Available online at www.sciencedirect.com

SciVerse ScienceDirect

Procedia Engineering 00 (2014) 000–000

Procedia Engineering

www.elsevier.com/locate/procedia

Humanitarian Technology: Science, Systems and Global Impact 2014, HumTech2014

Novel Situational Information in Mass Emergencies: What does Twitter Provide?

Haji Mohammad Saleem[a,*], Yishi Xu[a], Derek Ruths[a]

[a]School of Computer Science, McGill University, Montral, QC H3A 0G4, Canada

Abstract

It is widely held that social media is a valuable source of information for responders to mass emergencies. However, little is known about how much useful situational information social media data streams contain and how this content varies across types of events. In this paper, we conduct a thorough investigation of the novel situational information provided by public Twitter posts produced during and immediately following three fundamentally different disasters: the 2012 Newtown Connecticut shooting, 2013 Moore tornado, and 2013 Alberta floods. Our analysis reveals that these events yield data streams that contain quite different types and amounts of situational information, generated by different kinds of individuals with varying relationships to the event of interest. Furthermore, we find that, across all events, event-specific hashtags are rarely used in tweets that carry novel information, suggesting that event-specific hashtag/keyword filters miss much novel information. Collectively, these findings highlight a number of factors that can inform how social media is used to facilitate response and recovery efforts for specific mass emergencies.

Selection and peer-review under responsibility of the Organizing Committee of HumTech2014.

Keywords: Crisis Informatics; Humanitarian Data Mining; Social media; Situational Awareness; Twitter.

1. Introduction

During and in the immediate aftermath of a mass emergency, responders can have significant difficulty assessing conditions across affected areas and populations. This makes it hard to determine how best to prioritize, invest in, and deploy supplies (e.g., food, water, and shelter) and services (e.g., sanitation, medical, and law enforcement personnel), particularly to the hardest hit locations and communities. Countering this information vacuum requires the acquisition and interpretation of details about affected areas with the goal of assembling an accurate picture of the situational context across the region of interest. This heightened state of knowledge is called situational awareness [1].

The potential for social media to aid in the generation of situational awareness has attracted well-deserved attention in recent years. However, while much work has focused on building machinery to extract informational posts (e.g., [2–4]) and on characterizing the general kind of information such posts contain (e.g., [5–7]), the literature is silent on the question of how much information produced by social media will be useful to responders. In order to approach this topic, we introduce the concept of novel situational information (NSI).

* Corresponding author. Tel.: +1-438-831-1990.
E-mail address: haji.saleem@mail.mcgill.ca

1877-7058 © 2014 The Authors. Published by Elsevier Ltd.
Selection and peer-review under responsibility of the Organizing Committee of HumTech2014.

By novel information, we mean information that has not been previously reported elsewhere. We submit that one major (though certainly not the only) use of social media by responders is to obtain situational details that are novel: if social media is only telling responders what they already learned from other sources (e.g., the news), then the value of the data stream is marginalized (though certainly not nil since validation of other sources can be useful). For this reason, understanding the conditions under which social media will yield novel, actionable information is an exceedingly important question that will influence how much and what kind of attention responders should give social media data when handling a given mass emergency. Admittedly, we recognize that the concept of ”novel situational information” is not an earth-shattering insight. Nonetheless, while existing papers have shown that social media can provide situational information (e.g., [7–9]), to our knowledge, none have systematically investigated whether this information is strictly novel.

Thus, the goal of this study is two-fold: to establish (1) the extent to which novel situational information is present during a mass emergency in a social media data stream and (2) how the particular properties of an unfolding emergency influence the kind of NSI that is posted to social media platforms. Motivated by the prevalent use and availability of Twitter data to crisis response researchers and practitioners (see every reference save [10]), here we focus on the presence of situational information in the Twitter decahose. This said, while we focus on Twitter, many of our findings can be generalized to other social media platforms.

Our findings and conclusions are informed by case studies which investigate the NSI present in the Twitter feeds corresponding to three separate, recent crises: the 2012 Newtown shooting, 2013 Moore tornado, and 2013 Alberta floods. For each event, we identified a core set of situational details that would have been useful for responders. We then determined how many of these details were reported by social media before traditional news sources (a reasonable proxy for information already being widespread knowledge). This focus on discovering specific details in social media with an emphasis on novelty is a crucial distinction between the present and past work.

From the outset, we underscore that our intent in selecting these specific events was not to suggest that social media analysis would have aided responders to these particular events: the Oklahoma tornado was anticipated and was tracked by radar, law enforcement had a rapid and comprehensive response to the Newtown shooting, and flooding in Alberta could be somewhat quantified by satellite and local measurements. We have selected these events because they have well-documented timelines and damage reports, allowing us the ability to compare the Twitter record to real facts and events as they actually occurred.

Moreover, in analyzing these events, we treat them as proxies for similar kinds of events. The Newtown shooting proxies for violent events in which casualties are a foremost concern (e.g., the 2013 Westgate mall attack and the 2008 Mumbai attacks). The Oklahoma tornado proxies for events characterized by a hazard with focused and potentially several points or paths of destruction such as hurricanes, tornados, and bombings. The Alberta flooding proxies for events characterized by hazards with a large radius of destruction that is both diffuse and highly graduated (ranging from no damage to complete destruction) such as the 2011 Japanese tsunami and subsequent Fukushima nuclear disaster, the 2010 Haiti earthquake, and the 2011 England riots. Of course, it is unlikely that any one event can completely proxy for another event. Nonetheless, our analysis highlights a number of features and findings that enrich and, in some cases, challenge current wisdom concerning the use of social media in crisis response. More broadly, we consider the retrospective comparison of the Twitter record with the known record of a crisis to be an under-utilized, yet promising and generalizable paradigm for improving the relevance of social media to responders.

Our most fundamental finding is that available novel, situational information in social media is deeply affected by the way in which an emergency influences where people are and how willingly they use their cell phones (to access social media). Furthermore, novel situational information can be elusive: across all three events and the different aspects of situational awareness considered, we noted that the first tweets carrying situational information tended to lack the kind of identifying keywords and hashtags that would make them easy to discover in a full Twitter stream. Finally, the detailed analysis we conduct of each event provides insights into the kinds of NSI that particular classes of events may be expected to generate. An exciting direction of research that emerges from our work is the finding that there are clear, relatively easy-to-recognize features of emergencies that can be used to anticipate the kinds of novel situational information that a social media data stream could contain for a given event. We have identified several such features in this work and submit that flushing out these and discovering new distinguishing features is an important direction for future work.

2. Prior Work

Significant work has been done on the general topic of situational awareness obtained from social media in crises. Relevant to the findings of the present study, a notable limitation of most (if not all) existing studies is their dependence on datasets that were collected by sampling the Twitter feed using a small set of hashtags and keywords relevant to the event in question. As we will show throughout this study, many tweets with situational information carry no hashtags and often fail to even use keywords that might be expected (e.g., ”tornado” and ”forest fire”, etc.). Thus, despite some prior work on the topic of situational awareness, we expect that existing results may be hampered by the utter absence of the crucial content of interest (this is not to say that information may not be present, of course, only that it will not be detected when it is first reported, which is important insofar as timeliness of response and reconstruction of event progress is valuable).

Large-scale signatures of emergencies in social media. In recent years, a number of studies have conducted analyses of the large-scale signatures of crises in social media content (e.g., [6]). While these studies certainly offer insight into the cyber-projection of real-world disasters, they do not approach the core question of the present study: the extent to which this projection holds actionable information for responders.

Specific event analysis. Several studies have taken a case study-based approach, similar to that of the present work, in order to characterize mass emergency-relevant content posted to Twitter. A number of these studies considered the question of information propagation and sharing - irrespective of whether the information was both timely and actionable for responders ([5]; [7]; [11]). One study considered situational awareness in a relatively rural context, which was characterized by citizen tweets that significantly lagged news sources, confirming that, in population-sparse locations, social media is not a rich source of timely information ([9]). A recent study considered the information content of tweets made by medical teams responding to the 2010 Haiti earthquake ([12]). A key differentiator with the present study is that their analysis looked specifically at responders who were already in the field, which constitutes a very different population and timescale than we are interested in here. Another study related to the current work looked at community attempts to determine the victim list for the 2007 Virginia Tech shootings ([10]). A key distinction with our work, though, is the primary focus of this earlier work on the collective problem solving aspect of the victim list determination. As a result, very little is said about the situational information that the exercise produced.

Event detection and analysis. Another class of prior work has considered the problem of identifying tweets that pertain to a specific disaster (e.g., [2,13]) and that contain situational information (e.g., [3,4,8,14]). Our present study seeks to support such work by determining under what conditions (what kinds of crises and what kinds of situational information) such systems will yield actionable information for responders.

3. Events, Datasets and Methods

Our study considered the three distinct events. For each, our goal was to characterize how a specific kind of crucial and important information about (and somewhat unique to) the event was first reported in social media, comprising instances where social media would have provided useful situational information to responders. Thus, for each event, we required a Twitter stream spanning the event's time period. We used the Twitter decahose to collect data streams for each of the following events.

2012 Sandyhook Elementary School shooting in Newtown. A 20 year-old man forced his way into Sandyhook Elementary School and opened fire on some of the youngest pupils and their teachers, which ended in the deaths of 26 people including 20 children. The attack itself lasted only about 10-12 minutes, concluding with the shooter taking his own life. Crucially, the names of the victims were released 31 hours after the event itself and, in the intervening time, the names of victims were the basis for much speculation.

Given the short and intense nature of the event, our focus in this event was recovering the victim list. To be clear and to underscore our respect for the victims of this event, our intention here was not to determine whether voyeuristic reporting would have been possible. Rather, in violent events, it is important for responders to be able to establish the scope and, where possible, the specific victims of the attack. A social media-derived victim list would be an effective tool for assessing such information. Due to our focus on the reporting of victims, the dataset timeframe for this event spanned 9 AM Dec 14 to 4:35 PM Dec 15 (when the official victim list was posted by the police).

2013 Moore, Oklahoma tornado. On May 20, 2013, an EF5 tornado struck Moore, Oklahoma and nearby counties. The tornado touched down at 2:56 PM and remained on the ground for 40 minutes as it passed through the downtown core and residential areas of Moore. The tornado claimed 23 lives and injured 377. In this crisis, we focused on mapping the path of the tornado itself and points of the most extreme destruction, particularly through downtown Moore where most damage was done. The motivation behind this detail concerns the ability to provide responders with an accurate assessment of corridors of greatest destruction. Given our focus on the path, we collected a dataset that spanned 2-4 PM on May 20, encompassing the time the tornado was on the ground.

2013 Alberta, Canada floods. Several days of intense rainfall in regions surrounding Calgary resulted in a rapid and extreme increase in the flow of rivers in the area. Widespread flooding, including in the downtown core of Calgary, occurred as rivers overflowed their banks. In much of the region, mandatory evacuation orders were in effect over June 20 - 21st, the days of the most intense flooding. Property damage due to flooding was estimated at $1.7 billion in insurable losses and 5 people lost their lives. The situational information of interest here was on the reporting of damage due to flooding. Unlike a tornado, floodwaters rise over time, yielding damage that can worsen. We studied the extent to which Twitter data could be used to chart the progress of rising floodwaters, using a dataset that spanned from June 20 at 11:00 AM until June 21 at 2:40 PM.

4. Twitter Datasets and Methods

Our dataset for each event consisted of the complete contents of the Twitter decahose (10% of all public tweets) collected over the event's time range. As a result the vast majority of tweets in each event's dataset were unrelated to the event of interest. Removing or otherwise flagging these unrelated tweets is typically a necessary part of analyzing an event's record on Twitter. We did not perform this kind of curation, specifically to ensure that any and all information about the event was present in the tweets. As a result, our datasets were many orders of magnitude larger than previous studies of events (as reported below).

In order to perform this study, we inverted the traditional approach taken in the literature. Rather than deeply sub-sampling the Twitter stream using stringent hashtag and keyword filters (thereby decreasing false positives, but increasing false negatives), we narrowed our field of interest in the event to a very specific set of details for which we could search precisely in the collected Twitter datasets. By defining precisely the event features of interest in each case, we were able to pinpoint among the first tweets to report on each detail of interest (due to the 10% sampling of the decahose, we cannot guarantee seeing the first tweet, although in most cases it's clear we did and, moreover, this coverage is much more complete than previous studies). We considered such first tweets that arrived before news stories on the event to be a likely source of novel, situational information. Analysis of the specific details for each event was done as described below.

Victim reporting after Newtown shooting. We obtained the victim names from the official list published by the police on the day after the shooting. For each victim, we conducted searches for tweets in the dataset that mentioned the victim by name, nickname, relationship (e.g., cousin), or a range of other possibilities. Note that the decahose was essential to collecting this dataset since the correct victim name keywords could not have been known a priori for collection directly after the shooting. For this dataset, we included all english tweets in the range of 9:05 AM Dec 14 - 4:35 PM Dec 15. This yielded 49.4 million tweets.

Central corridor of destruction during Moore tornado. The path of the tornado was mapped and published soon after the event by a number of news and government sources. We used these sources to reconstruct the path of destruction, identifying along the path particular buildings that were singularly devastated. Among these were two schools, the Warren Theatres, and a medical centre. To identify tweets providing information on the path of the tornado, we searched for all tweets mentioning these locations as well as intersections along and near the reported path of the tornado. As with the Newtown shooting, the decahose was essential in building this dataset since the path of the tornado was not known at the time of the event - thus the necessary location keywords could not have been used as a filter query to the Twitter streaming API. For this dataset, we included all english tweets in the range of 2:00 PM May 20 - 4:00 PM May 20. This yielded 4.1 million tweets.

Gradual damage and flooding during Alberta floods. In order to identify flooding reports, we selected a number of locations (both precise buildings and neighbourhoods/boroughs) that sustained significant flooding. We searched for all tweets that made mention to these locations both directly and by nicknames that we learned through background

Table 1. The first tweets that identify individual Newtown victims. Out of respect for victims and their families, the text of the tweets have been omitted. *Rel* indicates the relationship of tweet author to the victim: N=News, EF=Extended Family, CF=Close Friend, CP=Close Non-family, P=Personal (generic). The generic *Personal* is used when the relationship was clearly non-news, but no additional details could be determined. Notable is the overwhelming lack of hashtags in tweets that first report victims.

Name	First Reference TweetID	Rel	Hashtags
DH	Dec.14 12:48PM 279652860818489344	N	*None*
CB	Dec.14 01:39PM 279792735223951360	EF	*None*
GM	Dec.14 02:57PM 279691710433488897	CP	*None*
CP	Dec.14 03:28PM 279684189375299585	CP	*None*
MS	Dec.14 03:34PM 279700887994236928	EF	*None*
BW	Dec.14 03:55PM 279736331532374016	EF	*None*
CK	Dec.14 05:35PM 279716259321491456	CF	#prayfornewtown
EP	Dec.14 01:39PM 279716259321491456	P	#praysfornewtown
AMG	Dec.14 08:58PM 279782455626567680	N	*None*
VS	Dec.14 11:04PM 279798979384967171	P	#rememberVS
CH	Dec.15 02:52AM 279992291056295936	N	#newtown
JL	Dec.15 04:22AM 279878989609332736	N	#prayfornewtown
DB	Dec.15 07:09AM 279921129769017345	P	#newtown
LR	Dec.15 11:38AM 279988813516922880	N	*None*
OE	Dec.15 01:17PM 280013647869247491	P	*None*
AR	Dec.15 04:20PM 280059823045877760	N	*None*
JR	Dec.15 04:20PM 280059823045877760	N	*None*
NP	Dec.15 04:20PM 280059823045877760	N	*None*
JP	Dec.15 04:27PM 280061639108227074	N	*None*
JG	Dec.15 04:29PM 280062096144740352	N	*None*
RD	Dec.15 04:29PM 280062096144740352	N	*None*
AMM	Dec.15 04:29PM 280062150884593664	N	*None*
DH	Dec.15 04:30PM 280063128056786944	N	*None*
JM	Dec.15 04:30PM 280062159558410240	N	*None*
MH	Dec.15 04:31PM 280062587088015360	N	*None*
AW	Dec.15 04:35PM 280063442805743616	N	*None*

research on the locations themselves (e.g., reading news reports and tweets about the locations). As before, the decahose was necessary here since the presence of flooding in particular locations could only be known after it was reported, which would have necessarily introduced a significant delay in arranging keyword-based monitoring of the Twitter feed for these locations. For this dataset, we included all english tweets in the range of 11:00 AM Jun 20 - 2:40 PM Jun 21. This yielded 58.9 million tweets.

5. Results

5.1. Newtown Shooting

We determined the timing (see Figure 1(a)) and text (see Table 1) of the first mentions of victims of the shooting. A number of features stand out as noteworthy. First, only 15 of the 26 victims were identified before the official police announcement on the afternoon of Dec 15th. Of these mentions, only 10 were personal tweets. The overwhelming majority of these posts were made by close family friends and extended family - crucially, none were made by individuals who would likely had been at the scene of the event. Moreover, the initial postings came well after the event, with the first personal mention arriving over 4 hours after the shooting.

When considering the problem of identifying these tweets as containing valuable situational information, several points are remarkable. First, these 15 early mentions were situated in a much larger corpus of tweets speculating about the names of the individuals, making discerning that these were the true victims much more difficult. Underscoring this fact, consider that the first mention of the principal, DH, includes incorrect speculation of a second gunman involved in the incident. Moreover, the majority of the first mentions (both overall and only personal) do not contain keywords or hashtags that might be used filter them out of the stream (e.g., ”newtown”, ”sandyhook” and ”school”) - again rendering these tweets harder to find at the time (and even quite arduous to find in retrospect).

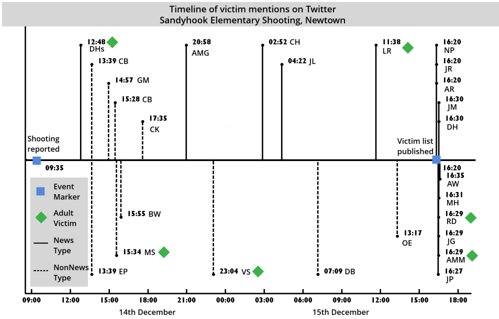

(a)

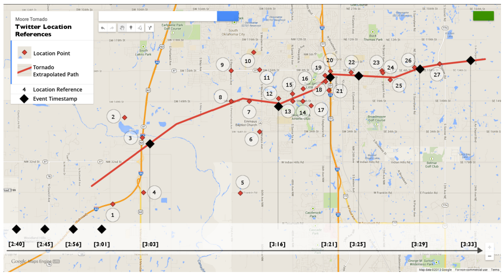

(b)

Fig. 1. (a) The timeline of the first Twitter mentions of victims of the Sandyhook Elementary Shooting. Out of respect, initials are used in place of full names. See Table 1 for the tweets. (b) The path of the tornado that struck Moore, Oklahoma on May 20, 2013. Overlaid are the tweets (see Table 2) that first reported critical points in the path.

5.2. *Oklahoma Tornado*

Our primary goal for this event was to evaluate timely reporting of the path of the tornado. Figure 1(b) provides an overlay of the path the tornado followed with the tweets (Table 2) that clearly reported on the path. As might be expected given the relative scarcity of geotagged content on Twitter, none of the early tweets that provided information on the path of the tornado were geotagged. In fact, the lack of geotagged content was so complete in the decahose that we were unable to find any tweets that (1) were geotagged, (2) mentioned the tornado in any way,and (3) were situated on the (known) path of the tornado. As a result, path-reporting tweets were identified by their mention of specific locations (e.g., place names, streets, and intersections) that clearly corresponded to locations along the tornado's path. Because here we were interested in being able to follow the path of the tornado, we focused on the first tweets to mention locations along the path. The correspondence between the time when a specific path-reporting post was made and when the tornado was known to have passed the point it mentioned provides the ability to know precisely how helpful Twitter data could have been in live tracking of the tornado (though admittedly, this benefit is somewhat theoretical since tracking would have required culling all mentions of locations in the area that were not along the path - again, the goal here is to determine just how good tracking could have been).

A number of features of the path-reporting tweets stand out as noteworthy. First, there is an ample and timely signature of the path of the tornado in Twitter. In general, tweets reporting about the tornado occurred within minutes of the time that it passed (or was about to pass) specific locations. However, many of these timely path-reporting tweets lacked keywords or hashtags that would facilitate their discovery by a standard Twitter filter (echoing a similar issue with the Newtown shooting tweets): 25% of tweets lacked both hashtags and keywords, 70% lacked hashtags. Second, not surprisingly and in agreement with prior work, coverage is much better in population-dense areas, in this case the Moore downtown core [9]. Of concern is the significant delay in the first reporting of the two elementary schools (Briarwood and Plaza Towers), which sustained direct hits. In both cases, nearly 30 minutes elapsed before either were reported (tweets 12 and 16 in Table 2), indicating that during emergencies, even in population dense areas, Twitter coverage does not always prioritize locations where the danger to human life is greatest. Somewhat ironically, more Twitter-based path reporting focused on the Warren Theatres (which were effectively abandoned during the storm) than on the schools, both of which had large student populations present when they were hit.

Also notable is the fact that most path-reporting tweets occur just after the tornado passes, possibly reflecting the fact that people were in hiding as the tornado passed and, therefore, were unable or unwilling to tweet.

5.3. *Alberta Floods*

In this part of the project, we focused on the way that gradual damage is reported in social media. To do this, we looked at personal Twitter posts that reported flooding progress in and around several Calgary neighbourhoods (Bowness and Inglewood) and landmarks (the zoo and stadium). For each location, tweets were identified which reported on the status of flooding at that location. These were ordered by time and a representative sample was taken (since, in all cases, too many location-specific tweets existed to include in the paper and analysis). Due to space

Table 2. Tweets from the Oklahoma Tornado that identify the location of damage along the path of the tornado. Quite notable is the significant delay in the mention of the two schools (Briarwood and Plaza Towers) which sustained direct hits and severe damage (tweets 12 and 16).

#	Location	Tweet	Time
1	Bailey Turnpike	I-44 on bailey turnpike. tornado on ground. watch:⟨*weblink*⟩	2:56
2	SW 156th Street	2:59 huge tornado on the ground near i-44 and sw 156.. take cover now!!!	2:59
3	Highway 37 and I-44	large tornado north of newcastle near hwy 37; i-44 headed toward moore. #okwx (3:00pm) ⟨*weblink*⟩	3:00
4	Newcastle Casino	here we go again. tornado on the ground. newcastle casino is not safe. watch out oklahoma.	2:46
5	Franklin Road	ef 2 and growing n. franklin road ready to move into sw okc and moore	3:00
6	SW 164th Street	debris ball a mile wide..i-44 bridge..south 164th..wedge, violent tornado.	3:05
7	B/W S Pennsylvania Ave. and S Western	police scans: the tornado is between s penn and western in southwest oklahoma city.#okwx	3:13
8	149th St. and S Penn Ave.	3:12 tornado doing terrible damage near sw 149th and penn...homes being blown away..	3:12
9	134th St. and S Penn Ave.	tornado, you stay away from 134th and penn!!!poor tatie and nate!!	3:03
10	Westmoore High	ef4 tornado on the ground with a mile wide debris ball.you live near westmoore school, get below ground! ⟨*weblink*⟩	3:06
11	McDonald's	major damage: massive tornado moving near i-44 near newcastle, okla. "just missed the mcdonald's!"	3:01
12	Briarwood Elementary	It hit briarwood	3:40
13	19th St. and Santa Fe Ave.	rain-wrapped tornado near sw 19th and santa fe in moore moving toward i-35. #okwx (3:18pm)	3:18
14	Moore Golf Course	@user no. down 19th street by the moore golf and athletic club. (part of conversation)	3:40
15	Santa Fe Ave.	oh my that huge tornado is going right towards s.f avenue full of cars that are in a traffic jam. this is horrible!	3:18
16	Plaza Towers Elementary	plaza towers elementary school has been hit.	3:51
17	Dick's Sporting Goods	19th street and more st near the dick sporting goods, home depot, you are in the path of tornado in moore oklahoma	3:17
18	Warren Theatres	it's heading to the warren theater	3:02
19	Moore Medical Center	moore, ok tornado is on intersect with moore medical center	3:15
20	SW 4th St. and I-35	the tornado is crossing i-35 in moore, okla. at the moment. very well-defined velocity couplet. please take shelter! #okwx — wrapped in rain plus debris - it is huge tornado now heading toward i-35. damage on 19th street in moore.	3:19
21	Freddy's Frozen Custard	warren theatre and freddy's frozen custard and steakburgers next to each other at i-35 and 19th st. in moore. kfor video shows some damage	3:58
22	Walmart	so the moore walmart and hospital are gone... everything is flattened damnn	3:50
23	4th St. and Bryant Ave.	4th, 12th, 19th streets - veterans memorial park - large wedge #tornado moving down 4th street has crossed train tracks take cover now #okwx	3:28
24	Veterans Memorial Park	veterans memorial park, moore administration building, has major damage along with the high school has damage. #okwx	3:45
25	Wimberley Creek	tornado approaching creeks of wimberly on 4th street in moore.(3:28 p.m.) #kocowx	3:28
26	Sunnylane Rd.	tornado crossing sunnylane road between 4th and 12th street in moore, ok	3:30
27	Sooner Rd.	#tornado about 1/8 mile from #sooner road by annex to lake. #moore #okwx #okc #oklahoma.grinding on ground 3:31 pm	3:31

limitations, in Figure 2 we show the Twitter content obtained for only one location (Bowness). As we found reporting styles consistent across all locations considered (discussed below), we consider the content shown to be representative.

Tweets reporting flooding shared several distinctive features that distinguished them from damage-reporting tweets posted during the Oklahoma tornado.

Use of photos. Foremost, as can be seen in the example tweets in Table 4, is the ubiquitous use of photos. While photos are not uncommon among tweets, flood-related tweets consistently showed dependence on photos to convey the damage that was being observed/reported. Furthermore, photos seemed to be primarily of two varieties: (1) canned photos taken from official news sites/tweets (this was particularly common with tweets about the zoo and stadium) and (2) photos taken from a convenient location near the person's house.

	*(NO IMAGE)*	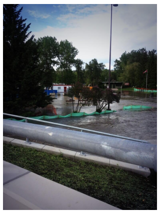	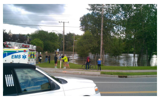		
it's high but not cresting. thinking you're ok for now. guess we'll see what happens #bowcabin #bowness #yyc ⟨*weblink*⟩	bowness park is swimming in a pool right now.	@ampcalgary #yycflood bowness park enterance by paddle boat pond ⟨*weblink*⟩	bowness already flooding. #yyc #abfloods ⟨*weblink*⟩	bowness road. #yyc #flood ⟨*weblink*⟩	Bowness park is a river ⟨*weblink*⟩
4:07 PM	4:49 PM	6:12 PM	7:08 PM	8:16 PM	8:53 PM

Fig. 2. Tweets reporting the status of flooding in and around the Bowness neighbourhood on June 20, 2013. Bowness is a suburban area near Calgary. Notable is the absence of flooding hashtags/keywords as well as the dependence on images to convey current conditions.

Uneven and uncertain coverage of an area. Presumably the reason for disproportionate numbers of ”porch” photos has to do with people's inability or unwillingness to move around during the flood. In fact, during the peak of the crisis, the Calgary municipal leadership regularly instructed citizens to stay in their houses for safety reasons. The combination of these and other factors produced a population of Twitter users who could only post pictures from a fixed viewpoint. Where situational awareness is concerned, this produced a systemic issue: for all the locations considered, individual tweets reported rather random perspectives - both in text and photos. Users did not have the luxury to survey their surrounding area and find particularly informative vantage points or collect particularly meaningful data. This was quite pronounced when looking at neighbourhoods, in which the aggregated tweet stream is primarily a sequence of pictures of users' backyards (which are by-and-large, spatially disjoint).

Use of imprecise language. From a situational awareness perspective, the tweet text (often simply accompanying photos) involved imprecise language that made judging the degree of flooding difficult. Figurative language was often used - in one tweet from Table 3, Bowness park was described as "swimming." Even tweets with valid literal interpretations provided little more than observations of locations being "flooded" or "underwater." Furthermore, tweets that attempted to express more detail often created unintentional ambiguity. For example, a tweet on June 21 at 12:11 PM reported that "the stampede[1] is underwater completely." Clearly, the author of this tweet is attempting to distinguish the flooding status from prior reports that flood water levels had reached the Calgary Stampede Park. However, how are we to interpret "underwater completely"? To be concrete, certainly the stampede audience risers (which stand well over 50 feet high) were not underwater completely. Has the user determined that all level ground in the stampede park is underground? Interpretive issues such as this created by imprecise descriptions of flooding damage, in combination with photos taken from different perspectives (and even of spatially distinct parts) of a location made proper assessment of flooding progress exceedingly difficult.

At small scale, human interpretation might be able to accommodate the factors identified above: humans can certainly interpret the contents of photographs, make reasonable conjectures about imprecise phrases like "underwater completely", and — with appropriate local knowledge — assess what a relatively random set of reports and pictures of flooding in a neighbourhood might mean for nearby streets and freeways. However, where large scale processing and computational assembly of situational awareness is concerned, these issues present serious challenges: computers do not handle arbitrary static or video footage well, have difficulty with even small ambiguities in human language, and typically lack the kind of local knowledge (and algorithmic capacity) to extrapolate environmental conditions from unstructured information from vague viewpoints. These challenges both flag the kinds of issues responders should

[1] The Calgary Stampede Park is a large building and fair-ground complex in the downtown core of Calgary.

expect to encounter when using social media data from such events and also identify important research directions in need of immediate attention.

6. Discussion

In this study, we have conducted the most comprehensive analysis to date of the novel situational information provided by Twitter, a social media platform that has often been used to track unfolding crises. Our findings yield a number of insights that both can improve the ways in which social media is used to inform mass emergency responders and also flag key methodological questions for future work on social media-based situational awareness systems.

Absence of keyword-based tagging. We were quite surprised to discover that many of the tweets that first reported actionable situational information (e.g., victim names and points of damage) did not include keywords or hashtags that would have made their discovery through standard Twitter filters possible. Admittedly, with the right choice of keywords, they would have been included, though presumably often along with many irrelevant tweets. Our point, however, is that the most obvious choice of keywords (e.g., the place, name of the disaster, and most common hashtags for the event) would not have been sufficient. As this is the way in which event tweets are typically harvested both in practice and in academic studies, such datasets will apparently lack many first reports of useful situational information. This strongly suggests that additional work is needed on developing more sophisticated filters for detecting emergency-relevant posts. We expect that this direction will involve computationally interesting challenges and also deliver valuable tools for academicians and responders.

Injuries and casualties are not reported quickly. Casualties - both injuries and deaths - occurred during all three of the events studied. With the benefit of hindsight knowledge of the names of the individuals killed, we looked for first mentions of the victims. In the case of the tornado and floods, we found no mentions at all during the time periods sampled, suggesting that Twitter users do not readily post this kind of information. A study of victim reporting in the Newtown shooting offers additional insight into why such information is omitted. In this event we found that victim names were reported by extended family or personal friends - individuals who were (at least in most cases) not at the scene during the crisis, did not receive the news firsthand, and were not directly involved in managing the aftermath of the event. Our hypothesis is that, to post such tragic information requires a certain amount of emotional distance which can be created either by time or, more likely, by degrees of separation from the victim themselves. As a result, we consider social media to be an unreliable source of information on emergency-related casualties.

Population mobility and access to cell phones impacts coverage. In both the floods and the tornado, we found evidence that the prevalence of populations with restricted mobility or limited access to cell phones impacted the quality of coverage of key features of the event. In the tornado, two elementary schools that sustained direct hits were not mentioned for nearly 30 minutes - whereas the destruction of the popular but (at the time, empty) Warren Theatres downtown was covered in nearly real time by Twitter posts. Two factors, which likely contributed to this disparity, are the relative rareness of smartphones (and Twitter accounts) among elementary school children and the fact that, at the time of the tornado, everyone in the school was likely preoccupied with surviving and then recovering from the damage done to the school. While obvious, it is important to remember that, when people are dealing with an immediate threat to their own safety, they will not post to social media.

The floods manifested this principle differently. Because movement was difficult, most individuals were trapped in a single location for an extended amount of time. Thus, while many individuals chose to post on their situation, they were unable to seek the most interesting or informative vantage point. Thus, most posts on the flood involved reports on quite random and unremarkable parts of the effected area. This kind of sampling made the task of obtaining broader situational awareness quite difficult. We suspect that only responders with expert knowledge of the areas in question might be able to use such information to achieve any degree of proper situational awareness, and this would be at small scale given the number of event-related tweets that might be mined for this information. It is reasonable to expect that this will be the issue whenever a mass emergency severely restricts the mobility of the affected population.

Tweets can ambiguously report degrees of damage. The floods also highlighted the general inadequacy of language (or, at least, tweeting conventions) to meaningfully quantify degrees of damage. While we know that flood levels changed in areas and that people continued to post about these areas over time - the language used to describe flooding conditions was not precise enough to allow fair assessment of these changes. Quite to the contrary, figurative and vague language often made even gross estimates of damage difficult. Furthermore, Twitter users often favoured using pictures to communicate their situation. While photos can be helpful, computers do not handle them well and even humans have difficulty determining degrees of damage from photos taken from different and irregular orientations.

As clearly illustrated by the Alberta floods case study, using social media to gain situational information about diffuse damage (damage with no clear epicentre) is quite difficult (e.g, floods and wildfires). Of course, there are some degrees of damage that are easy to parse (e.g., the "roof has been ripped off" a house). Thus, for some disasters and some forms of damage, social media will have significant information of this sort to provide. Formalizing our understanding of how different kinds of damage will be reported is a direction for future work, which could deeply inform effective use of social media in mass emergencies.

Factors affecting NSI in social media. Our analysis suggests a number of factors that influence the novel situational information that a social media stream can provide: (*danger*) people will not post purely informational content when their immediate safety is threatened; (*cell-equipped*) populations that have limited access to cell phones exist will yield reduced coverage of their environments; (*emotional distance*) casualties will only be reported by those who have achieved sufficient emotional distance from the victims' death or injury; (*mobility*) populations that have limited mobility will yield more opportunistic and less thorough coverage of key points of interest; (*gradual disaster*) damage that changes over time is difficult for populations to report in a systematic and interpretable way; (*diffuse damage*) disasters characterized by widespread, homogenous damage are difficult for social media to accurately report on.

Finally, we suggest that the inversion of event analysis — identifying known actionable ground truth and then searching for it in social media data — is a powerful strategy for overcoming the limitations of datasets collected using keyword-based filters and for focusing on specific pieces of situational information that hold particular strategic value for responders.

Acknowledgements

We thank Twitter for decahose access, which made this study possible. Most importantly, we recognize those who endured and, in some cases, perished, in the midst of the crises we have studied. To you, we humbly offer this work as discovering in your struggles some insights to better the fates of victims of future disasters.

References

[1] Endsley, M. R. 1995. Toward a theory of situation awareness in dynamic systems. Human Factors: The Journal of the Human Factors and Ergonomics Society 37(1):32-64.

[2] Sakaki, T.; Okazaki, M.; and Matsuo, Y. 2010. Earthquake shakes twitter users: real-time event detection by social sensors. In Proceedings of WWW, 851-860.

[3] Gao, H.; Barbier, G.; and Goolsby, R. 2011. Harnessing the crowdsourcing power of social media for disaster relief. Intelligent Systems, IEEE 26(3):10-14.

[4] Imran, M.; Elbassuoni, S.; Castillo, C.; Diaz, F.; and Meier, P. 2013. Practical extraction of disaster-relevant information from social media. In Proceedings of WWW, 1021-1024.

[5] Heverin, T., and Zach, L. 2010. Microblogging for crisis communication: Examination of twitter use in response to a 2009 violent crisis in the seattle-tacoma, washington area. In Proceedings of ISCRAM.

[6] Hughes, A. L., and Palen, L. 2010. Twitter adoption and use in mass convergence and emergency events. International Journal of Emergency Management 6(34):3-4.

[7] De Longueville, B.; Smith, R. S.; and Luraschi, G. 2009. Omg, from here, i can see the flames!: a use case of mining location based social networks to acquire spatio-temporal data on forest fires. In Proceedings of the International Workshop on Location-based Social Networks.

[8] Verma, S.; Vieweg, S.; Corvey, W. J.; Palen, L.; Martin, J. H.; Palmer, M.; Schram, A.; and Anderson, K. M. 2011. Natural language processing to the rescue? extracting situational awareness tweets during mass emergency. In Proceedings of ICWSM.

[9] Vieweg, S.; Hughes, A. L.; Starbird, K.; and Palen, L. 2010. Microblogging during two natural hazards events: what twitter may contribute to situational awareness. In Proceedings of CHI, 1079-1088.

[10] Vieweg, S.; Palen, L.; Liu, S. B.; Hughes, A. L.; and Sutton, J. 2008. Collective intelligence in disaster: An examination of the phenomenon in the aftermath of the 2007 virginia tech shootings. In Proceedings of ISCRAM.

[11] Mendoza, M.; Poblete, B.; and Castillo, C. 2010. Twitter under crisis: Can we trust what we rt? In Proceedings of the first workshop on social media analytics, 71-79.

[12] Sarcevic, A.; Palen, L.; White, J.; Starbird, K.; Bagdouri, M.; and Anderson, K. 2012. Beacons of hope in decen- tralized coordination: learning from on-the-ground medical twitterers during the 2010 Haiti earthquake. In Proceedings of CSCW, 4756. ACM.

[13] Chew, C., and Eysenbach, G. 2010. Pandemics in the age of twitter: content analysis of tweets during the 2009 h1n1 outbreak. PLoS ONE 5(11):e14118.

[14] Starbird, K.; Muzny, G.; and Palen, L. 2012. Learning from the crowd: Collaborative filtering techniques for identifying on-the-ground twitterers during mass disruptions. In Proceedings of ISCRAM.