

Evaluating Public Response to the Boston Marathon Bombing and Other Acts of Terrorism through Twitter

Cody Buntain, Jennifer Golbeck*, Brooke Liu†, and Gary LaFree‡
{cbuntain,golbeck}@cs.umd.edu, {bfliu,glafree@umd.edu}

Dept. of Computer Science, *College of Information Studies, †Dept. of Communication, ‡START Center
University of Maryland, College Park

Abstract

Twitter and other social platforms have become important communication channels during crises. While research into crisis informatics and social media is growing, the rarity of terrorist attacks in developed, Western countries complicates analysis of these specific events. To address this gap, we explore social media response to three terror events: the 2013 Boston Marathon Bombing, the 2014 Sydney Hostage Crisis, and the 2015 Charlie Hebdo Shooting. We show that, while these events do not significantly impact general Twitter usage, those users who are discussing the event behave in predictable ways across all three events. Such behaviors include increased references to the event and use of retweets, hashtags, and URLs. Furthermore, local news affiliates and law enforcement agencies (if present on social media) emerge as central actors in the networks.

Introduction

Social media has become an important tool during crises, as is apparent in recent surveys regarding the public's intent to use social media in response (Liu, Fraustino, and Jin 2013), public expectations that the government provide information through these channels (Cross 2010), and in federal efforts to improve communication during emergency situations (Lindsay 2011). Though such situations come in many forms and varying timeframes, terrorist attacks are of particular importance given their intentional destabilizing effects. Social media content during these crises might yield important insights, but the paucity of terrorist acts in developed countries with large Twitter populations complicates research. Some work has broached this topic, but much of it focuses on single terrorist acts with limited generalization (Liu, Fraustino, and Jin 2013; Gupta and Kumaraguru 2012; Sutton et al. 2014; Rogstadius et al. 2013). Olteanu, Vieweg, and Castillo enhance generalizability by investigating social media in the broader context of **all** disasters, including terrorist attacks (2015). Their work's broadness, while informative, precludes a deeper inspection into public response **specific to** terror attacks, and it is here that this paper makes its contribution.

We investigate social media use in developed, Western countries during three terror events: the 2013 Boston

Marathon Bombing, the 2014 Sydney Hostage Crisis, and the 2015 Charlie Hebdo Shooting. We analyze information sharing/seeking and social interactions before and after these events on Twitter, a popular microblogging site. This exploration characterizes public response to these events and identifies important accounts during these crises (spoiler: local police, if present on Twitter, and local news affiliates).

Contributions. Researchers seeking to leverage Twitter for crisis informatics, explore social media's informativeness during crises, or identify terror attacks should find this work of interest. Our primary contributions are to:

- Characterize effects of terror events on activity and social interactions (followers and mentions) on Twitter,
- Identify trends in actors who emerge as most important on Twitter during terror events, and
- Generalize these results across three terror events in industrialized, Western countries.

Related Work

A great deal of research explores the role of social media, especially Twitter, in disaster response (Vis 2013; Liu, Fraustino, and Jin 2013; Fraustino, Liu, and Yan 2012). Gupta and Kumaraguru (2012), Faustino et al. (2012), and Olteanu et al. (2015) have all shown social media contains important and useful information about these events. Olteanu et al. also sought to identify how the public's social media use adapts across a large collection of disasters, natural versus man-made and accidental versus intentional. Though this work included one terror event (the Boston Marathon Bombing), understanding the trends in social media *specific to terror events* was not part of their objectives. Our work builds upon these foundation by investigating a specific type of event (terror attacks) in a specific context (industrialized, Western countries). While others have explored social media and terror, this effort extends the state of the art by generalizing across a set of terror events. Social media has evolved over the past decade, so we also align our results with previously mentioned efforts to determine what trends still hold and what new trends have developed.

Methods

We characterize public response on Twitter during three terror-related events: the Boston Marathon Bombing and

resulting manhunt (15-19 April 2013); the 2014 Sydney Hostage Crisis (14-15 December); and the Charlie Hebdo Attack and manhunt (7-9 January 2015) and focus on the two weeks before and after each event. These events were selected because of their high coverage in Western media between April 2013 and May 2015 (the dates covered by our data set). We also constrain our data to only those tweets that mention target events (similar to Olteanu et al. (2015)) but also characterize trends in Twitter’s 1% sample around these events to gauge the general Twitter population’s responses and whether these events impact overall usage.

While Twitter restricts authors to 140 characters, users have a wide variety of content to share in this small space: website links (URLs), hashtags, multimedia, mentions of other users, and retweets. Trends in these artifacts evolve over time, and their patterns might inform us on information sharing. To this end, we determine whether each artifact experiences significant shifts in usage around target events, as determined through two tests: 1) calculating whether the activity’s frequency is more than twice the median absolute deviation (MAD) around the event, and 2) performing a Welch’s t-test of whether data before the event differs significantly from after (all tests are two-tailed and performed at $p < 0.05$). The following questions explore these activities:

- RQ1** Does relevant tweet volume change during the event?
- RQ2** Does the proportion of retweets, URLs, hashtags, mentions, or media change during the target event?
- RQ3** Do follower counts change for important accounts during the event?
- RQ4** How long do these changes persist?
- RQ5** What users emerge as important during terror attacks?

RQ1 is motivated by research that shows people’s intent to seek/share information increases in response to crises and asks whether this effect holds for Twitter (Jin, Fraustino, and Liu 2015). **RQ2** then identifies which specific artifacts are most affected by these events; e.g., significant increases in retweets might indicate higher information sharing as Twitter users rebroadcast information to their followers or join the conversation. **RQ3** then measures seeking behavior by quantifying users’ subscriptions to and thus seeking information from important Twitter accounts. **RQ4** is a natural follow-up on duration of these effects, and existing research is conflicting. Olteanu et al. (2015) showed changes on Twitter can persist for a few days to nearly two months, with the Boston Marathon Bombing persisting for 60 days, but Koutra, Bennett, and Horvitz demonstrated shocking events rarely influenced long-term user behavior in digital communities beyond social networks (2015). During such crises, it is also unclear which accounts become the center of attention. We investigate this question in **RQ5** by converting Twitter’s retweet and mention activity into a directed graph of interactions, where the vertices represent Twitter users, and the edges denote mentions/retweets. Research shows users with many followers or retweets often are not the most influential users (Cha et al. 2010), so we followed Kwak et al. and used a version of the PageRank algorithm to identify important accounts in this network (2010).

Data Collection This work leveraged a corpus gathered from Twitter’s 1% public sample stream (others have explored bias in this sample (González-Bailón et al. 2014; Morstatter et al. 2013)). We neither removed retweets nor filtered short tweets since retweets provide valuable insight into the network structure and help identify the central actors during these crises. To investigate the three terror events we target, we concentrate on the two weeks before and two weeks after each event:

- 1-30 April ’13: 134,245,610 tweets.
- 1-31 December ’14: 134,226,491 tweets.
- 23 December ’14-20 January ’15: 128,350,988 tweets.

To identify relevant content, we search for tweets containing “boston” for the Boston Marathon Bombing, “sydney” for the Sydney Hostage Crisis, and either “paris” or “hebdo” for the Charlie Hebdo Attack. This search is case-insensitive and matches keywords embedded in hashtags.

Results

A first step in understanding public Twitter response is to examine the 1% sample stream and its artifacts. Analysis of this data shows that general Twitter activity was unaffected by these events. Statistical analysis shows the only significant changes in activity were an increase in retweets on the day before and day of the Boston Marathon Bombing and a decrease in tweets containing URLs at the start of the Sydney Hostage Crisis, but tweet volume and other activities remain the same.

For **RQ1** and **RQ2** (relevant tweet volume and changes in activities), we refer to Figure 1, where the black asterisk curve shows tweet volume on the right axis, and remaining curves correspond to tweet proportions on the left. From these results, we see tweet volume increased on the date of each terror event. References to Boston increased from about 828 tweets per day to 138,000 on the day of the bombing. Sydney saw a smaller increase from an average 1,119 tweets per day to 16,000, and Paris saw an increase from 2,288 to nearly 78,000. Each event’s first day is between 34 and 70 times the MAD. In answer to **RQ1**, Twitter sees a significant increase in references to the target events on the day the terror attack starts.

For **RQ2**, in Boston, proportions of retweets, URLs, media, and hashtags exceed twice the MAD for 15 April, with retweets, media, and hashtags seeing increases while URLs decrease. Proportion of tweets containing mentions is not significantly affected. Of these activities, we see significant changes in retweets, URLs, and media sharing before and after the event ($p < 0.05$). Sydney sees a similar deviation from the MAD on 15 December in retweets, mentions, URLs, and tweets containing hashtags, with mentions being the only activity that decreases. Even though URL sharing did not initially deviate from the MAD, the Welch’s t-test suggests the day of the event to be a point in which sharing behavior changes with significant differences between average daily proportions. Media is unaffected, and retweet effects are insignificant. Finally, during the Charlie Hebdo attack on 7 January, only retweets, media sharing, and hashtags deviate from the MAD. URLs once again see a drop in




Figure 1: Relevant Twitter Activity Per Event

usage as in Boston, but the change is not significant on the first day of the attack. Average trends in all of these sharing activities, however, differ significantly before and after the event ($p < 0.05$ for retweets and mentions, and $p < 0.01$ for URLs, media, and hashtags). Furthermore, average daily mentions decrease following the event, whereas all other activities see an increase. Retweets and hashtags increase immediately following these events, while URL sharing responds more slowly. Tweets including media and mentions are inconsistently affected.

To answer **RQ3**, we examined the Boston Police Department (@bostonpolice), the New South Wales Police Department (@nswpolice), and Charlie Hebdo (@Charlie_Hebdo.). Results show the Boston Police Department and Charlie Hebdo see massive increases in followers. Prior to the event, the Boston Police Department had an average of 54k followers, which increased significantly to 264k followers on the day of the manhunt and peaked at over 300k through the end of the month. Similarly, followers of Charlie Hebdo increased from an average of 77k to 318k. The NSW Police Department also sees a significant increase, though not as substantial, from 61k to 80k followers. Follower increases for each account differ significantly from the average: Users referring to Boston, Sydney, and Paris saw a mean increase of 78, 136, and 253 followers ($\sigma = 2, 753, 989, \text{ and } 2, 731$) respectively. Therefore, all three accounts experience significant increases in followers ($p < 0.01$), with rapid increases that level off within 7 days.

For effect duration, retweets and hashtags experience significant but short-lived surges. Retweets referencing Boston retreat to 36% from a high of 59% on 15 April, and hashtags drop from 52% to 31%; hashtag usage before the bombing and seven days after is not significantly different ($p > 0.05$). Similarly for the Sydney Hostage Crisis, retweets and hashtags return to pre-event levels with no statistical difference ($p > 0.1$) within seven days. During the Charlie Hebdo attack, retweets are similar, but hashtags remain significantly higher than their pre-event levels for at least the next two weeks ($p < 0.01$). For URL sharing, after the Boston and Charlie Hebdo events, URL sharing increase and remain high, and while followers show a slight decrease after the event, they remain significantly higher than pre-event levels.

In answer to **RQ4**, changes in retweets and hashtags persist for only a few days, whereas URL sharing and follower counts see sustained increases for at least two weeks after the events.

Table 1 answers **RQ5** and depicts the ten most central accounts during each event. The Boston Police Department is the most central account during the Boston Marathon Bombing, closely followed by the Boston Globe. With the exception of @JFKLibrary and @bostonmarathon, all other accounts belong to news affiliates, three of which are local to the city of Boston (@BostonGlobe, @BostonDotCom, and @7News). Central accounts in Sydney also include government/law enforcement agencies (@nswpolice, @TonyAbbottMHR), news organizations (@abcnews, @BBCBreaking), and several unaffiliated accounts. Response to the Charlie Hebdo Attack is similar in that many central accounts belong to news organizations but differs noticeably with the absence of law enforcement.

Table 1: Central Accounts (Highest Rank to Least)

Rank	Boston	Sydney	Charlie Hebdo
1	BostonPolice	abcnews	itele
2	BostonGlobe	TonyAbbottMHR	Charlie_Hebdo..
3	JFKLibrary	nswpolice	AFP
4	AP	9NewsSyd	AFPphoto
5	BostonDotCom	CottonSydney	Le_Figaro
6	7News	sydneyharbert	jmdecugis
7	bostonmarathon	abit_wp	le_Parisien
8	ReutersUS	WolfSpirit2013	plantu
9	YourAnonNews	9NewsAUS	BFMTV
10	Reuters	BBCBreaking	ctxt.es

Discussion

An important observation is the limited response in Twitter’s sample stream; none of these events significantly altered overall Twitter activity. Given that the US and France account for more than 25% of Twitter’s user base (Richter 2013), we expected such national events to have a stronger effect. Gupta and Kumaraguru also suggest that sharing URLs increase during terrorist events (2012), but we see

this response only when constraining our analysis to relevant tweets. Despite absent overall response, Figure 1 shows a portion of Twitter respond to these events, with relevant tweets accounting for 3 – 4% of all tweet activity on those days. Taken together, surges in these activities suggest higher information sharing during these times even if these surges drop off quickly, which demonstrates Twitter’s short memory as relevant tweet volume dropped by 80% within three days. Follower counts are an exception as they remain high for at least two weeks after the events.

Central accounts also show an interesting result: In all three cases, several news affiliates emerge as leading sources of information. This result is unsurprising given media organizations report breaking news, but the presence of *local* news channels suggests users value information sourced close to the event. This result also corroborates Sutton et al. (2014), who showed local actors emerged as highly influential in Boston. Information from authoritative sources like local law enforcement have more influence on the public than random users and major media organizations (Liu, Fraustino, and Jin 2013), and the popularity of local police is consistent here. Central accounts during the Charlie Hebdo Attacks, however, lack such law enforcement presence, and after searching our data set, we were unable to find any evidence of Parisian law enforcement on Twitter. The resulting information vacuum was filled by media, which may be explained by the presence of a national news agency in France (something the US and Australia lack).

Limitations This work has several limitations we acknowledge. While we focus on Twitter, other platforms should also be considered (e.g., Facebook, reddit). We also likely under-sampled relevant tweets even though we attempted to mitigate this issue by extracting tweets from top actors, which did capture additional data. Finally, though our work is specific to developed, Western countries, results may not hold for terror attacks in less developed countries or where terrorism occurs more often. Future efforts should broaden topical and geographic scope and explore other major social media platforms.

Conclusions

This paper explores public Twitter response by characterizing trends in activities and accounts mentioned during terrorist attacks. Results show relevant tweets, retweets, and hashtags increase significantly and immediately and return to pre-event levels within days. Longer-lived responses appear in URL sharing and followers for central accounts. At the same time, the public coalesces around police/government (where available) and news organizations when sharing information. We see a surge of interest in police agencies during the Boston and Sydney events, but our results show surprisingly little government presence during the Charlie Hebdo Attack.

References

Cha, M.; Haddadi, H.; Benevenuto, F.; and Gummadi, K. P. 2010. Measuring User Influence in Twitter: The Million Fol-

lower Fallacy. In *Fourth International AAAI Conference on Weblogs and Social Media*.

Cross, A. R. 2010. Social Media in Disasters and Emergencies. *American Red Cross Disaster Newsroom* 1–19.

Fraustino, J. D.; Liu, B.; and Yan, J. 2012. Social Media Use during Disasters: A review of the Knowledge Base and Gaps. 1–39.

González-Bailón, S.; Wang, N.; Rivero, A.; Borge-Holthoefer, J.; and Moreno, Y. 2014. Assessing the bias in samples of large online networks. *Social Networks* 38(1):16–27.

Gupta, A., and Kumaraguru, P. 2012. Twitter explodes with activity in mumbai blasts! a lifeline or an unmonitored daemon in the lurking?

Jin, Y.; Fraustino, J. D.; and Liu, B. 2015. The scared, the outraged, and the anxious: How crisis emotions, involvement, and demographics predict publics conative coping. In *the Annual Convention of the International Communication Association*.

Koutra, D.; Bennett, P.; and Horvitz, E. 2015. Events and Controversies: Influences of a Shocking News Event on Information Seeking. In *Proceedings of the 24th International Conference on World Wide Web Companion*.

Kwak, H.; Lee, C.; Park, H.; and Moon, S. 2010. What is Twitter, a social network or a news media? In *Proceedings of the 19th international conference on World wide web*, 591–600. ACM.

Lindsay, B. R. 2011. Social Media and Disasters: Current Uses, Future Options and Policy Considerations. Technical report.

Liu, B.; Fraustino, J. D.; and Jin, Y. 2013. Social Media Use during Disasters: A Nationally Representative Field Experiment. Technical report, College Park, MD.

Morstatter, F.; Pfeffer, J.; Liu, H.; and Carley, K. 2013. Is the sample good enough? Comparing data from Twitter’s streaming API with Twitter’s firehose. *Proceedings of ICWSM* 400–408.

Olteanu, A.; Vieweg, S.; and Castillo, C. 2015. What to Expect When the Unexpected Happens: Social Media Communications Across Crises. In *In Proc. of 18th ACM Computer Supported Cooperative Work and Social Computing (CSCW15)*.

Richter, F. 2013. Twitter’s Top 5 Markets Account for 50% of Active Users.

Rogstadius, J.; Vukovic, M.; Teixeira, C. A.; Kostakos, V.; Karapanos, E.; and Laredo, J. A. 2013. CrisisTracker: Crowdsourced social media curation for disaster awareness. *IBM Journal of Research and Development* 57(5):4:1–4:13.

Sutton, J.; Spiro, E. S.; Fitzhugh, S.; Johnson, B.; Gibson, B.; and Butts, C. T. 2014. Terse Message Amplification in the Boston Bombing Response. *Proceedings of the 11th International ISCRAM Conference* (May):610–619.

Vis, F. 2013. Twitter As a Reporting Tool for Breaking News. *Digital Journalism* 1(1):27–47.