2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)

Finding Requests in Social Media for Disaster Relief

Tahora H. Nazer, Fred Morstatter, Harsh Dani, and Huan Liu
Arizona State University, Tempe, AZ
{tahora.nazer, fred.morstatter, harsh.dani, huan.liu}@asu.edu

***Abstract*—Natural disasters create an uncertain environment in which first responders face the challenge of locating affected people and dispatching aids and resources in a timely manner. In recent years, crowdsourcing systems have been developed to exploit the power of volunteers to facilitate humanitarian logistic efforts. Most of the current systems require volunteers to directly provide input to them and do not have the capability to benefit the large number of disaster-related posts that are published on social media. Hence, many social media posts in the aftermath of disasters remain hidden. Among these hidden posts are those that need immediate attention, such as requests for help. Hence, we have implemented a system that detects requests on Twitter using content and context of tweets.**

I. Introduction

Social media has become an important channel for crisis communication. During Hurricane Sandy in 2012, more than 20 million related tweets were published on Twitter. Social media is a fast-paced channel which is used by affected people to describe their situation and observations, seek information, specify their requests, and offer their voluntary assistance; providing actionable information [1], [2]. Officials have also used social media in disaster response. For instance, American Red Cross tweeted "You can text "HAITI" to 90999 to donate $10 to Red Cross relief efforts in #haiti" after the major earthquake in Haiti, in 2010[1].

To help actionable posts stand out among millions of others on Twitter, users take advantage of hashtags. In the aftermath of the Chennai Rains in 2015, #ChennaiRainsHelp was used by stranded people to request help and by volunteers to offer accommodations[2]. In the Paris Attacks in 2015, volunteers used #PorteOuverte to offer shelter and #StrandedInUS to help with French people whose flights have been canceled[3]. Actionable information is useful for first responders who face challenges such as the distribution of resources when their size and location are unknown and the environment is uncertain [3].

Although hashtags play an important role in finding the posts which contain actionable information, not all such posts are marked with proper hashtags. For instance, in the Boston Marathon Bombings in 2013, no consistent hashtag was offered by official organizations [4]. In such cases, valuable information could be neglected because of this sole dependence on hashtags. A trivial way to overcome this challenge is the manual inspection by the officials which is not feasible due to the fast pace of social media; millions of posts are published every hour. These issues have encouraged researchers to create systems that either benefits manual input from volunteers or are based on automatic methods to detect actionable information.

Several systems have been implemented to use crowdsourcing power in disaster relief. One example of such efforts is CrisisCamp which is managed by CrisisCommons. After the Haiti earthquake, CrisisCamp brought together volunteer programmers and developers who created digital maps for coordinating relief groups and applications for people to trade their supplies [4]. Another system is Ushahidi which has been launched in the 2008 Kenyan post-election crisis. In Ushahidi, volunteers can submit reports via various channels, SMS, MMS, or the online interface. These reports will be later categorized, geotagged, and put on interactive maps [5]. Manual annotation of these reports by volunteers was possible when the system started but not in later disasters. The number of reports grew from hundreds of daily reports in the Haiti Earthquake in 2010 to 300,000 tweets per minute in Japan Earthquake in 2011 and to 20 million tweets during Hurricane Sandy. Due to the immense volume, volunteers of Ushahidi were no longer able to annotate as they were overwhelmed by the data was produced by people on social media [6].

In the next generation of attempts, systems have been developed that use *automatic* methods to extract useful information from *social media* data for disaster relief. For example, in Artificial Intelligence for Disaster Response System, Imran et al. [7] extract crisis related tweets, ask the crowd to label a subset of them, and use this data to train a crisis-specific classifier. This classifier is later used to extract informative posts from Twitter. The system that we propose also benefits *social media* data and *automatically* finds requests for help on Twitter. Requests are a subset of informative tweets and hence, this problem is harder to tackle. Previous studies use the text of tweets as the main source. They extract n-grams and train a classifier for each specific disaster because the vocabulary in each disaster differs from others. N-grams were proven to be powerful features for extracting informative tweets but are expensive to achieve due to dependence on disaster-specific (manually) labeled datasets. Contextual features have been

[1]https://goo.gl/OdRDWa
[2]http://goo.gl/HoNaXB
[3]http://goo.gl/PgkcPn
[4]http://goo.gl/WlTww2

IEEE/ACM ASONAM 2016, August 18-21, 2016, San Francisco, CA, USA
978-1-5090-2846-7/16/$31.00

also used in previous studies but not for solving our specific problem and they have been exploited to slightly improve the performance of n-gram based methods. In this work, we suggest a set of contextual features that have the power to discriminate requests from other tweets with F-score 71%. We also show that the results will improve by adding n-grams, as expected. More details on our method and results can be found in Section II and Table III.

II. Learning Approach

To detect requests for help on Twitter, we calculate the probability of a tweet being a request using both its content and context. Content features are based on the text and include n-grams and topics (which are generated using LDA [8], see Section II-A). Contextual features are the metadata provided with each tweet such as entities (e.g., URLs, hashtags, and mentions), timestamp, retweet count, and author's information (e.g., username, location, number of friends and followers). In the rest of this section, we will provide the details of our process to create the model that has been used in our system.

A. Feature Extraction

To prepare the dataset for the learning step we have removed duplicate tweets, retweets, and tweets which are different in at most one word. Then, punctuation and stop words have been omitted from the text of tweets. After disasters, based on the possible outcomes, specific keywords will be used by affected people to report their problems and request for help. These keywords are important content features to distinguish requests from other types of tweets. We translate keywords to n-grams in our model by considering unigrams, bigrams, and trigrams. Another content-based feature is the topic which is the probability distribution over words. We have extracted 20 topics for each tweet using LDA.

The second group of features is contextual. *Source* is the software by which a tweet has been published such as "web" or "iPhone". The most popular sources are the Twitter website, Twitter web clients, and Twitter applications for iPhone and Android. *Location* of a tweet is a valuable information. Users can enable the geolocation service while tweeting to help other users and first responders verify that the author is actually located in the affected area [9]. Users can *mention* each other using "@" preceding a username. This feature is used to establish conversations on Twitter. Users who are publishing a request try to ask for resources or explain a situation which needs attention. Thus, getting involved in conversations is not aligned with their goal. Hence, the number of mentions in a post can distinguish it from discussions. In order to convey as much information as possible in the limited length of a tweet, requesters will use fewer *hashtags*. Twitter users exploit links to convey their message by forwarding users to external websites. Requests after a crisis will use *URLs* to introduce websites for humanitarian purposes such as collecting monetary donations (e.g., "Got a spare dollar? Give $1 to victims of Hurricane Sandy now! http://t.co/bvXfWxST #sandy #nyc"), collecting blood donations (e.g., "Sign up to donate blood to help victims of Hurricane #Sandy here: http://t.co/pqpn0oNX"), and encouraging charity activities. Table I shows the most discriminative contextual features of requests in comparison to other tweets.

TABLE I
Request vs. non-request tweets in terms of top three contextual features.

Feature	Request	Non-Requests
Tweets with at least one hashtag	38%	61%
Tweets with at least one URL	66%	44%
Average tweet length	118.2	94.9

B. Learning Method

To classify whether a tweet is a request, we have trained different classifiers to find the one which has the best results. Based on the results in Table II, the Decision Tree classifier outperforms other three methods based on F-Measure. Hence, we have used this method in building our system. Classifiers receive each tweet in the form of a vector of features (II-A) and calculate the probability of a tweet being a request.

TABLE II
Comparison of Classification Methods. We have tried four different supervised classifiers on our dataset and Decision Tree has the best performance.

Classifier	Precision	Recall	F-Measure
SVM	0.401	0.786	0.531
Decision Tree	0.904	**0.919**	**0.912**
Random Forest	0.904	0.832	0.867
AdaBoost	**0.913**	0.874	0.893

C. Data Set and Evaluation

The classifier has been trained on a dataset of 13,260 tweets on Hurricane Sandy with 3,261 requests and 9,999 normal tweets. The requests have been provided by Purohit et al. [10] and normal tweets are a random sample of our Hurricane Sandy dataset which was collected during the same period of time. We used Decision Tree classifier (which has the highest performance in our experiments) with a different subset of features. As reported in Table III, using all the features results in the best performance, as expected. However, the change is minor when changing n from 1 to 3 in n-grams, suggesting that unigrams have the most amount of information.

TABLE III
Decision Tree Classification Results in terms of Precision, Recall, and F-Measure.

Feature Set	Precision	Recall	F-Measure
Context	0.710	0.724	0.717
Context and unigrams	0.896	0.908	0.902
Context and uni/bigrams	**0.906**	0.916	0.911
Context and uni/bi/trigrams	0.904	**0.919**	**0.912**

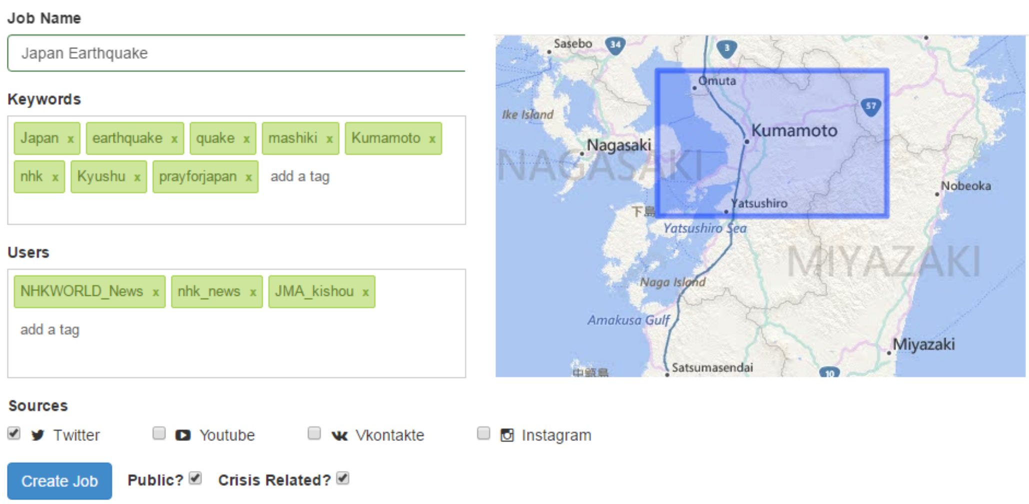

Fig. 1. Job creation in TweetTracker. When creating a job in TweetTracker, users can specify a set of keywords, users, and locations to filter the posts which are generated on any of the supported platforms, Twitter, YouTube, Vkontakte, and Instagram.

III. System Architecture

Our system has been implemented as part of TweetTracker which is a powerful tool for tracking, analyzing, and understanding activities on Twitter [11]. TweetTracker has been used by FEMA, The Red Cross, and others during crisis scenarios. In TweetTracker, tweets which are related to a job are crawled from the Twitter Streaming API as they are published. Each job in TweetTracker is defined to crawl the data related to an event from Twitter, Youtube, Vkontakte, and/or Instagram. Each job is a collection of keywords, users (authors), and locations. After creating a job, if a post is published on any of the selected social media sites and matches any of the job's criteria, it will be selected. A recent TweetTracker job, 2016 Japan earthquake in Kumamoto, is shown in Fig. 1. When defining a new job for the system, a user can flag the job as being related to a natural or man-made crisis such as an earthquake or a bombing. Tweets of the flagged jobs, after removing punctuations and stop words, will be labeled with the probability of being a request using the approach that was discussed in Section II-B. The probability of being a request is the output of our system (see Fig. 2).

A. Results

We periodically apply our model to new tweets for crisis related jobs and update the most probable requests. An example output of our system for a job on Hurricane Sandy is presented in Fig. 2. This figure shows a snapshot of TweetTracker Tweetalyzer in which there are several tabs to provide information of the selected job such as Tweets, Users, and Topics. Our system is shown in the rightmost tab, Requests. This tab is active when the selected job has been marked as crisis related. In Requests tab, text, author, and the probability of the most probable requests are provided to users.

Along with the probability of a tweet being a request, our system has the ability to show the geotagged most probable requests on the map. This feature is shown in Fig. 3.

IV. Conclusion

Social media is widely used in disaster response. Millions of tweets are published both during and in the aftermath of crises. Some of these tweets reflect observations, while many others contain offers and requests for help. On the other hand, a major issue that first responders face is locating victims and distributing resources in this uncertain environment. To help with these challenges, our system has the capability to distinguish requests and show their information and location. The underlying method of the proposed systems, along with content, benefit the context which has been used in previous studies but not regarding the problem of finding social media requests in the aftermath of disasters. In this paper, we have both proposed and implemented a method that can help first responders connect with people who are contributing and requesting aid, for the purpose of streamlining the process of providing aid to crisis-afflicted areas.

Acknowledgment

This work is sponsored, in part, by National Science Foundation grant 1461886 and Office of Naval Research grant N00014-16-1-2257.

References

[1] L. Palen and S. Vieweg, "The emergence of online widescale interaction in unexpected events: assistance, alliance & retreat," in *Proceedings of the Computer Supported Cooperative Work Conference*. ACM, 2008, pp. 117–126.

[2] S. Vieweg, A. L. Hughes, K. Starbird, and L. Palen, "Microblogging during two natural hazards events: what twitter may contribute to situational awareness," in *Proceedings of the Special Interest Group on Computer-Human Interaction Conference*. ACM, 2010, pp. 1079–1088.

[3] V. Zeimpekis, S. Ichoua, and I. Inis, "Humanitarian and relief logistics: Research issues, case studies and future trends," *Operations Research/Computer Science Interfaces Series*, vol. 54, pp. v–xi, 2013.

[4] J. Sutton, E. S. Spiro, S. Fitzhugh, B. Johnson, B. Gibson, and C. T. Butts, "Terse message amplification in the boston bombing response," in *Proceedings of the Conference for Information Systems for Crisis Response and Management*. Pennsylvania State University University Park, PA, 2014, pp. 612–621.

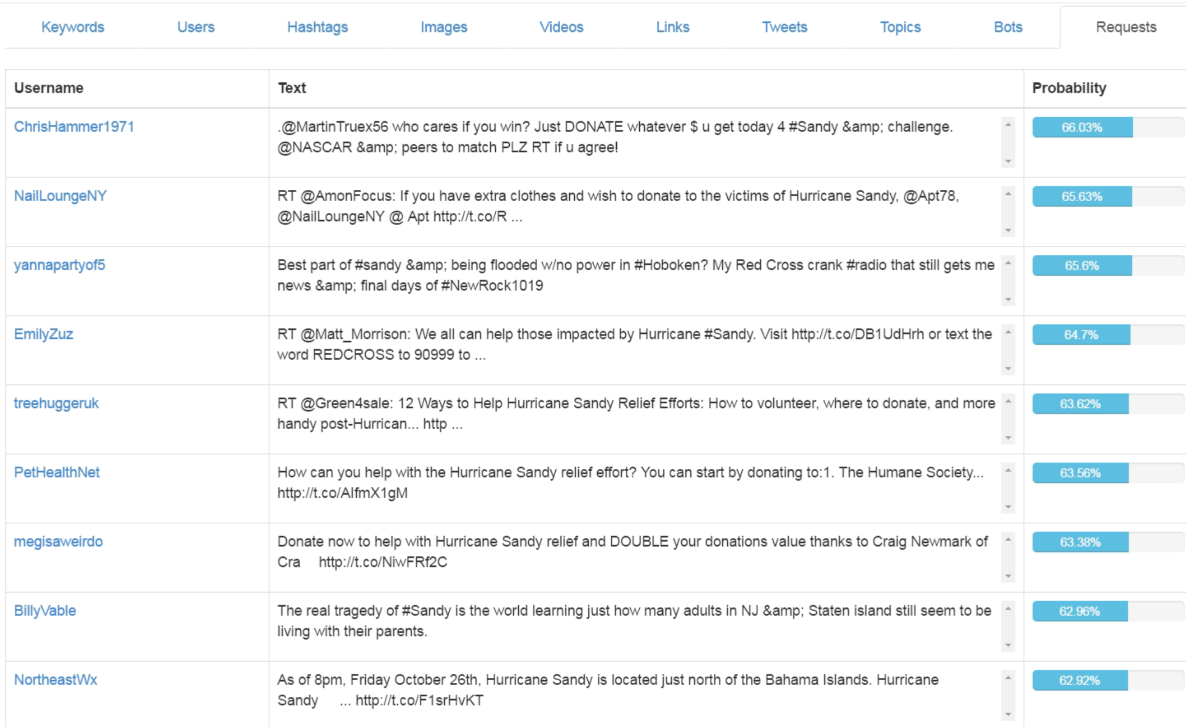
Keywords | Users | Hashtags | Images | Videos | Links | Tweets | Topics | Bots | Requests

Username	Text	Probability
ChrisHammer1971	.@MartinTruex56 who cares if you win? Just DONATE whatever $ u get today 4 #Sandy & challenge. @NASCAR & peers to match PLZ RT if u agree!	66.03%
NailLoungeNY	RT @AmonFocus: If you have extra clothes and wish to donate to the victims of Hurricane Sandy, @Apt78, @NailLoungeNY @ Apt http://t.co/R ...	65.63%
yannapartyof5	Best part of #sandy & being flooded w/no power in #Hoboken? My Red Cross crank #radio that still gets me news & final days of #NewRock1019	65.6%
EmilyZuz	RT @Matt_Morrison: We all can help those impacted by Hurricane #Sandy. Visit http://t.co/DB1UdHrh or text the word REDCROSS to 90999 to ...	64.7%
treehuggeruk	RT @Green4sale: 12 Ways to Help Hurricane Sandy Relief Efforts: How to volunteer, where to donate, and more handy post-Hurrican... http ...	63.62%
PetHealthNet	How can you help with the Hurricane Sandy relief effort? You can start by donating to:1. The Humane Society... http://t.co/AIfmX1gM	63.56%
megisaweirdo	Donate now to help with Hurricane Sandy relief and DOUBLE your donations value thanks to Craig Newmark of Cra http://t.co/NiwFRf2C	63.38%
BillyVable	The real tragedy of #Sandy is the world learning just how many adults in NJ & Staten island still seem to be living with their parents.	62.96%
NortheastWx	As of 8pm, Friday October 26th, Hurricane Sandy is located just north of the Bahama Islands. Hurricane Sandy ... http://t.co/F1srHvKT	62.92%

Fig. 2. Output of our system: most probable request for help tweets.

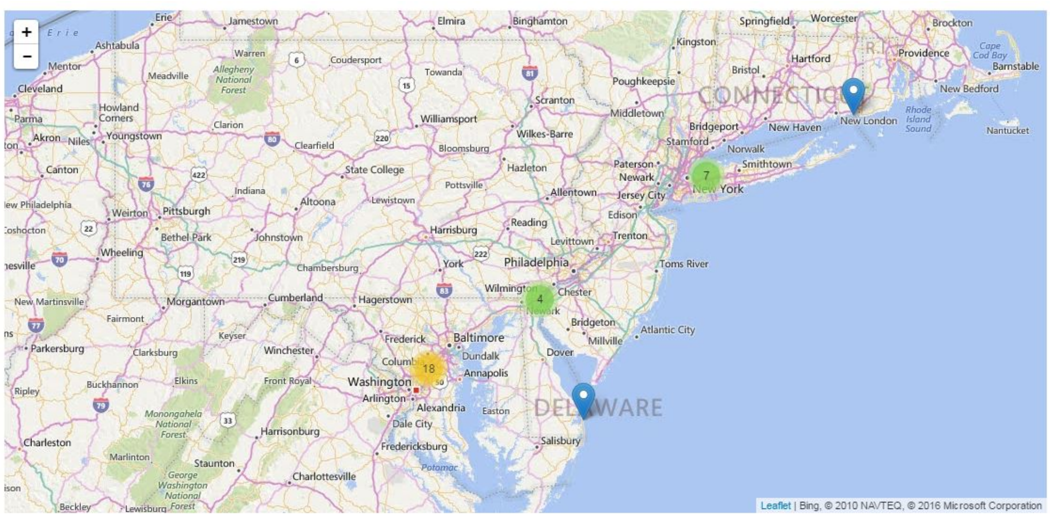

Fig. 3. Output of our system: location of the most probable geotagged requests on the map of the disaster-hit area.

[5] M. Zook, M. Graham, T. Shelton, and S. Gorman, "Volunteered geographic information and crowdsourcing disaster relief: a case study of the haitian earthquake," *World Medical & Health Policy*, vol. 2, no. 2, pp. 7–33, 2010.

[6] P. Meier, "Crisis maps: Harnessing the power of big data to deliver humanitarian assistance," 2013. [Online]. Available: http://goo.gl/1MF6lX, accessed 07.14.16.

[7] M. Imran, C. Castillo, J. Lucas, P. Meier, and S. Vieweg, "Aidr: Artificial intelligence for disaster response," in *Proceedings of the International World Wide Web Conference Committee*, 2014, pp. 159–162.

[8] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *the Journal of machine Learning research*, vol. 3, pp. 993–1022, 2003.

[9] F. Morstatter, N. Lubold, H. Pon-Barry, J. Pfeffer, and H. Liu, "Finding eyewitness tweets during crises," in *Proceedings of the ACL Workshop on Language Technology and Computational Social Science*, 2014, pp. 23–27.

[10] H. Purohit, C. Castillo, F. Diaz, A. Sheth, and P. Meier, "Emergency-relief coordination on social media: Automatically matching resource requests and offers," *First Monday Journal*, vol. 19, no. 1, 2013.

[11] S. Kumar, G. Barbier, M. A. Abbasi, and H. Liu, "Tweettracker: An analysis tool for humanitarian and disaster relief." in *Proceedings of the International Conference on Web and Social Media*, 2011, pp. 661–662.