition: Off the topic. Statements irrelevant to the rumor.	4
	Example: "Will this girl become the next Guo Meimei?"	

Table 2. Statement types and distribution (1,186 statements from 1,000 messages).

not seen in people's responses to other extreme events discussed on Sina Weibo (e.g., a major disaster [22]). Another 3% asked for information explicitly for clarifying the rumors.

We studied how posts of different statement types were reposted and found that the messages providing information (6.31 reposts in average) and disseminating opinions (8.82) were the most disseminated. All the other types, such as conveying emotions (0.98) and sensemakings (0.73) were rarely reposted in average. Overall, 80% reposts were made within the first 12 hours after the original messages were posted.

Our analysis of the statement trends revealed interesting patterns in the distribution of different conversation types over time. Figure 3 shows the combined trend, while Figure 4 gives a clearer view to the evolution of each statement type separately.

The rises and falls of messages *providing information* are quite similar to the overall waves of discussions. A closer look at the content of these posts also showed high relevance of the messages with the key phases illustrated in Figure 2. For example, one group of typical posts in this category is those to contrast and question Guo's wealth against her title in the early phase: "A 20-year-old young woman. Maserati, Lamborghini. Pure descent racehorse from Germany. Tens of Hermes handbags. First class airplane tickets. And...Commercial General Manager, Red Cross China." (@公益高剑)

" 'The role of Red Cross is to help the people in need.' Said Guo Meimei, the Commercial GM. This 20-year-old girl lives in a palatial villa, drives a Maserati." (@乐活山西)

Meanwhile, new materials were synthesized from various sources about Guo's hometown, family background, educations and personal life. Later on, another group of posts kept updating the Red Cross' reactions and news from other media. For example, "The Society says: Guo Meimei has no connection with the Red Cross [URL]", "The National Television calls Red Cross to prove itself innocent [URL]" and "Beijing Police: Guo Meimei has nothing to do with the Red Cross. [URL]"

Generally, the emotional curve reached its peaks a bit later than the informational curve in each stage, although not much. This is probably because people tend to express their emotions as intuitive reactions to unfamiliar situations. The surge of emotional statements could be roughly divided into three phases. The first phase was from the beginning of the rumor to the first two peaks of discussions around June 27 and June 30. The second phase followed the rumors about Guo's boyfriend and the police investigation during July 3 to 11. The last phase was generated around the Red Cross' donation to the railway accident victims in late July. Our preliminary qualitative analysis indicated people's different emotions in the different phases of rumor discussions. In the first phase, people expressed a lot of anxiety and worry. For example, "Guo Meimei has set a fire on the Internet!" "OMG! Guo Meimei! Explosive news! I can't go to work anymore!" "I can't bear this if this issue is not solved." Later, disappointment rose. "It has already been a week. I




Figure 3. Statement trend accumulated (1000 messages).

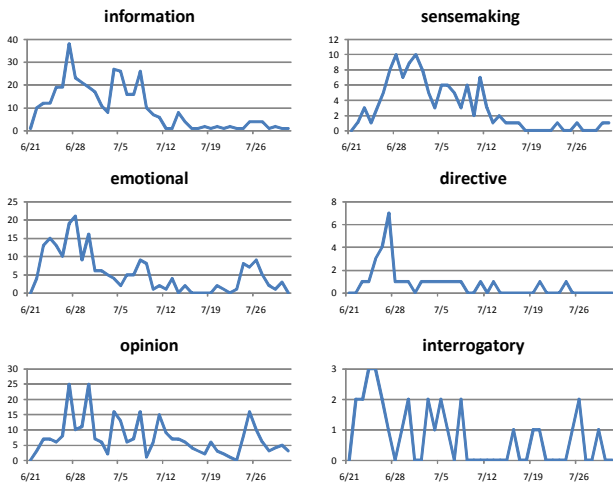


Figure 4. Statement trends separated (1000 messages)

am disappointed at the government.” “I am totally depressed.” In the second phase, we observed more anger and sarcasm. For example, “I feel like eating a worm in the apple.” “We all have to thank Guo Meimei. She is a hero!” During the last peak, some became indifferent and sneered down the Society’s small donation to the “7-23” railway accident victims: “Oh Red Cross! Save your donation for Guo Meimei’s handbags!”

Interestingly, it is observed that people’s emotional statements around the Guo Meimei scandal sometimes follow certain patterns of expression. For example, on the second day, a narrative format became very popular on Sina Weibo, which contains the pattern of “I am not in the mood”: “Guo Meimei’s identity remains unclear, I am not in the mood to sleep.” “Guo Meimei’s identity uninvestigated, I am not in the mood to write my paper”. This pattern appeared 20 times in the first phase of the emotional surge.

The opinion curve progressed differently. Both the informational and the emotional curves faded out gradually after the police investigation results came out on July 8th. Overall the opinion curve covered the whole span of the rumor discourse quite actively. Our preliminary observation suggested that the crowd tended to believe the rumor was true and question the credibility of the government. Even after the police released their investigation results which denied the rumor, the typical opinions on Weibo were

disbelief and criticism. For example: “The Guo Meimei Incident challenges transparency of the charity: the society has lost its credibility” (@G友 138159553), “The police is humiliating themselves! Yes Guo Meimei has no connection with the Red Cross Society but she has connections with the Red Cross officials!” (@冬去春来 9). Others challenged the government system: “We saw the fall of the whole society. Our philanthropy has become another source of corruption. Should we think more about the down side of our own system?” (@中国茶神)

Sensemaking statements, interestingly, peaked a little time later than informational, emotional and opinion-related statements. For example, the latter three reached their first peak on the same day of June 27, while sensemaking did that on June 28. Posts in this group increased quite slowly before its first peak. Similarly, while the other three had a second peak on June 30, sensemaking still peaked one day later. After that, it decreased, with several small peaks. One explanation for this latency could be that analyzing and interpreting the uncertain situation would depend on the understanding of information collected; so the sensemaking curve surged later than the informational curve. Compared to emotional expressions, which are more spontaneous, sensemaking utilizes one’s rational decision rules and could involve more efforts. Opinions, at the least, could be formed from one’s prior experience and thus surge before statements of sensemaking.

Lastly, directive and interrogatory statements showed novel trends. The directive messages emerged mainly with the first peak of the information curve and the first phase of the emotional surge. Posts in this category mainly asked for actions to put Guo Meimei into the *human flesh search* engine and to boycott the Red Cross. It is not surprising that the messages seeking information appeared immediately when the rumor was introduced. It is however notable that the interrogatory curve peaked a bit earlier than all other curves but disappeared after the information curve reached its peaks. This result could be explained by people’s needs to collect information before they try to understand the situation.

So far we have examined different statement types of the rumor discussions and their evolution along with the rumoring trend. Next we explore different user roles in the rumor transmission.

Who Led the Discussions?

To understand the different user roles in the rumor discussions around Guo Meimei, we classified the 38,221 users in our dataset into seven categories as shown in Table 1. We then analyzed the *contributions* and *influences* of each user group. Here contribution was measured by the number of participating users and messages *posted*. Influence was measured by the number of comments and reposts *received*. Table 3 summarizes our results.

Categories	Contribution		Influence	
	% users	% posts	% commented	% reposted
Celebrity	2.27	2.54	22.38	29.40
Certified	5.46	5.20	4.42	2.05
Mass Media	0.45	0.72	25.90	22.78
Organization	0.77	0.84	1.04	1.22
Website	0.71	0.88	1.20	1.55
Internet Star	2.74	4.21	18.23	28.74
Ordinary	87.62	85.61	26.85	14.25

Table 3. User categories and distribution in terms of contribution and influence.

As expected, the ordinary users were the most engaged. They occupied more than 85% share in both the number of participating users and messages posted during the rumor discussions. Their posts, however, were not often reposted or commented by others. In together, the ordinary users only received about 14% of the total reposts and 27% of the total comments. In contrast, three smaller groups invoked most discussions and occupied most reposts. They were, namely, celebrities, mass media and Internet Stars (iStar). The mass media, notably, while being the smallest group, received the second most comments. The Celebrities and iStars only posted a small number of messages related to the rumor (less than 7%); but their messages were the most broadcasted (received about 60% of all the reposts).

We then examined how users in different categories drove discussions along with the rumor's development. We analyzed the temporal trends of different user groups' influence in generating discussions and reposts with their original messages. As shown in Figure 5, in the early stage after the rumor broke out (June 21~25) and before the first peak (June 27), messages posted by *iStars* were the most widely disseminated (55%~83%). After that, around the first peak on June 27 when Guo Meimei made her online apologies (June 26~July 2), the *celebrity* took the flag and their posts were heavily forwarded (41%~65%). During the next two peaks around the Beijing police's investigation (July 4~July 10), the *mass media* made a strong presence, with their reports and investigations most circulated (39%~85%). Later, rumor discussions cooled down after more and more media reports came out trying to clarify the truth and people began to lose interests. It was not until July 24, the next day after the "7-23" high-speed train accident, that the Guo Meimei incident was recalled. In this stage, discussions were again led by celebrities and *iStars*.

Figure 6 shows different user groups' strength in attracting comments along with the rumor evolution. The overall trend revealed a similar pattern to the influence on reposts (Figure 5): *iStars* first took the leadership followed by *celebrities* and the *mass media* respectively. Notable differences were that the ordinary people were more

capable of generating discussions (receiving more comments) than producing wide-spread messages (receiving reposts), while *iStars* were the opposite.

In summary, our analysis of different user groups' influence over time reveals a 3-stage pattern of the changes of leading role in the rumor discussions on Weibo. In this pattern, the *iStars* first took the lead and generated the buzz. Then the celebrities broadcast the rumor to a larger audience. After that, the mass media take the leadership with their strong capabilities in providing substantial information.

DISCUSSION

In our case study, we categorized seven types of statements around rumor discussions and identified four major categories: providing information, expressing emotions, opinions, as well as analyzing and interpreting the situations. We also identified three leading user groups in the rumor discussions. While we do not generalize our findings to other microblogging communities (e.g., Twitter) or other extreme events, it's worth comparing our work with other case studies on the microblogging behavior in public crises (e.g., [19][22][27][29]) to understand the unique characteristics of rumor transmission.




Figure 5. The trend of user categories in receiving reposts.




Figure 6. The trend of user categories in receiving comments.

Rumor discussions are more dynamic in terms of its various statement patterns. First, the primary job after a natural disaster or public event is to take immediate responses for rescues. Messages on situational awareness and actions are the essential. Under a public rumor, the major goal is to reduce the uncertainty, where providing new evidence, expressing opinions and making interpretations could be the

driving force. Second, in all the previous studies on information spreading during extreme events, the posting trends largely follow the model proposed in [17][18]: the curve is shaped by a spike of breaking news and a power-law decay. Rumor discussions, however, are structured by various types of discussions that try to construct the truth or achieve a shared explanation. Next we discuss the contribution of our case study.

Conceptual Contribution to Rumor Theory

In this paper we take the social perspective that rumor transmission is a group explanation process for reducing collective anxiety in uncertain situations [7] [25]. Our study provided preliminary large-scale empirical evidence for the rumor theories that had been hindered by the difficulties in acquiring data under a natural observation and limited to small scale studies before the flourishing of SNS. As the next step, we plan to conduct deeper analysis on the interplay among the different types of statements. Conceptually we are interested in examining the relationship among uncertainty, anxiety and sensemaking activities. In our case, difficulties in finding credible evidences and the high public concern added to the complex of the rumors. Our case suggested that availability of reliable information sources, consistence in the information collected, variety of the materials as well as the scale of participants could be candidate indicators for measuring the uncertainty associated with rumors.

Furthermore, from our analysis we form a preliminary argument of the large-scale rumor transmission processes. We propose that at the early stage of rumor discourse, people post most messages for information seeking and emotion relieving, to address their intensive stress caused by the situation of great ambiguity. They form and share their opinions based on the situation and their past experience to give their initial explanations. Questions are asked for collective information gathering. When rumors continue to develop, with more details added to the situation, uncertainty increases. People put more effort in finding references, analyzing and making sense of the story. After that, emotional expressions decrease and interpretations increase. When people gradually build a collectively acceptable version of understanding, authenticating and emotional statements fade out.

While the contribution of our work is the rumor discourse analysis over a large scale of data, not much in-depth analysis has been included in this paper. Future work should be done to pin down and examine our above models proposed with more thorough analysis. It is also worth the effort to dig the roles and impact of different user groups while examining the model.

Microblogging during Massive Rumor Discourses

Our findings could inform the design and development of tools based on microblogs to respond to public rumors that may create collective stress in the society.

First, our results suggest that examining the combination of people's different statements may serve as an indicator for analyzing and predicting rumor development automatically, as well as measuring its impacts. For example, classifiers can be trained based on our analysis to classify the real-time large data on SNS that are related to rumor discussions. Our observation on different statement types has suggested many features that could be applied in machine learning technologies, e.g., statements in the "providing information" group often contain a URL link that points to other news websites, blog articles, or video sharing sites, and possess several patterns, e.g. "[entity] published an article". Emotional statements, as another example, often contain several exclamatory marks or mood-related keywords. Future work should be done to investigate the best features and measure their effectiveness for automatically identifying statement types accurately. Based on reliable large-scale statement classifications, portions and evolution trends of each statement types could be computed to build an index for understanding different rumor stages. Furthermore, the patterns of user behaviors for situation clarification and emotion relieving could be used to evaluate the levels of the importance and relevance of the topic, as well as the public anxiety induced by the uncertainty.

Second, technologies could be developed to analyze different user roles and identify key influencers that drive the direction of rumor discussions. In doing this, one should consider multiple dimensions for gauging different users' role in the microblogging community. Our analysis revealed different results on user roles between their contribution and influence. Interesting variances were also observed in different types of influences, e.g. capabilities for attracting comments versus reposts. While we did not include content analysis in measuring the user contribution and influence, future work could consider mining the text of users' discussions. For example, in our case, the mass media and celebrities have made the strongest presence in generating popular posts. However, it remains unclear *how* their posts are perceived and *how* their opinions influence the crowd. More in-depth research on the detailed interactions among different users should be done for identifying users that could help reduce uncertainty. With this understanding, future technology could be developed to detect opinion leaders, together with the analysis on the discourse patterns, to identify best strategies and channels for intervening and controlling the development of public rumors.

Social Media Transforming China's Civil Society

Lastly, our work also shed light on how social media is used in China under its social transitions. Recent years have witnessed the fast development of social media including microblogging in China since 2008, e.g. Kaixin001.com, Renren.com, weibo.com, etc. One of the reasons, we believe, lies in social media's great potential in driving public interests towards numerous controversial subjects

that have been emerging under China's social reforms. Bringing together the advantages of both blogging and social networking, Sina Weibo provides a great channel for the grassroots to voice their opinions towards public affairs, to collectively examine topics of shared concerns, and to form the power of the crowd in impacting government decisions of public interests. Our case study provides an example of how powerful an anonymous microblogger can be in the social transition that China has been going through. Rumor discussions on Guo Meimei and her rumored connection with the Red Cross urged the government-owned charity to be more transparent to the public. As a response to the furious crowd, the Red Cross has built a system to publicize how each donation is used and open itself to public supervisions. Future research could be conducted to obtain a deeper understanding of how social technologies change people's civic life in China.

CONCLUSION

In this paper, we have reported a case study on rumor transmissions in a Chinese microblogging community. Our content analysis on 1,000 messages identified seven different statement types and revealed their evolution dynamics as rumors developed over time. Another analysis of 38,221 unique users from 49,725 messages suggested different patterns of people taking the leading role in consecutive rumoring phases. Overall, our work provided preliminary evidence to social scientists' perspectives of viewing rumor transmission as a collective problem solving process. Our findings offered implications to unleashing the potential of using microblogging data to understand and control the public rumor development. Lastly, our study documented a case that added to the understanding of how social media is used in China for civil engagement.

ACKNOWLEDGMENTS

We would like to thank our colleagues John C. Thomas and Weijia Shen for their proofreading the paper and their valuable comments. Thank our anonymous reviewers for their insightful feedback, which helped us revise and improve the paper.

Lei Shi is supported in part by the National Basic Research Program of China (973 Program) under Grant No. 2010CB328105, National Science Foundation of China (NSFC) under Grant No. 60932003 and 60973144.

REFERENCES

1. Sina Weibo. <http://weibo.com>
2. Sina Weibo Wikipedia Entry. http://en.wikipedia.org/wiki/Sina_Weibo
3. Guo Meimei in Baidu Index. <http://index.baidu.com/main/word.php?word=%B9%F9%C3%C0%C3%C0>
4. Allport, G. W. and Postman, L. J. (1947). *The psychology of rumor*. New York: Holt, Rinehart and Winston.
5. Anthony, S. (1973). Anxiety and Rumor. *The Journal of Social Psychology* (89), pp. 91- 98.
6. Bordia, P. Studying verbal interaction on the Internet: The case of rumor transmission research, *Behavior Research Methods*, Vol. 28, No. 2 (1996), 149-151.
7. Bordia, P., and DiFonzo, N. Problem Solving in Social Interactions on the Internet: Rumor as Social Cognition. *Social Psychology Quarterly* 67(1), 2004: 33-49.
8. Bordia, P., DiFonzo, N., and Chang, A. Rumor as Group Problem Solving: Development Patterns in Informal Computer-Mediated Groups. *Small Research Group* 30(8), 1999, pp.8-28.
9. Bordia, P., and Rosnow, R.L. Rumor Rest Stops on the Information Highway: Transmission Patterns in a Computer-Mediated Rumor Chain. *Human Communication Research* 25(2), 1998, 163-179.
10. Buckner, H. T. A Theory of Rumor Transmission, *Public Opinion Quarterly*, Vol. 29, No. 1, Spring 1965, pp. 54-70.
11. Castillo, C., Mendoza, M., Poblete, B. Information Credibility on Twitter. In *Proc. WWW*, (2011).
12. DiFonzo, N., Bordia, P. and Rosnow, R. L. (1994). Reining in rumors. *Organizational Dynamics* 23, 47-62.
13. Hu, M., Liu, S., Wei, F., Wu, Y., Stasko, J., Ma, K., Breaking News on Twitter, In *Proc. CHI'12*.
14. Java, A., Song, X., Finin, T., and Tseng, B. Why we twitter: understanding microblogging usage and communities. In *Proc. WebKDD/SNA-KDD '07*, ACM (2007), 56-65.
15. Krull, D. S. and Craig A. A. (1997). The Process of Explanation. *Current Directions in Psychological Science* 6(1):1-5.
16. Kwak, H., Lee, C., Park, H., and Moon, S. What is twitter, a social network or a news media? In *Proc. WWW '10*, ACM (2010), 591-600.
17. Leskovec, J., Backstrom, L., and Kleinberg, J. Meme-tracking and the dynamics of the news cycle. In *Proc. KDD '09*, ACM (2009), 497-506.
18. Matsubara, Y., Sakurai, Y., Prakash, B. A., Li, L. and Faloutsos, C. Rise and Fall Patterns of Information Diffusion: Models and Implications. In *Proc. KDD '12*, ACM (2012), 6-14.
19. Mendoza, M., Poblete, B. and Castillo, C. Twitter Under Crisis: Can We Trust What We RT? KDD 2010 Social Media Analytics Workshop.
20. Morris M. R., Counts, S., Roseway, A., Hoff, A., Schwarz, J. Tweeting is Believing? Understanding Microblog Credibility Perceptions, In *Proc. CSCW 2012*.
21. Palen, L. and Vieweg, S. (2008). The Emergence of Online Widescale Interaction: Assistance, Alliance and Retreat. In *Proc. CSCW'08*.
22. Qu, Y., Huang, C., Zhang, P., Zhang, J. Microblogging after a Major Disaster in China: A Case Study of the 2010 Yushu Earthquake, In *Proc. CSCW 2011*.

23. Rosnow, R. L. (1980). Psychology of rumor reconsidered. *Psychological Bulletin*, 87, 578–591.
24. Rosnow, R. L. and Kimmel. A.J. (2000). Rumor. *Encyclopedia of Psychology*, Vol. 7, edited by Kazdin., A.E. New York: Oxford University Press. pp. 122-123.
25. Shibutani, T. *Improvised News: A Sociological Study of Rumor*. Indianapolis: Bobbs-Merrill, 1966.
26. Starbird, K., Palen, L. "Voluntweeters": Self-Organizing by Digital Volunteers in Times of Crisis. In *Proc. CHI 2011*.
27. Starbird, K., Palen, L., Hughes, A. L. and Vieweg, S., Chatter on The Red: What Hazards Threat Reveals about the Social Life of Microblogged Information. In *Proc. CSCW (2010)*.
28. Qazvinian, V., Rosengren, E., Radev, D. R., and Mei. Q. Rumor has it: identifying misinformation in microblogs. In *Proc. EMNLP '11*, (2011), 1589–1599.
29. Vieweg, S., Hughes, A. L., Starbird, K. and Palen, L., Microblogging during two natural hazards events: what twitter may contribute to situational awareness. In *Proc. CHI*, (2010), 1079-1088.
30. Wu, D. K. Aligning a parallel English-Chinese corpus statistically with lexical criteria. In *Proc. ACL '94*, 80-87.
31. Wu, S., Hofman, J. M., Mason, W. A., and Watts, D. J. Who says what to whom on twitter. In *Proc. WWW '11*, ACM (2011), 705–714.
32. Yang, J. and Leskovec, J. Patterns of temporal variation in online media. In *Proc. WSDM '11*, ACM (2011).
33. Bernardi, D., Cheong, P. H., Lundry, C. and Ruston, S. W. *Narrative Landscapes: Rumors, Islamist Extremism, and the Struggle for Strategic Influence*. Rutgers University Press, 2012.