

A Behavior Analytics Approach to Identifying Tweets from Crisis Regions

Shamanth Kumar, Xia Hu, Huan Liu
Computer Science & Engineering, School of CIDSE, Arizona State University
{shamanth.kumar, xia.hu, huan.liu}@asu.edu

ABSTRACT

The growing popularity of Twitter as an information medium has allowed unprecedented access to first-hand information during crises and mass emergency situations. Due to the sheer volume of information generated during a disaster, a key challenge is to filter tweets from the crisis region so their analysis can be prioritized. In this paper, we introduce the task of identifying whether a tweet is generated from crisis regions and formulate it as a decision problem. This problem is challenging due to the fact that only $\sim 1\%$ of all tweets have location information. Existing approaches tackle this problem by predicting the location of the user using historical tweets from users or their social network. As collecting historical information is not practical during emergency situations, we investigate whether it is possible to determine that a tweet originates from the crisis region through the information in the tweet and the publishing user's profile.

Categories and Subject Descriptors

H.2.8 [Database Management]: Database Applications—*Data mining*

Keywords

crisis tweets; user behavior in tweets; situational awareness

1. INTRODUCTION

Social media, particularly Twitter, has been a popular medium to discuss and report disaster related information. As an information dissemination platform Twitter has been used with varying success in several recent crises and mass emergency situations. For example, during Hurricane Sandy in 2012, people published videos and images of the damage caused by the storm. The continued usage of Twitter as a platform to submit crisis related information motivates us to identify relevant information during a crisis. As Twitter is a globally visible information medium the high volume of tweets generated during a crisis makes this a challenging

task. It is reasonable to assume that first-hand reports on a crisis such as reports of damage, originate in the crisis region. Therefore, we propose to study the problem of identifying information generated from crisis regions.

Twitter facilitates the geotagging of tweets, which is a process by which precise location information can be appended to a tweet. However, a recent investigation [14] has shown that only a small fraction ($\sim 1\%$) of all tweets contain location information. Thus, it is necessary for us to find alternate methods to identify a tweet's location. Recent approaches to tweet location identification leverage the geographic bias in the language of the tweet [2]. However, these techniques do not consider the topic bias in the information stream during a crisis. Thus, it is much harder under these circumstances to determine whether a tweet is generated inside the crisis region. These challenges motivate us to identify tweets from crisis region.

Identifying the user's location is another alternative, which can be accomplished using the social network information [19] and the user's historical tweets [3, 11]. Typically, identifying and extracting additional information such as a user's network, or his tweet history during a crisis is not practical due to the API constraints imposed by Twitter. These approaches also fail when there is insufficient network, or content history. Hence, we cannot apply these techniques to identify tweets from crisis region. Recent studies have also shown that a user's location may not necessarily correspond to the location of the tweet [4], due to user mobility. Thus, it is more reasonable to identify the location of a tweet.

In this paper, we conduct a study of crisis tweets to gain deeper insight into their characteristics and to specifically answer the following questions: 1) Do tweets inside crisis region express different behavioral patterns and can these patterns be used to identify tweets from crisis region when explicit location information is unavailable. Our contributions are the following:

- We formally define the novel problem of identifying tweets from a crisis region and highlight the challenges;
- We conduct a study of tweets from major crises to discover behavioral patterns in Section 2; and
- We propose an approach to identify tweets from crisis region in Section 3.

Problem statement: Given a crisis C associated with a crisis region R , and a collection of tweets relevant to the crisis T , where each tweet $t \in T$ contains tweet data t_d , including the tweet text and the user's profile information t_u , but does not include the location information. For each such tweet $t \in T$, we need to decide whether the tweet is

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

HT'14, September 1–4, 2014, Santiago, Chile.

Copyright 2014 ACM 978-1-4503-2954-5/14/09 ...\$15.00.

<http://dx.doi.org/10.1145/2631775.2631814>.

Table 1: Dataset Characteristics

| Dataset | # Tweets | # Retweets | # Geotagged Tweets | # Inside Tweets | # Outside Tweets |
|----------------------|------------|------------|--------------------|-----------------|------------------|
| EQ Japan | 2,734,431 | 1,223,609 | 105,669 | 44,119 | 28,953 |
| FL Mississippi River | 157,435 | 72,377 | 3,042 | 944 | 1,355 |
| HUR Sandy | 4,344,308 | 2,203,262 | 58,092 | 36,324 | 15,455 |
| SP OWS | 10,722,020 | 5,039,152 | 95,313 | 43,489 | 38,557 |
| WF AZ | 8,679 | 2,865 | 213 | 85 | 35 |

generated from inside the crisis region, i.e., $t \in R$ or the tweet is generated from outside the crisis region, i.e., $t \notin R$.

2. BEHAVIORAL PATTERNS IN TWEETS

We begin with a study of the characteristics of crisis tweets to identify distinct behavioral patterns.

2.1 Datasets

To conduct this study, we collected tweets pertaining to major crises in 2011 and 2012. All the data was collected using our tweet monitoring platform TweetTracker [10] through parameters specified by virtual volunteers from the NGO Humanity Road¹ and analysts from various governmental agencies monitoring the crises. Specifically, three types of parameters were used to collect the tweets: keywords/hashtags, geographic regions, and userIDs.

Preparing the dataset: To study the characteristics of the tweets, we must first identify tweets originating in the crisis region. The affected region for each crisis was decided based on the nature and scale of the crisis. For example, as Hurricane Sandy affected the entire East coast of the United States, we consider the region extending from Florida to New York as the crisis region R . Once a crisis region was determined, tweets from crisis region were identified through the geotagging information explicitly provided by the users. As the location information is voluntarily provided, we assume that it is accurate. Geotagged tweets which contained only a link or no content at all were removed. From Table 1, it can be observed that geotagged tweets typically comprised a small fraction of the data. The distribution of the tweets is presented in Columns 5 and 6 in Table 1.

2.2 Tweets From Crisis Regions

Previous studies have shown that the primary application of Twitter during a crisis involves information dissemination [9, 8]. In [18], the authors investigated the characteristics of tweets generated using mobile devices. In [22], the authors investigated the utility of linguistic features to detect tweets with situational awareness. Motivated by these studies, we propose to conduct a study of the characteristics of crisis tweets along the following dimensions: device usage behavior, attention seeking behavior, resource sharing behavior, and the originality of information.

Below we investigate the differences in these behaviors in crisis tweets. For each identified behavior, we will follow the following procedure to compare the behavior in the tweets. Comparison of the likelihood of a behavior is performed by computing the *Likelihood Ratio*. The *Likelihood Ratio* is indicative of how likely is the behavior to be exhibited by tweets inside crisis region when compared to tweets from outside the crisis region. Given a behavior b , its *Likelihood Ratio* can be computed as:

$$LR_b = \frac{P(b|inside)}{P(b|outside)}, \quad (1)$$

where $P(b|inside)$ is the likelihood of the behavior to be exhibited in tweets from inside crisis region and $P(b|outside)$ is the likelihood of the behavior to be exhibited in tweets from outside crisis region. If $LR_b > 1$, then the tweets inside the crisis region are more likely to express the behavior and the magnitude of the ratio indicates how likely this is. Further, to establish the significance of the differences in the observed behavior, we employ statistical tests.

Testing statistical significance: To establish that the difference in the behavioral patterns is statistically significant, we will employ the two tailed t-test. For all the comparisons, we set the significance level $\alpha = 0.05$. Let μ be the mean of the number of tweets exhibiting the behavior inside crisis region and μ_0 be the mean of the number of tweets exhibiting the behavior outside crisis region. The null hypothesis H_0 for the test can be defined as:

H_0 : the tweets inside the crisis region and tweets outside the crisis region demonstrate similar behavior, i.e., $\mu = \mu_0$.

If the p-value of the test is below the chosen significance level, then we can accept the alternate hypothesis i.e., $\mu \neq \mu_0$ and conclude that the difference in the behavior observed in the two types of tweets is statistically significant. Otherwise, we accept the *null hypothesis* and conclude that the behavior in the two types of tweets is similar.

2.2.1 Is the usage of mobile devices more prevalent inside crisis region?

Mobile devices are ubiquitous. The usage of capable mobile devices such as smartphones is increasing rapidly. In the United States alone, the usage of smartphones has exceeded 60% of all mobile subscribers². Therefore, we begin with an investigation of the usage of mobile device during a crisis.

Each tweet is associated with a client or an application which was used to publish it. From the client, it is possible to determine whether the tweet was published using a mobile device. As there are no explicit resources to distinguish between mobile and non-mobile clients, below we present a strategy to perform this task.

Procedure to identify mobile clients: In our datasets, we observed that the popularity of clients followed the power law distribution, where only a few clients were used to publish most tweets. As mobile clients need to be identified manually, we focus our effort on the most popular 100 clients from each dataset. Our investigations show that the top 100 clients account for more than 94% of the tweets in all the datasets. Among these clients, some clearly indicate their mobile nature. For example *Ubertwitter for Android*. For others, we visited the associated homepage where additional information was available. If a client indicated that it was an API, provided a desktop tool, or it was a bot, then it was

¹<http://www.humanityroad.org>

²<http://bit.ly/1bgXMLX>

Table 2: Behavioral characteristics in tweets: I

| Dataset | LR_{mobile} (p-value) | $LR_{hashtag}$ (p-value) | $LR_{resource}$ (p-value) | $LR_{retweet}$ (p-value) |
|----------------------|-------------------------|--------------------------|---------------------------|--------------------------|
| EQ Japan | 2.35 (0.0) | 0.06 (0.0) | 2.37 (0.0) | 0.11 (3.14E-85) |
| FL Mississippi River | 0.86 (6.046E-4) | 0.72 (4.53E-8) | 0.79 (1.44E-4) | 0.89 (0.55) |
| HUR Sandy | 1.07 (5.25E-56) | 1.34 (6.37E-102) | 1.50 (3.26E-222) | 0.40 (7.51E-23) |
| SP OWS | 0.96 (3.99E-6) | 0.86 (9.27E-131) | 1.03 (0.0) | 0.81 (3.856E-10) |
| WF AZ | 0.46 (0.01) | 2.59 (5.04E-6) | 1.51 (0.04) | 2.47 (0.30) |

classified as a non-mobile source, otherwise it was considered to be a mobile client.

Using the annotated clients, we investigated the tweets published using mobile clients. We set the behavior b in Equation 1 to *mobile* and compute LR_{mobile} . We found that the tweets from crisis region were more likely to be generated using mobile devices in two datasets. In the case of *EQ Japan*, tweets from crisis region were more than twice as likely to be published using a mobile device. Natural disasters are typically associated with the failure of utilities and increased mobility of users, such as during Hurricane Sandy³ and we observed this behavior in the two biggest events in our data: EQ Japan and HUR Sandy. A summary of LR_{mobile} is presented in column 2 in Table 2. The results were found to be statistically significant (p-values from the test are presented next to the ratio).

2.2.2 Are tweets from crisis region more likely to seek visibility?

Hashtags are typically used to indicate the topic of a tweet and for organizational purposes. However, studies have shown that the use of multiple hashtags indicates an effort to seek visibility [16]. The usage of multiple hashtags allows a tweet to be indexed under these hashtags and increases the chances of a tweet to be seen by a wider audience. To investigate this behavior in crisis tweets by setting the behavior b in Equation 1 to *hashtag*. The results show that tweets inside crisis region are less likely to include multiple hashtags in the tweet. We observed that this pattern was consistent across most datasets. Crises typically have specific popular hashtags associated with them. In addition, the usage of multiple hashtags reduces the amount of information that can be included in the tweet, which might explain this behavior. The $LR_{hashtag}$ for the datasets are presented in Column 3 of Table 2.

2.2.3 Are tweets from crisis region more likely to share an external resource?

Twitter messages are restricted to 140 characters. Thus, longer content and media such as images and videos can only be shared through external references. Therefore, we investigate whether tweets inside crisis region are more likely to share external resources. We measure $LR_{resource}$ by setting the behavior b in Equation 1 to *resource*. The $LR_{resource}$ for the datasets is summarized in Column 4 of Table 2. In this study, we only consider original tweets. Our study shows that the tweets from crisis region were more likely to contain URLs in most of the datasets. During a crisis, we expect the tweets from crisis region to contain links to resources such as videos and images describing the impact. For example, during Hurricane Sandy images of the flooding in the streets and the subway were widely shared by the local residents⁴.

³<http://ti.me/1neNTZB>

⁴<http://bit.ly/1sgSj9h>

Table 3: Behavioral characteristics in tweets: II

| Dataset | LR_{action} (p-value) | $LR_{entities}$ (p-value) |
|----------------------|-------------------------|---------------------------|
| EQ Japan | 0.34 (0.0) | 0.43 (0.0) |
| FL Mississippi River | 1.10 (3.08E-10) | 0.86 (9.63E-13) |
| HUR Sandy | 0.93 (1.28E-55) | 0.93 (9.38E-76) |
| SP OWS | 1.08 (4.50E-112) | 0.98 (7.07E-5) |
| WF AZ | 0.97 (0.83) | 1.06 (0.55) |

2.2.4 Are tweets from crisis region more likely to be a retweet?

There are several ways to publish tweets. When tweets posted by another user are forwarded by a user, the tweet is called a retweet. Retweets are characterized by the inclusion of the symbol “RT” at the beginning of the tweet. Typically retweets constitute a large part of crisis related tweets. Therefore, we study the originality of tweets from crisis region. We measure $LR_{retweet}$ by setting the behavior b in Equation 1 to *retweet*. The results in Column 5 of Table 2 show that the tweets from crisis region are more original and less likely to be a retweet. As first hand information is more readily available inside crisis region, it is reasonable to expect tweets from crisis region to contain original information and this is confirmed by the results.

2.2.5 Are tweets from the crisis region more likely to indicate an action?

Action words imply that the user publishing the tweet is performing an action. Verbs are typically used to indicate an action, such as leaving, moving etc. in the context of a crisis. To detect verbs in a tweet, we use the Ark NLP Part-Of-Speech (POS) tagger [15]. To study this behavior, we measure LR_{action} . The results are presented in Column 2 of Table 3. Our study shows that the tweets from crisis region are less likely to contain action words in a majority of the datasets. The results were also found to be statistically significant in most of the datasets studied.

2.2.6 Are tweets from crisis regions more likely to reference entities?

Entities typically refer to the names of people, buildings, or specific locations. During a crisis, such entities may refer to people and landmarks in the crisis region etc. Proper nouns are typically associated with names of people and places. Therefore, we employ the Ark POS tagger to identify the proper nouns in tweets. By analyzing proper nouns contained in the tweets we can determine whether the tweets reference entities. We set the behavior b in Equation 1 to *entity* and describe the results in Column 3 of Table 3. We expect that the first hand reports from crisis regions would be more likely to reference local entities who may be affected/involved in the crisis, thus providing situational awareness to first responders and other responding agencies. However, LR_{entity} for the datasets shows that tweets from crisis regions are less likely to reference entities. This might be due to the quick propagation of the information regard-




Figure 1: Avg. probability of tweets in various crises data using local entities outside the region after a crisis leading to frequent references in tweets from outside the region.

2.2.7 Are tweets from the crisis region novel?

Novel content indicates original information and not merely a tendency to publish popular content. To answer this question, we construct a unigram language model assuming the tweets are constructed using the bag-of-words strategy. The likelihood of each word $P(w)$ is estimated using the Maximum-Likelihood approach as

$$P(w) = \frac{c(w \in W)}{\sum_w c(w)}, \quad (2)$$

where $c(w)$ is the count of the word in the corpus. The probability of a tweet ($P(t)$) is computed as

$$P(t) = \prod_{w \in W} P(w). \quad (3)$$

In Figure 1, we present a comparison of the average probability of a tweet in each dataset. The experiment shows that tweets inside crisis region are generally more novel.

From the above study we have three key insights on tweets from crisis regions: 1) they are associated with original content and are more likely to discuss novel topics, which reaffirms previous findings on the information dissemination behavior of tweets during crisis, 2) they are less likely to seek attention, and 3) they are more likely to use external resources to convey their message.

3. USING BEHAVIORAL PATTERNS

Using the insights from the above study, we extract the following different types of features to model a tweet:

Mobile Features: Using the annotated list from Section 2.2.1 we identify if a tweet is published using a mobile client.

Resource Features: We constructed features indicating the presence of a URL and the number of URLs. In addition, we consider if the URLs point to an image or a video. References to Foursquare location are considered separately as an indication of location information.

Textual Features: Patterns in tweets, such as whether a tweet is a retweet, a directed message, or contains a user mention as well as the usage of hashtags. Positive and negative emoticons are distinguished using a list of popular happy and sad emoticons⁵. The quality of tweets is measured through punctuations, character, and word length.

Linguistic Features: References to actions and entities in tweets is measured using both the presence and the frequency of part-of-speech tags including proper nouns, verbs, and pronouns. Additionally, given a tweet t with vocabulary w , we compute and use the probability of the most probable word ($\max_w P(w)$), the least probable word ($\min_w P(w)$),

⁵http://en.wikipedia.org/wiki/List_of_emoticons




Figure 2: A comparison of the F_1 score




Figure 3: A comparison of the weighted AUC score

and the probability of the tweet ($P(t)$) computed using the language model previously discussed.

User Features: From the user’s profile, the user’s familiarity with the medium is extracted through the number of status messages and the user’s network activity.

Baseline Approaches: As the we only consider the tweets collected during crises and the profile of the publishing user as the available information, we cannot directly compare with other location prediction approaches, such as [11] and [2] as they require historical data. Therefore, we propose to use the following baselines:

- **Majority** - Since there is class imbalance, we consider predicting the majority class as a baseline.
- **Verma et al. (LF)** identified various linguistic features such as the unigrams and their raw frequency and the POS tags and demonstrated their effectiveness in identifying tweets containing situational awareness. Thus, we apply this model to our problem as a baseline. It should be noted that this strategy leads to a very large number of features which was generally proportional to the number of instances.

All experiments were performed using implementations in Weka [6] and the results were generated using 5-fold cross validation with 80% training data and 20% test data.

3.1 Evaluating the Performance

We evaluated our approach using the naive Bayes and *random forest* classifier. Comparisons with the baseline approaches are made using the F_1 score computed as

$$F_1 = \frac{2 * (precision * recall)}{precision + recall}. \quad (4)$$

Figure 2 summarizes the performance of the naive Bayes classifier applied to the baseline (NB_LF) and the proposed model (NB_TP). We find that our approach performs better than the baseline with significantly greater efficiency due to the compact representation. In addition, the random forest classifier (RF_TP) using the proposed features performed considerably better than the naive Bayes classifier.

To account for class imbalance, we also present the weighted AUC score in Figure 3. We omit the “majority” baseline

Table 4: Arizona wildfires summary created using Z_i

| Summary Tweets |
|---|
| it's really smokey and hazy today. #wallowfire |
| smoke near eagar #wallowfire http://twitpic.com/5ci9i7 |
| wildfire info: wallow fire pm update 6/19/11 (wallow wildfire) #azfire #wallowfire |
| #wallow fire swept thru greer. |
| glenwood gazette - breaking news: #wallowfire 06/10/11 map http://t.co/xr8e23b |

here as it was significantly worse than the other approaches. The proposed approach outperformed the baseline in almost all datasets. We verified the statistical significance of the results using the Wilcoxon-Signed Rank test. We compare the baseline (NB_LF) and our approach (NB_TF), treating them as a paired sample. The observed p-value was 0.0019 and the improvement was statistically significant at $\alpha = 0.05$.

3.2 Importance of Features

To evaluate the importance of different feature classes we constructed a *logistic regression* classifier. This classifier learns a weight for each feature, which can be interpreted as a measure of the feature's importance for the prediction task. The overall rank of a feature across different datasets was determined by the sum of its rank on each dataset. Lower rank values indicate that a feature is important across the datasets. Our investigation revealed that linguistic features were the most important class of features. This can be attributed to the novelty of tweets from crisis region. Textual features such as the use of punctuations were found to be more useful than reference to entities and action words. User related features were the least important class of features for the task, thus suggesting that prediction can be reasonably performed with just the information contained in tweets.

4. CASE STUDY: ARIZONA WILDFIRES

Tweets from crisis region can be used to obtain situational awareness and can be used to generate post-crisis summarization of the event from the perspective of tweets. In the task of event summarization, the goal is to identify a small number of representative tweets from the entire corpus, which can describe the event or a crisis. In this case study, we will use the task of event summarization to demonstrate that the application of our approach enables the generation of a more meaningful summary of the crisis.

Extracting representative tweets from topics derived from the tweets is a commonly used approach to event summarization [5]. To illustrate the differences between the two sets of tweets, we will summarize the Arizona Wildfires (*WF AZ*). First, the proposed model is used to classify all tweets whose location information is unknown. Then, the following procedure is applied:

- Extract 10 topics Z_i from tweets inside and Z_o from tweets outside crisis region.
- For each detected topic, rank tweets t with vocabulary w by its perplexity score defined as $perplexity(t) = \exp\left(\frac{-\log P(t|z)}{|w|}\right)$.
- Create a summary of the crisis by picking the 5 most relevant tweets from the top 10 tweets in the topic.

The extracted summaries in Tables 4 and 5 show that the summary created using Z_i has more relevant information and it highlights the relevance of our approach.

Table 5: Arizona wildfires summary created using Z_o

| Summary Tweets |
|---|
| wildfires wreaking havoc in arizona. http://bit.ly/jsgwpv |
| #arizona - y su bonito glowing bird suena en radio paranoia :) |
| rt @radionoisefm cel mai devastator incendiul din a.. #12 #ore #arizona #devastator #dublat #incendiul |
| hello #arizona, #bringit :] http://instagr.am/p/f4ife/ |
| 1600 quadratkilometer wald durch brand vernichtet #arizona |

5. RELATED WORK

Social media services have been extensively studied as social sensors to monitor important events occurring in the real world. In particular, recent research has focused on the analysis of the use of social media during emergencies [20], including earthquakes [12], riots [17], wildfires [21], etc. Seeking high-quality social media data pertaining to crisis serves as the basis of these studies and motivates this study to identify tweets from crisis regions.

Identifying a user's home location using social media data[7] is an interesting and important problem. The existing research on this topic can be divided into two groups. The first set of research methods assume that a user's tweets might contain distinct features due to their proximity to the region. Cheng et al. [2] estimated a Twitter user's home city based on the content of their tweets. Mahmud et al. [11] used an ensemble of statistical and heuristic classifiers to infer the home location of Twitter users at different granularities by using the content information and their tweeting behavior. However, topic specific variation of content has not been investigated. These approaches also rely on the availability of a user's tweet history, which is not readily available during a crisis. The second set of research methods assume that a user's home location is strongly correlated with his friends' home location. Backstrom et al. [1] estimated the home location of Facebook users using user-supplied address data and the network of associations between members. But, due to the API limitations it is not practical to extract network information during a crisis under time constraint. Therefore, these approaches cannot be directly applied to our data.

The problem of recognizing eyewitness tweets was independently investigated in [13]. While the authors evaluated whether linguistic features could be used to identify such tweets, here we analyzed several kinds of behavioral patterns in tweets from crisis regions.

6. CONCLUSIONS AND FUTURE WORK

Identifying tweets from crisis regions is becoming increasingly important due to information overload on Twitter. In this paper, we used tweets from several crises to conduct a study of the tweet characteristics and behavior and used the observations to build a novel method to detect tweets from crisis regions. Through experiments, we demonstrated that our approach is successful in identifying such tweets. As part of our future work, we will investigate the impact of the size of training data on the performance of the model as the process of filtering tweets from crisis regions should be initiated as soon as possible. We also plan to study the temporal effects on patterns in crisis tweets.

Acknowledgments

This work was supported in part by the Office of Naval Research grants N000141110527 and N000141410095.

7. REFERENCES

- [1] L. Backstrom, E. Sun, and C. Marlow. Find Me If You Can: Improving Geographical Prediction with Social and Spatial Proximity. In *Proceedings of the 19th international conference on World wide web*, pages 61–70, 2010.
- [2] Z. Cheng, J. Caverlee, and K. Lee. You Are Where You Tweet: A Content-Based Approach to Geo-Locating Twitter Users. In *CIKM*, pages 759–768, 2010.
- [3] Z. Cheng, J. Caverlee, and K. Lee. A Content-driven Framework for Geolocating Microblog Users. *ACM Trans. Intell. Syst. Technol.*, 4(1):2:1–2:27, 2013.
- [4] E. Cho, S. A. Myers, and J. Leskovec. Friendship and mobility: User Movement in Location-Based Social Networks. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1082–1090, 2011.
- [5] F. C. T. Chua and S. Asur. Automatic Summarization of Events From Social Media. *Technical Report*, 2013.
- [6] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The WEKA Data Mining Software: An Update. *ACM SIGKDD Explorations Newsletter*, 11(1):10–18, 2009.
- [7] B. Hecht, L. Hong, B. Suh, and E. H. Chi. Tweets From Justin Bieber’s Heart: The Dynamics of the Location Field in User Profiles. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 237–246, 2011.
- [8] T. Heverin and L. Zach. *Microblogging for Crisis Communication: Examination of Twitter Use in Response to a 2009 Violent Crisis in the Seattle-Tacoma, Washington, Area*. ISCRAM, 2010.
- [9] A. L. Hughes and L. Palen. Twitter Adoption and Use in Mass Convergence and Emergency Events. *International Journal of Emergency Management*, 6(3):248–260, 2009.
- [10] S. Kumar, G. Barbier, M. A. Abbasi, and H. Liu. TweetTracker: An Analysis Tool for Humanitarian and Disaster Relief. In *ICWSM*, 2011.
- [11] J. Mahmud, J. Nichols, and C. Drews. Where Is This Tweet From? Inferring Home Locations of Twitter Users. In *ICWSM*, 2012.
- [12] M. Mendoza, B. Poblete, and C. Castillo. Twitter Under Crisis: Can We Trust What We RT? In *Proceedings of the first workshop on social media analytics*, pages 71–79, 2010.
- [13] F. Morstatter, N. Lubold, H. Pon-Barry, J. Pfeffer, and H. Liu. Finding Eyewitness Tweets During Crises. In *Workshop on Language Technology and Computational Social Science*, 2014.
- [14] F. Morstatter, J. Pfeffer, H. Liu, and K. M. Carley. Is the Sample Good Enough? Comparing Data from Twitter’s Streaming API with Twitter’s Firehose. *ICWSM*, 2013.
- [15] O. Owoputi, B. O’Connor, C. Dyer, K. Gimpel, N. Schneider, and N. A. Smith. Improved Part-Of-Speech Tagging for Online Conversational Text with Word Clusters. In *Proceedings of NAACL-HLT*, pages 380–390, 2013.
- [16] R. Page. The Linguistics of Self-Branding and Micro-Celebrity in Twitter: The Role of Hashtags. *Discourse & Communication*, 6(2):181–201, 2012.
- [17] P. Panagiotopoulos, A. Z. Bigdeli, and S. Sams. “5 Days in August”—How London Local Authorities Used Twitter during the 2011 Riots. In *Electronic Government*, pages 102–113. 2012.
- [18] M. Perreault and D. Ruths. The Effect of Mobile Platforms on Twitter Content Generation. In *ICWSM*, 2011.
- [19] D. Rout, K. Bontcheva, D. Preotiuc-Pietro, and T. Cohn. Where’s @Wally?: a Classification Approach to Geolocating Users Based on their Social Ties. In *HT*, pages 11–20, 2013.
- [20] T. Sakaki, M. Okazaki, and Y. Matsuo. Earthquake Shakes Twitter Users: Real-Time Event Detection by Social Sensors. In *Proceedings of the 19th international conference on World wide web*, pages 851–860, 2010.
- [21] S. Sinnappan, C. Farrell, and E. Stewart. Priceless Tweets! A Study on Twitter Messages Posted During Crisis: Black Saturday. 2010.
- [22] S. Verma, S. Vieweg, W. J. Corvey, L. Palen, J. H. Martin, M. Palmer, A. Schram, and K. M. Anderson. Natural Language Processing to the Rescue? Extracting “Situational Awareness” Tweets During Mass Emergency. In *ICWSM*, 2011.