2015 48th Hawaii International Conference on System Sciences

Combating Rumor Spread on Social Media: The Effectiveness of Refutation and Warning

Pinar Ozturk
Howe School
Stevens Institute of Technology
pozturk@stevens.edu

Huaye Li
Howe School
Stevens Institute of Technology
hli21@stevens.edu

Yasuaki Sakamoto
Howe School
Stevens Institute of Technology
ysakamot@stevens.edu

Abstract

Twitter and other social media are now a major method of information exchange and dissemination. Although they can support rapid communication and sharing of useful information, they can also facilitate the spread of rumors, which contain unverified information. The purpose of the work reported here was to examine several design ideas for reducing the spread of health-related rumors in a Twitter-like environment. The results have shown that exposing people to information that refutes rumors or warns that the statement has appeared on rumor websites could reduce the spread of rumors. These results suggest that social media technologies can be designed such that users can self correct and inactivate potentially inaccurate information in their environment.

1. Introduction

Developments in information technology and changes in the media landscape have fundamentally influenced the ways in which we communicate. Social media and other kinds of computer-mediated communication are now a major method of information exchange and dissemination.

In micro-blogs like Twitter, users tweet up to 140 characters. While users can generate original tweets by providing first-person information, observations and thoughts; they can also re-tweet someone else's tweets. In social networking sites like Facebook, users can also share any links, status updates, and photos posted on the site. Especially if the content was posted publicly, the sharing of it exposes the friends and the followers of the sharer to the content who can again share the content as well. This low marginal cost of online information exchange has allowed information to circulate at a faster pace and in greater amounts than ever before. As a result, any claim can find expression and can spread quickly on social media, providing a rich substrate for rumor propagation [1].

A rumor is commonly defined as a statement whose informational truthfulness is unverified [2]. Rumors may spread misinformation in the absence of verifiable information regarding uncertain circumstances [3] or deliberate false information (disinformation). The misinformation or disinformation carried by rumors can quickly spread within a network.

Rumors, as part of our social communication, also exist in our offline lives. However, online, the ability to participate pseudonymously, low levels of entry barrier, social presence, and lack of gatekeeping mechanisms existing in traditional mainstream media, create a setting of low accountability and uncertainty. In such environments, rumors thrive [4] and acquire an air of legitimacy through being shared by numerous people. So much so that they can easily shape public opinion and cause harm to others [5].

When false rumors find traction in sizable segments of the population, governments intervene through issuing announcements or setting up rumor controls (http://www.fema.gov/hurricane-sandy-rumor-control). However, rumors might already be widely spread when official rebuttals are released in a temporally disjointed format similar to retractions or corrections in traditional media. In addition, selective exposure to the official rebuttals might not prevent people from absorbing misinformation carried by rumors. Since misinformation and rumors rarely come with a warning label, people usually cannot recognize that the information is incorrect until they refer to outside sources or receive a refutation. Therefore, most of the time they are left to their own tools to assess the accuracy of the information and to either disseminate or counter the rumor spread.

In the absence of established official sources, people counter the rumor spread by sending posts/tweets questioning or criticizing the rumor directly. Such citizen participation was observed on social media during responses to disasters [5, 6]. Some information is not easily verifiable through traditional media and official sources. Social media can react more quickly through user-generated counter rumors.

1530-1605/15 $31.00 © 2015 IEEE
DOI 10.1109/HICSS.2015.288

However, it is unclear if counter rumors act as important signals for the veracity of the information, affecting people's perception of the information and their spreading behavior. Thus, the following question naturally arises: Can counter rumors on social media reduce the spread of rumors?

To answer this question, we conducted an experiment to examine the effectiveness of the design ideas in reducing the spread of health-related statements containing rumors. Unverified health-related information can do harm to people if it is found to be inaccurate. Thus the ability to reduce the spread of such information is important. Using known health related rumors identified as either incorrect or unverified (debatable) by health professionals and their counterparts rumors, selected from the "health myths" sections of Discovery, Food Networks and National Institute of Health websites, we measured participants' sharing intentions of these rumors and counter rumors, presented through different design architectures by which the likelihood of rumor sharing might be reduced.

The purpose of the work reported here is to evaluate several design ideas for reducing the spread of rumors on social media like Twitter. The designs, which are motivated by work in rumor psychology and observation in Twitter, involve displaying information that refutes or warns about a rumor along with the rumor.

The contributions of the current work are as follows. It extends past work on rumor transmission in face-to-face settings to social media environments. In addition, its results suggest a way to use user-generated content to reduce the spread of rumors on Twitter and similar social media. By better understanding how people spread information online and how to reduce the spread of rumors online, we wish to contribute to improving the quality of information on social media.

The remainder of this paper is organized as follows: We first introduce related past work in rumor research. We then describe the experimental design in detail, present the results, and conclude with a discussion of our findings and their implications for future research.

1.2 Literature Review

Although most online social networks has been in our lives for less than ten years, rumor research investigating rumor generation, transmission and refutation date back over seventy years in the field of psychology [3]. The two early studies [7, 8] that are frequently referenced across this literature have found that rumor generation and transmission are usually attempts at sense making and finding answers in unclear situations, especially when the level of anxiety is at peak. Since those studies, others have also investigated the psychological factors that affect rumor generation and transmission. They have found additional factors such as accuracy and importance that impact rumor behavior [9, 10].

Rumor rebuttal and refutation has also been the subject of considerable scholarly attention in the field of psychology. A number of studies have shown rumor criticisms and rebuttals decrease the belief in rumors [8, 9] and eliminate reliance on misinformation to a degree [11]. However, some studies in psychology also found that when these rebuttals and retractions are presented in a temporally and contextually disjointed format, they proved to be ineffective and didn't have the intended effect of eliminating reliance on misinformation carried by a rumor. The persistent reliance on this misinformation, even when later people are presented with a correction or retraction, has been label as the continued influence effect. Johnson and Seifert [12] suggested the concept of mental event models to explain the continued influence effect. They proposed that people build mental models of information and when no alternative information exists or a criticism/correction is presented after the model is already built, they may keep their model despite knowing it is false. According to the mental model approach, the initial integration of information when the model is being built is more readily performed than is its updating after a retraction [13].

Research has identified only few factors that can increase the effectiveness of corrections including (a) warnings at the time of the initial exposure to misinformation and (b) repetition of the correction and retraction [14]. Based on these factors, in this study we have investigated different design interfaces through which we presented rumor and counter rumors to participants and measured their sharing intentions of the rumor.

While studies in psychology have extensively studied the propagation of rumors and the effectiveness of counter rumors in offline world, very few of the past work have examined rumor and counter rumor dissemination in online communication. Prior research examining online communication has mostly concentrated on information diffusion [15, 16, 17] and the importance of social ties in disseminating information. However, the veracity of information, whether it'd be false or questionable, is generally not included in these studies.

In some recent work investigating rumors in online communication through social media, the main area has been identifying rumors and the effects of psychological factors in the spread of rumors. Oh et al. [18] investigated tweets posted during the Haiti earthquake of 2010 and found that anxiety and information ambiguity are the key variables that affect rumor transmission under extreme events. In Castillo et al. [19], Amazon Mechanical Turk was used to crowdsource credibility evaluation of trending topic tweets and crowd's credibility judgments were used to train a machine learning system that rates the credibility of tweets on a particular topic.

Ratkiewicz et al. [20] created a web service (Truthy) that visually represents the diffusing of misleading political memes on Twitter using tweet features, including hashtags, links, and mentions. Canini et al. [21] designed an automated method to identify and rank credibility of information sources on Twitter according to their relevance and expertise on a given topic. Gupta et al. [22] investigated tweets containing images that were posted during Hurricane Sandy and developed classification models to predict whether images being transmitted on Twitter were real or fake. The current work focuses on what to do after identifying rumors in order to minimize the spread of rumors.

Closely related to the work reported here, Friggeri et al. [1] examined the spread of rumors on Facebook and found that the spread of counter rumor didn't dominate the rumor cascade until rumors had achieved wide distribution. This is one of the reasons why rumors keep spreading after counter rumors appear on social media. Rumors are already spread so far that people are more likely to encounter rumors than counter rumors. Following this line of thoughts, Tanaka et al. [6] examined the effect of exposure to counter rumors on people's likelihood to spread rumors and found that if people were exposed to counter rumor before rumors rather than after, rumor spread was significantly reduced. The present work extends Tanaka et al.'s work by showing rumors and counter rumors at the same time and by measuring intent to retweet instead of number of people to share. In other words, we made the current work more realistic.

2. Experiment

2.1 Method

Platform. The experiments were conducted on Amazon Mechanical Turk (AMT), an online crowdsourcing platform where requesters post jobs/tasks and workers choose which jobs/tasks to perform for pay (https://www.mturk.com/mturk/).

Past researchers found that the AMT platform can assist researchers to collect high quality data and replicate classic psychological phenomena [23, 24, 25, 26, 27]. The design and procedure of current work followed their recommendations.

Participants. In return for a nominal fee, 259 workers of AMT completed the experiment.

Materials. Ten health-related statements were hand-selected from the "health myths" sections of the Discovery, Food Networks, and National Institute of Health websites. These health-related statements were selected with two constraints:

(i) Each statement was either unverified or false according to health professionals, and
(ii) The information carried by each statement was not too specialized.

Each statement and associated information was read and evaluated by the authors to ensure the two constraints were satisfied. Ten counter rumors stated that the rumor statements were not true. See Figure 1 for an example of a rumor and counter rumor.

Design and Procedure. All experiments were conducted through AMT. Figure 1 shows an example screen presented to subjects. Each screen contained one "ReTweet" button to simulate the real Twitter environment. All subjects were instructed to imagine reading a health-related statement on Twitter, and then to answer if they would retweet the statement. Ten statements were presented sequentially in a random order. Each participant was randomly assigned to one of only five conditions.

(i) **Only Rumor:** Ten rumor statements appeared by themselves. The left side of Figure 1 shows the task shown to the subjects in this condition.
(ii) **Only Counter:** Ten counter statements appeared by themselves. The right side of Figure 1 highlights the difference between this condition and the only rumor condition.
(iii) **Rumor and Counter:** Ten rumor statements appeared with corresponding counter statements. The right side of Figure 1 shows how the rumor and counter was presented together to the subjects.
(iv) **Rumor with Checkbox:** Ten rumor statements appeared by themselves as in the only rumor condition. However, subjects clicked a checkbox instead of a radio button used in the non-exposure condition. Whereas a checkbox allows users to unselect their selection, a radio button does not. We wanted to test if the results would change if the subjects could change their minds.
(v) **Rumor with Information:** Ten rumor statements appeared with a warning stating that the content had appeared on a rumor website. The right side of Figure 1 shows how the warning appeared with the statement.

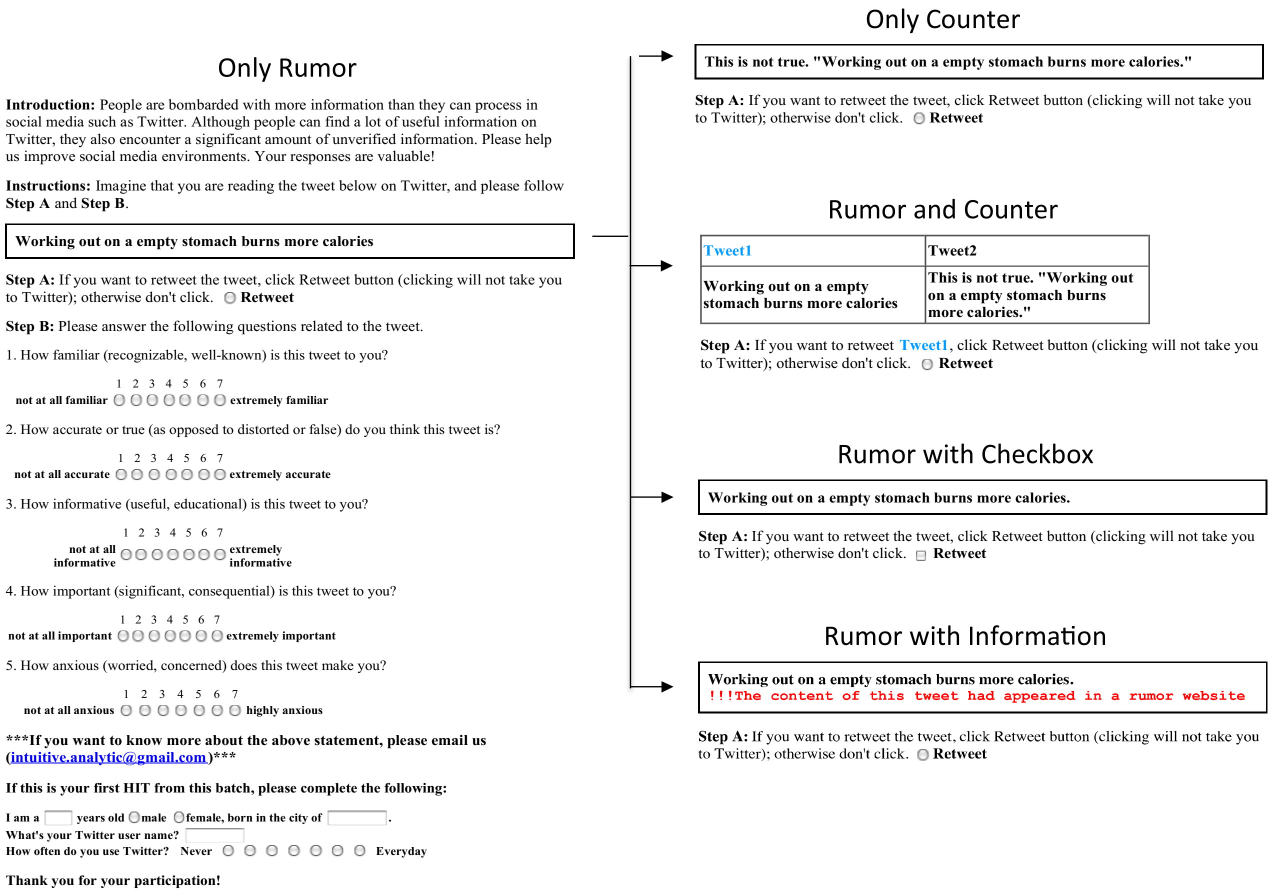

Figure 1. Experimental design

2.2 Results

The main interest of the current work was to test the effectiveness of counter statement and warning in reducing the spread of rumors. Thus, our analyses focus on the only rumor condition, only counter condition, rumor and counter condition, and rumor with information condition.

To assess whether the retweeting behavior differ among these four conditions, we conducted a one-way analysis of variance (ANOVA), which calculate F statistic and the associated probability, p, that the given differences will take place by chance. The ANOVA results revealed that the proportion of retweeting in the only rumor (0.275), only counter (0.205), rumor and counter (0.165), and rumor with information conditions (0.17) differed statistically, $F(1, 36) = 10.14$, $p = .003$. Figure 2 displays the probabilities of sharing in these four conditions. Error bars indicate 95% confidence intervals.

As shown in figure 2 and Table 1, subjects were statistically more likely to share rumor statements in the only rumor condition than in the rumor and counter condition as well as the rumor with information condition. Thus people tend to reduce their sharing of rumor statements in the presence of counter statements or a warning. In addition, subjects were more likely to share rumor statements than counter statements when rumor statements appeared alone.

These results suggest that pairing rumors with counter rumors or a warning message can reduce the spread of rumors. Developing social media technologies that can search counter tweets and display them together with corresponding rumors could help combat the spread of rumors.

Although we found statistical differences between the only rumor condition and rumor and counter condition, the difference was only 11%. To transfer this difference into reach in Twitter, we simulated the number of people who would receive the tweets based on our results. We report the results of this simulation next.

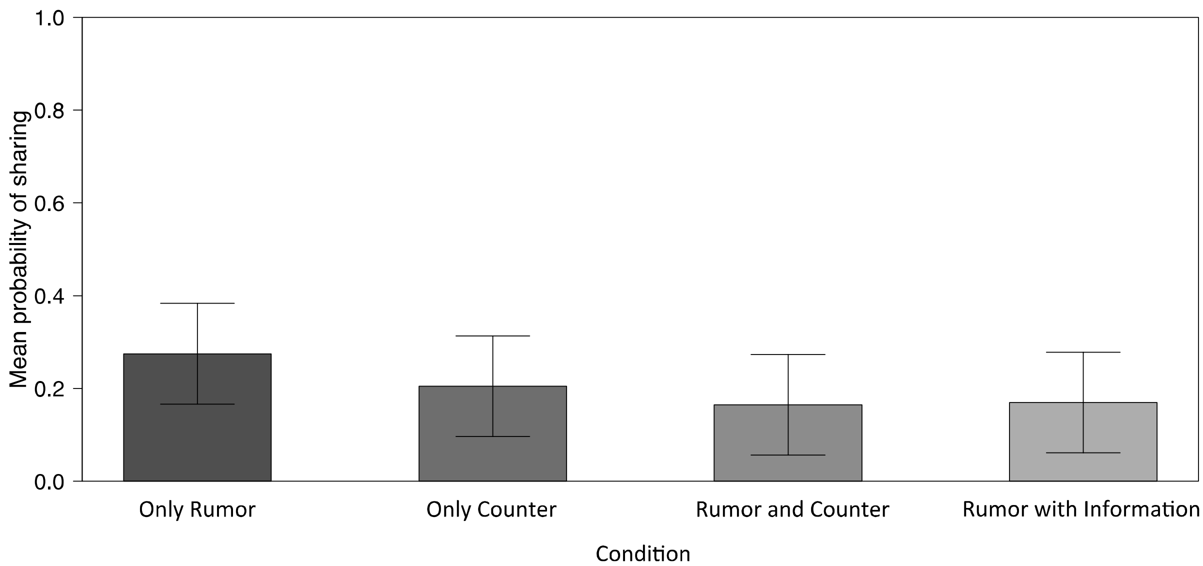

Table 1. The results of t-tests comparing pairs of conditions.

	Only counter	Rumor and counter	Rumor with information
Only rumor	t = -2.689, p = 0.0249	t = 3.161, p = 0.012	t = 3.042, p = 0.014
Only counter		t = 1.206, p = 0.259	t = 1.024, p = 0.333
Rumor and counter			t = -0.165, p = 0.872

3. Simulating the Spread

In order to simulate rumor and counter rumor spread, in each condition we have asked the subjects to share their twitter handle, if they have a Twitter account and willing to share. 132 workers stated having a Twitter account, of which 4 with no followers. Twitter follower numbers of the subjects ranged from 0 to over 100,000. Because the distribution was heavily skewed due to several outliers, we created bins to place the follower numbers instead of using the actual number. The horizontal axis of Figure 3 shows the bins. The use of bins remove the bias introduced by the outliers. For example, if a subject in one condition happens to have ten million followers, and the largest follower count in another condition was a million, the wide spread of tweets in the former condition could be simply due to one exceptional person. Subjects were randomly assigned to each condition and thus the number of Twitter followers was in general distributed evenly across the conditions. However, the follower count was not controlled in the experiment. Thus it was unavoidable to have outliers in one condition but not others.

Figure 3 shows the potential rumor spread by those workers who had indicated they would share the presented rumors. Regardless of the follower numbers, in each bin the potential rumor spread is much wider without counter rumor and warning. Figure 3 also shows that if these rumors were indeed to be circulating on Twitter, they would reach to almost 60,000 people through 132 initial rumor spreaders. The breakdown of the rumor reach in different conditions is as follows: 26,055 people in Only Counter; 19,890 people in Only Rumor; 760 people in Rumor and Counter; and 5,005 people in Rumor with Information.

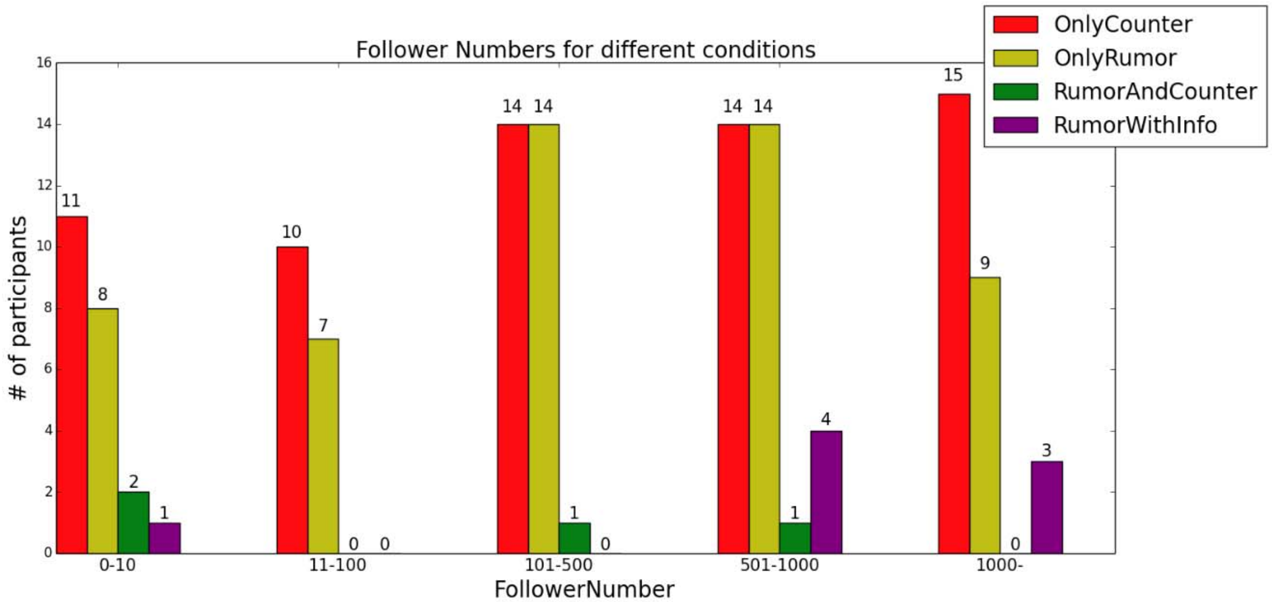

Figure 3. Potential rumor and counter rumor spread based on follower numbers self-reported by subjects

The simulation results are promising. First, people try to share counter rumor. They try to police themselves. Second, exposing people to counter rumors when they see rumors can greatly reduce the spread of rumors relative to seeing rumors only. Simply providing a warning message can also greatly reduce the spread of rumors but not as much as paring rumors with counters.

4. Discussion

With the new developments in information technology we rely more and more on online platforms such as Twitter for our information needs. As these platforms become increasingly significant as our news resources, in particular during emergency and crisis situations, it becomes critical to validate the veracity of the online information we receive [28]. However, the clues that we have in the real world to make those judgments do not exist in these online platforms. In fact, the intensification of information through being shared by hundreds of other people further undermines our ability to assess the accuracy of the information.

As the users of the platforms, we can only make judgments based on available information presented through the features of the immediate user interface. Hence, how and what information is presented to us with each post/tweet becomes important on how it will be perceived and spread subsequently. While most information received online cannot be verified and rebutted instantaneously by official sources or through traditional media, criticisms and counter information from other users will be available and can act as a signal for the accuracy of the information received.

Our study has shown that presenting counter information from other users or outside sources along with the original rumor did indeed decrease the likelihood of rumor sharing behavior. Consistent with previous findings in the psychology literature, giving the warning of “tweet may contain misinformation” at the initial exposure reduced the rumor spread but was not as effective as the repetition of the counter information. The results we obtained in this study suggest that matching counter rumors with rumors would help users to assess the accuracy of the information they receive and possibly decrease far-reaching rumor spread.

With respect to interface design, we highlight the issue that users are dependent on what is prominent in the user interface when making credibility and veracity judgments and deciding on whether to disseminate that piece of information. The current design of Twitter environment only allows the tweets posted by users followed to appear in the one’s timeline. But the tweets published by other users who are not followed will not appear in the timeline even if they contain the exact piece of text. One’s social network tends to exhibit homophily [29] and have likeminded individuals with shared viewpoint. This, coupled with the current design, reduces the chances of an individual ever seeing a counter rumor if members of his/her Twitter network decide not to share it. In this fashion, current design of the Twitter environment fosters selective exposure and limited distribution, resembling the cases of official corrections and rebuttals in the traditional media. In addition, a user can follow as many other users as he/she wants to; so there are many tweets displayed in one's timeline and a high possibility that one may miss some of them. Because there is no existing mechanism that matches rumors and counter rumors, users may be subjected to either of two at any time but it is very unlikely to receive both at the same time. A similar matching mechanism does exist in Facebook News feed, aggregating posts with similar content together even if they are posted by different friends of the user. Such mechanism allows users to see the comments of different people and a chance to see if any of their friends contradicted the shared content.

Currently, Twitter Alerts are available only for the use of select local, national and international institutions that can provide critical information relevant to an unfolding event, such as public safety warnings and evacuation instructions during emergency situations [30]. Tweets sent under this alert program contain a warning image (depending on the kind) and are pushed to the subscribers as a notification or text message. However, even the rumor tweets that are falsified through outside sources are not considered under this program, leaving them to get lost in the millions of others. Since the possibility of sending a warning image along with the tweets is clearly demonstrated in the alert program, we think rumor tweets should also be considered critical as they facilitate the dissemination of false information, potentially creating widespread harm and panic. Matching rumor and counter rumor tweets may not be technically possible right now. Our results show that presenting a warning may lower rumor spread, improve the quality of information circulating on online platforms, and may save lives in crisis situations where people rely on social media platforms make sense of the situation.

It is ultimately up to people to decide whether to believe and/or to spread the information they receive online, but we believe that they can use some help from the platforms to assess the veracity of the information. This is especially true for a micro blog platform such as Twitter where the information given is limited but the stream of tweets flows rapidly.

5. Limitations and Future Work

In this study, we only examined 10 rumor tweets and their associated counter tweets. Each tweet was selected from outside websites (Discovery, Food Networks and National Institute of Health) and was analyzed by the researchers. However, in the real Twitter environment this process would need to be automated.

In addition, we only were able to simulate the rumor spread using the twitter follower numbers of those participants who were willing to share their account handle with us. Still, simulated spread clearly demonstrated a decrease in the conditions when counter rumor or a warning was presented. However, it is possible that if we were able to include those who didn't share, the simulated spread would be different.

Our current experiment had included questions representing psychological factors that might affect the transmission and spread of rumor. In a future study, we will analyze those factors to investigate the effect of presenting counter rumors further and understand the logic behind the sharing behavior of individuals.

There are many factors that influence the spread of information on social media, including social network structure [31, 32, 33, 34, 35, 36, 37, 38], the content of the information [39, 40, 41], and how people perceive the information [2, 7, 8, 9, 10, 42, 43, 44]. Future work should examine how these factors influence the effectiveness of counter rumors and warnings in reducing the spread of rumors.

6. Conclusions

Many people turn to social media to seek information and improve their decisions and predictions. One challenge, however, is that unverified messages spread on social media that are later found to be false, which are often followed by other messages that question their accuracy. These messages, coupled with irrelevant messages, make it difficult to discover accurate and useful information from social media.

One way to reduce the amount of false information on social media might be to facilitate the self-correcting nature of social media. Although people notice questionable information on social media and tries to correct it by posting denial messages, doing so will not be effective unless people see the denial messages before they spread the questionable messages.

We have studied the effect of presenting counter rumor or a warning along with the rumor to help reduce the transmission of rumor on Twitter. We demonstrated that such counter methods did indeed decrease participants' likelihood of sharing rumor information. Our results suggest that a mechanism to match rumor and counter rumor on Twitter and display the two sides of information together will be helpful for reducing the spread of rumors on Twitter.

We hope that further research extending the current work could help improve the quality of information on social media.

7. Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. IIS-1138658 and Grant No. BCS-1244742.

8. References

[1] Friggeri, A., Adamic, L. A., Eckles, D., & Cheng, J. (2014). Rumor cascades in social networks. Presented at the *8th International AAAI Conference on Weblogs and Social Media* (*ICWSM*).

[2] DiFonzo, N., & Bordia, P. (2007). Rumor, gossip and urban legends. *Diogenes*, *54*(1), 19-35.

[3] Dunn, H. B., & Allen, C. A. (2005, March). Rumors, urban legends and Internet hoaxes. In *Proceedings of the Annual Meeting of the Association of Collegiate Marketing Educators* (p. 85).

[4] Marett, K., & Joshi, K. D. (2009). The decision to share information and rumors: Examining the role of motivation in an online discussion forum. *Communications of the Association for Information Systems*, *24*.

[5] Mendoza, M., Poblete, B., & Castillo, C. (2010, July). Twitter under crisis: Can we trust what we RT? In *Proceedings of the First Workshop on Social Media Analytics* (pp. 71-79). *ACM.*

[6] Tanaka, Y., Sakamoto, Y., & Matsuka, T. (2013, January). Toward a social-technological system that inactivates false rumors through the critical thinking of crowds. In *Proceedings of the 46th Hawaii International Conference on* System Sciences (pp. 649-658). *IEEE.*

[7] Prasad, J. (1935). The psychology of rumour: A study relating to the great Indian earthquake of 1934. *British Journal of Psychology. General Section*, *26*(1), 1-15.

[8] Allport, G. W., & Postman, L. J. (1947). *The Psychology of Rumor*. New York, NY: Holt, Rinehart, & Winston.

[9] Bordia, P., & DiFonzo, N. (2004). Problem solving in social interactions on the Internet: Rumor as social cognition. *Social Psychology Quarterly*, *67*(1), 33-49.

[10] Rosnow, R. L., & Foster, E. K. (2005). Rumor and gossip research. *Psychological Science Agenda*.

[11] Seifert, C. M. (2002). The continued influence of misinformation in memory: what makes a correction effective? *Psychology of learning and motivation*, *41*, 265-294.

[12] Johnson, H. M., & Seifert, C. M. (1994). Sources of the continued influence effect: When misinformation in memory affects later inferences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*(6), 1420.

[13] Ecker, U. K., Lewandowsky, S., Swire, B., & Chang, D. (2011). Correcting false information in memory: Manipulating the strength of misinformation encoding and its retraction. *Psychonomic Bulletin & Review*, *18*(3), 570-578.

[14] Lewandowsky, S., Ecker, U. K., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and its correction continued influence and successful debiasing. *Psychological Science in the Public Interest*, *13*(3), 106-131.

[15] Bakshy, E., Rosenn, I., Marlow, C., & Adamic, L. (2012, April). The role of social networks in information diffusion. In *Proceedings of the 21st international conference on World Wide Web* (pp. 519-528). *ACM.*

[16] Romero, D. M., Meeder, B., & Kleinberg, J. (2011, March). Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. In *Proceedings of the 20th International Conference on World Wide Web* (pp. 695-704). *ACM.*

[17] Goel, S., Watts, D. J., & Goldstein, D. G. (2012, June). The structure of online diffusion networks. In *Proceedings of the 13th ACM Conference on Electronic Commerce* (pp. 623-638). *ACM.*

[18] Oh, O., Kwon, K. H., & Rao, H. R. (2010). An exploration of social media in extreme events: Rumor theory and twitter during the Haiti earthquake 2010. *ICIS* 2010.

[19] Castillo, C., Mendoza, M., & Poblete, B. (2011, March). Information credibility on twitter. In *Proceedings of the 20th International Conference on World Wide Web* (pp. 675-684). *ACM*.

[20] Ratkiewicz, J., Conover, M., Meiss, M., Gonçalves, B., Patil, S., Flammini, A., & Menczer, F. (2011, March). Truthy: Mapping the spread of astroturf in microblog streams. In *Proceedings of the 20th International Conference Companion on World Wide Web* (pp. 249-252). *ACM.*

[21] Canini, K. R., Suh, B., & Pirolli, P. L. (2011, October). Finding credible information sources in social networks based on content and social structure. In *Privacy, Security, Risk and Trust (passat), 2011 IEEE Third International Conference on Social Computing (socialcom)* (pp. 1-8). *IEEE.*

[22] Gupta, A., Lamba, H., Kumaraguru, P., & Joshi, A. (2013, May). Faking Sandy: Characterizing and identifying fake images on Twitter during Hurricane Sandy. In *Proceedings of the 22nd International Conference on World Wide Web* (pp. 729-736). *ACM*.

[23] Paolacci, G., Chandler, J., and Ipeirotis, P. G. (2010). Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making*, *5*, 411-419

[24] Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, *6*, 3-5.

[25] Crump, M. J. C., McDonnell, J. V., and Gureckis, T. M. (2013). Evaluating Amazon's Mechanical Turk as a tool for experimental behavioral research. *PLoS ONE* 8(3): e57410. doi:10.1371/journal.pone.0057410

[26] Mason, W., and Watts, D. (2009). Financial incentives and the 'performance of crowds.' *HCOMP*, 77-85.

[27] Mason, W., & Suri, S. (2011). A Guide to Behavioral Experiments on Mechanical Turk. *Behavior Research Methods*, 44, 1-23.

[28] Zubiaga, A., & Ji, H. (2013). Tweet, but Verify: Epistemic Study of Information Verification on Twitter. *arXiv preprint arXiv:1312.5297*.

[29] McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual review of sociology*, 415-444.

[30] Retrieved on June 15, 2014 from: https://about.twitter.com/products/alerts/guidelines-faq

[31] Aral, S., Muchnik, L., and Sundararajan, A. (2009). Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. In *Proceedings of the National Academy of Sciences*, *106*, 21544-21549.

[32] Aral, S. and Walker, D. (2011). Creating social contagion thought viral product design: A randomized trial of peer influence in networks. *Management Science*, *57*, 1623-1639.

[33] Aral, S. and Walker, D. (2012). Identifying influential and susceptible members of social networks. *Science*, *337*, 337-341.

[34] Cha, M., Haddadi, H., Benevenuto, F., and Gummadi, K.P. (2010). Measuring user influence in Twitter: The million follower fallacy. In *Proceedings of the 4th International Conference on Weblogs and Social Media* (*ICWSM*). *AAAI*.

[35] Fang, X., Hu, P. J.-H., Li, Z., & Tsai, W. (2013). Predicting adoption probabilities in social networks. *Information Systems Research*, *24*, 128-145.

[36] Huberman, B., Romero, D., and Wu, F. (2009). Social networks that matter: Twitter under the microscope. *First Monday*, *14*, 1-5.

[37] Kwak, H., Lee, C., Park, H., and Moon, S. (2010). What is Twitter, a social network or a news media? In *Proceedings of the 19th International Conference on World Wide Web WWW*, *10*, 591.

[38] Leskovec, J., Adamic, L. A., & Huberman, B. A. (2007). The dynamics of viral marketing. *ACM Transactions on the Web*, *1*, 1-39.

[39] Asur, S. and Huberman, B. A. (2010). Predicting the future with social media, in *Proceedings of WIIAT*.

[40] Bandari, R., Asur, S., and Huberman, B. A. (2012) The pulse of news in social media: Forecasting popularity. In *Proceedings of the 6th International Conference on Weblogs and Social Media* (ICWSM). *AAAI*.

[41] Ha, S. and Ahn, J. (2011). Why are you sharing others' tweets?: The impact of argument quality and source credibility on information sharing behavior. In *Proceedings of the International Conference on Information Systems*.

[42] Chen, R., & Sakamoto, Y. (2013). Perspective matters: Sharing of crisis information in social media. *HICSS* 46.

[43] Chen, R., & Sakamoto, Y. (2014). Feelings and perspective matter: Sharing of crisis information in social media. *HICSS* 47.

[44] Li, H., Sakamoto, Y., Tanaka, Y., & Chen, R. (2014). The psychology behind people's decision to forward disaster-related tweets. In *Proceedings of the 18th Pacific Asia Conference on Information Systems*.