#pray4victims: Consistencies In Response To Disaster on Twitter

CODY L. BUNTAIN, University of Maryland, USA
JUNG KYU (RHYS) LIM, University of Maryland, USA

This paper studies commonalities in response across disasters in online social networks (OSNs) and Twitter specifically. After presenting an algorithm for extracting vocabularies across disasters, we extract type-specific vocabularies for terrorist attacks, earthquakes, and climate-related disasters between 2012 and 2017. Within similar disasters, commonalities emerge: terrorism responses reference the "attack" and law enforcement, earthquake responses mention the quake and its magnitude, and climate-related responses include safety and requests for aid. Across disaster types, tweets regularly mention victims/affected and prayer, consistent with communal coping and social support in crisis aftermath. Using these disaster-type vocabularies, we study Twitter as an alternate measure for severity, correlating casualties to Twitter volume. These vocabularies better correlate with casualties than baseline crisis lexica, especially in western countries. Twitter response and casualties diverge at the extreme, and Twitter response is stronger in Western countries, suggesting perceived severity is driven by additional factors.

Additional Key Words and Phrases: social media, twitter, crisis, disaster response, terrorism, earthquakes, climate change

ACM Reference Format:
Cody L. Buntain and Jung Kyu (Rhys) Lim. 2018. #pray4victims: Consistencies In Response To Disaster on Twitter. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW, Article 25 (November 2018), 18 pages. https://doi.org/10.1145/3274294

1 INTRODUCTION

An increasing majority of Americans are consuming news from OSNs [45], and recent reports show the public expects emergency responders to use OSNs in their decision-making processes [21]. These trends are boosting the role OSNs play in crisis response infrastructure [11], where access to timely information, social support, and family/friends is critical [15, 51]. Numerous efforts have studied ways to capitalize on these information sources during crises [23, 40, 49], but skepticism surrounds OSNs, with open questions about information quality [9, 46], emotional-but-uninformative content [3, 36, 49], and how information seekers and crisis responders can sift through this data [15].

This paper tackles these questions through an observational, quantitative study of Twitter and the textual similarities in user responses to terror attacks, earthquakes, and climate-related disasters. Such commonalities are touched upon in other research and recent media references to "disaster Mad Libs", or the "scripts" in which disaster details are filled in around otherwise standard responses about resilience, thoughts, and payers [52]. Despite the prevalence of such responses, this content is often removed during analysis to identify more useful information, but questions

Authors' addresses: Cody L. Buntain, University of Maryland, USA, College Park, Maryland, 20742, cbuntain@cs.umd.edu; Jung Kyu (Rhys) Lim, University of Maryland, USA, College Park, Maryland, 20742, jklim@umd.edu.

2573-0142/2018/11-ART25 $15.00
https://doi.org/10.1145/3274294

remain about 1) how consistent these disaster scripts are across tragedies, and 2) where these scripts manifest. To answer these questions, we develop an algorithm to identify keywords making up these disaster scripts and apply this algorithm to seven terror attacks, seven earthquakes, and seven climate-change-related disasters between 2012 and 2017. Several consistencies then emerge: "attack" and "police" are referenced in the majority of terrorist attacks; "magnitude" and "quake" are common in response to earthquakes; and "safe" and "help" are often used in climate-related events. A common sentiment across all disaster types is also mention of the "affected", "victims", and "pray", consistent with work by Olteanu et al. [36] and the common "thoughts and prayers" sentiment referenced in the media [52]. We posit this popular type of emotional-but-uninformative content is a manifestation of communal coping through social support, which drives indirectly affected populations to participate in the conversation. These messages may therefore have value to the public beyond their information content. This paper's first contribution then is widening the bridge between crisis informatics and communications literature on communal coping and social support [2, 27, 44].

We then extend these results and prior research to investigate whether use of these disaster scripts correlate with disaster severity. As a proxy for severity, we use casualty metrics provided in the Global Terrorism Database (GTD) [32] and the National Geophysical Data Center's Significant Earthquakes Database (SED) [33]. We also highlight complications in measuring disaster severity via a single online platform, as many factors may bias online data; e.g., geopolitical issues have led to prohibitions against Twitter in several countries (e.g., China and Iran), language- and region-specific platforms are more popular in some countries (e.g., Sina Weibo in China), and socioeconomic issues raise barriers to Twitter (and the larger Internet) in rural areas [34, 53]. To illuminate some of these potential biases, we decompose our comparative study between disaster-script vocabularies and severity across several geographic regions, provided by the GTD and SED. Results show our disaster scripts for terrorism and earthquakes both significantly correlate with disaster severity, but this correlation is strong primarily in Western, English-speaking countries. We then contrast days in which our terror and earthquake disaster scripts see the most use to days with highest casualties, showing divergence between casualties in non-Western countries and disaster-script Twitter volume. Instead, days with the most frequent use of terror-related disaster scripts tend toward Western countries. This divergences suggests volume of these scripts may yield an alternate method to compare societal impact of different crises, distinct from casualties but limited to Western societies. Additionally, a better understanding of these cross-disaster scripts could facilitate *information and user response specific to a single crisis event* like where to seek shelter, what areas are heavily damaged, and other data contributing to situational awareness [22, 49, 50].

To establish the foundations for this work, we begin with an overview of relevant literature on OSN-based crisis informatics before describing our disaster vocabulary construction algorithm and Twitter dataset. Descriptions of these studies and our results on consistent vocabularies within and across disasters and correlations between these vocabularies and casualties follow before a discussion on implications, to include Twitter as a source of social support, regional biases, and evaluating crisis severity. We close with a discussion of this work's limitations and future work.

Researchers studying Twitter for crisis informatics, crisis responders interested in leveraging OSN data, and aid organizations wishing to engage better in the aftermath of disaster will find the following contributions:

- An algorithm for constructing disaster-type vocabularies,
- A set of language commonalities within and across disaster types,
- A bridge between emotional "disaster scripts" and communal coping/social support,
- A regional decomposition of emotional response to disasters, and
- An alternate, non-casualty based method for comparing societal impacts of disasters.

2 RELATED WORK

It is well established that OSNs contain useful and actionable information in a disaster's wake, as studied in the burgeoning crisis informatics literature. Examples include an automated classification system to identify tweets that contribute to situational awareness during crises [49], a system to identify potential eyewitnesses during mass unrest using Twitter data [12], and others (e.g. [15, 20, 26, 36, 51]). Despite quality issues plaguing OSNs (especially after crises) [8, 9, 46], the public wants emergency responders to use this data [21, 29], and recent research has shown platforms like Twitter are becoming more integrated into crisis response infrastructure [11].

As a result, much of the research in this area tends to focus on methods for removing uninformative content (e.g., Verma et al. [49]), as much of the social media data posted during high-impact crises are more sentimental than informative [24, 36]. Such efforts tend to fall on a spectrum between deep, qualitative content analysis across a few crises (e.g., Glasgow et al. [17, 18]) and broad, quantitative analysis of sharing behaviors/propagation around popular content (e.g., Sutton et al. [48]). Results on OSNs' role in supporting victims, while providing insight into the social solidarity networks formed online, study a single event [17, 18]. Investigations into similarities in structural response show increases in retweets and a preference for local news and law enforcement organizations as information sources, but these studies deal with only one *type* of disaster: natural disasters in Starbird and Palen [47]/Verma et al. [49] and terrorism in Buntain et al. [4]. In contrast, this paper studies not what is informative about a specific crisis but instead identifies the potentially uninformative commonalities across many crises *and* distinct disaster types. These uninformative trends are of particular interest because common sentiments like thoughts and prayers (seen even in early work by Starbird and Palen [47] and measured in Olteanu et al. [36]), are often filtered out during analyses, but these messages generally comprise the majority social media content during national or global coverage of disasters. Their prevalence, however, suggests some utility to the broader social media population; this paper then presents ways in which these sentiments could be constructively used rather than simply removed.

This paper is not the only study of social media across disasters, as Verma et al. studied earthquakes, wildfires, and floods [49], and Sutton et al. studied terrorism, wildfires, blizzards, a hurricane, and a flood [48]. Rather than differentiating responses across disaster types, these works focus on common methods for extracting information and stimulating information propagation. Most similar to our work is the study by Olteanu et al., where they present a characterization of information types and sources shared across 26 disasters of various types [36]. While the most similar in goals, Olteanu et al. perform a qualitative study using crowdsourcing to characterize information types and sources across several disaster dimensions. In contrast, the investigation described herein focuses on the tokens used and how they correlated with severity and region, yielding new insight into disaster-type-specific language and regional effects. Both our work and that by Olteanu et al. builds on the CrisisLex effort, which maintains a lexicon of crisis-related keywords across many types of disaster [35]. Our paper is distinct in its objectives, however, as CrisisLex maintains a dataset of keywords for improving information retrieval tasks in the wake of crises, whereas we study the lexical similarity across disasters to characterize trends in language and region.

Beyond crisis informatics, researchers are increasingly aware of potential validity issues in connecting social media to the real world as we better identify ways in which platforms like Twitter are biased and not representative of real-world populations [53]. Researchers must therefore exercise caution when designing sampling methods for social media-based studies [34]. Recent efforts have shown Twitter's geographic limitations are part of this bias, with Twitter populations often centered on heavily populated urban areas[6, 14]. Hence, solutions like earthquake alert systems as proposed in Sakaki et al. [41], while performant in densely populated countries like

Japan, may fail in more rural or sparsely populated regions (especially unfortunate as these areas likely have limited infrastructure and would benefit most from such systems). This work is similarly affected, as disasters strike regardless of whether sufficient Internet infrastructure or Twitter populations are present. While correcting for such biases is difficult, we highlight how they may affect this work, and we add to the literature on bias by identifying regional differences in response and where researchers should take care in studying response to disaster.

3 METHODS

This section lays out our data sources, disaster-script algorithm, and the study for comparing disaster scripts and measuring correlation between their use and disaster severity.

3.1 Social Media Data Source

This effort uses Twitter as a model OSN given its public-by-default context, which makes it ideal for information dissemination, and the large archives available for it. Our dataset is constructed from Twitter's 1% public sample stream, covering 1 January 2012 to 31 March 2017, containing 7,584,676,020 tweets. This collection **includes retweets and short tweets** as users express and share common sentiment via these behaviors, and retweets tend to increase in response to crises [4]. While others have explored bias in Twitter's public sample, research shows popular content is relatively well represented in the data [19, 31], and disasters tend to become popular topics. This method should be equally applicable to Reddit, Facebook, or other textual OSNs.

3.2 Disaster Selection

This investigation studies commonalities both within and across disaster types. Using the idea of "disaster madlibs" as inspiration, we first investigate whether Twitter response to the same types of disasters share similar vocabulary. Second, we explore whether these similarities hold regardless of disaster type. We therefore select a diverse set of disasters, borrowing from the hazard categories presents in CrisisLex: terrorist attacks; earthquakes; and a combination of meteorological, hydrological, and climatological disasters via hurricanes, typhoons, floods, and wildfires (referred to as "climate-related" events). Table 1 describes the events of interest, keywords, and analysis timeframes. For each event, we study the day of the event and the six days following, as prior research indicates event-related discussion generally returns to pre-event levels within one week [4, 36]. Search keyword are also selected in consultation with subject-matter experts.

For terrorist attacks, we select seven events based on their significant coverage in Western media and span across countries and continents. We select earthquakes by first using a list of high-fatality and high-magnitude earthquakes occurring between 2012-2016 [54] and filtering quakes that induce a large response on social media following Sakaki et al. [41]. Narrowing the list of earthquakes is essential because many earthquakes occur in remote locations where OSN access is limited, and some level of Twitter response is required for our analysis. Lastly, for our climate-related events, we leverage events from 2013 used in CrisisLex and their associated search terms [36].

Keywords shown in Table 1 are also purposefully general. As our algorithm considers the *intersection* of content across multiple crises, response at the disaster-type level is robust against false positives (i.e., irrelevant Twitter content) this generality may introduce. That is, even if the tweet set for a specific crisis event contains many irrelevant tweets, our algorithm identifies content that exists across *multiple* crises, which is likely to exclude irrelevant, non-disaster-related tweets.

3.3 Identifying Consistencies in Disaster Response

This section presents Algorithm 1, which takes a set of short event descriptions E and an unfiltered OSN dataset T as input and produces a set of keywords V used across the majority of these events.

Table 1. Selected Disasters By Type. "Search Keywords" are queries used to identify relevant tweets in Twitter's Public Sample Stream during the given timeframes. These keywords are intentionally general, as false positives will be excluded when taking the intersection of content across these crises.

Disaster	Search Keywords	Timeframe
Terrorist Attacks		
Boston Marathon Bombing	boston	15-22 April 2013
Westgate Mall Attack	westgate	21-27 September 2013
Kano Mosque Bombing	kano, bombing	11 Nov.-3 Dec. 2014
Charlie Hebdo Attack	paris, hebdo, jesuischarlie	6-12 January 2015
Paris November Attacks	paris	12-18 November 2015
Brussels Bombing	brussels, zaventem	15-22 March 2016
Pulse Nightclub Shooting	orlando, pulse, nightclub	11-17 June 2016
Earthquakes		
2012 Indian Ocean	sumatra, indian ocean, indonesia, tsunami	11-17 April 2012
2014 Iquique, Chile	tarapaca, chile, tsunami, iquique	17-23 March 2014
2014 Napa Valley	napa, california, francisco, bayarea, american canyon	11 Nov.-3 Dec. 2014
2015 Nepal	nepal, everest, kathmandu	25 Apr.-1 May 2015
2015 Nepal Aftershock	nepal, everest, kathmandu	12-18 May 2015
2015 Hindu Kush	afghanistan, pakistan, kashmir, south asia	26 Oct.-1 Nov. 2015
2016 Perugia, Italy	italy, perugia, amatrice, pescara	24-30 August 2016
Climate-Related Disasters		
Hurricane Sandy	frankenstorm, hurricane, sandy	25-31 October 2012
Alberta Floods	#abflood, alberta flood, alberta flooding, alberta floods, calgary flood, calgary flooding, calgary floods, canada flood, canada flooding, canada floods, #yycflood, #yycflooding, #yycfloods	20-26 June 2013
Manila Monsoon Floods	maringupdates, #phalert, philippine flood, philippine floods, #prayforthephilippine, #reliefph, #rescueph, #safenow, #trafficph	18-24 August 2013
Colorado Floods	#boulderflood, #coflood, #cofloodrelief, #coloradoflood, colorado flooding, colorado floods, #longmont, #opcoflood	11-17 September 2013
Australia Wildfires	australia #bushfire, australia fire, australia fires, #bushfire sydney, #faulconbridge, #nswbushfire, nsw #bushfire, #nswbushfires, nsw fire, #nswfires, nsw fires, #nswrfs, #nswrfs, #prayforaustralia, prayfornsw, #prayfornsw, #redoctober, #sydneybushfire, #sydneybushfires, #sydneyfire, sydney fire, #sydneyfires, sydney fires	15-21 October 2013
Typhoon Yolanda	#floodph, #haiyan, #philip-, pines, #redcrossphilippines, #reliefph, #rescueph, rescue ph, #safenow, #supertyphoon, super typhoon, #tracingph, #typhoon, #typhoonaid, #typhoonhaiyan, typhoon haiyan, typhoonph, typhoon philippines, #typhoonyolanda, typhoon yolanda, #yolanda, #yolandaactionweekend, #yolandaph	5-11 November 2013
Sardinia Floods	#cleopatra alluvione, #cyclonecleopatra, cyclone cleopatra, #sardegna, sardegna alluvione, #sardinia, sardinia flooding, sardinia floods	16-22 November 2013

To avoid requiring researchers to provide exhaustive event descriptions, this algorithm uses a data-driven approach to identify shared vocabularies via pseudo-relevance feedback (PRF), a common query expansion technique [28] and used in CrisisLex [35].

At a high level, researchers select a set of related disaster events and provide a small set of keywords and a date for each event; our algorithm then searches the OSN dataset for tweets containing these keywords in the seven-day period after each event. Messages matching the event comprise a pseudo-relevant set; we tokenize this set and rank the resulting tokens by frequency. For each event, the algorithm saves the top 100 most relevant keywords as the primary response vocabulary for that event. Our algorithm then ranks keywords by the number of events in which they appear, returning those that are used in a super-majority of events as the shared, cross-event vocabulary. These cross-event vocabularies are specific to one event type, however; one could run this algorithm across all event types simultaneously at the expense of a type-specific vocabulary.

At a more detailed level, lines 1-4 of Alg. 1 set up the event descriptions, and line 5 queries the OSN dataset for documents matching the given query q for each event. This search uses substring matching to identify pseudo-relevant documents, allowing query terms like "paris" to match "#pray4paris" and other common hashtag patterns. This `match()` function then returns a set of pseudo-relevant messages T_{rel} containing the query patterns.

By tokenizing these messages, we find common keywords in each event. To expand this keyword set, we use a standard probabilistic query expansion model based on Kullback-Leibler divergence (KLD) [7]. While CrisisLex is similarly data-driven, it uses keyword co-occurrence networks to identify event-specific keywords, but we opt for KLD here as it can provide higher-signal expansion terms [7]. KLD scores for each token are calculated via a foreground/background model, in which the foreground model is the distribution of keywords in the pseudo-relevant set, and the background model is the distribution in the unfiltered set. This divergence metric gives more weight to keywords that are both frequent in the foreground model and more frequent in the foreground than in the background (i.e., keywords that are both common in the set of documents matching the event and diverge from their typical use across all documents). We then rank each keyword by its KLD score via the `rank()` function and take the top 100 keywords (line 7) as the most relevant for this event and add them to the overall vocabulary *vocabMap*.

After scoring these top 100 keywords for each event, *vocabMap* contains event-specific keywords and the number of events in which they occur. To identify keywords that occur in the majority of these disaster events, we remove stopwords, search terms, and any keyword that occurs in fewer than two-thirds of the events. An alternate approach would be to pool all relevant tweet sets T_{rel} together and compute KLD on this pool, but such an approach would bias keywords towards events with larger OSN response volumes. We then return the resulting filtered keyword set as the disaster-type vocabulary V, which contains keywords that are both frequently used in each event and across the majority of events in the event set.

We apply this algorithm to the disaster classes described above to study the following hypotheses:

- **H1** – Crises of similar type will exhibit similar keywords across events.
- **H2** – Regardless of crisis type, references to social support and social solidarity will be common.

H1 is motivated by research into event detection via OSNs, as in prior work on detecting earthquakes via frequencies of "quake" and "shaking" [41]. Successful event detection, as CrisisLex enables, suggests similar events should exhibit similar responses. **H2**, however, is based on the "disaster Mad Libs" concept and the "thoughts and prayers" sentiment shown in prior work. This sentiment is supported by both Glasgow et al., who show significant response to social solidarity following a mass shooting [18], and Olteanu et al., who show "sympathy and emotional support" is the second most common type of information shared in the wake of disaster [36].

Algorithm 1: Process for extracting "disaster scripts", or keywords oft-used across events

input : A set of event descriptions E, and OSN dataset T
output: A set of keywords V shared across events

$vocabMap \leftarrow Dict()$; `/* Cross-event vocabulary */`
foreach $event \in E$ **do**
  `/* Find event-specific keywords */`
  $q \leftarrow event.query$; `/* Event-specific query */`
  $d \leftarrow event.date$; `/* Event start date */`
  $T_{rel} \leftarrow match(T, q, d, d + 7)$; `/* Match msgs */`
  $v_e \leftarrow rank(T_{rel}, T, d, d + 7)$; `/* Score keywords */`
  $v_{top} \leftarrow topK(v_e, 100)$; `/* Take top 100 keywords */`
  foreach $kw \in v_{top}$ **do**
    `/* Increment top-k keywords in vocabMap */`
    $vocabMap[kw] = vocabMap.get(kw, 0) + 1$
`/* Keep keywords shared by event majority */`
$V \leftarrow filter(vocabMap, ceil(1/2 * |E|))$;
return V

3.4 Disaster Severity and Regional Twitter Response

Constructing disaster scripts is valuable for understanding why the public is using these platforms, but how these vocabularies are used in the interim between disasters is equally important for understanding whether these responses are driven by the disaster or by some other factor. In this effort's second study, we investigate these intervening periods by comparing frequencies of disaster-category-specific vocabularies in Twitter to time series of disaster severity, operationalized via daily casualty rates. We test these correlations both at the global and regional levels to evaluate impact of disasters in different locations and their affects on different populations, as motivated by the following hypotheses:

H3 – For specific disaster types, normalized frequencies of the relevant disaster script correlate with disaster severity.
H4 – Correlations between OSN response and disaster severity are modulated by region.

H3 should be a consequence of our vocabulary construction, but these vocabularies may be used more generally than in response to disaster (e.g., references to the "dead" may be confounded by popular cultural phenomena like season premieres of "The Walking Dead", a popular television show). For **H4**, we expect to see changes based on the region where a disaster occurs, as a result of our English-centric language, socioeconomic factors like Internet access, or other bias-inducing issues. To this end, we construct a set of time series datasets for terrorist attacks and earthquakes, each consisting of normalized daily tweet frequencies that include tokens from the corresponding disaster script, and a time series of the daily casualties reported in the GTD and SED respectively. Climate-related disasters are excluded here because such disasters have less-well-defined event boundaries, as hurricanes are often foreseen several days in advance, and an analogous dataset on climate-related casualties is unavailable. Using these time series, we then compare the top five days by Twitter response to casualties to evaluate whether these two measures of severity are similar. The following two sections detail how we compare terror- and earthquake-specific casualties to Twitter.

3.4.1 Twitter and Terrorism. For terrorism, the 2015 version of the GTD [32] contains over 170,000 records of terrorist attacks between 1970 and 2015, replete with information on deaths, injuries, attack types, and regions. One cannot simply compare Twitter frequencies in 2013 with frequencies in 2015, however, as the Twitter platform has grown considerably in users. To account for these fluctuations, we normalize disaster-relevant tweet frequency by dividing through by the daily tweet frequency. This normalization yields a percentage of Twitter conversation that discusses the disaster, which is more comparable over time. We compare this Twitter percentage to casualties (deaths plus injuries) and calculate Pearson's correlation coefficient between these time series. We also filter the GTD to remove duplicate terror attacks, failed terror-related events, assassinations, hijackings, kidnappings, and infrastructure attacks and align to our Twitter dataset. This filtering leaves 24,967 "successful"[1] bombings and armed assaults between 1 January 2012 and 31 December 2015.

To provide benchmarks for this correlation, we include two baselines: a time series of tweets mentioning "terror", and a dataset of tweets mentioning a string in the CrisisLex dataset. One would expect our terrorism-specific vocabulary to outperform the CrisisLex lexicon given its more general construction (i.e., it covers disasters in general rather than a specific class of disaster), and our terrorism-specific vocabulary should outperform the simple terror baseline given expressions of sympathy and condolences common to disasters. Lastly, to address issues of access and language in Twitter, we also compare these lexica across the twelve regions covered in the GTD (e.g., North/South America, Western/Eastern Europe, East/South/Southeast/Central Asia, the Middle East, etc.) and combinations of regions.

3.4.2 Twitter and Earthquakes. For earthquakes, the SED [33] provides ground truth on lives lost, injuries, missing, and cost of damage for significant earthquakes between 2150 BC and the present. The SED defines an earthquake as significant if it causes "damage (approximately $1 million or more), 10 or more deaths, Magnitude 7.5 or greater[2], Modified Mercalli Intensity X or greater, or ... generated a tsunami" [33]. Given the more rigorous and data-driven definition for an earthquake, we perform no de-duplication in this dataset.

As in the terrorism study, we normalize the frequency of tweets including keywords from our earthquake vocabulary by dividing by the daily tweet frequency in our dataset. We also compare vocabulary against a baseline of "earthquake" and "shaking", as suggested by Sakaki et al. [41], and the CrisisLex lexicon. The earthquake covariate here is the daily casualties reported by the SED, which we also divide by region. To support comparisons against terrorism, we also coalesce the regions into groups consistent with GTD region definitions; that is, in the GTD, the North America and Western Europe regions cover the SED's North America, West Europe, and South Europe regions.

4 RESULTS

4.1 Consistency Findings

After applying Algorithm 1 to the event sets listed in Table 1, several keywords appear across each disaster class. Table 2 shows these commonalities, with 17 keywords present in the majority of terror attacks, 11 are shared by the majority earthquakes, and nine are shared by climate events.

For **H1**, we find the following similarities within each disaster type: Following terrorist attacks, references to the attack, victims (e.g., "victims", "killed", "families", and "dead"), religion ("prayers"

[1] "Success" here is defined by the GTD in contrast to attacks prevented by law enforcement, technical issues, or other factors before it occurred and should not be taken as an endorsement or measure of efficacy.

[2] "Magnitude" varies by event date, with early events measured by a "felt area" scale, more recent quakes measured using the well-known Richter scale, and modern events measured via the surface-wave magnitude.

Table 2. Tokens Occurring in the Majority of Disasters ($|E|$ represents the number of events in which the keyword occurred in the top 100 divergent tokens)

Terror		**Quake**		**Climate**	
Token	**|E|**	**Token**	**|E|**	**Token**	**|E|**
attack	7	quake	6	affected	6
victims	7	magnitude	6	safe	6
killed	7	earthquake	6	stay	5
police	6	terremoto	6	victims	5
terror	6	#earthquake	6	evacuation	4
families	5	affected	5	thoughts	4
terrorist	5	sismo	4	help	4
prayers	5	pray	4	prayers	4
dead	5	gempa	4	damage	4
heart	4	prayers	4		
thoughts	4	victims	4		
affected	4				
praying	4				
injured	4				
world	4				
tragedy	4				
people	4				

and "praying"), law enforcement, and terror are common. In response to earthquakes, we see similar references to affected populations and prayer as well as multiple references to "earthquake" across several languages ("terremoto","sismo", and "gempa" are Italian, Spanish, and Indonesian for quake respectively). In the more diverse climate-related events, mentions of the affected population ("affected", "victims", "damage", and "evacuation") and religion are still common along with requests to stay safe and provide assistance.

Across these three disaster types, we find two common themes: references to victims and religions invocations of prayer, with "victims", "affected", and "pray" present across all three disaster types. We extract a few samples of this usage and see these mentions of victims and prayers are consistent with **H2** and generally represent instances of social solidarity expressed through hashtags like "#pray4boston", "#pray4paris", or "#pray4nepal".

4.2 Twitter and Terrorism

Comparisons between Twitter and affected populations, as shown in Table 3, demonstrate Twitter response exhibits positive, weak correlation with casualties in North America and Western Europe. In both North America and Western Europe, our terrorism vocabulary is more correlated with the affected population size than CrisisLex (as expected given CrisisLex's generality). Our terrorism response vocabulary is a better predictor of casualties in North America than the "terror" baseline, while "terror" performs better in Western Europe. Among all pairs of regions, we also find our terror vocabulary has the strongest correlation with North America and Western Europe together, where it outperforms both baselines.

To further investigate this difference between Western and non-Western countries and the difference between the perceived impact on social media and casualties, we compare Twitter's most terror-filled days versus to the GTD's most dangerous days between 2012 and 2015. A common criticism of Western media is the limited coverage of deaths in non-Western countries, which is reflected by the regional differences between Twitter and the GTD. Similarly, given the modern

Table 3. Pearson Correlation with Twitter and Terrorism Affected Population

Region	"Terror"	CrisisLex	Terror Vocabulary
Australia	0.009472	0.01965	-0.001151
Central America	0.02683	0.01688	0.01713
Central Asia	0.002054	0.03481	-0.006242
East Asia	-0.01843	-0.007652	-0.01144
Eastern Europe	0.02376	0.05021	0.03446
Middle East	-0.007476	0.02216	0.00063
North America	0.1065***	0.2508***	0.2743***
South Africa	0.0621	0.0724	0.0183
South America	-0.03845	-0.04122	-0.01998
South Asia	-0.01333	-0.03445	-0.01863
Southeast Asia	-0.02797	-0.02355	0.01274
Western Europe	0.2247***	0.1627***	0.2131***
North America and Western Europe	0.1955***	0.2973***	0.3408***
Global	0.01849	0.04701	0.02191

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 4. Pearson Correlation with Twitter and Earthquakes

Region	Sakaki et al.	CrisisLex	Quake Vocabulary
Central America	0.006708	-0.01533	0.01417
Central Asia	0.1474***	0.0211	0.1098***
East Asia	0.03244	-0.01548	0.03763
Middle East	0.01887	-0.02366	0.0114
North Africa	-0.009628	0.008498	-0.006536
North America	0.09811***	0.01971	0.0727**
Pacific	0.0404	-0.0194	0.0293
S. and SE. Asia	0.4411***	0.0077	0.3871***
South Africa	0.009019	0.016642	0.007566
South America	0.06121**	0.008598	0.09988***
Southern Europe	0.2041***	0.0235	0.3777***
Western Europe	0.005261	0.0495*	0.03206
North America, West & South Europe	0.2132***	0.0232	0.3633***
Global	0.4342***	0.0018	0.3985***

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

media climate, a terrorist attack in a Western country need not injure many to have a significant impact on Western society.

Results show difference between Twitter's most terror-filled days and GTD's most dangerous days. Of the five days with the highest proportion of terror-related tweets, four cover the Paris November attacks between 13-16 November 2015, in which 130 people were killed and 413 were injured; the remaining day covers the Boston Marathon Bombing on 15 April 2013, in which 3 people were killed and 264 were injured. The five most deadly days in the GTD do not overlap with these events. Instead, they are days with large death tolls in Syria, Sudan, Afghanistan, and Libya, which averaged 652 casualties each.

4.3 Twitter and Earthquakes

Table 4 presents these comparisons and shows both the Sakaki et al. baseline and our earthquake vocabulary have moderate positive correlation to global earthquake severity. Regionally, the Sakaki et al. baseline and our vocabulary exhibit strongest correlations with earthquakes in South/Southeast Asia, Southern Europe, and Central Asia. CrisisLex, however, is only very weakly correlated with all regions (again, as one would expect given its more general nature). As we observed from the terrorism data analysis, our disaster-specific vocabulary dominates in the combined North American and Western/Southern Europe regions.

To further study the difference between intensity of public response on Twitter and the human toll of disaster, we again compare the five most intense disasters based on Twitter response and casualty. Results show the difference in intensity between Twitter responses and disaster impact. For Twitter, these five dates cover three earthquakes: one in Amatrice, Italy on 24 August 2016, resulting in 705 casualties; one in Salango, Ecuador on 16 April 2016, resulting in 6,936 casualties; and the three days following the Nepalese earthquake in Kathmandu on 25 April 2015, causing 26,066 casualties.

Sorting based on casualties in the SED, the top five earthquakes by casualties include the Nepalese and Ecuadorian earthquakes but also covers the 2013 Lushan earthquake in China on 20 April 2013 (11,666 casualties), the 2012 East Azerbaijan earthquake in Iran on 11 August 2012 (3,306 casualties), and the 2014 Ludian earthquake in China on 3 August 2014 (2,989 casualties).

5 DISCUSSION

These studies advance our understanding of the OSN-using public's response to various categories of disaster, which in turn illuminate several broader topics in communications, OSN analysis, crisis informatics, and machine translation. Exploring these topics below, we discuss (i) similarities across languages in response to disasters, (ii) social solidarity mechanisms during disaster and its consistency with communications theory on communal coping; (iii) the disconnect between human casualty and interest on OSNs; and (iv) Twitter's regional biases and consequences for future OSN-oriented research.

5.1 Cross-Disaster Language Similarities

Our results find strong commonalities of language within each disaster type, as discussed above, which suggest differing information needs across the contexts. Common references to police during terrorist attacks are absent during earthquakes and climate-related events, indicating higher interest in law enforcement and police action in response to the attack, consistent with prior work [25]. In reviewing consistencies in earthquake response, we see four different languages represented in this vocabulary, all of which refer to earthquakes: English, Italian, Spanish, and Indonesian. This event-specific similarity is consistent with prior research into sporting events, in which Twitter response to dramatic moments of play also crosses languages (e.g., in response to a goal during the 2014 World Cup, references to the semantic concept of "goal" saw a spike in usage in Spanish, Portuguese, German, English, and other languages) [5]. Two questions emerge from this result: First, why are earthquakes unique in their multi-lingual response, as neither the terrorism nor climate-related events show this similarity? Earthquakes are rather unique in the disaster hierarchy because they affect large populations nearly simultaneously and without warning, while hurricanes are more easily foreseen, and terrorist attacks generally only affect small geographic areas. Second, can these temporal and event-response similarities connect syntactically different but semantically related concepts for machine translation? For research into low-resource languages (languages that have limited datasets or computational language models), leveraging this temporal correlation may support bootstrapping translation from the low-resource language by providing cross-language pairs with which to seed the translation process.

5.2 #Pray4___: Coping and Social Solidarity

Prior research into OSN usage during 26 disasters in 2012-2013 found expressions of sympathy, support, and references to affected individuals were ubiquitous [36]. Glasgow et al. similarly found supportive expressions of solidarity to be an important aspect of OSN grief response in the aftermath of the Newtown, CT school shooting. Our results are consistent with these past works, as references to prayer and mentions of the "victims" and the "affected" are present in 16 of 21 disasters we study and in all three disaster categories. Our results further indicate sympathetic expression and solidarity often have religious connotations. In fact, during the 15 April 2013 and the Boston Marathon Bombing, 9,572 of the 20,850 disaster-vocabulary tweets include a hashtag starting with "pray", often of the form "#pray4boston" or "#prayersforboston". We see similar patterns in "#pray4paris". One can see this language even in early studies of Twitter, where "#pray4" appeared in response to flooding in North Dakota in 2009-2010 [47, 49].

Twitter's emotional response to disaster appears common beyond the most well-represented regions as well. For example, many Twitter users responded to the 16 April 2014 sinking of a South Korean ferry with "#Pray4SouthKorea", with similar responses from the Philippines during other crises via "#pray4philipines". Other research has also touched on this result, with variants of "#prayforchile" present in the top hashtags and references to victims with "missing" and "death" trending around the 2010 Chilean earthquake [30]. In Turkey, researchers have studied Twitter's response to the 2014 Soma mining disaster, in which 301 people perished [13], and these researchers identified a common religious response in Turkish as well[3].

While prayer is generally seen as religious in nature, it is unclear whether this use in Twitter is truly a religious response; instead, it may be more of a communal coping strategy, or an expression of social support and solidarity. This pseudo-religious language may be more about establishing a shared connection through common religious convictions or vernacular and may instead serve an important purpose in helping OSN users cope with *seeing* these tragedies even if they do not directly experience them. Communications literature suggests such social support/coping manifests when the responsibility to address the crisis is also shared [2], and shared concerns in these disaster contexts are aligned with the social support literature [44]. Given the population of users who are geographically unaffected but still engage in the event is significantly larger than the directly affected population [24, 47, 49], and the commonality of thoughts-and-prayers-style emotional content despite the limited resharing it receives [3, 48], these patterns perhaps deserve additional consideration beyond the noise they inject into OSNs. This contribution extends our understanding about coping for the general online public and provide a possible theoretical bridge to the communications literature for why this emotional but uninformative content is so prevalent in the aftermath of crises.

5.3 Tragedy and Statistics: Impact vs. Casualty

In general, our results for terrorist attacks and earthquakes demonstrate **H3** holds in that our disaster scripts for these two event types outperform a simple baseline and more general information retrieval-based keyword sets. Despite this result, measuring a disaster's severity is a complex task, with casualties or property damage often used as standard indicators, but these data points may not match the disaster's salience to or impact on the public. For instance, despite terror attacks in the Middle East and North Africa and earthquakes in China and Iran having high casualty

[3]Gunes Ertan, correspondence with the author, 27 July 2017.

rates, our results show these events did not stimulate the largest responses on Twitter. Instead, the Boston Marathon Bombing and the Paris November Attacks were the two highest-impact terror events in 2012-2015 on Twitter, although they had comparatively limited casualties. In addition, the Paris November attacks had nearly twenty times more fatalities than the Boston Marathon Bombing, so even in the highest impact attacks by Twitter, we see an imbalance in the number of affected. Social media may therefore provide an alternate means for measuring a disaster's impact through a means separate from raw casualty numbers. By providing a standard vocabulary of crisis response for a particular type of disaster, we can better compare these events' impact on the public by comparing frequencies and durations of these vocabularies as a measure for the crisis salience. This information is particularly important for terror-like events where impact is the goal.

This additional axis for salience also suggests an opportunity to enhance disaster databases like the GTD. While we show tracking terrorism vocabularies in Twitter can support GTD construction by identifying terror attacks as they occur, traditional news media can likely fulfill this role as well in Twitter's well-represented regions. Instead, our study suggests including public impact/salience as an additional axis in such a database could be a useful ranking mechanism. In quantifying this salience, however, one must take care about the population of interest: Our work also suggests terrorism salience on Twitter is likely a proxy for salience in Western countries rather than global salience while salience of earthquakes is better distributed across the globe.

5.4 Twitter's Regional Biases

An important result from Tables 3 and 4 is that the disaster scripts we have identified are regionally biased, as anticipated in **H4**. Correlations between our disaster scripts and crises are strongest in North America and Western/Southern Europe, with earthquakes also strong in Central Asia.[4] The nearly non-existent global correlation with terror (Table 3) and limited earthquake response in South America, Africa, and East Asia (Table 4) reinforce these limitations. While disasters in other countries like the South Korean ferry sinking do receive coverage in Twitter, with comments about religion and victims, they may be potentially overwhelmed out by discussions of The Walking Dead or topics more germane to Western audiences. Ideally, researchers could focus on specific regions of interest and remove tweets from Western regions that could drown out the target signal, but tweets with geolocation data are rare, and triangulating individual Twitter users is difficult at scale [10]. Researchers should therefore exercise caution when studying disaster response in other regions via Twitter despite its appeal as an easily accessible data source.

While Twitter's Western-centric response in our results may partially be explained by our focus on terrorism in Western countries and geopolitical/infrastructure issues (e.g., Twitter usage in China is low), several countries in Africa have English as an official language (e.g., Nigeria) or tweet primarily in English (e.g., South Africa), and yet terrorism and earthquake disasters in sub-Saharan Africa are uncorrelated with Twitter response, despite the relatively large number of events. Similarly, Saudi Arabia has a significant Twitter population, accounting for 40% of all tweets from Arab world, with Egypt being the next largest source [43], and yet the Middle East and North Africa are not well-represented by our disaster vocabularies. To address language issues, we have explored using Google Translate to convert the terrorism vocabulary into the top five languages on Twitter and rerunning our experiments, but this translation yielded poorer results. An alternate explanation for Twitter's regional biases, however, may better explain the moderate global correlation for earthquakes versus the non-existent correlation with terrorism and the poor

[4] Limited correlation between Twitter and crises in Australia is notable but likely a result of the rarity of these disaster types (the GTD shows only four terrorist attacks, with a sum of six casualties in 2012-2015, and the SED has no significant earthquakes in Australia in 2012-2016).

results with Google Translate: Perhaps non-Western regions may exhibit different coping patterns to terrorism. Natural or unintentional disasters like earthquakes or hurricanes may better stimulate religious coping, as such events are often referred to as "acts of God," but terrorism's intrinsically political nature may elicit different responses from the affected population. Such a result would provide a counterpoint to prior work on Western coping mechanisms [1, 37]. Google Translate may also be insufficient to capture the nuances of public coping, so future research should explore non-English disaster responses, as we mention above with the Soma mining disaster in Turkey.

As a third point in studying Twitter's regional biases, a large volume of recent crisis informatics work has explored using Twitter to detect crises [41, 42], study population movements [16, 55], and similar research avenues. This work presents a paradox between Twitter use and Twitter-based detection needs: In Western countries, Twitter can provide insight into disaster detection and response, but these countries already have significant disaster response infrastructure in place. On the other hand, countries or regions where a Twitter-based detection/response mechanism would be most useful, the infrastructure or population likely is insufficient for Twitter to provide useful information. For instance, if a terrorist attack or earthquake were to occur in New York City or London, mainstream media will pick up the event rapidly [38], but when an earthquake occurred in Chile, it took several hours for Twitter to respond in a significant way even though geological instruments identified the quake and estimated its magnitude rapidly [30]. As a counterpoint, however, recent work by Dailey and Starbird on OSNs as crisis infrastructure find these platforms are still valuable in less-densely populated areas, suggesting a more targeted and locally oriented approach is needed [11]. Furthermore, even in regions with high correlations (e.g., the northeastern United States during Hurricane Sandy or the earthquake in Haiti), we know much of the content posted to Twitter about a disaster tends to come from areas that are not geographically affected by the disaster [24, 49]. While recent work has studied identifying eyewitnesses on Twitter [12] during riots in London, our results suggest more research should separate Twitter's global response from the locally affected population and focus more on how the local population's OSN response differs or enhances existing infrastructure (as was the claim during Hurricane Harvey in Houston, TX [39]).

6 LIMITATIONS AND FUTURE WORK

While our results are consistent with many prior efforts across several fields, this work has several limitations that future work could address. Foremost in these limitations is our selection of terrorist attacks that occurred in mostly Western countries. This selection was motivated partially by the Twitter platform to ensure we had sufficient response data from which to draw to identify similar keywords across events and partially by familiarity with these events. Future work could expand this selection to identify major terror attacks across the globe. This limitation is mitigated by including multiple disaster categories, however, as our earthquake and climate-related event selections are less Western-centric, so results on consistent references to prayer and the affected should still hold.

Related to limitations introduced by our event selection, we also intrinsically bias our results towards the English language as all of our search keywords in Table 1 are English. While English is the most common language on Twitter, accounting for other major languages like Arabic or Japanese could provide better results in regional Twitter-disaster correlations. Though machine translation tools did not increase this performance, future work could employ linguists with expertise in these other high-frequency languages and better explore whether our results stem from differences in language, cultural coping mechanisms, or Twitter access. Our use of substring search rather than token matching partially addresses linguistic issues, however, as patterns like "#Pray4..." seem conserved across languages, as we found many instances with such a hashtag embedded in tweets that were otherwise in Indonesian, Korean, or another language.

Finally, a third major limitation is our reliance on Twitter and Twitter's public sample stream, which only gives access to 1% of Twitter's full content stream. Prior work [31] has identified biases present in this supposedly random sample, which could impact our results. Much of these issues affect events in the "long tail" of Twitter coverage, however, so for particularly high-impact events like disasters, these issues should not affect our results in terms of coverage. Our cross-disaster vocabulary may be affected though as words that are used relatively rarely in a single event but appear across all events may be omitted when relying on the 1% sample. Future work could contract with expensive third-party data resellers and apply our vocabulary construction algorithm, but a sampling methodology would still be necessary as acquiring all tweets posted in a 7-day period after a single disaster (approximately 2.8 billion tweets) would be prohibitively expensive. Integrating data from Facebook or Snapchat or other major OSNs would also reduce our biases, but these sources are notoriously hard to access.

7 CONCLUSION

This paper explores public use of Twitter in response to disasters. We develop an algorithm for extracting commonly used vocabularies across a set of related disaster events and show several patterns are shared within each disaster category: responses to terrorism events often mention an "attack" and law enforcement; earthquake responses mention the quake and its magnitude in various languages; and climate-related disaster responses tend to include safety and requests for aid. Across all three types of disaster, Twitter users also regularly mention victims, the affected, and invocations of prayer, consistent with communication theory on communal coping and social support in the aftermath of crises. Using these disaster scripts, we test Twitter's regional biases, results of which suggests Twitter response favors Western countries, and the most tragic days in terms of casualty do not align with impact on Twitter. While other work has studied Twitter response to disasters, this work provides a novel analysis into the lexical form and locations of these responses. These results have implications for emergency responders, aid organizations, and crisis informatics researchers, illuminating potential regional biases in Twitter and suggesting social support and solidarity behaviors are common across disasters and disaster types. This commonality suggests an underlying motivation may be present in driving indirectly affected populations to respond and engage with each other during times of stress.

8 ACKNOWLEDGEMENTS

This research was supported by an appointment to the Intelligence Community Postdoctoral Research Fellowship Program at the University of Maryland, administered by Oak Ridge Institute for Science and Education through an interagency agreement between the U.S. Department of Energy and the Office of the Director of National Intelligence. The authors would also like to thank Brooke Liu, Michael Egnoto, and Holly Roberts for sharing their insights on crisis coping mechanisms.

REFERENCES

[1] Terri Adams, Leigh Anderson, Milanika Turner, and Jonathan Armstrong. 2011. Coping through a disaster: Lessons from Hurricane Katrina. *Journal of Homeland Security and Emergency Management* 8, 1 (2011).

[2] Tamara D Afifi, Susan. Hutchinson, and Stephanie Krouse. 2006. Toward a Theoretical Model of Communal Coping in Postdivorce Families and Other Naturally Occurring Groups. *Communication Theory* 16, 3 (2006), 378–409. https://doi.org/10.1111/j.1468-2885.2006.00275.x

[3] Melissa Bica, Leysia Palen, and Chris Bopp. 2017. Visual Representations of Disaster. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing - CSCW'17 (CSCW '17)*. ACM, New York, NY, USA, 1262–1276. https://doi.org/10.1145/2998181.2998212

[4] Cody Buntain, Jennifer Golbeck, Brooke Liu, and Gary LaFree. 2016. Evaluating Public Response to the Boston Marathon Bombing and Other Acts of Terrorism through Twitter. http://www.aaai.org/ocs/index.php/ICWSM/ICWSM16/paper/view/13038
[5] Cody Buntain, Jimmy Lin, and Jennifer Golbeck. 2016. Discovering Key Moments in Social Media Streams. In *Consumer Communications and Networking Conference (CCNC), 2016 13th Annual IEEE.*
[6] Fink Clayton C. 2012. Mapping the Twitterverse in the Developing World: An Analysis of Social Media Use in Nigeria. *Social Computing, Behavioral - Cultural Modeling and Prediction* 7227 (2012), 164.
[7] Claudio Carpineto, Renato de Mori, Giovanni Romano, and Brigitte Bigi. 2001. An information-theoretic approach to automatic query expansion. *ACM Transactions on Information Systems* 19, 1 (2001), 1–27. https://doi.org/10.1145/366836.366860
[8] Carlos Castillo, Marcelo Mendoza, and Barbara Poblete. 2011. Information credibility on twitter. In *Proceedings of the 20th international conference on World wide web (WWW '11)*. ACM, New York, NY, USA, 675–684. https://doi.org/10.1145/1963405.1963500
[9] Carlos Castillo, Marcelo Mendoza, and Barbara Poblete. 2013. Predicting information credibility in time-sensitive social media. *Internet Research* 23, 5 (2013), 560–588. https://doi.org/10.1108/IntR-05-2012-0095
[10] Ryan Compton, David Jurgens, and David Allen. 2014. Geotagging one hundred million twitter accounts with total variation minimization. In *Big Data (Big Data), 2014 IEEE International Conference on*. IEEE, 393–401.
[11] Dharma Dailey and Kate Starbird. 2017. Social media seamsters: Stitching platforms & audiences into local crisis infrastructure. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing - CSCW'17*. 1277–1289. https://doi.org/10.1145/2998181.2998290
[12] Nicholas Diakopoulos, Munmun De Choudhury, and Mor Naaman. 2012. Finding and assessing social media information sources in the context of journalism. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 2451–2460. https://doi.org/10.1145/2207676.2208409
[13] Burcak Erkan, Gunes Ertan, Jungwon Yeo, and Louise K. Comfort. 2015. Risk, profit, or safety: Sociotechnical systems under stress. *Safety Science* 88 (2015), 199–210. https://doi.org/10.1016/j.ssci.2016.02.002
[14] Clay Fink, Nathan Bos, Alexander Perrone, Edwina Liu, and Jonathon Kopecky. 2013. Twitter, public opinion, and the 2011 Nigerian Presidential election. In *International Conference on Social Computing (SocialCom)*. 311–320. https://doi.org/10.1109/SocialCom.2013.50
[15] Julia Daisy Fraustino, Brooke Liu, and Jin Yan. 2012. Social Media Use during Disasters: A review of the Knowledge Base and Gaps. (2012), 1–39.
[16] Song Gao, JA Yang, Bo Yan, Yingjie Hu, Krzysztof Janowicz, and G McKenzie. 2014. Detecting Origin-Destination Mobility Flows From Geotagged Tweets in Greater Los Angeles Area. *Geog.Ucsb.Edu* (2014), 0–4. http://www.geog.ucsb.edu/{~}sgao/papers/2014{_}GIScience{_}EA{_}DetectingODTripsUsingGeoTweets.pdf
[17] Kimberly Glasgow, Clayton Fink, and Jordan Boyd-Graber. 2014. "Our Grief is Unspeakable": Automatically Measuring the Community Impact of a Tragedy. *Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media* (2014), 161–169. http://www.aaai.org/ocs/index.php/ICWSM/ICWSM14/paper/view/8086
[18] Kimberly Glasgow, Jessica Vitak, Yla Tauczik, and Clay Fink. 2016. Grieving in the 21st Century: Social Media's Role in Facilitating Supportive Exchanges Following Community-Level Traumatic Events. In *Proceedings of the 7th 2016 International Conference on Social Media & Society*, Vol. 1691. ACM, 4. https://doi.org/10.475/123 arXiv:arXiv:1602.05561v1
[19] Sandra González-Bailón, Ning Wang, Alejandro Rivero, Javier Borge-Holthoefer, and Yamir Moreno. 2014. Assessing the bias in samples of large online networks. *Social Networks* 38, 1 (2014), 16–27. https://doi.org/10.1016/j.socnet.2014.01.004 arXiv:1212.1684
[20] Aditi Gupta and Ponnurangam Kumaraguru. 2012. Twitter explodes with activity in mumbai blasts! a lifeline or an unmonitored daemon in the lurking? (2012).
[21] Amanda Lee Hughes and Leysia Palen. 2009. Twitter adoption and use in mass convergence and emergency events. *International Journal of Emergency Management* 6, 3/4 (2009), 248. https://doi.org/10.1504/IJEM.2009.031564
[22] Muhammad Imran, Carlos Castillo, Fernando Diaz, and Sarah Vieweg. 2015. *Processing Social Media Messages in Mass Emergency: A Survey*. Vol. 47. 1 – 38 pages. https://doi.org/10.1145/2771588 arXiv:arXiv:1407.7071v1
[23] Muhammad Imran, Carlos Castillo, Ji Lucas, Patrick Meier, and Sarah Vieweg. 2014. AIDR: Artificial intelligence for disaster response. *Proceedings of the companion publication of the 23rd international conference on World wide web companion* October (2014), 159–162. https://doi.org/10.1145/2567948.2577034
[24] Marina Kogan, Leysia Palen, and Kenneth M. Anderson. 2015. Think Local, Retweet Global. In *Proceedings of the 2015 ACM Conference on Computer Supported Cooperative Work and Social Computing - CSCW'15*. 981–993. https://doi.org/10.1145/2675133.2675218
[25] Gary LaFree and Amy Adamczyk. 2015. *Change and Stability In Attitudes Toward Terrorism: the Impact of the Boston Marathon Bombings*. Preprint. START Center, University of Maryland.

[26] Brooke Liu, Julia Daisy Fraustino, and Yan Jin. 2013. *Social Media Use during Disasters: A Nationally Representative Field Experiment.* Technical Report. College Park, MD. 1–53 pages.

[27] Weixu Lu and Keith N Hampton. 2017. Beyond the power of networks: Differentiating network structure from social media affordances for perceived social support. *New Media & Society* 19, 6 (2017), 861–879. https://doi.org/10.1177/1461444815621514

[28] Christopher D. Manning, Prabhakar Ragahvan, and Hinrich Schutze. 2009. *An Introduction to Information Retrieval.* Cambridge University Press, Cambridge. https://doi.org/10.1109/LPT.2009.2020494 arXiv:0521865719 9780521865715

[29] Richard McCreadie and Ian Soboroff. 2017. *TREC Incident Streams Track Proposal.* Technical Report. University of Glasgow, Glasgow. 1–4 pages. http://dcs.gla.ac.uk/{~}richardm/TREC{_}IS/TREC-IS-Proposal.pdf

[30] Marcelo Mendoza, Barbara Poblete, and Carlos Castillo. 2010. Twitter Under Crisis: Can We Trust What We RT?. In *Proceedings of the First Workshop on Social Media Analytics (SOMA '10).* ACM, New York, NY, USA, 71–79. https://doi.org/10.1145/1964858.1964869

[31] Fred Morstatter, J Pfeffer, H Liu, and Km Carley. 2013. Is the sample good enough? Comparing data from Twitter's streaming API with Twitter's firehose. *Proceedings of ICWSM* (2013), 400–408. https://doi.org/10.1007/978-3-319-05579-4_10 arXiv:arXiv:1306.5204v1

[32] National Consortium for the Study of Terrorism and Responses to Terrorism (START). 2016. *Global Terrorism Database [Data file].* Technical Report. https://www.start.umd.edu/gtd

[33] NOAA National Geophysical Data Center. 2017. *National Geophysical Data Center / World Data Service (NGDC/WDS): Global Significant Earthquake Database.* Technical Report. https://data.nodc.noaa.gov/cgi-bin/iso?id=gov.noaa.ngdc.mgg.hazards:G012153

[34] Alexandra Olteanu, Carlos Castillo, Fernando Diaz, and Emre Kiciman. 2016. Social Data: Biases, Methodological Pitfalls, and Ethical Boundaries. *SSRN Preprint (8 March 2017)* (2016), 1–44.

[35] Alexandra Olteanu, Carlos Castillo, Fernando Diaz, and Sarah Vieweg. 2014. CrisisLex: A Lexicon for Collecting and Filtering Microblogged Communications in Crises. *Proc. of the 8th International Conference on Weblogs and Social Media* (2014), 376. http://www.aaai.org/ocs/index.php/ICWSM/ICWSM14/paper/download/8091/8138

[36] Alexandra Olteanu, Sarah Vieweg, and Carlos Castillo. 2015. What to Expect When the Unexpected Happens: Social Media Communications Across Crises. In *In Proc. of 18th ACM Computer Supported Cooperative Work and Social Computing (CSCW'15).*

[37] Kenneth I Pargament. 2001. *The psychology of religion and coping: Theory, research, practice.* Guilford Press.

[38] Saša Petrovic, Miles Osborne, Richard McCreadie, Craig Macdonald, Iadh Ounis, and Luke Shrimpton. 2013. Can Twitter replace Newswire for breaking news?. In *Proceedings of the 7th International AAAI Conference on Weblogs and Social Media*, Vol. 2011.

[39] Maya Rhodan. 2017. 'Please Send Help.' Hurricane Harvey Victims Turn to Twitter and Facebook. *Time* (aug 2017). http://time.com/4921961/hurricane-harvey-twitter-facebook-social-media/

[40] J Rogstadius, M Vukovic, C A Teixeira, V Kostakos, E Karapanos, and J A Laredo. 2013. CrisisTracker: Crowdsourced social media curation for disaster awareness. *IBM Journal of Research and Development* 57, 5 (sep 2013), 4:1–4:13. https://doi.org/10.1147/JRD.2013.2260692

[41] Takeshi Sakaki, Makoto Okazaki, and Yutaka Matsuo. 2010. Earthquake shakes Twitter users: real-time event detection by social sensors. In *Proceedings of the 19th international conference on World wide web (WWW '10).* ACM, New York, NY, USA, 851–860. https://doi.org/10.1145/1772690.1772777

[42] Takeshi Sakaki, Makoto Okazaki, and Yutaka Matsuo. 2013. Tweet analysis for real-time event detection and earthquake reporting system development. *IEEE Transactions on Knowledge and Data Engineering* 25, 4 (2013), 919–931. https://doi.org/10.1109/TKDE.2012.29

[43] Fadi Salem. 2017. Social Media and the Internet of Things: Towards Data-Driven Policymaking in the Arab World: Potential, Limits and Concerns. *Dubai: MBR School of Government* 7 (2017).

[44] Irwin G Sarason and Barbara R Sarason. 1985. *Social support: Theory, research, and applications.* Martinus Nijhoff Publishers.

[45] Elisa Shearer and B Y Jeffrey Gottfried. 2017. News Use Across Social Media Platforms 2017. *Pew Research Center* (2017), 17. http://www.journalism.org/2017/09/07/news-use-across-social-media-platforms-2017/

[46] Kate Starbird, Jim Maddock, Mania Orand, Peg Achterman, and Robert M. Mason. 2014. Rumors, False Flags, and Digital Vigilantes: Misinformation on Twitter after the 2013 Boston Marathon Bombing. *iConference 2014 Proceedings* (2014), 654–662. https://doi.org/10.9776/14308

[47] Kate Starbird and L Palen. 2010. Pass it on?: Retweeting in mass emergency. *Proceedings of the 7th International ISCRAM Conference* December 2004 (2010), 1–10. https://doi.org/10.1111/j.1556-4029.2009.01231.x

[48] Jeannette Sutton, C. Ben Gibson, Nolan Edward Phillips, Emma S. Spiro, Cedar League, Britta Johnson, Sean M. Fitzhugh, and Carter T. Butts. 2015. A cross-hazard analysis of terse message retransmission on Twitter. *Proceedings of the National Academy of Sciences* 112, 48 (2015), 14793–14798. https://doi.org/10.1073/pnas.1508916112

[49] Sudha Verma, Sarah Vieweg, William J Corvey, Leysia Palen, James H Martin, Martha Palmer, Aaron Schram, and Kenneth M Anderson. 2011. Natural language processing to the rescue? Extracting "Situational Awareness" tweets during mass emergency. In *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media*. 385–392.

[50] Sarah Vieweg, Amanda L. Hughes, Kate Starbird, and Leysia Palen. 2010. Microblogging During Two Natural Hazards Events: What Twitter May Contribute to Situational Awareness. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 1079–1088. https://doi.org/10.1145/1753326.1753486

[51] Sarah Elizabeth Vieweg. 2012. *Situational Awareness in Mass Emergency : A Behavioral and Linguistic Analysis of Microblogged Communications*. Ph.D. Dissertation. University of Colorado.

[52] PJ Vogt. 2016. Stolen Valor. 70 (jul 2016). https://gimletmedia.com/episode/70-stolen-valor/

[53] Katrin Weller and Katharina E Kinder-Kurlanda. 2015. Uncovering the Challenges in Collection, Sharing and Documentation: The Hidden Data of Social Media Research? *Proceedings of the Ninth International AAAI Conference on Web and Social Media* 2014 (2015), 28–37.

[54] Wikipedia. 2017. List of 21st-century earthquakes — Wikipedia{,} The Free Encyclopedia. https://en.wikipedia.org/w/index.php?title=List{_}of{_}21st-century{_}earthquakes{&}oldid=807624566

[55] Chao Zhang, Guangyu Zhou, Quan Yuan, Honglei Zhuang, Yu Zheng, Lance Kaplan, Shaowen Wang, and Jiawei Han. 2016. GeoBurst: Real-Time Local Event Detection in Geo-Tagged Tweet Streams. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '16)*. ACM, New York, NY, USA, 513–522. https://doi.org/10.1145/2911451.2911519

Received April 2018; revised July 2018; accepted September 2018