

Integrating Social Media Communications into the Rapid Assessment of Sudden Onset Disasters

Sarah Vieweg, Carlos Castillo, and Muhammad Imran

Qatar Computing Research Institute, Doha, Qatar
svieweg@qf.org.qa, chato@acm.org, mimran@qf.org.qa

Abstract. Recent research on automatic analysis of social media data during disasters has given insight into how to provide valuable and timely information to formal response agencies—and members of the public—in these safety-critical situations. For the most part, this work has followed a bottom-up approach in which data are analyzed first, and the target audience’s needs are addressed later. Here, we adopt a top-down approach in which the starting point are information needs. We focus on the aid agency tasked with coordinating humanitarian response within the United Nations: OCHA, the Office for the Coordination of Humanitarian Affairs. When disasters occur, OCHA must quickly make decisions based on the most complete picture of the situation they can obtain. They are responsible for organizing search and rescue operations, emergency food assistance, and similar tasks. Given that complete knowledge of any disaster event is not possible, they gather information from myriad available sources, including social media.

In this paper, we examine the rapid assessment procedures used by OCHA, and explain how they executed these procedures during the 2013 Typhoon Yolanda. In addition, we interview a small sample of OCHA employees, focusing on their uses and views of social media data. In addition, we show how state-of-the-art social media processing methods can be used to produce information in a format that takes into account what large international humanitarian organizations require to meet their constantly evolving needs.

Keywords: Crisis informatics; Microblogging; Humanitarian computing

1 Introduction

The role of social media as a conduit for useful information during emergencies is increasingly acknowledged and accepted by formal response and humanitarian agencies [11]. We focus on the information gathering processes of a large, diverse, international humanitarian relief agency. We explain how the United Nations Office for the Coordination of Humanitarian Affairs (UN OCHA, or OCHA), views social media data, which are considered a legitimate source of information during the data-gathering process OCHA goes through when they respond to a crisis (though concerns exist) [27]. In addition, we discuss apprehensions some have about incorporating social media data in rapid assessment procedures. Informed by these observations, we perform both human and automatic analyses of tweets broadcast during Typhoon Yolanda, and provide the

findings to OCHA. Each of these steps research leads to suggestions for how humanitarian agencies can use social media communications, and critically, on how methods to process social media data can effectively support these agencies.

1.1 Related Work

Much research on processing social media data during emergencies has focused on applying computational methods—such as Information Retrieval, Natural Language Processing, and/or Machine Learning—to the creation of systems for filtering, classifying, and summarizing messages (for a recent survey, see [12]).

In addition, an interdisciplinary line of work looks at how the end-users of these systems—including various formal response organizations and agencies—use social media during disasters to understand the situation, coordinate relief efforts, and manage information. For instance, several articles analyze how social media was used by response agencies during the 2010 Haiti earthquake. Starbird [25] studies how Twitter was used by hospitals to broadcast availability to care for victims. Sarcevic et al. [23] shows how it was used to report on the relief activities of various medical groups. Goggins et al. [9] analyzes an online discussion forum used by the US Navy to coordinate with NGOs during the same event, and finds that forum discussions correspond with “on the ground” activities during the earthquake.

Beyond the 2010 Haiti earthquake, Deneff et al. [5] describe how two different police departments used Twitter in response to the 2011 riots that took place in England, and find that disparate adoption styles led to different images and relationships with the public. Cobb et al. [4] characterize the use of social media in disasters by “digital volunteers,” and provide insight into how to best support the needs of geographically-dispersed volunteers when they respond to mass emergencies. Hughes [10] studies the usage of social media by Public Information Officers, who handle the public relations aspects of emergency response in the United States.

We contribute to this literature by showing how a large humanitarian organization can adapt its internal procedures—particularly those related to rapid assessment of an emergency situation—to incorporate social media communications. In particular, we focus on the information needs of OCHA, and give insight into the process of providing organizations with particular data that fits within their defined procedures and established information needs. By providing OCHA with Twitter data that contain information they specifically require to better oversee the management of dozens of organizations who perform myriad tasks related to disaster relief, the hope is that we ease the burden on those responsible for collecting and analyzing in the immediate post-impact period of a disaster. In addition, we show how existing methods can lead to more productive uses of social media data by stakeholders, and affected populations.

1.2 Background

We focus on Typhoon Yolanda (internationally known as Typhoon Haiyan), one of the strongest tropical cyclones ever recorded. The typhoon made landfall in the central Philippines on November 8, 2013, affecting over 14 million people, killing over 6,000, leaving over 4 million displaced, and causing billions of US dollars in property damage.

Though it was predicted and residents in some areas were able to take precautions, the typhoon brought about devastating effects.¹ Many national and international aid organizations were deployed to respond to Typhoon Yolanda, including OCHA.

According to the International Monetary Fund, the Philippines is an emerging market.² However, despite 24% of the population being classified as “poor” [2], telecommunications infrastructure is widespread. Nearly every Filipino adult has access to a mobile phone, and the cellular network covers almost the entire country.³ Regarding Twitter use, the Philippines is ranked 10th in the world for number of Twitter accounts,⁴ and English is widely spoken due to its former occupation by the United States. When Typhoon Yolanda struck the Philippines, the combination of widespread network access, high Twitter use, and English proficiency led to many located in the Philippines to tweet about the typhoon in English. In addition, outsiders located elsewhere tweeted about the situation, leading to millions of English-language tweets that were broadcast about the typhoon and its aftermath.

2 OCHA Procedures in Disaster

When a sudden-onset disaster happens and government capabilities are exceeded, OCHA is mobilized.⁵ OCHA is tasked with quickly assessing the situation “on the ground” and coordinating response efforts. The UN uses a framework for division and organization of needs during humanitarian crises based on eleven different *clusters*, which are “groups of humanitarian organizations (UN and non-UN) working in the main sectors of humanitarian action ... [c]lusters provide a clear point of contact and are accountable for adequate and appropriate humanitarian assistance.”⁶ The clusters are: Logistics, Nutrition, Emergency Shelter, Camp Management and Coordination, Health, Protection, Food Security, Emergency Telecommunication, Early Recovery, Education, and Water, Sanitation and Hygiene (WASH).

We can think of the role of OCHA as overseeing the organization of humanitarian response in disasters. Their responsibilities include: assessing the situation, understanding the needs of the affected population and of responding organizations, deciding on priorities, obtaining access to affected areas (which may have political as well as logistical implications), ensuring sufficient funding and resources, consistently and clearly communicating with the public, and monitoring progress.⁷ Our goal in this research is to use OCHA’s needs as the lens through which we examine Twitter communications

¹ <http://reliefweb.int/disaster/tc-2013-000139-phl>

² <http://www.imf.org/external/country/phil/index.htm?type=9988>

³ <http://www.infoasaid.org/guide/philippines/telecommunications-overview>

⁴ <http://business.inquirer.net/111607/telcos-report-record-number-of-customers> http://www.mediabistro.com/alltwitter/twitter-top-countries_b26726

⁵ The UN differentiates between “slow-onset” and “sudden-onset” disasters. Sudden-onset disasters occur quickly and with little to no warning. Slow-onset disasters have a longer period of buildup, and extend for greater periods of time. [3].

⁶ <http://www.unocha.org/what-we-do/coordination-tools/cluster-coordination>

⁷ <http://www.unocha.org/what-we-do/coordination/overview>

broadcast during Typhoon Yolanda. In addition, we examine data compiled from interviews with OCHA employees and consider how information communicated via Twitter may (or may not) be used by aid organizations, and how advocates within the UN who strive to incorporate social media data into their assessments and decision-making process can further their cause.

In response to disasters, the international humanitarian community undertakes a series of actions called the “Humanitarian Program Cycle.” OCHA manages the cycle while working with additional agencies, NGOs, and other technical bodies. One of the goals of the program cycle is to issue a MIRA (Multi Cluster/Sector Initial Rapid Assessment) report two weeks after disaster onset.

2.1 The MIRA Framework

The MIRA framework specifies how to quickly assess the needs of affected populations, and assign responsibilities to various response agencies soon after disaster impact.⁸ The goal of MIRA is threefold: 1) to systematically collate and analyze secondary data to provide an accurate-as-possible understanding of the situation; 2) to perform a “community level assessment,” where aid workers and volunteers interview and talk with affected populations to gather primary data, and; 3) to bring data together into a coherent picture that provides decision makers with a current report that incorporates information from the various clusters and allows them to have a common understanding of the situation. The MIRA process starts during the *immediate post-impact period* ([7, page 8]) of sudden-onset disasters.

The first step in the MIRA process is to produce a “Situation Analysis” report, written and released within the first 48 hours after impact. It starts with an overview that includes crisis severity, priority needs, and government capacity to respond. This is followed by the “humanitarian profile,” which describes the population, estimated number of affected people, and casualties, among other details.

The Situation Analysis is comprised of “secondary data,” and “primary data.” Secondary data describe the population of concern under typical circumstances (i.e. poverty level), as well as data that have been collected from sources such as mainstream media and satellite imagery. The secondary data form an up-to-date picture of the situation that provides a common understanding among stakeholders. Primary data are collected on the ground in the immediate aftermath of the disaster. Teams of aid workers and volunteers conduct interviews with “key informants,” i.e. those who are most likely to know the current state of the population, what are the most pressing needs, and how vulnerable populations are affected.⁹ Due to the quick turnaround necessary to produce the Situation Analysis, and the potential for Twitter communications to provide data that is potentially useful as well as broadly representative of population needs in those first 48 hours of response, we focus on this aspect of the MIRA process.

⁸ https://docs.unocha.org/sites/dms/Documents/mira_final_version2012.pdf

⁹ OCHA recruits key informants knowing that they will obtain a purposive sample of data sources that are not necessarily representative of the affected population.

2.2 Situation Analysis of Typhoon Yolanda

The Situational Analysis for Typhoon Yolanda [2] was released on November 10, 2013. It starts with a list of six priority needs, and goes on to provide a high-level overview of the current situation in hard-hit cities, explaining where people are without food and water, where electricity is not available, and what areas are inaccessible. Sources include the Philippines' government, Red Cross, Atmospheric, Geophysical and Astronomical Service Administration, and Department of Social Welfare and Development, in addition to existing demographic and census data.

This Situation Analysis also includes points such as “typical assistance needs,” how impact “may” affect the local population, and how the typhoon “can” cause additional complications and problems. i.e. much of this report is based on previous knowledge of similar events, and describes what is *likely* to happen, and where supplies and services will *typically* be needed, but it does not include first-hand accounts of the situation. This is *not* a criticism of OCHA—it is not possible for any individual nor organization to grasp the situation on the ground so soon after a large-scale disaster that affects millions of people, across a large geographical area. OCHA's reliance on past experience and expected needs is necessary in the first days after disaster impact, when getting substantial data about the situation on the ground is so difficult.

Given the difficulty of assessing a disaster situation in such a short time period, OCHA is open to using new sources—including social media communications—to augment the information that they and partner organizations so desperately need in the first days of the immediate post-impact period. As these organizations work to assess needs and distribute aid, social media data can potentially provide evidence in greater numbers than what individuals and small teams are able to collect on their own.

2.3 Social Media Experiment in Typhoon Yolanda

OCHA attempted to gain information from Twitter communications during the immediate post-impact period of Typhoon Yolanda, to triangulate and/or augment information they had from other sources. The social media data OCHA employees had access to during the crisis were gathered and analyzed by MicroMappers.¹⁰

MicroMappers is a digital volunteer organization devoted to annotating and mapping tweets (and other data) produced during disasters. In the case of Typhoon Yolanda, OCHA contacted MicroMappers to see what information they could gain from Twitter communications. Starting on November 7, 2013 at 19:28 (GMT) engineers began to collect tweets that contained specific keywords that described the typhoon and/or relief efforts. Data collection continued until November 13, 2013 at approximately 12:00 (GMT). Details are provided in Section 4.1.

A set of tweets sent between the start of the collection and November 10, 2013 (at 07:00 GMT) were first sampled for quick response efforts on behalf of MicroMappers. These tweets were uploaded to the MicroMappers platform for volunteers to read and label based on the information they contained. A total of 3,678 tweets were labeled by volunteers. The categories volunteers used to label tweets at this early stage

¹⁰ <http://www.micromappers.org/>

in the response efforts are detailed below. These categories were quickly identified during the typhoon response by MicroMappers volunteers working around the clock to label and organize information communicated via Twitter; they do not claim these categories were all-encompassing, nor that they were representative of all tweets about the typhoon. The three categories that appeared to be most salient in the immediate post-impact period, and which MicroMappers used to start their tweet-labeling procedure were:

- *Infrastructure Damage*: Information about destruction and/or damage of roads, bridges, buildings; disruptions to basic services, e.g. hospitals.
- *Community Needs*: Information about shelters, food, location of missing persons, water, and hygiene.
- *Humanitarian Support*: Information about deployment of aid, recovery services, and in-kind donations and contributions of goods and services.

In addition, volunteers produced a map including 600 of these tweets that were associated to a location based on geographical references contained in them. All the labeled tweets were then sent to OCHA for further analysis.

The Situation Analysis for Typhoon Yolanda mentions the Twitter data provided by MicroMappers, and includes approximately one page about social media data. This section on social media shows a general map containing data produced by MicroMappers volunteers. However, no specific information about tweet content is provided.

As expressed by the OCHA staff interviews we describe in the following section, OCHA's *hope* was that social media data could contribute to a better understanding of the situation on the ground in these three specific areas. However, at the time the Situation Analysis was published, they did not have sufficient data. Later analysis of the Typhoon Yolanda MIRA report and additional documents do not indicate that OCHA was able to garner additional information about these topics from social media communications. However, OCHA is hopeful that future disaster response efforts will successfully involve social media data.

3 Information Management in Practice

To better understand disaster response from OCHA's perspective, and how they perceive social media data in these situations, we interviewed OCHA employees about their experiences during the response to Typhoon Yolanda, and about the potential for social media data *vis-à-vis* response efforts going forward. OCHA staff described the procedures they follow, spoken candidly about the challenges they face when responding to disasters, and worked with us to formulate ideas on how to best use social media data when responding to disasters.

3.1 Information Management Roles

The OCHA employees tasked with "information management" during disaster situations are "information management officers" or "IMOs," as well as "humanitarian affairs officers," or "HAOs."

The duties of IMOs and HAOs in disaster response are to drive coordination. Different IMOs and HAOs have different skill sets, but as a group, they are tasked with gathering data, liaising with various cluster leaders, communicating with volunteers, updating databases and common data repositories, and producing a variety of documents. In the immediate aftermath of a disaster, they often experience “ad-hoc craziness” brought on by a need to complete myriad tasks in a short period of time [27]. Additionally, they answer requests from all manner of stakeholders, and are responsible for writing reports that provide up-to-the-minute information.

When IMOs and HAOs collect data in the field, they focus on eight themes that guide the MIRA framework: 1) drivers of the crisis and underlying factors; 2) scope of the crisis and humanitarian profile; 3) status of populations living in affected areas; 4) national capacities and response; 5) international capacities and response; 6) humanitarian access; 7) coverage and gaps; 8) strategic human priorities. The MIRA framework includes specific questions that coincide with each of these themes to guide OCHA employees, as well as others who may be working with them, to collect data.

3.2 Interviews with OCHA Staff

As yet, social media data are a somewhat amorphous source of information for OCHA. The population of OCHA staff who can speak to the use of social media data in disaster response is relatively small; thus, we were able to secure interviews in person, and via phone, with four OCHA staffers. While we recognize the small sample size, we nevertheless stress that—together with the documentation we analyzed—the insight and firsthand knowledge we gained by speaking to these interviewees provides a sufficient backdrop regarding the potential benefits and hindrances to using social media data, particularly in the large, multi-organization coordination efforts that OCHA undertakes.

Our first interview was with an OCHA staffer in New York, NY, United States, who we refer to as O1. O1 provided us with a big picture view of what OCHA staff are responsible for when they deploy in disaster situations, and also gave insight into the Typhoon Yolanda response effort. S/he laid out the initial background information regarding the role of OCHA, the types of information they need to collect and organize when assessing a situation, and how they usually perform the myriad tasks for which they are responsible. Our discussion with O1 provided us with the foundational information we needed to understand what OCHA does, and helped us frame questions and points for discussion for our subsequent interviews.

Our next interview was with an HAO based in Geneva, Switzerland (who we refer to as O2). His/her job in the first 48 hours after impact is to quickly compile as much information as possible about the area of impact and the current situation, organize it into a coherent narrative, and present it in a Situation Analysis.

For Typhoon Yolanda, in addition to the traditional sources that OCHA turns to, O2 and his/her colleagues were also open to seeing what information they could gain from Twitter data. They received the dataset of 3,678 labeled tweets from MicroMappers, in which each tweet was associated with one of three categories of information described in the previous section. O2 and colleagues looked at the content of these tweets; their impression was that around 200-300 of them provided what they considered relevant

information. In addition, they found that reading and analyzing tweets was an interesting exercise, but it was very time consuming. During those initial hours of disaster response, so much work needs to happen so quickly that the OCHA employees who responded to Typhoon Yolanda are not sure the social media data augmented what they already knew. Overall, in this case, they felt “the time investment was too high.” O2’s experience speaks to the need for a way to process social media data that addresses and centers on their specific information requirements. In other words, O2 is implying that s/he needs to get Twitter data that are processed from a “top-down” perspective.

O2 was working from UN headquarters in Geneva during the days after Typhoon Yolanda made impact, and produced the Situation Analysis from there. Subsequently, O2 traveled to the Philippines after the Situation Analysis was released, and continued to work on the MIRA report, which is published two weeks after impact. During the time in the Philippines, s/he had access to the primary data that were collected from in-person interviews with key informants, and additional sources of local knowledge.

O2 observed that in comparing social media data to primary data, there seemed to be a considerable bias in the social media data toward those located in urban areas, with access to telecommunications networks.¹¹ However, despite the (well-understood and often inherent) bias of social media data, eyewitness accounts, first-hand knowledge and additional useful information captured via social media can still augment situational awareness, and OCHA is aware of this.

Subsequent discussions with HAOs and IMOs have provided further insight into the difficulty OCHA employees face when they perform a rapid needs assessment and write the Situation Analysis; getting a handle on the needs of a large population (i.e. millions of people) that has been severely affected by a disaster is a monumental task. IMOs, HAOs and other UN employees are ready and willing to use any viable source of information available to help them better understand and assess these situations. They *want* to use social media data; the question is how to provide them with these data in a timely, easily understandable format that they can use to triangulate and/or augment other sources within the immediate post-impact period—and which correlate to the specific information they are seeking.

Another interview with a OCHA employee (an IMO based in Geneva who we refer to as O3) revealed further difficulties with including social media data into disaster response procedures. Though O3 sees great potential in incorporating social media data in disaster response, s/he points out that using social media as a bona fide source of information in crises is a tough sell to UN management. O3 points to the notion among many at the UN is that social media data are more likely to contain “bad,” “false,” or “unverified” information persists. S/he also pointed out the problem of information expiration—information that is posted on social media sites often has a short period of time during which it is “true,” or “actionable.”

In further discussion, O3 also stressed that the role of UN agencies in disaster is not to respond to individual requests. If OCHA staff see a tweet about a trapped family, or where someone needs medical attention—regardless of whether the information is verified—they are not in a position to act upon such information. Rather, OCHA seeks

¹¹ This bias has been observed in other domains, particularly politics, see e.g. [8, 20].

a collective view of the situation, particularly with respect to the eleven clusters, and with respect to the location of various needs.

In OCHA's view, social media data could contribute to this type of assessment by e.g. counting how many tweets are being sent (or not) from particular areas, how many tweets mention the need for food, water, or other supplies, and to locate tweets containing specific information about macro-level population needs, e.g. "*2,000 people in <village> are affected by the typhoon—all need shelter.*" This assessment is provided by another interviewee—O4, an IMO also based in Geneva—who points to the potential for Twitter data to "complete the picture" when OCHA is trying to gain an overview of the situation and ascertain how to coordinate and activate the various cluster agencies.

Equipped with this understanding of how social media data can be of most use to OCHA, we show how current technologies can be used to develop reports of social media data that could be readily incorporated into OCHA processes, with a focus on humanitarian clusters and regional location of needs.

4 Data Analysis

Having spent time with OCHA staff who are open to using social media data, our next step was to perform an analysis of a separate dataset of tweets collected during Typhoon Yolanda (i.e. different from the dataset that was labeled by MicroMappers.) Our express goal was to identify information that coincides with the UN humanitarian clusters. We then determined how Twitter data compares to and/or augments the information the IMOs and HAOs are typically able to collect within the first 48 hours of a disaster.

4.1 Data Collection and Pre-Processing

To obtain a set of tweets sent during Typhoon Yolanda that were likely to include information about the event, we performed a keyword search using Twitter's Streaming API; keywords included: "YolandaPH," "Yolanda," "RescuePH," "TyphoonHaiyan," and many more that were identified during the typhoon by colleagues who were closely monitoring the Twitter stream as the event unfolded.¹² The keyword search resulted in a dataset of 2,302,569 tweets from November 7, 2013 19:28 (GMT) to November 13, 2013 12:00 (GMT), as shown in Table 1. Though many tweets about the typhoon were posted weeks after this time period, we stopped data collection on November 13 at 12:00 (GMT) because our OCHA interviewees stated that this six-day period would be of most interest to them. Further, we divided the dataset into two periods.

The first period represents the time frame OCHA considered for the Typhoon Yolanda Situation Analysis report. This period within our dataset consists of 1,173,850 tweets. In addition to the first set of tweets, we also consider tweets posted during the next 48 hours after the first period. The second period is from November 10, 2013 20:31 (GMT) to November 13, 2013 12:00 (GMT), and contains 1,128,719 tweets. We include tweets from this second time period in our analysis to determine: (i) to what extent the information posted on Twitter changes after the initial period, and (ii) to what extent those

¹² More details on the data collection, including sampled keywords, are included in the data release, see URL at the end of this paper.

Period	Start (GMT)	End (GMT)	# of Tweets
First	Nov. 7, 2013 19:28	Nov. 10, 2013 20:30	1,173,850
Second	Nov. 10, 2013 20:31	Nov. 13, 2013 12:00	1,128,719
Total			2,302,569

Table 1. Breakdown of tweets into two time periods.

changes may affect OCHA’s ability to gain situational awareness information that may be included in later reports (i.e. the MIRA report, and other reports that are generated after the Situation Analysis is released.) It was only after all tweets were collected and we had done some preliminary analysis that we spoke with OCHA employees.

4.2 Automatic Classification by Region

The Philippines are divided into seventeen different regions, or administrative divisions. The UN breaks down its analysis per region, as shown in Table A1 (in the Appendix). The information in this table is based on the most complete data provided by the UN, which they released on November 23, 2013. Regarding the UN-provided data, some regions have no data available; these regions were either unaffected, or no data were provided about them. Previous work points to social media activity increases in regions affected by a disaster [6, 24]. However, this is not always the case, as frequency of social media postings can increase in areas that are not strongly affected by a disaster.

We measure to what extent the number of tweets sent from particular regions correlate with amounts of damage or number of affected people. We classify tweets by region using two strategies. First, by *geolocation*, for tweets that include GPS coordinates, which are added by mobile clients when the user enables this functionality. Second, by *keyword*, i.e. we considered all tweets that mentioned the name of a region or the name of any municipality in that region. Table A1 shows the results. We note that while the activity on Twitter was in general more significant in regions heavily affected by the typhoon, the correlation is not perfect. For instance, there are more tweets from the National Capital Region and from CALABARZON, which were not among the most affected, than from the Bicol Region, where more than half a million people were affected. Though our results show that classifying tweets by region was not a reliable undertaking in this particular case, we maintain that it is a worthwhile exercise, as it can prove useful in some circumstances [21, 22, 24].

We also attempted to measure these correlations in relative terms, e.g. by expressing the affected people as a percentage of the population, and/or by expressing the number of tweets in proportion to the tweets “normally” present in each region (using a data sample from one month prior to the crisis). Results were similar, in terms of showing some correlation but not a perfect one. We did not expect that tweets would predict the number of affected people per region, for the same reason that they do not predict winners in political elections [8]. Again, we see an example that likely points to bias in the Twitter data; urban, affluent, tech-savvy people are more capable of posting to microblogging services than rural, poor populations. Knowing that these biases exist and are likely to continue is critical for OCHA to take into account as they work to incorporate social media data into future response efforts.

Category	Human Labeled	Automatically Labeled
Informative	845	42% 1,109,480 48%
Not informative	613	31% 661,228 29%
Not related	542	27% 531,861 23%

Table 2. Initial classification task on a sample of tweets posted during Typhoon Yolanda.

4.3 Automatic Classification

We implement supervised machine learning to perform automatic classification of tweets. In this approach, a relatively small number of human-labeled items are used to train an automatic classification system (this is the “training set”). Then, this automatic classification system that has been trained on human-labeled data is used to classify the remaining tweets.

Automatic Classification of Informative Tweets. We filtered the datasets to identify messages that might contain useful information using the supervised classification approach of Random Forests. Tweets were first converted to binary feature vectors in which each word (unigram), or a sequence of two consecutive words (bigram), is a coordinate in the vector ([15, 28]). A random sample of 2,000 tweets was used as a training set. The choice of the learning approach, features types, and training set size was based the empirical evidence presented in [16]. We used crowdsourcing services from CrowdFlower, which provided us with workers who labeled tweets with the appropriate category. Workers were given the following instructions:¹³

Indicate if the item contains information that is useful for capturing and understanding the situation on the ground:

- A. Informative: contains useful information that helps you understand the situation.*
- B. Not informative: refers to the crisis, but does not contain useful information that helps you understand the situation.*
- C. Not related to this crisis.*

Two out of three workers’ agreement was required to finalize a label.

Results are in Table 2. The resulting classifier has an AUC of 0.89, measured using 10-fold cross validation, which indicates fairly high classification accuracy.¹⁴

Automatic Classification into the Cluster Framework. Next, we again turned to crowdsource workers to perform manual labeling, and used the output to train an automatic classifier.¹⁵

To reduce the amount of false positives—i.e. messages automatically classified as informative, but not containing useful information—we imposed the constraint that classification confidence must be higher than 0.8. The first data period yielded 270,781

¹³ CrowdFlower is an online crowdsourcing service that allows clients to upload tasks with instructions, which Crowdfower workers are then paid to complete: <http://crowdfower.com/>

¹⁴ AUC is Area Under Receiver Operating Characteristic curve, 50% means a random classifier and 100% means a perfect classifier. We do not use accuracy, as it is misleading in this context.

¹⁵ For this Crowdfower labeling task, we grouped “camp management” and “shelter” clusters together, and “food security” and “nutrition” clusters together for clarity, which gave us a total of 9 cluster categories from which workers could choose.

Cluster	Human Labeled		Automatically Labeled	
	(period 1)	(period 1)	(period 1)	(period 2)
Food and nutrition	54	4,712	39,448	↑
Camp and shelter	41	1,870	8,470	↑
Education and child welfare	50	18,076	22,198	↓
Telecommunication	90	8,002	5,899	↓
Health	57	1,008	2,487	
Logistics and transportation	51	2,290	3,259	
Water, sanitation, and hygiene	31	1,210	82,568	↑
Safety and security	87	7,884	4,970	↓
Early recovery	216	14,602	46,388	
None of the above	1,323	382,906	451,122	↓
Total	2,000	442,560	666,809	

Table 3. Classification of informative tweets posted during Typhoon Yolanda, according to the Humanitarian Clusters Framework. Up/down arrows indicate relative increase/decrease of 50% or more in period 2, proportional to the total number of tweets in each period.

tweets (23% of tweets during that period), from which 2,000 tweets were sampled uniformly at random and labeled according to the humanitarian clusters (the same constraint in the second data period yields 351,070 tweets: 53% of tweets during that period, which suggest an increase in informative content, consistent with Table 2).

Results of both the manual and automatic classification are shown in Table 3, where we also indicate whether there is an increase or decrease of 50% or more in the proportion of messages in each cluster. In the first time period (roughly the first 48 hours), we observe concerns focused on early recovery and education and child welfare. In the second time period, these concerns extend to topics related to shelter, food, nutrition, and water, sanitation and hygiene (WASH). At the same time, there are proportionally fewer tweets regarding telecommunications, and safety and security issues.

In general, Table 3 shows a significant increase of useful messages for many clusters between period 1 and period 2. It is also clear that the number of potentially useful tweets in each cluster is likely on the order of a few thousand, which are swimming in the midst of millions of tweets. This point is illustrated by the majority of tweets falling into the “None of the above” category, which is expected and has been shown in previous research [29].

4.4 Drilling Down into Clusters: Topic Models

OCHA staff indicated their preference for being presented with aggregate information, as opposed to a list of individual tweets. In this section, we examine how information relevant to each cluster can be further categorized into useful themes. We employ topic modeling using Latent Dirichlet Allocation (LDA) [18]; a common method used to analyze datasets of thousands or millions of documents, and whose application to disaster-related tweets is described in [17].

Results of the topic modeling, including example tweets and representative words according to the LDA algorithm, are in Table A2. Due to a small number of items in the clusters, two themes were generated for most of them. However, some clusters

e.g. “telecommunications, safety and security” resulted in only one theme because the majority of tweets in that cluster mention the same words/information.

Topic models allow us to quickly group thousands of tweets, and to understand the information they contain. In the future, this method can help OCHA staff gain a high-level picture of what type of information to expect from Twitter, and to decide which clusters or topics merit further examination and/or inclusion in the Situation Analysis.

Feedback from the UN. To find if we were on a helpful path regarding our post-hoc analysis of Typhoon Yolanda Twitter data, we asked OCHA staff to look at the information we present in Table A2. We provided a description of the data, and explained that though the data are from a past event, we were concerned with whether they could use this type of information in future events.

Feedback was positive and favorable. Regarding the information in Table A2, O4 said: “it could potentially give us an indicator as to what people are talking most about—and, by proxy, apply that to the most urgent needs.” O4 goes on to say “There are two places in the early hours that I would want this: 1) To add to our internal “one-pager” that will be released in 24-36 hours of an emergency, and 2) the Situation Analysis: [it] would be used as a proxy for need.” Another UN staffer, who works for a non-OCHA sector in disaster response, stated: “Generally yes this [information] is very useful, particularly for building situational awareness in the first 48 hours.” One staffer (O1) did express concern that this level of analysis may be too general for some applications, saying that “the [topic] words seem to general.” However, s/he went on to say that the table gives a general picture of severity, which is an advantage during those first hours of response. This validation from UN staff supports our continued work on collecting, labeling, organizing, and presenting Twitter data to aid humanitarian agencies with a focus on their specific needs as they perform quick response procedures.¹⁶

5 Discussion

Twitter is established as a place to communicate, gather and disperse information, and gain situational awareness during disasters. Furthermore, research suggests that there is abundant useful information broadcast on Twitter during mass emergencies [14, 21, 26, 29]. This has led many within OCHA to view Twitter communications as a way to triangulate what they know from other, more conventional, sources.

5.1 Obstacles

Emergency responders face technological and organizational barriers to the adoption of social media in their processes, including a growing need for institutional change [10]. OCHA has an overwhelming amount of work to do when tasked with assessing a crisis, identifying needs, and distributing reports that provide an overview of the situation. Social media communications are yet another item on the lengthy list of sources for them to consider when attempting to gain an accurate understanding of a crisis situation.

¹⁶ We were unable to get feedback from all staff we interviewed earlier due to field deployments.

This is an obstacle noted by others: “Even when good data is available, it is not always used to inform decisions. There are a number of reasons for this, including data not being available in the right format, not widely dispersed, not easily accessible by users, not being transmitted through training and poor information management. Also, data may arrive too late to be able to influence decision-making in real time operations, or may not be valued by actors who are more focused on immediate action.” [1]

Concerns about veracity of social media information were also voiced. These issues are not unlike those faced by Public Information Officers (PIOs) in the United States who also wrestle with knowing if they can trust information that is found on social media. [10]. However, regardless of the questions around “truth,” and “trust,” it is clear that social media data can be used to augment situational awareness. [26, 29].

5.2 Recommendations

Providing social media data to humanitarian organizations requires, first and foremost, an understanding of how those humanitarian organizations work. Organizations that have existed for decades will rarely re-invent themselves around a new technology. However, they can be guided toward making new tools and data an established feature of their processes. In this sense, OCHA staff cited the need to know what they are likely to find—and not find—on social media when they are in the midst of a response.

The next consideration is to present the information in a format that answers target users’ questions. OCHA staff are supportive of incorporating social media in their processes, but they need data to be presented in a format that is easily consumable. This echoes concerns expressed by Public Information Officers interviewed by Hughes [10], who also note the complexity of social media as an information space. OCHA does not want to read thousands of tweets; they require a high-level snapshot that explains the Twittersphere, and which they can use to augment their assessment of the situation.

This research has shed light on the fact that providing the “big picture” of a crisis situation via an analytic view of tweets is helpful to OCHA, and potentially other aid agencies. While we do not deny the value of information found in individual tweets, organizations such as OCHA require a higher-level overview of the activities and behaviors that play out on Twitter in the immediate post-impact period of a disaster. Therefore, we suggest presenting results in *multiple levels*. For example, a higher level shows the number of tweets per geographical region, followed by the number of tweets per cluster, and the topics inside each cluster (the scheme we have followed in this paper).

Finally, it is important to have the right systems in place. Given the consensus among OCHA staff that social media data are particularly valuable during the early hours of a disaster, real-time acquisition and analysis of data is critical. This involves large amounts of time and effort on behalf of many people, so in addition to digital volunteer platforms such as MicroMappers—which employs humans as a sole source of information processing—we have pointed to other systems that perform real-time social media analysis using supervised machine learning, and which incorporate humans in the process when required [13, 19].

Data release. The data we collected is available for research purposes at <http://crisislex.org/>.

Acknowledgments. The authors wish to thank Patrick Meier and the UN participants.

Bibliography

- [1] Promoting innovation and evidence-based approaches to building resilience and responding to humanitarian crises. Tech. rep., Department for International Development, UK AID (2012)
- [2] Situation Analysis Philippines, Typhoon Haiyan (Yolanda). Tech. rep., United Nations Office for the Coordination of Humanitarian Affairs (Nov 2013)
- [3] Capelo, L., Chang, N., Verity, A.: Guidance for collaborating with volunteer and technical communities. Tech. rep., Digital Humanitarian Network (Aug 2012), <http://digitalhumanitarians.com/content/guidance-collaborating-volunteer-technical-communities>
- [4] Cobb, C., McCarthy, T., Perkins, A., Bharadwaj, A., Comis, J., Do, B., Starbird, K.: Designing for the deluge: Understanding & supporting the distributed, collaborative work of crisis volunteers. In: Proc. of CSCW. ACM (2014)
- [5] Deneff, S., Bayerl, P.S., Kaptein, N.A.: Social media and the police: tweeting practices of British police forces during the August 2011 riots. In: Proc. of CHI. pp. 3471–3480. ACM (2013)
- [6] Evnine, A., Gros, A., Hoffleitner, A.: On Facebook when the earth shakes... Tech. rep., Facebook Data Science team (Aug 2014)
- [7] Fischer, H.W.: Response to disaster: fact versus fiction & its perpetuation—the sociology of disaster. University Press of America (1998), <http://www.worldcat.org/isbn/0761811826>
- [8] Gayo-Avello, D.: Don't turn social media into another "Literary Digest" poll. Commun. ACM 54(10), 121–128 (Oct 2011), <http://doi.acm.org/10.1145/2001269.2001297>
- [9] Goggins, S., Mascaro, C., Mascaro, S.: Relief work after the 2010 Haiti earthquake: leadership in an online resource coordination network. In: Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work. pp. 57–66. ACM (2012)
- [10] Hughes, A.L.: Participatory design for the social media needs of emergency public information officers. In: Proc. of ISCRAM (2014)
- [11] Hughes, A.L., Peterson, S., Palen, L.: Social Media in Emergency Management. FEMA in Higher Education Program (2014)
- [12] Imran, M., Castillo, C., Diaz, F., Vieweg, S.: Processing social media messages in mass emergency: A survey (Aug 2014), <http://arxiv.org/abs/1407.7071>
- [13] Imran, M., Castillo, C., Lucas, J., Meier, P., Vieweg, S.: AIDR: Artificial intelligence for disaster response. In: Proc. of WWW (Companion). pp. 159–162. International World Wide Web Conferences Steering Committee (2014)
- [14] Imran, M., Castillo, C., Lucas, J., Patrick, M., Rogstadius, J.: Coordinating human and machine intelligence to classify microblog communications in crises. In: Proc. of ISCRAM (2014)
- [15] Imran, M., Elbassuoni, S., Castillo, C., Diaz, F., Meier, P.: Practical extraction of disaster-relevant information from social media. In: Proc. of Workshop on Social Media Data for Disaster Management. pp. 1021–1024 (2013)

- [16] Imran, M., Elbassuoni, S.M., Castillo, C., Diaz, F., Meier, P.: Extracting information nuggets from disaster-related messages in social media. Proc. of ISCRAM, Baden-Baden, Germany (2013)
- [17] Kireyev, K., Palen, L., Anderson, K.: Applications of topics models to analysis of disaster-related Twitter data. In: NIPS Workshop on Applications for Topic Models: Text and Beyond. vol. 1 (2009)
- [18] Mark, G., Bagdouri, M., Palen, L., Martin, J., Al-Ani, B., Anderson, K.: Blogs as a collective war diary. In: Proc. of CSCW. pp. 37–46. ACM (2012)
- [19] Melville, P., Chenthamarakshan, V., Lawrence, R.D., Powell, J., Mugisha, M., Sapra, S., Anandan, R., Assefa, S.: Amplifying the voice of youth in africa via text analytics. In: Proc. of KDD. pp. 1204–1212. ACM (2013)
- [20] Mitchell, A., Hitlin, P.: Twitter reaction to events often at odds with overall public opinion. Tech. rep., PEW Research Center (Mar 2013), <http://www.pewresearch.org/2013/03/04/twitter-reaction-to-events-often-at-odds-with-overall-public-opinion/>
- [21] Olteanu, A., Castillo, C., Diaz, F., Vieweg, S.: CrisisLex: A lexicon for collecting and filtering microblogged communications in crises. In: Proc. of ICWSM. pp. 376–385 (Jun 2014), <http://chato.cl/papers/olteanu.castillo.diaz.vieweg.2014.crisis.lexicon.disasters.twitter.pdf>
- [22] Sakaki, T., Okazaki, M., Matsuo, Y.: Earthquake shakes Twitter users: real-time event detection by social sensors. In: Proc. of WWW. pp. 851–860. ACM, New York, NY, USA (Apr 2010), <http://dx.doi.org/10.1145/1772690.1772777>
- [23] Sarcevic, A., Palen, L., White, J., Starbird, K., Bagdouri, M., Anderson, K.: Beacons of hope in decentralized coordination: learning from on-the-ground medical twitterers during the 2010 Haiti earthquake. In: Proc. of CSCW. pp. 47–56. ACM (2012)
- [24] Shelton, T.: Visualizing the relational spaces of hurricane Sandy. <http://www.floatingsheep.org/2013/08/visualizing-relational-spaces-of.html> (Aug 2013)
- [25] Starbird, K.: Delivering patients to Sacré Coeur: collective intelligence in digital volunteer communities. In: Proc. of CHI. pp. 801–810. ACM (2013)
- [26] Starbird, K., Palen, L., Hughes, A.L., Vieweg, S.: Chatter on the red: what hazards threat reveals about the social life of microblogged information. In: Proc. of CSCW. pp. 241–250. ACM, New York, NY, USA (2010), <http://dx.doi.org/10.1145/1718918.1718965>
- [27] Verity, A., Mackinnon, K., Link, Y.: OCHA information management guidance for sudden onset emergencies. Tech. rep., UN OCHA (Feb 2014)
- [28] Verma, S., Vieweg, S., Corvey, W.J., Palen, L., Martin, J.H., Palmer, M., Schram, A., Anderson, K.M.: K.m.: Natural language processing to the rescue?: Extractingsituational awareness tweets during mass emergency. In: Proc. of ICWSM (2011)
- [29] Vieweg, S.: Situational Awareness in Mass Emergency: A Behavioral and Linguistic Analysis of Microblogged Communications. Ph.D. thesis, University of Colorado at Boulder (2012)

Appendix

In this appendix, we include details of people and houses affected by a disaster, compared with geo-located tweets by region (Table A1). We also include details of the results of topic modeling for each UN Cluster (Table A2).

Region designation and name	OCHA Information		Number of Tweets	
	Affected people	Affected houses	By geolocation	By keywords
I Ilocos Region	–	–	228	189
II Cagayan Valley	–	–	344	1,905
III Central Luzon	–	–	705	575
IV-A CALABARZON	27,076	840	2,034	2,524
IV-B MIMAROPA	425,903	33,499	150	1,339
V Bicol Region	695,526	12,129	1,372	1,214
VI Western Visayas	2,694,031	476,844	14,110	6,329
VII Central Visayas	5,180,982	101,789	19,075	7,938
VIII Eastern Visayas	4,156,612	504,526	1,110	19,224
IX Zamboanga Peninsula	–	–	25	165
X Northern Mindanao	19,592	20	381	2,174
XI Davao Region	5,175	40	847	1,217
XII SOCCSKSARGEN	–	–	74	39
XIII Caraga	45,063	549	198	660
ARMM Autonomous Region in Muslim Mindanao	–	–	175	442
CAR Cordillera Administra- tive Region	–	–	353	1,428
NCR National Capital Region	–	–	2,211	15,909

Table A1. A chart showing the number of affected people and houses per region, compared with both the number of tweets geo-located in that region (“by geolocation”), as well as the number of tweets that contain a region or municipality name (“by keywords”). Source: United Nations Typhoon Yolanda data reports, November 2013.

Cluster	Number of Tweets		Topic Words	Example Tweet
	(period 1)	(period 2)		
Food and nutrition	2,340	17,559	food, need, please, goods, relief, help, volunteer	<i>Multi-climate ration packs or healthy army food. Folks need practical food specially the kids @sarah-meier @USArmy #YolandaPh #urgentneed</i>
	2,372	21,889	donate, food, wfp, families, water	<i>RT @radikalchick: Red Cross asks for help from police / military. their trucks w/ food and water for 25000 families are stopped in Tanauan</i>
Camp and shelter	846	3,447	homes, destroyed, areas, relief, moving, many	<i>Roxas says many homes in Leyte's coastal areas destroyed: They're like matchsticks that were flung inland & talagang sira</i>
	1,024	5,023	shelter, seek, millions, apart, super, rise, super	<i>Super typhoon Haiyan slams central Philippines millions seek shelter Read more: http://.../</i>
Education and child welfare	14,153	12,275	suspended, today, classes, work	<i>RT @AdamsonUni: Classes and work at all levels are suspended today Nov 8 in anticipation of Typhoon Yolanda. Stay safe Adamsonians. #wala</i>
	3,923	9,923	relief, kids, help, support, emergency	<i>Support UNICEF. emergency relief efforts for kids in the #Philippines. How to help:http://.../ #Haiyan http://.../</i>
Telecommunications	8,002	5,899	satellite, call, image, mtsat, officials, countries	<i>MTSAT enhanced-IR satellite image of #YolandaPH as of 2:30 am 09 November 2013: http://.../ via @dost_pagasa RT @govph</i>
Health	542	1,030	medical, doctors, help, volunteer, charities, team	<i>MSF emergency & medical teams continue to closely monitor the #Typhoon #Haiyan situation and are ready to respond to needs</i>
	466	1,457	supplies, red cross, hospital, medical, send	<i>@KarloPuerto: Davao City 911 sends rescue and medical equipment and personnel to Tacloban City #YolandaPH</i>
Logistics and transport	1,138	1,649	goods, help, repack	<i>RT @DepEd.PH: DSWD needs volunteers to help repack relief goods. Call DSWD-NROC at 851-2681 to schedule your shift. #YolandaPH http://.../</i>
	1,153	1,609	roads, river, affected, debris	<i>Debris on roads in Tacloban is blocking delivery of aid from airport to victims of Typhoon #Haiyan in #Philippines http://.../</i>
Water, sanitation, and hygiene	613	34,825	water, clean, need, food, supply	<i>Heard from @ExtremeStorms who is still in Tacloban. Desperate need for drinking water. Need for military ship & supplies #haiyan #yolanda</i>
	596	47,743	donate, clean, water, millions, appeal	<i>No potable water supply power outage & impassable roads in Leyte. Immediate needs r clean water food & shelter-staff in OrmocMai #haiyan</i>
Safety and security	7,884	4,970	safe, dead, killed, ridiculous	<i>7000 kid's parents have been killed by the storm in the Philippines and #StayStrongJustin is trending... Ridiculous http://.../</i>
Early recovery	14,602	46,338	donate, relief, efforts, support, donations, goods	<i>Doing relief efforts now for #YolandaPH. Need free shipping line info. @indayearona @juanxi @kwittiegirl</i>

Table A2. Results of topic models with two topics per cluster. We include representative topic words generated by the topic model algorithm, and one example tweet per topic.