Hyperlocal Toponym Usage in Storm-Related Social Media

Rob Grace
Texas Tech University
rob.grace@ttu.edu

ABSTRACT

Crisis responders need to locate events reported in social media messages that typically lack geographic metadata such as geotags. Toponyms, places names referenced in messages, provide another source of geographic information, however, the availability and granularity of toponyms in crisis social media remain poorly understood. This study examines toponym usage and granularity across six categories of crisis-related information posted on Twitter during a severe storm. Findings show users often include geographic information in messages describing local and remote storm events but do so rarely when discussing other topics, more often use toponyms than geotags when describing local events, and tend to include fine-grained toponyms in reports of infrastructure damage and service disruption and course-grained toponyms in other kinds of storm-related messages. These findings present requirements for hyperlocal geoparsing techniques and suggest that social media monitoring presents more immediate affordances for course-grained damage assessment than fine-grained situational awareness during a crisis.

Keywords

Volunteered Geographic Information, Twitter, Information Behavior, Crisis Informatics, Emergency Management.

INTRODUCTION

To effectively monitor social media, crisis responders require geographic information in the metadata or content of social media messages to locate events within an operational jurisdiction and gather actionable information to support response operations. The nature of the response, in turn, determines the granularity of geographic information required by crisis responders. Emergency managers, for instance, can make use of course-grained information (e.g., city name), to assess geographically dispersed impacts of disasters such as floods, tornados, and hurricanes, while emergency dispatchers and first responders require fine-grained information (e.g., building name) to respond to localized incidents.

However, despite crisis responders' need to locate and characterize events reported on social media, geographic information is often scarce in social media messages. In the case of Twitter, responders often rely on geographic metadata tagged to tweets (i.e., geotags) to collect and analyze crisis social media despite consistent observations that only 1% of tweets include geotags (Carley, Malik, Landwehr, Pfeffer and Kowalchuck, 2016; Morstatter, Pfeffer, Liu and Carley, 2013). Looking "beyond the geotag" reveals other sources of geographic information (Crampton et al., 2013), such as the place names, or toponyms, users include in social media messages. While the availability of geotagged social media data has been extensively studied (Carley et al., 2016; Huang and Carley, 2019), far less attention has been devoted to toponym usage during a crisis, especially usage of fine-grained place names that refer to events at hyperlocal levels of analysis, e.g., buildings, street intersections, etc. Moreover, automated techniques to identify and locate events reported on social media often rely on techniques that extract only coarse-grained, municipal-level place names while overlooking fine-grained, common, and colloquial place names users often use when describing local events (Middleton, Kordopatis-Zilos, Papadopoulos and Kompatsiaris, 2018; Wang and Ye, 2018). Consequently, beyond the contributions of a few studies (Avvenuti, Cresci, Del Vigna and Tesconi, 2016; Middleton et al., 2018; Middleton, Middleton and Modafferi, 2014), characteristics of toponym usage and granularity in crisis social media remain poorly understood.

This study examines toponym usage and granularity on Twitter during a severe storm and tornado that struck the

Northeastern United States to examine the extent to which users include toponyms in tweets posted during the storm (toponym usage), reference places inside or outside a geographic jurisdiction (toponym location), and provide course and fine-grained geographic information (toponym granularity). Toponym usage, location, and granularity are compared between storm-related and non-storm-related messages, as well as across six categories of storm-related information: warnings, forecasts, experiences, updates, damages, and disruptions.

The study presents three primary findings with implications for research in crisis informatics (Reuter, Hughes and Kaufhold, 2018). First, analysis of toponym usage shows that crisis-related social media includes disproportionate amounts of geographic information compared to social media messages about other topics. Second, analysis of toponym location shows that while crisis social media includes reports of geographically dispersed events, local events tend to be described using toponyms, not geotags. Third, analysis of toponym granularity shows that the granularity of place names in crisis social media varies by information category: warnings, forecasts, and accounts of personal experience mostly include course-grained toponyms referring to cities and counties, while reports of infrastructure damage and service disruption often include fine-grained toponyms referring to places within cities and counties.

The following sections discuss research on crisis responders' geographic information requirements and the availability of geographic information in crisis social media, including prior research on toponym usage and granularity in social media. The next sections outline methods of data collection and analysis, present findings, and discuss their implications for crisis informatics research.

GEOGRAPHIC INFORMATION REQUIREMENTS FOR CRISIS SOCIAL MEDIA

Studies have shown that, in nearly all cases, crisis responders require useful geographic information to assess and act on information reported on social media (Kropczynski et al., 2018; Zade et al., 2018). Responders' initial priority is to assess if an event is located inside or outside an area of operations or jurisdiction. As Zade et al. (2018) explain, "Information about the location of a disaster event places it either within a responder's geographic range of operability and helps her assess its actionability, or allows her to redirect the information to a co-located responder authority" (p. 10). Responders can then use course or fine-grained geographic information for situational awareness during response operations. However, as jurisdictions and operational roles vary widely among crisis responders, the granularity of geographic information responders require varies too.

For emergency responders typically serving in municipal-level jurisdictions, Kropczynski et al. (2018) observe that prospective uses of social media depend on the availability of precise geographic information that allow 911 dispatchers to route first responders to the scene of an emergency. During participatory design workshops in which dispatchers, police officers, firefighters, and paramedics composed examples of social media messages providing actionable information, each hypothetical tweet included hyperlocal references to "intersections, buildings... and in some cases, specific street addresses" (p. 5). As Kropczynski et al. conclude, "Simply put, without an incident location it is very difficult to provide a response" (p. 6).

In contrast, interview and survey studies with emergency managers and disaster responders highlight uses of social media that require relatively coarse-grained geographic information to conduct rapid impact assessments of infrastructure damage, flooding, and other geographically dispersed events at the outbreak of large-scale emergencies or disasters (Grace, Kropczynski and Tapia, 2018; McCormick, 2016; Reuter, Ludwig, Kaufhold and Spielhofer, 2016). These officials seek to collect and categorize social media messages to measure and compare crisis impacts across a geographic jurisdiction to inform decision-making related to resource allocation, emergency declarations, and requests for assistance (Giacobe and Soule, 2014; Hiltz et al., 2019; Plotnick and Hiltz, 2018).

These studies, in turn, provide general (i.e., location) and specific (i.e., granularity) geographic information requirements for the design of social media monitoring systems and workflows that can support crisis responders in different roles. Moreover, these requirements suggest that the affordances of social media for crisis response — which responders can use it and for what — hinge on the availability and quality of geographic information users post on social media platforms.

VOLUNTEERED GEOGRAPHIC INFROMATION IN CRISIS SOCIA MEDIA

The availability of geographic information in social media shapes opportunities for monitoring crisis events in a geographic area. Twitter, for instance, offers two types of geographic information: geotags, geographic coordinates that users and devices assign to tweets, and toponyms, place names users include in the content of tweets. The scarcity of geotagged tweets is well evidenced, with studies routinely observing that approximately 1% of all tweets include geotags (Carley et al., 2016; Crampton et al., 2013; Morstatter et al., 2013). This scarcity also characterizes tweets posted during crises. Middleton et al. (2014) observe that geotagged tweets account for only 1% of crisis-related messages posted during hurricane, earthquake, and blackout events. Moreover, studies highlight the possibility for systematic bias when gathering information only from geotagging users (who are more likely to be younger and live in urban centers compared to the general user population) (Hecht and Stephens, 2014), or inaccuracy when naively using geo or placetags to locate events described, in tweet content, as remote (Shelton, Poorthuis, Graham and Zook, 2014; Tasse, Liu, Sciuto and Hong, 2017). Consequently, Crampton et al. (2013) have led the call to move "beyond the geotag" by utilizing alternative geographic information available on Twitter, namely toponyms.

Toponym Usage in Crisis Social Media Data

Toponyms are place names that are either specific, referring to a unique place (e.g., 126 Belmont Road), or generic, referring to a class of place or instance of that class (e.g., my house). Toponyms can be further distinguished by location, referring to local or remote places inside or outside a geographic area of interest, respectively. Lastly, toponyms differ by granularity. References to buildings, street intersections, and other hyperlocal places provide fine-grained location information, while the names of towns, regions, and countries provide course-grained location information.

Motivated by the lack of geotagged data (Carley et al., 2016; Kitamoto and Sagara, 2012; Middleton et al., 2018), researchers have turned to toponyms to locate and characterize events reported on social media platforms. These studies focus on toponym extraction techniques to identify place references in social media messages and toponym resolution techniques to match extracted place entities with unique geographic areas, often using databases, or gazetteers, of place names and associated latitude/longitude coordinates. However, the effectiveness of toponym extraction and resolution methods remains contingent on the availability of toponyms in crisis social media: the extent to which users reference places in social media posts (toponym usage), refer to places of interest to crisis responders (toponym location), and provide detailed geographic information that can support "actionable" decision-making (toponym granularity).

Toponym Usage

Studies hint at differences in toponym usage between crisis-related and non-crisis-related tweets. Cheng et al. (2010) collected 29,479,600 tweets to find 4,124,960, or 14%, referencing a city listed in the 2000 U.S. Census Gazetteer. Kitamoto and Sagara (2012) observe that while the number of tweets posted during weather events typically increases, the percentage of tweets including toponyms remains stable at approximately 5-10% of all posted tweets. Elsewhere, Salas et al. (2018) find that, of 13,410 tweets collected using traffic-related keywords, 1,161, or 8.6%, referred to a roadway or roadside location.

Among crisis-related tweets, however, toponym usage appears to increase significantly. Of 55.1 million tweets including flood-related keywords, de Brujin et al. (2018) found 19.2 million, or approximately 35%, mention specific places at the country, regional (e.g. province), and municipal levels. Manually coding a subset of the flood-related tweets, de Bruijin found 53% to include at least one toponym (p. 9). Elsewhere, Middleton et al. (2014) used event-related keywords to collect 92,300 tweets posted after a 2013 tornado struck Moore, Oklahoma, of which 42,434, or 46%, were found to contain some kind of location reference. Similarly, Avvenuti et al. (2016) observe that 35-58% of crisis-related tweets posted during two Italian disasters contain at least one toponym.

However, these studies do not report differences in toponym usage among crisis-related messages— a diverse yet generally consistent set of information ranging from messages of support and prayer to situational reports of infrastructure damage and requests for help (Olteanu, Vieweg and Castillo, 2015). Dunkel et al. (2019), for example, note that place names frequently appear in descriptions of damage tagged to photographs uploaded to Flickr during the 2013 St. Jude Storm, but do not consider the extent to which toponyms appear in other kinds of crisis-related posts during the storm. Understanding how toponyms are distributed across crisis-related messages would reveal the events that prompt users to communicate geographic information and, in turn, requirements for identifying and locating events reported on social media during crisis response operations.

Toponym Granularity

Toponyms are further distinguished by spatial granularity, referring to fine-grained locations such as individual buildings and street intersections, or coarse-grained locations such as municipalities, regions, and countries. The granularity of toponyms in crisis social media create opportunities for social media use in crisis response: fine-grained toponyms can support emergency response, while course-grained (and fine-grained) toponyms can inform, for example, impact assessments in emergency management. Whether crisis responders can exploit these opportunities, however, hinges on effective techniques to identify toponyms in crisis social media and associate extracted toponyms with hyperlocal places within an operational jurisdiction. Although studies develop toponym extraction and resolution techniques for crisis management (Avvenuti, Cresci, Nizzoli and Tesconi, 2018; Dredze, Paul, Bergsma and Tran, 2013; Halterman, 2017; Middleton et al., 2014), there remains need to understand the diversity and distribution of generic and fine-grained toponyms in crisis social media in order to inform requirements for geolocation techniques that meet the needs of crisis responders, particularly those serving in municipal-level jurisdictions.

Generic toponyms, including vernacular and ambiguous place references often used among local people, pose an obstacle for automated geolocation techniques. Despite their prevalence, generic toponym usage has not yet been examined:

> More importantly, compared with the accurate and widely used geo-coordinates and toponyms, some place names that have vague boundaries but are largely communicated by social media users have been rarely explored in disaster contexts… 'downtown' and 'city center' are often used to represent the core area of a city… These 'vernacular place names'... could be utilized to better enrich the useful geographic information for natural disaster management. (Wang and Ye, 2018, p. 51)

Moreover, geolcation methods often overlook generic toponyms such as "home," "building," or "office" to identify unique, relatively rare words and proper nouns that prove useful for toponym recognition. De Brujin et al. (2018), for instance, not only employs a gazetteer listing only the names of villages, towns, and greater administrative areas (e.g. counties, provinces, etc.), but discards uni and bi-grams containing any of the 1000 most frequently used words in a language (according to a corpus of Wikipedia articles) when identifying candidate locations in a tweet's text (p. 5).

Moreover, toponym resolution techniques perform poorly when relying on available, coarse-grained gazetteers to geolocate place names extracted from crisis social media. Middleton et al. (2018) compare state-of-the-art geoparsing techniques using tweets posted during three terrorist attacks— 13 November 2015 Paris shootings, 22 March 2016 Brussels airport bombing, and the 8 May 2016 Turkish police station bombing— to find best performance when recognizing regional toponyms and the poorest when recognizing ground-truth building and street names assembled for each incident. For the latter, Middleton et al. observed best performance among approaches making use of detailed street and building data for the geographic areas of the terror attacks.

The global and regional-level collection and analysis of social media data and routine use of gazetteers composed of municipal and regional toponyms— the names of towns, cities, counties, and larger administrative areas— across toponym extraction and resolution research often results in the exclusion of fine-grained and generic toponyms within these studies: specific and generic places within a town or city that social media users mention during a crisis. As a result, we know little about the usage of hyperlocal toponyms among social media users during crises, and, in turn, design requirements for extraction and resolution techniques that can help responders detect and locate information about events occurring within operational jurisdictions.

RESEARCH QUESTIONS

Given crisis responders' geographic information requirements when using social media, the inadequacy of geotag-based analysis methods, and the gaps that appear in literature examining the usage and granularity of toponyms posted on Twitter during crises, we pose three research questions: (RQ1) To what extent do messages include toponyms in crisis social media? (RQ2) To what extent do messages including toponyms refer to local or remote places? (RQ3) What is the granularity of toponyms included in crisis social media?

METHODS

Tweets were collected using three data collection methods— location, keyword, and network filtering— to gather 22,343 unique tweets (after removing 343 duplicates) during a severe storm and F1 tornado that struck Centre County, Pennsylvania, on May 1st, 2017. First, location filtering employed Twitter's Streaming API to collect 9,098 tweets within a bounding box[1] covering Central Pennsylvania, and encompassing Centre County, during a six-hour period (3pm-9am) before and after the peak of a severe storm and tornado that struck the area just after 6pm on May 1st. Second, keyword filtering employed Twitter's Streaming API to collect 4,566 tweets including at least one of 46 place names, including "Centre County" and the names of the county's 45 municipalities, boroughs, and census-designated places. Third, using network filtering to infer 185,176 Twitters users with social network ties (i.e. following) to ground truth accounts within Centre County, 8,679 tweets were collected via Twitter's Streaming API during the 6-hour period on May 1st. In a previous study (Grace, Kropczynski, Pezanowski, et al., 2018), this "wide net" inferencing approach was evaluated, finding that among 80K users with self-entered profile locations, 68% entered a location inside the county. Furthermore, during the qualitative coding

[1] SW 40.170196, -79.1214656; NW 41.3434653, -76.5996462

process every tweet from all three datasets was manually examined to determine if the post provided storm-related or off-topic information and described events inside or outside the area of observation. Opportunities for future work comparing toponym usage, location, and granularity across the three datasets are discussed in the conclusion.

Data Analysis

In a previous study (Grace, Halse, Aurite, Montarnal, & Tapia, 2019), three of the authors manually coded each of the 22,343 tweets to understand the information users posted on Twitter during the storm. The qualitative content analysis involved two rounds of coding. First, to determine relevance, tweets were coded as "storm-related" if referring to weather or its consequences (e.g., damage caused by high winds, etc.), and "off-topic" if referring to other topics (e.g., sporting events, political issues, etc.). A random set of 1000 tweets were coded by all three authors and a Cronbach alpha test was run yielding $\alpha = 0.92$. Coding differences were deliberated and reconciled, and then the remaining data was subdivided and coded for relevance.

Second, storm-related tweets were coded to understand the types of information reported during the storm. Together, the three authors engaged in a grounded, iterative process of open coding that involved assigning meanings, in the form of emergent codes, to all 2,364 storm-related tweets (Campbell, Quincy, Osserman and Pedersen, 2013; Glaser and Strauss, 1967; Hsieh and Shannon, 2005). During this process the authors consulted and unpacked codes developed in previous literature (Olteanu et al., 2015; Starbird, Palen, Hughes and Vieweg, 2010). For what Olteanu et al. (2015) code as "Infrastructure & Utilities," for example, five new codes emerged in the data: property damage, road damage, power line damage, Internet outage, and power outage. Ultimately, the coding process resulted in 19 types of information organized into six categories to account for the diversity of tweets posted about the storm (Table 1).

Table 1. Categories and Types of Storm-related Information

Category	Information	Tweets	% Storm-Related	% Total
Disruptions	Power outage (140), Internet outage (8)	148	6%	1%
Experiences	Humor (99), Admiration (80), Complaint (54), Fear (29), Appreciation (10)	272	12%	1%
Forecasts	Forecast (278)	278	12%	1%
Damages	Road damage (224), Property damage (112), Power line damage (18)	354	15%	2%
Updates	Current conditions (383), Automated update (74), Information request (39), Event information (14)	510	22%	2%
Warnings	Storm watch/warning (463), Tornado watch/warning (206), Flood advisory (96), Advice (37)	802	34%	4%
Total		2364	100%	11%

The present study involved three further rounds of coding, performed by the author, to examine toponym usage and granularity across both storm-related and off-topic tweets gathered during the storm. First, all 22,343 tweets were coded for toponym usage, distinguishing tweets *with* specific or generic toponyms from tweets *without* toponyms. Second, the 3,497 tweets with toponyms were coded for references to *local* places within Centre County, references to *non-local* places outside Centre County, and references to *generic* places, such as "house," for which a specific location cannot be determined. If a tweet included both local and non-local toponyms it was coded as local. Third, tweets with toponyms were further coded for granularity, distinguishing references to *hyperlocal* places including names of landmarks (i.e., addresses, buildings, and road intersections) and areas (i.e., neighborhoods, parks, campuses, and street blocks), *municipal* places including names of towns, villages, cities, and counties, and *regional* places including names of states and other regional markers (i.e., Upstate New York, Central Pennsylvania, etc.). In this coding scheme, hyperlocal refers to places within municipalities which refer to places within regions. If a tweet included multiple toponyms with different granularities, it was assigned a code based on the most granular toponym available.

FINDINGS

Examining toponym usage, referent location, and granularity in tweets posted during the storm that struck Centre County on May 1st reveals three initial findings (Table 2). First, users frequently make explicit references to places in storm-related tweets but do so rarely in off-topic tweets. In the dataset, 68% of storm-related tweets include toponyms compared to only 9% among off-topic tweets.

Second, among tweets with toponyms, 74% refer to non-local places outside the geographic area targeted for data collection while only 17% of tweets refer to specific places within or encompassing the area, i.e., Centre County. Local toponyms referring to places within the county or the county itself constitute only 25% of storm-related information and 9% of off-topic information. Generic toponyms referencing unspecified landmarks, areas, towns, and regions appear in 4% of on-topic and 14% of off-topic tweets.

Third, among all tweets posted during the storm, 38% include hyperlocal toponyms— names of specific or generic landmarks and areas within a municipality. Although relatively uncommon, toponyms in off-topic tweets tend to be more granular (48%) than toponyms in storm-related tweets (25%). However, both storm-related and off-topic tweets nearly always refer to places at the municipal level or below, 94% and 92% respectively.

Table 2. Toponym Usage, Location, and Granularity in Storm-related and Off-topic Tweets

		Storm-related Tweets		Off-Topic Tweets		Totals	
Usage	*With Toponyms*	1600	68%	1897	9%	3497	16%
	Without	764	32%	18082	91%	18846	84%
	Total	2364	100%	19979	100%	22343	100%
Location	*Local*	406	25%	176	9%	582	17%
	Non-Local	1130	71%	1460	77%	2590	74%
	Generic	64	4%	261	14%	325	9%
	Total	1600	100%	1897	100%	3497	100%
Granularity	*Hyperlocal*	407	25%	918	48%	1325	38%
	Municipal	1105	69%	832	44%	1937	55%
	Regional	88	6%	147	8%	235	7%
	Total	1600	100%	1897	100%	3497	100%

These finding can be compared with the 6% of tweets that include geotags in the dataset, a relative abundance compared to the 1% of geotagged data in the global Twitter stream. Surprisingly, 24% of storm-related tweets include geotags, compared to only 4% of off-topic tweets. However, as will be shown below, the preponderance of geotagged, storm-related information consists of weather warnings, forecasts, and weather updates referring to places outside the area of observation but posted with geographic coordinates inside Centre County. When excluding geotagged tweets with non-local toponyms, only 2% of storm-related tweets are geotagged.

Toponym Usage

Further insight is possible by examining toponym usage across the six categories of storm-related information: disruptions to electrical and internet service, weather forecasts, user experiences such as fear and humor, damages to powerlines, buildings, and roadways, updates on current weather conditions, and flood, storm, and tornado warnings and watches. Comparing toponym usage across these categories shows that while users generally refer to places in storm-related tweets, they do so much more frequently when posting some kinds of information than others (Figure 1).

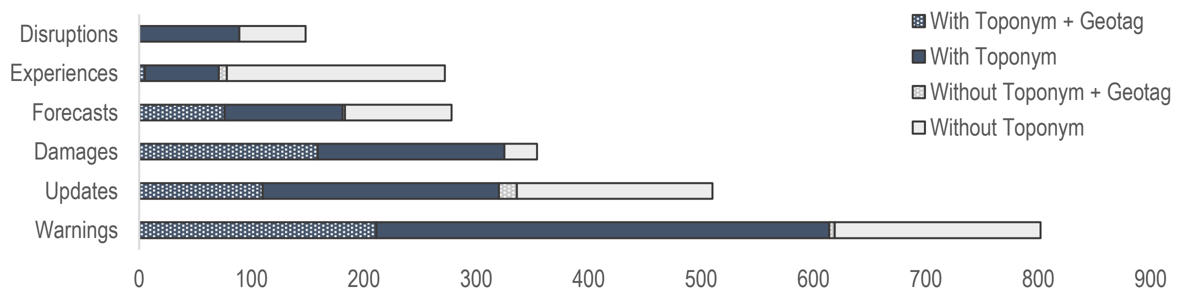

Figure 1. Toponym Usage and Geotagging (Dotted) across Categories of Storm-related Information

Twitter users refer to places in most tweets providing warnings (77%) and, strikingly, nearly every tweet describing damages (92%). These tweets describe places where damage has occurred to power lines (100%), property (78%), and roadways (98%). Comparatively, users providing weather updates (63%), forecasts (65%), or reports of electrical and Internet disruptions do so less often (60%). In contrast, far fewer tweets describing users' experiences of the storm include toponyms (26%).

Furthermore, users similarly geotag some types of information more than others. Tweets describing experiences

(2%) and disruptions (0%) include few geotags, whereas updates (22%), warnings (26%), forecasts (27%), and damage (45%) reports include relatively many. Moreover, storm-related tweets including toponyms tend to also include geotags. For instance, 92% of damage reports include toponyms and nearly half of those include geotags. Overall, the distributions of toponym usage and geotagging across storm-related information, in stark contrast to the dearth of volunteered geographic information available in off-topic tweets, highlight how Twitter users actively report the locations of storm-related events.

Toponym Location

Overall, most toponyms in the dataset refer to remote places outside the geographic area— Centre County, Pennsylvania— targeted for data collection and analysis (Figure 2). This holds true for each category of information although, again, some types of information include slightly more toponyms than others. Generic toponyms constitute a fraction of toponyms in each category, ranging from ≤1% for forecasts and warnings to 14% of toponyms included all reports of disruption (e.g., "The power is out in the whole building").

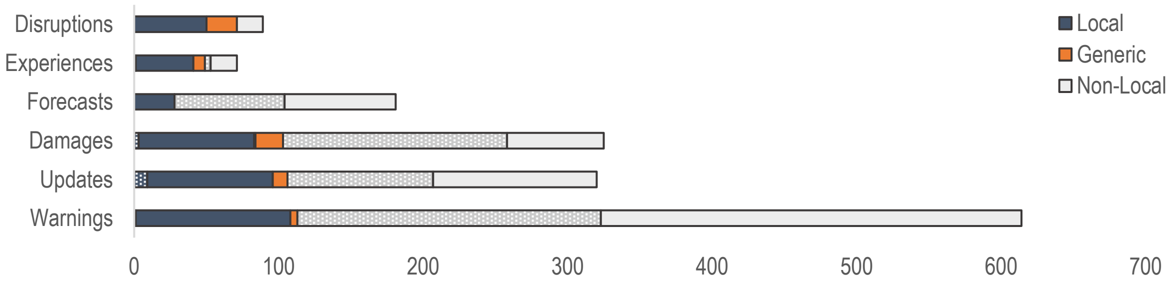

Figure 2. Toponym Location and Geotagging (Dotted) across Categories of Storm-related Information

However, comparing the distribution of geotags across tweets with local, generic, and non-local toponyms reveals that 92% of geotagged, storm-related tweets refer to events occurring in remote places. Figure 2 illustrates the imbalance of geotagging between tweets referencing local and non-local places. Among warnings, for example, 97% of geotagged tweets refer to remote places, while only a single geotagged tweet refers to a location within Centre County. This imbalance likely results from users' generally selective geotagging behaviors (Tasse et al., 2017), and the inverse tendency of media, government, and emergency management Twitter accounts. Most forecasts and warnings, unsurprisingly, are posted by local news stations and meteorologists, as well as the Twitter accounts of municipal governments (e.g. "A Tornado Warning has been issued for parts of Elk County until 4:15 this evening"). Furthermore, as the storm front moved east to west, first impacting areas across eastern Pennsylvania before striking Centre County (in central Pennsylvania) and, later, cities such as Pittsburg in western Pennsylvania, local media tweeted a steady stream of reports about events in nearby counties before and after the storm peaked in the area of observation (e.g., "Traffic along Loyalsock Ave in Montoursville snarled because of trees and power lines down between Turkey Hill and Indian Park").

Toponym Granularity

Lastly, while users typically post storm-related tweets without geotags and refer to either fine-grained, hyperlocal places (38%) or course-grained, municipal (55%) places in the area of observation, the granularity of toponyms varies according to the type of information users post about the storm (Figure 3). Warnings (76%), updates (58%), experiences (55%), and forecasts (54%) mostly include municipal-level toponyms— names of towns within Centre County and references to Centre County itself— while reports of infrastructure damage (71%) and service disruption (63%) include mostly hyperlocal toponyms— names of landmarks and areas within these municipalities. While most common in updates (8%), geotags are rarely present in any of the six information categories. Consequently, due to the lack of geotags and frequent use of course-grained toponyms, most tweets about the storm lack precise geographic information.

Not shown in Figure 3 are generic toponyms for each category, references to common nouns, such as "house" or "road," that make geolocation difficult. Use of generic toponyms varies by information type, constituting a notable minority of local toponyms included tweets describing disruptions (30%), damage (19%), and experiences (16%), while less present in updates (9%), warnings (4%), and forecasts (0%). These findings again point to different behaviors among users when sharing geographic information in posts describing the storm.

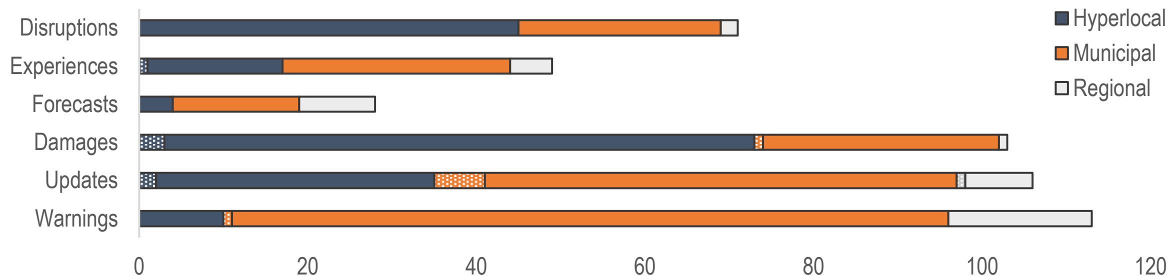

Figure 3. Toponym Granularity and Geotagging (Dotted) across Categories of Local, Storm-related Information

DISCUSSION

These findings reveal toponym usage and granularity patterns in crisis social media that hold significant implications for crisis informatics research. This study extends findings of prior content analyses that survey the types of information users post on social media during a crisis by revealing patterns of toponym usage and granularity across these information types (Grace, Halse, Aurite, Montarnal and Tapia, 2019; Olteanu et al., 2015). Furthermore, by illustrating patterns of generic and fine-grained toponym usage at hyperlocal-levels of analysis, this study informs requirements for geolocation techniques that can support crisis responders monitoring social media in municipal-level jurisdictions (Avvenuti et al., 2016; Middleton et al., 2018, 2014). Overall, this study makes three primary contributions.

First, the frequent use of toponyms and geotagging in storm-related messages (68% and 25% respectively), compared to off-topic messages (9% and 4% respectively), suggest that *crisis social media includes disproportionate amounts of geographic information*. Overall, via toponym or geotag, 1,630 (69%) of 2,364 storm-related messages in the dataset include geographic information compared to just 2,397 (12%) of 19,979 off-topic messages. However, the availability of geographic information in crisis social media varies by information category, with reports of infrastructure damage including the most toponyms (92%) and descriptions of personal experience including the least (26%). Similarly, and extending findings by Avvenuti et al. (2016), tweets with damage information include the most geotags (45%) in the dataset, although few describe local events. These findings, in turn, suggest that users tend to communicate the locations of events that present community risks or demand immediate assistance from crisis or volunteer responders.

Second, analyzing whether toponyms reference places inside or outside the area of observation shows that *crisis social media includes reports of geographically dispersed events, but that social media users tend to describe local events using toponyms, not geotags, during a crisis.* The majority of storm-related (92%) and off-topic (33%) messages describe events in remote places. This tendency is especially apparent among storm-related messages geotagged within the area of observation, of which only 3% describe events occurring within that area. In contrast, 25% of storm-related messages including toponyms describe local events. Consequently, while crisis-related messages may include disproportionate geographic information compared to non-crisis-related messages, this study finds that social media users more often employ toponyms than geotags when reporting crisis-related events. These findings contribute to understanding uneven geotagging behaviors among Twitter users (Crampton et al., 2013; Huang and Carley, 2019; Tasse et al., 2017), and highlight the frequency of "sensor-subject displacement" (Robertson and Feick, 2018; Shelton et al., 2014): local users[2] can and often do talk about non-local events. Consequently, collecting and geolocating situational awareness information using only geographic metadata is likely inadequate for most crisis management applications.

Third, findings show that *toponym granularity in crisis social media varies by information category*: warnings, forecasts, and accounts of personal experience tend to include course-grained toponyms referring to municipalities, while reports of infrastructure damage and service disruption often include fine-grained toponyms referring to places within municipalities. Across every category, however, course-grained toponyms are common, suggesting that the availability of geographic information in crisis social media more readily afford rapid impact assessments useful to emergency managers responsible for relatively large geographic areas than emergency response activities which require more precise geographic information or external information sources (e.g., geographic information gathered from emergency telephone numbers such as 911 and 112).

Together, these findings on toponym usage, location, and granularity provide requirements for the design of social media monitoring techniques that can provide crisis responders with situational awareness in hyperlocal operational jurisdictions (Imran, Castillo, Diaz and Vieweg, 2015). The observation that local, storm-related

[2] More precisely, tweets geotagged within a geographic area can refer to places outside that area.

information tends to include toponyms rather than geotags suggests affordances for systems that can leverage toponym usage to filter non-local information and geolocate local, actionable information (Avvenuti et al., 2018). That most existing toponym extraction and resolution techniques perform poorly with hyperlocal or generic place names such as "downtown" (de Bruijn et al., 2018; Middleton et al., 2018; Wang and Ye, 2018), suggest that such systems require context-sensitive toponym extraction and resolution techniques to exploit the kinds of colloquial and fine-grained toponyms included in most messages describing infrastructure damage and service disruption, and over a quarter of all storm-related tweets examined in this study. Developing such techniques will likely require community-specific gazetteers that include hyperlocal place names, especially of critical infrastructure, and that can grow by adding location entities identified during ongoing social media monitoring efforts.

Limitations and Future Work

This study is limited to the analysis of toponym usage on Twitter during a single crisis event. The findings are, necessarily, contingent to contextual information behaviors among Twitter users, such as geotagging, that may vary from the behaviors of Twitter users in different locales or crisis conditions (Huang and Carley, 2019). Comparative studies of toponym usage and geotagging for different crisis events and locations would inform general conclusions on the availability of volunteered geographic information on social media during crises. Furthermore, contextual factors— local, national, and international events occurring on May 1st, 2017— shaped what users discussed in off-topic tweets and, perhaps, storm-related tweets. Outside the scope of the present study, the comparative analysis of storm-related and off-topic information provides an opportunity for future research.

The lack of geotagged Twitter data also points to the importance of innovative data collection methods that can collect information posted by users in crisis-impacted areas without resorting to geotag-based filtering techniques (Grace et al., 2017). The use of three separate data collection methods in this study provides an opportunity for comparing, for example, the relative affordances of location, keyword, and network-based methods for collecting relevant information (e.g., infrastructure damage) and geographic information (e.g., fine-grained toponyms) during a crisis. Beyond the scope of the present study, these questions will be addressed in future research.

CONCLUSION

This study analyzes toponym usage and granularity across types of storm and non-storm-related information posted on Twitter. Findings show that users report the locations of hyperlocal storm impacts using place names rather than geotags, and frequently include fine and course-grained toponyms in storm-related information but do so rarely when discussing other topics. These findings reveal patterns of toponym usage and granularity in crisis social media and hold implications for the design and use of social media monitoring applications in crisis management.

REFERENCES

Avvenuti, M., Cresci, S., Del Vigna, F. and Tesconi, M. (2016) Impromptu Crisis Mapping to Prioritize Emergency Response, *Computer*, *49*, 5, 28–37. https://doi.org/10.1109/MC.2016.134

Avvenuti, M., Cresci, S., Nizzoli, L. and Tesconi, M. (2018) GSP (Geo-Semantic-Parsing): Geoparsing and Geotagging with Machine Learning on Top of Linked Data, *Lecture Notes in Computer Science*, *10843 LNCS*, 17–32. https://doi.org/10.1007/978-3-319-93417-4_2

Campbell, J. L., Quincy, C., Osserman, J. and Pedersen, O. K. (2013) Coding In-depth Semistructured Interviews: Problems of Unitization and Intercoder Reliability and Agreement, *Sociological Methods & Research*, 42, 3, 294–320. https://doi.org/10.1177/0049124113500475

Carley, K. M., Malik, M., Landwehr, P. M., Pfeffer, J. and Kowalchuck, M. (2016) Crowd sourcing disaster management: The complex nature of Twitter usage in Padang Indonesia, *Safety Science*, 90, 48–61. https://doi.org/10.1016/J.SSCI.2016.04.002

Cheng, Z., Caverlee, J. and Lee, K. (2010) You are where you tweet, *Proceedings of the 19th ACM International Conference on Information and Knowledge Management - CIKM '10*, 759. https://doi.org/10.1145/1871437.1871535

Crampton, J. W., Graham, M., Poorthuis, A., Shelton, T., Stephens, M., Wilson, M. W. and Zook, M. (2013) Beyond the geotag: situating 'big data' and leveraging the potential of the geoweb, *Cartography and Geographic Information Science*, 40, 2, 130–139. https://doi.org/10.1080/15230406.2013.777137

de Bruijn, J. A., de Moel, H., Jongman, B., Wagemaker, J. and Aerts, J. C. J. H. (2018) TAGGS: Grouping Tweets to Improve Global Geoparsing for Disaster Response, *Journal of Geovisualization and Spatial Analysis*, 2, 1, 2. https://doi.org/10.1007/s41651-017-0010-6

Dredze, M., Paul, M. J., Bergsma, S. and Tran, H. (2013) Carmen: A Twitter Geolocation System with Applications to Public Health, In *Workshops at the Twenty-Seventh AAAI Conference on Artificial Intelligence*. Retrieved from https://github.com/mdredze/carmen

Dunkel, A., Andrienko, G., Andrienko, N., Burghardt, D., Hauthal, E. and Purves, R. (2019) A conceptual framework for studying collective reactions to events in location-based social media, *International Journal of Geographical Information Science*, 33, 4, 780–804. https://doi.org/10.1080/13658816.2018.1546390

Giacobe, N. A. and Soule, P. J. (2014) Social media for the emergency manager in disaster planning and response. *Proceedings of the 11th International Conference on Information Systems for Crisis Response and Management (ISCRAM)*, 570–574.

Glaser, B. and Strauss, A. (1967) *The discovery of grounded theory: Strategies for qualitative research*, Hawthorne, NY: Aldine de Gruyter.

Grace, R., Halse, S., Aurite, W., Montarnal, A. and Tapia, A. (2019) Expanding awareness: Comparing location, keyword, and network filtering methods to collect hyperlocal social media data, *Proceedings of the 52nd Hawaii International Conference on System Sciences (HICSS), 2699-2708. http://hdl.handle.net/10125/59707*.

Grace, R., Kropczynski, J., Pezanowski, S., Halse, S., Umar, P. and Tapia, A. (2017) Social Triangulation: A new method to identify local citizens using social media and their local information curation behaviors, *Proceedings of the 14th International Conference on Information Systems for Crisis Response and Management (ISCRAM)*, 902–915.

Grace, R., Kropczynski, J., Pezanowski, S., Halse, S., Umar, P. and Tapia, A. (2018) Enhancing emergency communication with social media: Identifying hyperlocal social media users and information sources, *International Journal of Information Systems for Crisis Response and Management (IJISCRAM)*, 10, 3, 20-41. https://doi.org/10.4018/IJISCRAM.2018070102

Grace, R., Kropczynski, J. and Tapia, A. (2018) Community coordination: Aligning social media use in community emergency management, *Proceedings of the 15th International Conference on Information Systems for Crisis Response and Management (ISCRAM), 609-620*.

Halterman, A. (2017) Mordecai: Full Text Geoparsing and Event Geocoding, *The Journal of Open Source Software (JOSS)*, 2, 9, 91. https://doi.org/10.21105/joss.00091

Hecht, B. and Stephens, M. (2014) A Tale of Cities: Urban Biases in Volunteered Geographic Information, *Eighth International AAAI Conference on Weblogs and Social Media (ICWSM)*, 197–205.

Hiltz, S. R., Hughes, A. L., Imran, M., Plotnick, L., Power, R. and Turoff, M. (2019) Exploring the usefulness and feasibility of software requirements for social media use in emergency management, *International Journal of Disaster Risk Reduction*, 101367. https://doi.org/10.1016/j.ijdrr.2019.101367

Hsieh, H.-F. and Shannon, S. E. (2005) Three Approaches to Qualitative Content Analysis, *Qualitative Health Research*, 15, 9, 1277–1288. https://doi.org/10.1177/1049732305276687

Huang, B. and Carley, K. M. (2019) A Large-Scale Empirical Study of Geotagging Behavior on Twitter. Retrieved from http://arxiv.org/abs/1908.10948

Imran, M., Castillo, C., Diaz, F. and Vieweg, S. (2015) Processing Social Media Messages in Mass Emergency, *ACM Computing Surveys*, 47, 4, 1–38. https://doi.org/10.1145/2771588

Kitamoto, A. and Sagara, T. (2012) Toponym-based geotagging for observing precipitation from social and scientific data streams, *Proceedings of the 2012 ACM International Workshop on Geotagging and Its Applications in Multimedia*, 23–26. https://doi.org/10.1145/2390790.2390799

Kropczynski, J., Grace, R., Coche, J., Obeysekare, E., Bénaben, F., Halse, S., … Tapia, A. (2018) Identifying Actionable Information on Social Media for Emergency Dispatch. *Proceedings of ISCRAM Asia Pacific 2018*, 1–11. Retrieved from https://prioritydispatch.net/discover_proqa/

McCormick, S. (2016) New tools for emergency managers: an assessment of obstacles to use and implementation, *Disasters*, 40, 2, 207–225. https://doi.org/10.1111/disa.12141

Middleton, S. E., Kordopatis-Zilos, G., Papadopoulos, S. and Kompatsiaris, Y. (2018) Location extraction from social media: Geoparsing, location disambiguation, and geotagging, *ACM Transactions on Information Systems*, 36, 4. https://doi.org/10.1145/3202662

Middleton, S. E., Middleton, L. and Modafferi, S. (2014) Real-Time Crisis Mapping of Natural Disasters Using Social Media, *IEEE Intelligent Systems*, 29, 2, 9–17. https://doi.org/10.1109/MIS.2013.126

Morstatter, F., Pfeffer, J., Liu, H. and Carley, K. M. (2013) *Is the Sample Good Enough? Comparing Data from Twitter's Streaming API with Twitter's Firehose*, 400–408. https://doi.org/10.1007/978-3-319-05579-4_10

Olteanu, A., Vieweg, S. and Castillo, C. (2015) What to Expect When the Unexpected Happens, *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW)*, 994–1009. https://doi.org/10.1145/2675133.2675242

Plotnick, L. and Hiltz, S. R. (2018) Software Innovations to Support the Use of Social Media by Emergency Managers, *International Journal of Human–Computer Interaction*, 34, 4, 367–381. https://doi.org/10.1080/10447318.2018.1427825

Reuter, C., Hughes, A. L., & Kaufhold, M.-A. (2018) Social Media in Crisis Management: An Evaluation and Analysis of Crisis Informatics Research, *International Journal of Human–Computer Interaction*, 34, 4, 280–294. https://doi.org/10.1080/10447318.2018.1427832

Reuter, C., Ludwig, T., Kaufhold, M.-A. and Spielhofer, T. (2016) Emergency services' attitudes towards social media: A quantitative and qualitative survey across Europe, *International Journal of Human-Computer Studies*, 95, C, 96–111. https://doi.org/10.1016/j.ijhcs.2016.03.005

Robertson, C., & Feick, R. (2018) Inference and analysis across spatial supports in the big data era: Uncertain point observations and geographic contexts, *Transactions in GIS*, 22, 2, 455–476. https://doi.org/10.1111/tgis.12321

Salas, A., Georgakis, P. and Petalas, Y. (2018) Incident detection using data from social media, *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*, 751–755. https://doi.org/10.1109/ITSC.2017.8317967

Shelton, T., Poorthuis, A., Graham, M. and Zook, M. (2014) Mapping the data shadows of Hurricane Sandy: Uncovering the sociospatial dimensions of 'big data,' *Geoforum*, 52, 167–179. https://doi.org/10.1016/J.GEOFORUM.2014.01.006

Starbird, K., Palen, L., Hughes, A. L. and Vieweg, S. (2010) Chatter on the red what hazards threat reveals about the social life of microblogged information, *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work (CSCW)*, 241. https://doi.org/10.1145/1718918.1718965

Tasse, D., Liu, Z., Sciuto, A. and Hong, J. I. (2017) *State of the Geotags: Motivations and Recent Changes*. Retrieved from www.aaai.org

Wang, Z. and Ye, X. (2018) Social media analytics for natural disaster management, *International Journal of Geographical Information Science*, 32, 49–72. https://doi.org/10.1080/13658816.2017.1367003

Zade, H., Shah, K., Rangarajan, V., Kshirsagar, P., Imran, M. and Starbird, K. (2018) From Situational Awareness to Actionability, *Proceedings of the ACM on Human-Computer Interaction* (CSCW), 1–18. https://doi.org/10.1145/3274464