Towards Real-time Remote Social Sensing via Targeted Advertising

Masoomali Fatehkia
Qatar Computing Research Institute, HBKU
mfatehkia@hbku.edu.qa

Muhammad Imran
Qatar Computing Research Institute, HBKU
mimran@hbku.edu.qa

Ingmar Weber*
Saarland Informatics Campus, Saarland University
iweber@cs.uni-saarland.de

ABSTRACT

Social media serves as an important communication channel for people affected by crises, creating a data source for emergency responders wanting to improve situational awareness. In particular, social listening on Twitter has been widely used for real-time analysis of crisis-related messages. This approach, however, is often hindered by the small fraction of (hyper-)localized content and by the inability to explicitly ask affected populations about aspects with the most operational value. Here, we explore a new form of social media data collected through targeted poll ads on Facebook. Using geo-targeted ads during flood events in six countries, we show that it is possible to collect thousands of poll responses within hours of launching the ad campaign, and at a cost of a few (US dollar) cents per response. We believe that this flexible, fast, and affordable data collection can serve as a valuable complement to existing approaches.

Keywords

remote social sensing, real-time polling, flood mapping, Facebook advertising

INTRODUCTION

During crisis response, information is hardest to obtain during the early phase after the disaster onset when the on-the-ground presence of response teams is limited. At the same time, it is exactly during this early phase where a better situational awareness and needs assessment could have the biggest impact on mitigating the crisis impact, and on saving lives (S. E. Vieweg 2012).

To help provide more real-time insights, social listening-based approaches, such as real-time analysis of tweets, have been proposed (S. Vieweg et al. 2010; Castillo 2016). These approaches can provide valuable, complementary insights, but they often lack fine-grained geographic resolution, as only a tiny fraction of tweets is geo-coded (Carley et al. 2016). Additionally, they do not provide a way to guide the data collection by actively inquiring about certain dimensions which might have the biggest operational impact, such as whether roads in certain parts of an area are still navigable.

In this paper, we explore a novel type of social media use for collecting real-time information during disasters. Specifically, we offer a proof-of-concept study on targeted poll advertisements on Facebook as a means of remote sensing in real-time in the aftermath of flood disasters. Poll ads are a type of ads that support collecting data, through polls, natively within the social media feed. See Figure 1 for an example. Such ads can be selectively shown in small geographic regions.

We tested this approach in different flood events across several countries as a feasibility test concerning (i) the geographic targeting capabilities, (ii) the real-time nature of such data, and (iii) the cost-effectiveness. Our results

*corresponding author

demonstrate that within a few hours after launching an ad campaign, thousands of poll responses can be collected across small geographic units at a cost of a few (US dollar) cents per response. At the same time, our analysis also shows that there are data quality challenges related to noticeable baseline rates of "yes" answers, even in implausible settings.

In the following, we first situate this particular type of data within a broader set of "big data" sources that have been used during disaster response. We then give details about Facebook's advertising platform and the particular type of poll ads we used. Next, we describe our experiments with collecting data around flood events in six countries. In particular, we include a number of "lessons learned" for other researchers to build on. Finally, we discuss potential ways forward for this methodology.

RELATED WORK

The need for real-time, highly localized information during disasters has been widely recognized as a challenge for effective disaster response (Grace 2021). In the following, we give a brief overview of some of the non-traditional data sources which have been proposed in this context, including (i) satellite imagery, (ii) mobile phone data (iii) Twitter data, and (iv) surveys through Facebook ads.

Earth Observation (EO) satellite data is being used more and more widely for disaster mapping and damage assessment tasks (Voigt, Giulio-Tonolo, et al. 2016), as well as for mapping the environmental impact of crises. This has included the use of nighttime light intensities for monitoring power outages after a disaster (Molthan and Jedlovec 2013), and studying armed conflicts (Li et al. 2013). The use of daytime satellite imagery has ranged from monitoring the human and environmental effects of violent conflicts (Witmer 2015), mapping and monitoring floods (Schumann et al. 2018; Moumtzidou et al. 2020) to rapid mapping and damage assessment for a variety of other sudden-onset[1] disasters (Voigt, Kemper, et al. 2007). Satellite imagery is clearly an important data source, and likely to grow in importance, but (i) its availability can be limited due to cloud cover, cost constraints, or revisit rates, and (ii) it lacks the *social*, i.e., human dimension as similar disaster images might be linked to different social impacts.

Mobile phone Call Detail Records (CDR) are another source of remote sensing data which have been utilized in a variety of disaster management applications. These applications have included the use of CDR data to study changes in communications patterns during and after floods (Ghurye et al. 2016; Hong et al. 2018), to estimate the size and trends in population mobility following earthquakes (Bengtsson et al. 2011; Wilson et al. 2016) and to infer spatial distribution of displaced populations after a cyclone (Cumbane and Gidófalvi 2021). In particular, the mobility aspect of CDR data has huge potential value for mapping disaster displacement. At the same time, limitations include (i) difficulty of access due to considerable privacy risks, and (ii) challenges around sensing real-time changes in well-being merely from, say, changes in calling behavior.

On the social media side, there is extensive work on the use of Twitter for disaster management and recovery efforts. Earlier work has looked at information sharing behavior on Twitter after flood events (Palen et al. 2010). Information extracted from Twitter messages has been used to detect a variety of sudden-onset disasters events ranging from floods (Arthur et al. 2018; Bono et al. 2022), wild-fires (Boulton et al. 2016; Papadimos et al. 2022) to predicting the center and trajectory of earthquakes (Sakaki et al. 2010). Beyond incident detection, Wang and Taylor 2016 used geo-located tweets to study the impact of disasters on urban mobility patterns while Kryvasheyeu et al. 2016 used Twitter activity during Hurricane Sandy for damage assessment. Social media data have also been used in combination with other data sources such as satellite imagery for flood mapping (Akhtar et al. 2021). Other efforts have focused on building systems to collect, classify and filter Twitter messages in real-time to assist with disaster response activities (Imran et al. 2014; Castillo 2016).

A limitation of Twitter mining is that there is no way to actively ask about a topic as one has to rely on passive listening from messages posted by users. Moreover, the percentage of geo-located content at the sub-city level is very, very small (Kruspe et al. 2021). In order to collect data actively from Twitter users, Avvenuti et al. 2017 proposed a hybrid crowdsensing approach which they tested in earthquake emergencies. In this approach, relevant tweets were first filtered out in order to identify potential users to approach. In the subsequent phase, these users were messaged via a tweet requesting more information. The responses were then recorded and analyzed. The active engagement of targeted Twitter users offers a more guided data collection, compared to passively listening. However, this approach cannot be spatially targeted, and risks getting blocked by Twitter due to potential violation of the Terms of Service (ToS), which limited unsolicited and automated interactions between bots and users.

[1]Based on reviewer feedback, we have decided to avoid the term "natural disaster" in this publication (see https://www.nonaturaldisasters.com/).

Closer to our work, previous studies have explored the use of targeted advertising on Facebook to cost-effectively recruit participants for survey research (Grow et al. 2022; Neundorf and Öztürk 2023). These surveys administered through Facebook ads have been used to study perceptions and responses to the COVID-19 pandemic (Perrotta et al. 2021), to survey samples of migrant populations (Pötzschke and Braun 2017) as well as for cost-effectively surveying populations in the Global South (Rosenzweig et al. 2020). However, none of the aforementioned studies explored data collection in a real-time context, such as in the immediate 24-hour period after a sudden-onset event, and at sub-city spatial resolutions.

Another relevant effort led by Facebook is the Facebook Safety Check feature which is activated during disaster events by sending notifications to Facebook users in the area asking them to mark themselves as safe (Lee 2019); users can then see a list of their Facebook friends who have also marked themselves as safe and can invite other friends to receive and respond to the notification. The data gathered through this feature is also integrated into Facebook's disaster maps[2] which provide information for humanitarian agencies. While Facebook poll ads provide somewhat related functionality, they offer more flexibility on the choice of questions asked and can be demographically targeted to specific demographic subgroups of interest such as women, for example. Moreover, poll ads can be utilized independently of the Facebook Safety Check feature which has not always been activated for all crisis events[3].

Against this backdrop of existing data sources, our proposed methodology offers a complementary approach by (i) offering sub-city geographic resolution, (ii) providing data within hours of starting the collection, and (iii) allowing an open-ended set of questions to be asked directly to the affected population.

THE FACEBOOK ADVERTISING PLATFORM

As a social media platform that is free to its users, Facebook generates revenues from targeted advertising. Advertisers are provided with a rich array of targeting options to enable them to reach their desired audience. These targeting options include targeting users based on geographic location and demographics such as age and gender. Furthermore, in order to suit the needs of different advertisers, Facebook provides a variety of ad formats ranging from images and videos to more interactive content with links to external sites. While these features make the platform valuable for businesses advertising their goods and services, they also serve as a useful tool for research by enabling researchers to gather information from the Facebook user base through advertising surveys (Grow et al. 2022). In this paper, we use a specific feature, namely poll ads, as a means of remote social sensing from Facebook users. We chose Facebook for our experiments as it uniquely offers (i) native within-app poll ads, (ii) fine-grained geographic ad targeting, and (iii) has high levels of market penetration in many low-and-middle-income countries. This section provides more details about poll ads and how we set up our advertising campaigns on Facebook.

Facebook Poll Advertisements

Poll advertisement[4] is one of the several supported ad formats on Facebook. They provide advertisers with a means of engaging their target audience and receiving feedback through a poll. This ad format consists of a video with a poll appearing underneath. The poll consists of a single question along with two possible answer choices. Any user who sees the ad can participate by clicking on one of the two answer choices. The responses are then recorded and reported to the advertiser in aggregate, showing the total number of users who participated as well as the number who voted for each of the two answer choices. The advertiser is not provided with information on the identity of the individual users who participated or how they voted.

Poll ads are displayed in the user's Facebook Feed alongside other content. As they are an advertising format natively supported by the platform, a user seeing the ad can respond to the poll while browsing through other content without the need to leave the platform. This is in contrast with survey-based advertising where the user is required to click on a link directing them to an external online survey. This feature makes poll ads ideal for rapid remote social sensing in a disaster context as they provide a simple and convenient way to gather basic information from users in a particular location.

Setting up a Facebook poll ad

The first required step before running any advertisement is to create a Facebook page with which the ad will be associated. For our experiments, we created a Facebook page called "Disaster Polling"[5] consisting of a short

[2] https://research.facebook.com/blog/2017/6/facebook-disaster-maps-methodology/
[3] https://www.wired.co.uk/article/the-inside-story-of-how-facebook-is-transforming-disaster-response
[4] https://www.facebook.com/business/help/215415943324822?id=603833089963720
[5] https://www.facebook.com/DisasterPolling

introductory post and a disaster-related profile picture. The introductory post contained a brief overview of the research project in order to provide some information for users who might look at the page after seeing the ad. Once the page had been set up, the next step was to launch an advertising campaign. This was done through the Ads Manager[6], Facebook's online platform for creating and managing advertising campaigns.

Figure 1 (Panel A) depicts an example poll advertisement that we ran as part of our experiments. The top portion of the ad displays the name and profile picture of the associated Facebook page. For the ad content, we used a short video consisting of two images each shown for a duration of two seconds. There was no sound. The first image shown in the video was an image of a flooded location. In order to make the image more relevant to the specific geographic location where the ad was shown, we varied this image across different ad campaigns. Three different images were used across all of our advertising campaigns, as shown in the sample Ad in Panel A and the images (i) and (ii) in Panel B. Image (iii) in Panel B was the second image shown in the video; this second image was the same for all the ads.

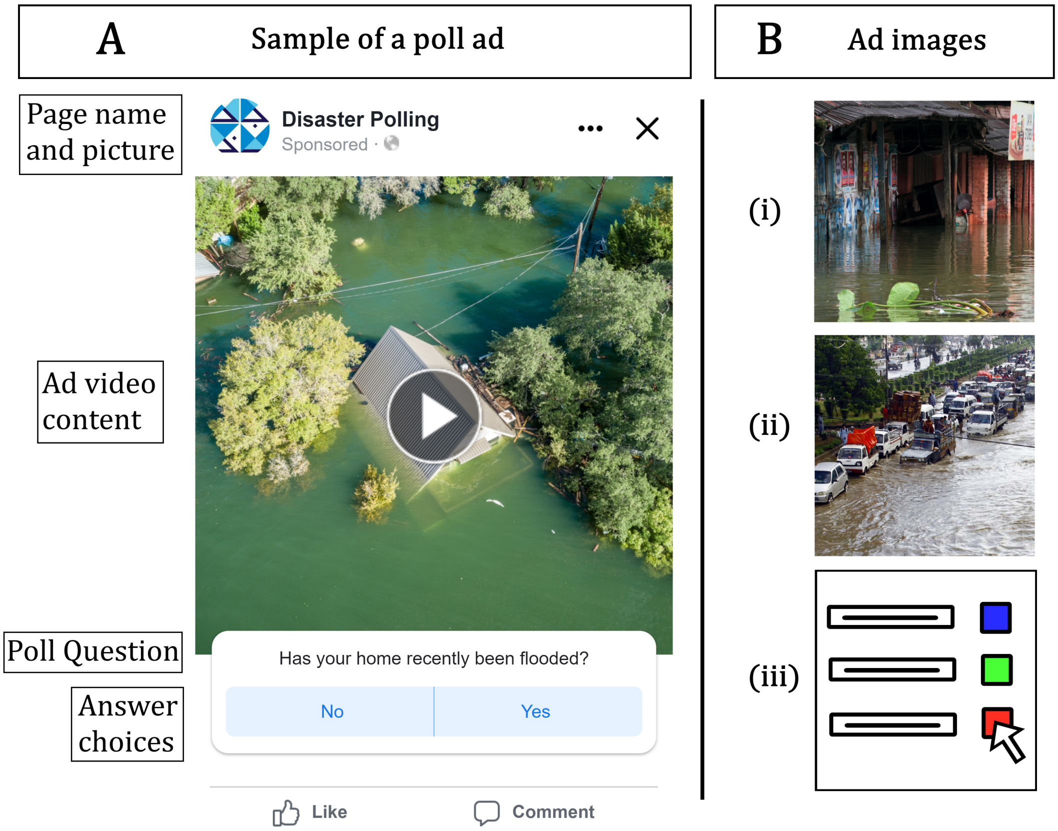

Figure 1. Panel A shows an example of a poll advertisement. The Ad consists of a video with a poll question and two answer choices. Panel B shows other images that were used in the ad videos.

EXPERIMENTS WITH FLOOD EVENTS

Table 1 shows a summary of the advertising campaigns that were run as part of our experiments. We ran poll advertisements for a total of ten different flood events across six countries. The flood events were identified by monitoring news sources and social media such as Twitter for mentions of ongoing flood disasters. Once we had identified an event, the next step was to identify locations of interest that may have potentially been affected by floods. We attempted to gather a list of such locations by searching news articles and social media reports for names of affected locations. In addition to the list of potentially affected areas, we also tried to identify one or two control location(s). The control location was another city or area in the same country for which we could not find any news mentions of recent floods. The control location was included for comparison in order to assess what poll responses looked like in areas that were not affected by floods. The total number of locations, including control locations, that were targeted for each flood event are indicated in Table 1.

[6] https://www.facebook.com/business/tools/ads-manager

The smallest geographic resolution for which ads can be targeted on Facebook are custom locations which consist of a radius around a given latitude and longitude coordinate. The smallest permitted radius for custom locations is one mile. In order to gather data at finer spatial resolutions, we used custom locations with a one mile radius to target the locations of interest. The exception to this were the ads for the control locations which were targeted at the city level. Based on experience from earlier campaigns, in the last two experiments (E9 and E10) the radius of the custom locations were allowed to vary. Covering a larger area by allowing a bigger radius was necessary in order to ensure a sufficient number of poll responses for locations with few users.

Across rounds of advertising campaigns, we experimented with different elements of the ad content. Table 1 provides details about the ad content for each experiment. The poll question was slightly varied in the initial campaigns, after which we continued to use the more concise 'Has your home recently been flooded?' prompt for the poll. Another varying element of the poll was the order of the response choices, i.e whether the *Yes* button appeared on the right or the left. Both variants were explored in the third campaign, the results of which are reported later in the paper. In most campaigns, we explored two variants of the ad, one in English and one in the local language. The local language version of the ad was translated from English using the Google translate[7] tool. Where possible, the translated version was verified with a speaker of that language. In the next sections we provide more information on the time and cost of running these ad experiments.

Table 1. Summary of the advertising campaigns that were run as part of our experiments. The poll question was slightly varied across different campaigns: (a) Did you have to leave your home due to flooding?, (b) Did you have to leave your home in the past week due to flooding?, (c) Has your home recently been flooded?, (d) Has your neighborhood recently been flooded?; Where the ad was in a language other than English, a translated version of the above questions was used. The poll response layout refers to the order in which the Yes and No answer choices were presented from left to right, eg: YN indicates Yes on the left and No on the right. The first image in the ad videos was also varied as shown in figure 1: (I) figure 1 panel A, (II) figure 1 panel B(i), (III) figure 1 panel B(ii). *The campaign for Australia did not run to completion. **The Hindi ads did not run for very long as the campaign was nearly over by the time the ads passed review.

	Flood event	*Ad language(s)*	*Poll question*	*Poll response layout*	*Ad video image*	*Number of locations targeted*
E1	Maguindanao, Philippines, Flash floods	English	a	YN	I	3
E2	Bangalore, India floods	English, Hindi	b	YN	I	2
E3	Bangladesh flooding in North and North-East regions	English	c	YN, NY	I	2
E4	Bangladesh floods	English, Bengali	c	NY	II	14
*E5**	Western Australia flooding	English	c	NY	I	26
*E6***	Manikaran, India flash floods	English, Hindi	c	NY	II	15
E7	Flash floods in several villages of Kinnaur district, India	English, Hindi	c	NY	II	10
E8	Karachi floods, Pakistan	English, Urdu	c	NY	III	26
E9	Odish floods, India	English, Hindi	c	NY	II	13
E10	Kinshasa floods, DRC Congo	French	c, d	NY	II	7

Time

Collecting information in the immediate aftermath of a disaster is crucial for planning and allocating humanitarian resources. In this section we present information about the temporal aspect of collecting data through poll advertisements based on our experiments. There are several factors that affect the speed with which information can be collected through poll advertisements after a flood event. These include (i) how long it takes to prepare the ad after the event, (ii) the time it takes for the published ads to undergo ad review and be approved to run, and (iii) how long the campaign runs to get responses. This information is provided in Table 2, which reports the date(s) of each flood event, the scheduled start date/time, and the duration of the ad campaigns.

[7] https://translate.google.com/

Table 2. Details about the time and responses received for the ad campaigns. The ad review time is the range of observed review times across all ads in each campaign. For campaigns with ads in more than one language, we report the review times separately by language only when there was a significant difference; otherwise, when the review times were similar for the different language ads, we report the overall range of review times across all ads (this was the case for campaigns E2, E4 and E8). Reach is defined as 'the number of people who saw your ads at least once'. Responses is the number of poll responses. Response rate is the ratio of Responses to the ad Reach multiplied by 100. *The ad campaign for Australia did not run to completion (the ads ran for 3 hours). **The Hindi ads did not run for very long as the campaign was nearly over by the time the ads passed review.

	Event date(s) (2022)	*Ad campaign start date*	*Campaign duration (hours)*	*Ad review time*	*Reach*	*Responses*	*Response rate*
E1	27 Apr	27 Apr 4:30pm	24	16-16.5 hours	7,134	78	1.09 %
E2	17 May	18 May 1:00pm	24	6.3-7 hours	26,159	83	0.32 %
E3	16-18 Jun	19 Jun 10:20am	24	3.2 hours	71,088	866	1.22 %
E4	22 Jun	22 Jun 10:00pm	24	5-10 min	29,415	680	2.31 %
*E5**	3 Jul	4 Jul 2am	6	5-35 min	2,705	21	0.78 %
*E6***	6 Jul	6 Jul 5:20pm	6	English: 5-30 min; Hindi: 5-7 hours	24,735	69	0.28 %
E7	18 Jul	19 Jul 1pm	9	English: 5-30min; Hindi: 5-6 hours	18,152	98	0.54 %
E8	24-25 Jul	26 Jul 8am	6	5-25 min	65,024	1,632	2.51 %
E9	17 Aug	17 Aug 6pm	6	English: 5-20min; Hindi: 6-6.5 hours	124,190	1,209	0.97 %
E10	13 Dec	14 Dec 1pm	6	11 hours	31,696	578	1.82 %

Among the factors discussed above, the time taken to prepare and publish an ad after an event is likely to be minimal once a flood event and locations of interest are identified. Using the Facebook Ads Manager, it is possible to create drafts of campaigns and ads in advance. These could then be populated with the necessary information about the target locations and the scheduled run time of the ads for that flood event before being published. The ad creation process could also be potentially automated by using Facebook's Marketing Application Programming Interface (API), which provides functionality for creating and managing Facebook ads programmatically.

All published ads on Facebook must undergo an ad review process before they can go live and start running[8]. This is done to ensure compliance with Facebook's advertising policies and community standards. According to Facebook, the ad review process is generally completed within twenty-four hours and is mostly automated, though in some cases a manual review may be required. The length of the ad review process could therefore present a bottleneck to how rapidly poll ads can be launched in a time-sensitive context. Table 2 provides a summary of the ad review times that we observed in our campaigns. We experienced the longest review time (16.5 hours) in our first campaign which was also the first advertising campaign under that account. Thereafter, in subsequent campaigns, the ad review time dropped significantly to as low as five minutes for some ads.

We also observed differences in ad review times across ad languages and geographic locations. Despite multiple campaigns in the Hindi language, we continued to experience long ad review times of up to seven hours for ads in that language while the equivalent English ads for the same locations were approved within half an hour. Our ad for the floods in Kinshasa (E10) which was in a new location and language also took significantly longer to be approved (11 hours). Overall, our empirical observations suggest that, barring possible exceptions for some languages/locations, generally ad review times are minimal and should not present a significant bottleneck to data collection after a disaster. Moreover, it appears that ad review times might reduce over time as an advertising account builds 'credibility' after multiple successful campaigns with no violations.

Once an ad is running, the time it takes to collect responses will influence how rapidly information can be gathered. Table 2 provides information about how long each campaign was scheduled to run as well as information on the reach and number of poll responses aggregated across all ads in each campaign. The initial campaigns were scheduled for twenty-four hours while in later campaigns we reduced this to six hours. Note that in some cases, the actual campaign may have run for a shorter duration than scheduled due to the ad review time; for example, in the first campaign the ads only started running about sixteen hours into the scheduled duration.

[8]https://www.facebook.com/business/news/facebook-ad-policy-process-and-review

As shown in the table, the ads were viewed by a large number of users and attracted responses with response rates as high as 2.5%. This was the case for the day-long campaigns as well as for the shorter campaigns. Please note that the absolute numbers for the reach and responses are not directly comparable across campaigns as they differed in terms of duration, the number of available Facebook users in the targeted locations and budget. In the next section, we provide more information on the budget and cost of the campaigns.

Cost

Table 3 presents information on the cost of running the advertising campaigns. Each ad campaign had a budget which was specified at the ad set level. In our case, each ad set corresponded to one targeted location. The total budget for the campaign is the budget per ad set multiplied by the number of ad sets (targeted locations). The budget per ad set varied across our advertising campaigns. The table also reports the final cost of running the campaign. The final cost might be lower than the budgeted amount, for example when there were too few users in a location to display the ads.

The table also presents information about the cost per people reached and the cost per poll response. Information about the reach and number of responses received was presented earlier in Table 2. The ad campaign in Australia (E5) was by far the most expensive with a cost per poll response of close to a dollar. The cost per poll response was considerably lower for the other campaigns. The higher cost could be due to differences across geographic markets which may be more competitive and hence more expensive in some countries. Nonetheless, for many of the other countries in our campaign we were able to acquire poll responses for as low as a few cents per response.

Table 3. Summary of the cost of running the advertising campaigns. All monetary figures are quoted in US Dollars. The cost per 1000 people reached is the amount spent divided by the Reach (see Table 2) multiplied by 1000. The cost per response is the amount spent divided by the number of poll responses.

	Budget per Ad set	*Total budget*	*Amount spent*	*Cost per 1000 people reached*	*Cost per response*
E1	5	15	11.69	1.64	0.150
E2	5	10	10	0.38	0.120
E3	5	10	10	0.14	0.012
E4	2	28	24.32	0.83	0.036
E5	2	52	19.91	7.36	0.948
E6	2	30	12.49	0.5	0.181
E7	2	20	7.93	0.44	0.081
E8	2	52	38.96	0.6	0.024
E9	4	52	50	0.4	0.041
E10	4	28	28	0.88	0.048

Some lessons learned

Baseline rate of yes responses

In addition to the potentially flood affected locations, all of our advertising campaigns included a control location for which we could not find any recent reports of floods. The responses from the control locations can serve as a baseline for comparison between affected and not-affected locations. In an ideal setting, one would expect that all respondents from the control locations would reply *No* when asked about flooding in their homes. However, our empirical observations suggest that this is generally not the case as a certain percentage of respondents from control locations do answer *Yes* to the flood poll question.

Figure 2 shows a box-plot of the distribution of the percentage of respondents who replied *Yes* (their home/area was flooded) to the poll question, across all of our ads, disaggregated by whether the locations was a control or not. With one exception, the percentage of respondents who replied *Yes* was non-zero for all the control locations, ranging from 9.8% to 47.2% with a median of 26.6%. For the other potentially flood affected locations, the percentage of *Yes* responses was more variable with a median of 32.6%.

Based on the above results, the percentage of *Yes* responses is non-zero and in some cases quite high for the control locations where one would have expected no floods to be reported. For some of the flood events, large swathes of the country were affected which may have influenced the observed responses from the control location. Experimenting with more specific questions, e.g., adding location or time information about the event of interest, could shed light on whether this is part of the reason behind such *Yes* responses. It is also plausible that there may be a certain percentage of error responses due to users accidentally hitting a response while scrolling their phone[9]. Overall, these results highlight the importance of including some baseline measures from other not-affected locations. These baseline rates should then be taken into account when interpreting the poll results.

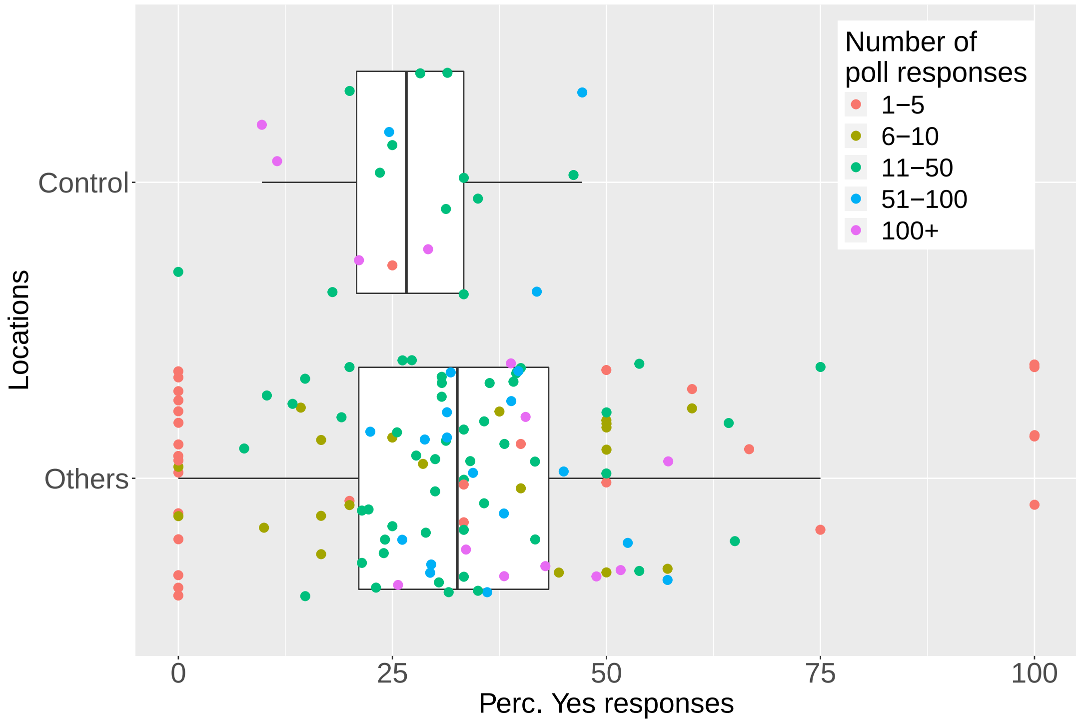

Figure 2. Distribution, across all campaigns, of the percentage of respondents who answered "Yes" to the flood poll question, indicating that their home had recently been flooded. The plot is shown separately for ads targeted to the control locations (top) versus other potentially affected locations (bottom). The individual dots show the aggregate "Yes" percentages for each poll ad with the number of responses indicated by the color of each point. The ads for E5 and Hindi ads for E6 are excluded as they did not run to completion.

Layout of the response choices

Table 4 displays the results from E3 campaign where two variants of the ad were shown in each location with the position of the *Yes* and *No* response choices swapped: either *Yes* on the left (YN) or *No* on the left (NY). As can be seen in the table, for the flooded location the percentage of yes responses was roughly the same for both ad variants. However, in the control location, the percentage of yes responses decreased by about half when the *No* option was placed on the left-hand side; this is more consistent with results that would be expected in a control location where there were no floods. Based on these observations, we used the NY response layout for subsequent campaigns thereafter.

Table 4. Results for ad campaign E3. Two variants of the ad were launched, differing in terms of the order of the response choices, with either Yes on the left (YN) or No on the left (NY).

Location name	*location*	*Ad variant*	*Poll responses*	*Yes responses*	*Perc. of Yes responses*
Sunamganj	Flood	YN	98	56	57.14 %
Sunamganj	Flood	NY	332	190	57.22 %
Khulna	Control	YN	313	66	21.09 %
Khulna	Control	NY	123	12	9.76 %

[9]After hitting a response choice on the poll, users are shown the poll results and would not be able to change their responses: `https://www.facebook.com/business/help/437605740190763?id=603833089963720`

User comments/reactions to the ads

Facebook users can interact with an ad by posting a comment or sharing a reaction from a set of available choices. The available reactions are Like, Love, Care, Haha, Wow, Sad and Angry. We received a total of five comments across all the ads, three of which were single-word comments expressing grief and in one case a comment replying No in response to the poll. Two users posted images that were not relevant to the ad or the flood context. Some ads also received reactions with Like as the most common reaction, while we also observed the other reaction types. While the user reactions and comments are not directly relevant to gathering data through the poll responses, they could potentially be relevant for emergency responders in case users share more information through a comment during a disaster. A different poll design might also explicitly try to encourage such commenting behavior.

DISCUSSION

In this paper, we explored the use of poll advertisements on Facebook as a tool for remote social sensing within the context of flood disasters. We tested this approach for ten different flood events across six countries. The results from our ad campaigns highlight the timeliness and cost-effectiveness of this approach for gathering data on flood-affected areas at sub-city spatial resolution. Using this approach, it is possible to gather information within several hours of an event. Moreover, it is possible to do so with small budgets as we observed costs as low as a few cents per poll response.

An important consideration in using this approach for remote social sensing is the quality of the poll responses. Here the question of interest would be the extent to which the spatial patterns of positive responses to a poll asking about floods in the user's area reflect the actual spatial variations in the extent of flooding. One challenge with validating the poll responses, which also motivates this work, is the lack of available ground truth, especially in the immediate aftermath of a disaster. Nonetheless, a potential approach to validation could include comparing the poll responses with other data sources where available. These could include flood extent maps derived from satellite imagery or on-the-ground reports collected through surveys after the disaster. We see such in-depth validation as the next step for this current proof-of-concept work.

As with any study that uses social media data sources, one has to account for potential sources of bias and how these may affect the results. With poll ads on Facebook, we are only able to probe a sample of the population who are digitally connected and use Facebook. To limit the impact of selection bias on the results, we suggest focusing on questions where a small set of respondents can answer on behalf of a geographic area. For example, flooding or the accessibility of roads is likely to affect the whole neighborhood, but personal well-being might vary even from neighbor to neighbor. Similarly, we believe that platform-specific response biases might be more pronounced for personal-level concerns, such as questions around gender-based violence, and less pronounced for externally visible neighborhood-level concerns.

Our results here highlight the feasibility, in terms of time and cost, of using Facebook poll ads to gather information at fine-grained spatial resolutions during a flood disaster. An in-depth validation of these poll responses and an investigation of sources of bias are important avenues for future work. Also, more work can be done on experimenting with various ad settings as well as the ad content (for example, the images used, how the poll question is phrased) in order to better understand how these influence the performance of the ads and the data that is obtained. Finally, this approach to remote social sensing is not limited to flood events and can be tested and deployed in a variety of other disaster settings.

REFERENCES

Akhtar, Z., Ofli, F., and Imran, M. (2021). "Towards Using Remote Sensing and Social Media Data for Flood Mapping". In: *ISCRAM 2021 Conference Proceedings – 18th International Conference on Information Systems for Crisis Response and Management*. Blacksburg, VA (USA): Virginia Tech, pp. 536–551.

Arthur, R., Boulton, C. A., Shotton, H., and Williams, H. T. P. (Jan. 2018). "Social sensing of floods in the UK". en. In: *PLOS ONE* 13.1, e0189327.

Avvenuti, M., Bellomo, S., Cresci, S., La Polla, M. N., and Tesconi, M. (Apr. 2017). "Hybrid Crowdsensing: A Novel Paradigm to Combine the Strengths of Opportunistic and Participatory Crowdsensing". In: *Proceedings of the 26th International Conference on World Wide Web Companion*. WWW '17 Companion. Republic and Canton of Geneva, CHE: International World Wide Web Conferences Steering Committee, pp. 1413–1421.

Bengtsson, L., Lu, X., Thorson, A., Garfield, R., and Schreeb, J. v. (Aug. 2011). "Improved Response to Disasters and Outbreaks by Tracking Population Movements with Mobile Phone Network Data: A Post-Earthquake Geospatial Study in Haiti". en. In: *PLOS Medicine* 8.8, e1001083.

Bono, C. A., Pernici, B., Fernandez-Marquez, J. L., Shankar, A. R., Mülâyim, M. O., and Nemni, E. (2022). "TriggerCit: Early Flood Alerting using Twitter and Geolocation - A Comparison with Alternative Sources". In: *ISCRAM 2022 Conference Proceedings – 19th International Conference on Information Systems for Crisis Response and Management*. Tarbes, France, pp. 674–686.

Boulton, C., Shotton, H., and Williams, H. (2016). "Using Social Media to Detect and Locate Wildfires". en. In: *Proceedings of the International AAAI Conference on Web and Social Media* 10.2, pp. 178–186.

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". In: *Safety science* 90, pp. 48–61.

Castillo, C. (2016). *Big crisis data: social media in disasters and time-critical situations*. Cambridge University Press.

Cumbane, S. P. and Gidófalvi, G. (June 2021). "Spatial Distribution of Displaced Population Estimated Using Mobile Phone Data to Support Disaster Response Activities". en. In: *ISPRS International Journal of Geo-Information* 10.6, p. 421.

Ghurye, J., Krings, G., and Frias-Martinez, V. (June 2016). "A Framework to Model Human Behavior at Large Scale during Natural Disasters". In: *2016 17th IEEE International Conference on Mobile Data Management (MDM)*. Vol. 1, pp. 18–27.

Grace, R. (2021). "Toponym usage in social media in emergencies". In: *International Journal of Disaster Risk Reduction* 52, p. 101923.

Grow, A., Perrotta, D., Del Fava, E., Cimentada, J., Rampazzo, F., Gil-Clavel, S., Zagheni, E., Flores, R. D., Ventura, I., and Weber, I. (2022). "Is Facebook's advertising data accurate enough for use in social science research? Insights from a cross-national online survey". en. In: *Journal of the Royal Statistical Society: Series A (Statistics in Society)* n/a.n/a ().

Hong, L., Lee, M., Mashhadi, A., and Frias-Martinez, V. (2018). "Towards Understanding Communication Behavior Changes During Floods Using Cell Phone Data". en. In: *Social Informatics*. Ed. by S. Staab, O. Koltsova, and D. I. Ignatov. Lecture Notes in Computer Science. Cham: Springer International Publishing, pp. 97–107.

Imran, M., Castillo, C., Lucas, J., Meier, P., and Vieweg, S. (Apr. 2014). "AIDR: artificial intelligence for disaster response". In: *Proceedings of the 23rd International Conference on World Wide Web*. WWW '14 Companion. New York, NY, USA: Association for Computing Machinery, pp. 159–162.

Kruspe, A., Häberle, M., Hoffmann, E. J., Rode-Hasinger, S., Abdulahhad, K., and Zhu, X. X. (Nov. 2021). "Changes in Twitter geolocations: Insights and suggestions for future usage". In: *Proceedings of the Seventh Workshop on Noisy User-generated Text (W-NUT 2021)*. Online: Association for Computational Linguistics, pp. 212–221.

Kryvasheyeu, Y., Chen, H., Obradovich, N., Moro, E., Van Hentenryck, P., Fowler, J., and Cebrian, M. (Mar. 2016). "Rapid assessment of disaster damage using social media activity". In: *Science Advances* 2.3, e1500779.

Lee, K. S. (Oct. 2019). "Explicit Disaster Response Features in Social Media: Safety Check and Community Help Usage on Facebook during Typhoon Mangkhut". In: *Proceedings of the 21st International Conference on Human-Computer Interaction with Mobile Devices and Services*. MobileHCI '19. New York, NY, USA: Association for Computing Machinery, pp. 1–12.

Li, X., Chen, F., and Chen, X. (Oct. 2013). "Satellite-Observed Nighttime Light Variation as Evidence for Global Armed Conflicts". In: *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 6.5, pp. 2302–2315.

Molthan, A. and Jedlovec, G. (2013). "Satellite Observations Monitor Outages From Superstorm Sandy". en. In: *Eos, Transactions American Geophysical Union* 94.5, pp. 53–54.

Moumtzidou, A., Bakratsas, M., Andreadis, S., Karakostas, A., Gialampoukidis, I., Vrochidis, S., and Kompatsiaris, I. (2020). "Flood detection with Sentinel-2 satellite images in crisis management systems". In: *ISCRAM 2020 Conference Proceedings – 17th International Conference on Information Systems for Crisis Response and Management*. Blacksburg, VA (USA): Virginia Tech, pp. 1049–1059.

Neundorf, A. and Öztürk, A. (Feb. 2023). "How to improve representativeness and cost-effectiveness in samples recruited through meta: A comparison of advertisement tools". In: *PLOS ONE* 18.2, pp. 1–18.

Palen, L., Starbird, K., Vieweg, S., and Hughes, A. (2010). "Twitter-based information distribution during the 2009 Red River Valley flood threat". en. In: *Bulletin of the American Society for Information Science and Technology* 36.5, pp. 13–17.

Papadimos, T., Pantelidis, N., Andreadis, S., Bozas, A., Gialampoukidis, I., Vrochidis, S., and Kompatsiaris, I. (2022). "Real-time Alert Framework for Fire Incidents Using Multimodal Event Detection on Social Media Streams". In: *ISCRAM 2022 Conference Proceedings – 19th International Conference on Information Systems for Crisis Response and Management*. Tarbes, France, pp. 623–635.

Perrotta, D., Grow, A., Rampazzo, F., Cimentada, J., Fava, E. D., Gil-Clavel, S., and Zagheni, E. (Dec. 2021). "Behaviours and attitudes in response to the COVID-19 pandemic: insights from a cross-national Facebook survey". en. In: *EPJ Data Science* 10.1, p. 17.

Pötzschke, S. and Braun, M. (Oct. 2017). "Migrant Sampling Using Facebook Advertisements: A Case Study of Polish Migrants in Four European Countries". en. In: *Social Science Computer Review* 35.5, pp. 633–653.

Rosenzweig, L., Bergquist, P., Pham, K. H., Rampazzo, F., and Mildenberger, M. (Oct. 2020). *Survey sampling in the Global South using Facebook advertisements*. en-us.

Sakaki, T., Okazaki, M., and Matsuo, Y. (Apr. 2010). "Earthquake shakes Twitter users: real-time event detection by social sensors". In: *Proceedings of the 19th international conference on World wide web*. WWW '10. New York, NY, USA: Association for Computing Machinery, pp. 851–860.

Schumann, G. J.-P., Brakenridge, G. R., Kettner, A. J., Kashif, R., and Niebuhr, E. (Aug. 2018). "Assisting Flood Disaster Response with Earth Observation Data and Products: A Critical Assessment". en. In: *Remote Sensing* 10.8, p. 1230.

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 the SIGCHI conference on human factors in computing systems*, pp. 1079–1088.

Vieweg, S. E. (2012). "Situational awareness in mass emergency: A behavioral and linguistic analysis of microblogged communications". PhD thesis. University of Colorado at Boulder.

Voigt, S., Giulio-Tonolo, F., Lyons, J., Kučera, J., Jones, B., Schneiderhan, T., Platzeck, G., Kaku, K., Hazarika, M. K., Czaran, L., et al. (July 2016). "Global trends in satellite-based emergency mapping". In: *Science* 353.6296, pp. 247–252.

Voigt, S., Kemper, T., Riedlinger, T., Kiefl, R., Scholte, K., and Mehl, H. (June 2007). "Satellite Image Analysis for Disaster and Crisis-Management Support". In: *IEEE Transactions on Geoscience and Remote Sensing* 45.6, pp. 1520–1528.

Wang, Q. and Taylor, J. E. (Jan. 2016). "Patterns and Limitations of Urban Human Mobility Resilience under the Influence of Multiple Types of Natural Disaster". en. In: *PLOS ONE* 11.1, e0147299.

Wilson, R., Erbach-Schoenberg, E. zu, Albert, M., Power, D., Tudge, S., Gonzalez, M., Guthrie, S., Chamberlain, H., Brooks, C., Hughes, C., et al. (Feb. 2016). "Rapid and Near Real-Time Assessments of Population Displacement Using Mobile Phone Data Following Disasters: The 2015 Nepal Earthquake". In: *PLoS Currents* 8, ecurrents.dis.d073fbece328e4c39087bc086d694b5c.

Witmer, F. D. W. (May 2015). "Remote sensing of violent conflict: eyes from above". In: *International Journal of Remote Sensing* 36.9, pp. 2326–2352.