

Social Media in Disaster Relief

Usage Patterns, Data Mining Tools, and Current Research Directions

Peter M. Landwehr and Kathleen M. Carley

Abstract. As social media has become more integrated into peoples' daily lives, its users have begun turning to it in times of distress. People use Twitter, Facebook, YouTube, and other social media platforms to broadcast their needs, propagate rumors and news, and stay abreast of evolving crisis situations. Disaster relief organizations have begun to craft their efforts around pulling data about where aid is needed from social media and broadcasting their own needs and perceptions of the situation. They have begun deploying new software platforms to better analyze incoming data from social media, as well as to deploy new technologies to specifically harvest messages from disaster situations.

1 Introduction

In this chapter, we review the ways in which individuals and organizations have used social media in past disaster events and discuss ways in which the field will progress. In the first section, we cover the how both individuals and organizations have used social media in disaster situations. Our discussion emphasizes how both types of groups focus on searching for new information and disseminating information that they find to be useful. In general, facts about disasters collected from the small number of individuals located near the scene of a disaster are the most useful when dealing with specific disaster situations. Unfortunately, this data is rare and difficult to locate within the greater sea of social media postings related to the disaster.

We follow this by discussing a framework for considering how to analyze and use social media. This framework consists of several different use cases and

Peter M. Landwehr · Kathleen M. Carley
Carnegie Mellon University,
Pittsburgh, Pennsylvania
e-mail: {plandweh, kathleen.carley}@cs.cmu.edu

analytic steps: collecting social media data; managing a workflow for analyzing social media data; constructing a narrative from social media data; processing social media data to find relevant information; working with geolocation data; analyzing the text of social media postings; and broadcasting information using social media. Along with each step we provide reference examples of tools and libraries that can be used by analysts and first responders.

The chapter concludes with a section looking at current research into how we can better analyze and understand social media. Our discussion centers on methods for automatically classifying text and using visual analytics to gain new insights.

2 Usage Patterns in Disaster

The questions confronting people in a disaster are almost always the same: What happened? Are my friends and possessions safe? How can we remain safe? Social media is a new resource for addressing these old needs.

Locals at the site of the disaster who are posting information about what they are witnessing are in many ways the gold of the social media world, providing new, actionable information to their followers. They are few in number, and while their messages are sometimes reposted they often don't circulate broadly. Locating their content is an ongoing challenge akin to finding a needle in a haystack. Such local information can serve as an early alert system, leading traditional news sources [1].

While non-local users cannot provide reportage on the disaster, they can propagate local stories across the network and help them gain traction. By simply discussing a disaster or using hashtags associated with it they can contribute to other users' perception that the disaster is relevant. They can also collate data from other media sources, ferret out local users, and debunk false rumors as they begin to propagate. They can also serve as a workforce to sort through postings for the few that are cries for help, identify locations based on photos, find missing people in scanned video, and create maps where there are none.

Organizations fill a different role in the social media ecosystem. While individuals seek out information to preserve their own well-being, news media and aid groups use social media to help carry out their missions. Reporters look to social media to find stories and get feedback on their coverage. They and their parent organizations also often post links to breaking stories hosted on their own websites or being broadcast in the traditional media. Relief groups post requests for resources, announcements about their activities, and monitor social media for information they can use in their relief work.

Individuals within organizations are often charged with monitoring social media for any and all content that might be relevant to understanding the disaster as it relates to the organization's mission. This is a free-form search for information, conditioned only on the organization's role. It's similarly difficult to

constrain what an organization might post to their account beyond that it be relevant to this mission.

Along with press-releases and information about services, organizations may engage in “beaconing” behaviors by trying to solicit particular information or resources from individuals in their communities. Organizations will also carry out immediate dialogue with users, responding to their comments directly and via the same medium. Such efforts can help with the success of things like beaconing by making it clear that the organization takes the medium and its users seriously.

It is important to note that not all usage during a disaster is benign. Some social media users will start spreading false information to create additional panic. The occurrence of a cascade of damaging rumors (a “virtual firestorm”) during a crisis can serve to undermine a first responder, hamper the relief effort, and lead to innocent victims being harmed [2]. Organizations have set up “fake” meeting places to identify those whom they wished to contain.

As we explore these different ways in which social media have been used in disaster, bear in mind that this framing is not necessarily crucial for any of these activities to occur. Individuals and organizations use social media to find and share information in standard contexts as well. In a disaster and its aftermath these activities are heightened in particular ways but should not be construed as necessarily restricted to it.

2.1 Individuals

Disasters rarely end instantaneously. Aftermaths can drag on for days, weeks, months, or years. (As we write this, three years after Haiti was struck by a devastating earthquake, thousands of individuals remain in tent cities [3].) Disaster researchers often divide disasters and disaster response into four phases: preparedness; response to the event; recovery, including rebuilding after the response; and mitigation, including enacting changes to minimize the impact of future events [4].

When people are confronted with a disaster they don’t just seek to preserve their lives at a single critical moment. Users actively seek out information that can help them understand what’s happening for a prolonged period of time. They try and connect with other members of local communities for support, aid, and understanding. Often, they will use technology to do so. ([5] as cited by [6].)

Shklovski et al. documented this process for individuals in California who were afflicted by wildfires in 2007 [6]. These fires dragged on for weeks, covering large swathes of rural countryside. Californians in at-risk areas found the news media unhelpful, citing a focus on stories about damage to celebrity homes. What locals wanted was general information about where fires were occurring and who was in danger. To combat this lack of knowledge, the Californians being studied had set up two different online forums for posting news and warnings. At the end of wildfire season, one of the subject forums was closed because it was no longer

useful. The other remained open as a community hub and remained part of its users' lives.

These researchers later witnessed a similar phenomenon among a community of musicians in New Orleans in the aftermath of Hurricane Katrina [7]. The musicians adopted SMS messaging, more regular cell phone use, and posting to online forums in order to stay in touch during the disaster. Like the Californians living in range of the wildfires, these New Orleans natives felt that the television media focused on the most dramatic aspects of the disaster while ignoring the majority of the afflicted. The victims used satellite images and message boards set up by the local newspaper to find information that was relevant to them. They turned to previously unused technologies to socialize in disaster, and in many cases adopted these new practices into their regular lives.

In both California and New Orleans, individuals turned to technological resources carry out established information seeking patterns via new media. Since these studies were carried out, we have seen the advent of Web 2.0 and the plethora of social media platforms that exist today. It is easier than ever to search the web for information about disaster, but filtering out rumor, falsehood, and off-topic discussion from the ocean of online content remains difficult. The outstanding research challenge remains helping people to find information they need and to post information so that it can be found.

While people don't intentionally confine themselves to a particular medium, they naturally favor those with which they are comfortable and those from which they believe they can gain more information. Since its introduction in 2007, Twitter has benefitted from generally positive media coverage [8]. Thanks to both this positive portrayal and its widespread adoption, the microblogging platform has become seen as an important source for disaster information. Leading up to Superstorm Sandy in 2012, blogs published guides for how to best search Twitter for data [9]. In the storm's wake, blogs and news media published stories about how much Twitter had been used [10, 11].

Despite the press coverage, Twitter isn't the dominant means of electronic communication. Its usage is the barest fraction of SMS and email [12]. While a personal email is often rich in meaningful content, Twitter's broadcast nature has meant that the relevant tweets sent during any event are buried in a sea of off-topic noise. Nonetheless, the ready availability of data, as well as the perception that the service is the "new thing" has made it a popular choice for academic research. Twitter is by no means insignificant – its millions of users are real- but it is perhaps overvalued. Even as we focus heavily on it in this chapter, we advise that you consider the platforms relative position and situate your findings correspondingly.

The ready availability of data from the platform has also made it a popular choice for academic research. This doesn't mean that Twitter is a particularly dominant communications platform: its usage is the barest fraction of SMS and email and it suffers as a data source from a great deal of noise generated by third party users. This also doesn't mean that Twitter should be dismissed as

insignificant. Rather, it highlights the fact that other platforms, especially SMS, should be given additional research more in keeping with their usage patterns.

The vast pool of research on how Twitter has been used outside of disaster is generally beyond this chapter's scope. However, it is useful for understanding the how the service has been generally used, so we provide a brief overview here.

Kwak et al. collected a very large corpus of Twitter users, tweets, trending topics, and social relations between users, and provide a large collection of summary statistics for each. The researchers make a variety of observations, not least of which is that there is little overlap in their data between the most followed users on Twitter and the users who are most retweeted. They also find that following has a low degree of reciprocity, and that users who follow each other tend to be in the same time zone [13]. Java et al. have used network methods to analyze Twitter to try and identify meaningful user communities. In the process, they categorized the bulk of twitter interactions as consisting of "Daily Chatter" (descriptions of routine life), conversations, information sharing, and reporting news. They also characterized users as primarily being defined as information sources, information seekers, and friends [14]. Naaman et al. collected tweets from approximately 125000 users over a prolonged period, developed nine overlapping categories for the tweets, and then identified two clusters of users: reformers, who often broadcast personal information, and informers, who generally shared different types of information [15]. Bakshy et al. tried to identify how one could successfully inject a particular idea into Twitter by influencing a particular user. The researchers consider a user to have "influence" based on when users retweet a URL that they have posted; the researchers caution that this requires a relatively strong signal to detect influence, but is also precisely measurable. While they identified certain users as possessing influence and causing cascades of information, they found it difficult to predict when a cascade would occur or which of these potentials might cause a cascade. The researchers concluded that the most cost-effective for propagating a particular URL or idea on twitter would be to seed many non-influential users. These users would have the potential to create many small information cascades which might then add up to one of relatively rare large cascades [16].

Research on how Twitter is used in disaster often takes the form of looking at data collected from a particular subset of users commenting on a disaster and looks at the particular features of their discussion. For example, Starbird et al. attempted to understand usage patterns during the 2009 Red River flood by qualitatively analyzing tweets collected during the flood period that used the terms "red river" and "redriver". The researchers identified two overlapping types of useful tweets by users: generative and synthetic. Generative tweets introduce new information via description of lived experience or factual commentary on an extant tweet [17]. Synthetic tweets pull in a variety of outside information and repackage it specifically for Twitter: a 140-character summary of a news story, for example, as might be produced by a news organization. While the authors noted other types of tweets, the generative and synthetic made up the kernel of the useful data that arrived during the disaster. Original tweets are also hard to find. They

made up less than 10% of the sample used by the researchers, and more than 80% of that small number were produced by individuals located within 6 hours driving time of the afflicted area.

Similarly, Sinappan et al. attempted to categorize tweets broadcast by Australians during the 2009 Black Saturday brush fires. Using another search-based approach, the authors coded the tweets using a modified version of Naaman et al.'s general tweet categorization scheme specifically for disasters. When looking at 1684 tweets captured, the researchers found that only 5% contained directly actionable information [18]. Similarly, only 4% of messages posted to the Chinese microblog service Sina Weibo after the Yushu Earthquake in 2010 related to actions that individuals could or needed to take [19]. Roughly 25% of the messages were tied to situation updates about Yushu, but a large number of them were from secondary sources, something also true for the data analyzed by Sinappan et al.

In her thesis research, Sarah Vieweg developed a new categorization system for the subset of tweets that contain useful information. Synthesizing tweets from four disasters and referencing the disaster research literature, she created three overarching categories (social, built, and physical environment) for useful tweets. These categories are themselves split into 35 subcategories that capture the message's content [20]. Sample categories include "Status – Hazard", "Advice – Information Space", and "Evacuation".

These phenomena (a small number of actionable tweets, a small number of tweets from locals providing primary source data) play out repeatedly in analyses of different disasters. The non-local tweets often play secondary roles that are important in the broader context of the disaster. Sutton witnessed this when researching Twitter discussions of the 2008 spill of 5.4 million cubic yards of coal ash into the Tennessee River [21]. While many of the Twitterers were local, Sutton describes them as using the medium as a "grassroots mechanism" for getting national media attention aimed at the disaster. They are the non-influential users trying to start local cascades.

While demanding that a retweet must include a particular URL is stringent, the basic idea of using retweets as a measure of endorsement is natural and useful. Starbird & Palin found this to be true in the tweets broadcast during the 2011 Egyptian uprising [22]. (Bear in mind that an uprising differs from conventional disasters as it features two opposing forces, not simply people in distress.) The researchers draw the same lines that they have before between locals and non-locals and the relative importance of these tweets for knowing the condition on the ground. However, they also note that retweets make up 58% of the corpus they collected, and that the most circulated tweets were all variants of a particular "progress bar" meme about uninstalling a dictator or installing democracy. The meme originated with Twitterers outside of Cairo but eventually made its way into the city proper, getting picked up by other Twitterers nearer the heart of the protest. The researchers characterize the meme as the "complex contagion" described by Centola & Macy, arguing that the remixing of the different meme elements "show some degree of shared understanding of its purpose". ([23] as

cited by [22].) Meme retweeting and remixing kept the protesters involved, and can be seen as a way that even those outside a developing crisis situation can try to connect themselves to it, possibly as a precursor to additional action.

In addition to trying to raise awareness of the disaster, the Twitterers responding to the coal ash spill in Tennessee also tried to debunk false rumors about the disaster's scope. Indeed, the segment of the twitter community affiliated with any particular disaster has taken on the job of suppressing rumors relating to it. NPR Reporter Andy Carvin, who gained acclaim covering global news events solely on Twitter, has likened his many followers to the staff of a news room: "rather than having news staff fulfilling the roles of producers, editors, researchers, etc., I have my Twitter followers playing all of those roles [24]." Carvin relies on the platform to eventually provide him with access to domain experts who can verify content or help him debunk it. For example, Carvin was able to work with his followers to determine that a prominent blog ostensibly written by a Syrian lesbian documenting the local unrest was actually a hoax [25]. Similarly, during Superstorm Sandy reporter Jack Stuef exposed user @comfortablysmug as spreading false information about what was happening in New York City. Many of @comfortablysmug's tweets were identified as false by other Twitter users, while Stuef found images from @comfortablysmug's Twitter profile in his YouTube and was able to determine the user's true identity [26, 27].

Mendoza et al. attempted to systematically analyze the practice of individual Twitterers debunking and supporting the various rumors that can arise as a disaster progresses [28]. The researchers identified tweets sent in the wake of the 2010 Chilean earthquake that been retweeted at least one thousand times and that were promulgating ideas externally verified as either true or false. They then looked at the responses that these tweets had elicited. None of the verified truths were substantially contested by Twitterers, while all of the falsehoods saw a number of tweets denying their accuracy. Additionally, the falsehoods were generally affirmed as true in other tweets more rarely than were the genuine truths. The exception to this was the widespread reporting of looting in certain areas of Santiago; tweets about this topic performed similarly to the other true tweets. The study suggests that while generally rumors can be expected to be called out on Twitter, particular types of rumor will still fly under the radar and be hard to detect. The study suggests that true reports of disasters will not be regarded as controversial, which may be useful in automatically confirming their accuracy from social media data.

Contra the Mendoza et al. study, however, we emphasize that even if eventually corrected, falsehoods have been propagated on social networks for long enough to enter the mass media. @comfortablysmug's stories of flooding at the NYSE were rebroadcast by several major news outlets before Stuef outed him. In the aftermath of the 2013 Boston Marathon Bombings, Twitter users and Redditors incorrectly identified a missing Brown University student and individual mentioned on a police scanner as the bombing suspects [29, 30]. This caused a brief but potent online witch-hunt for which Reddit administration apologized [31]. The Boston

Police, which has made extensive use of Twitter before and after the bombing, published the facts of the case to the platform to counter the rumors [32].

We have mentioned that Twitter has gotten a great deal of research attention relative to other media used in disaster. While it remains the focus of this chapter, it is important that we acknowledge the ways that individuals are leveraging other social media in these circumstances. To simply focus on Twitter, when the reality is that an individual equipped with a smartphone can already function on any number of social media platforms at once. Technologies for analyzing social media will not remain confined to a single platform but will exploit as many as possible for data. They will leverage not just Twitter, but also RSS feeds, Facebook, SMS, Sina Weibo, Four Square, and others from a variety of different nations.

As mentioned at the start of this section, musicians in New Orleans adopted SMS messaging in the aftermath of Katrina to stay in touch. SMS's ability to directly connect individuals and the widespread availability of the technology on low-tech cellphones has made it critical in emergency situations. In the immediate aftermath of the 2010 Haiti Earthquake, a small group of actors from relief organizations and the US Government got DigiCel, Haiti's main cellular service provider, to reserve the SMS short code 4636 as a dedicated number for processing distress messages. These messages were archived and translated into English by Haitian expatriates mobilized by the organizers of "Mission 4636". Both expatriates and Haitians still on the island worked to promote the short code as a useful resource. By Week 3, Mission 4636 was dealing with such a volume of messages that it began working with the CrowdFlower and Samasource crowdsourcing platforms to better coordinate message translation [33].

From our perspective on how individuals use social media, SMS was key in this disaster because significant numbers of Haitians used low-tech cellphones that could access SMS services in the wake of the earthquake, and because the SMS infrastructure itself was still working. In that sense, it was the right medium for the time. AS has occurred with Twitter data in other disasters, the SMS data was rife with falsehoods despite being sent to a dedicated help line. According to the Harvard Humanitarian Initiative's (HHI's) study of relief organization responses to the Haiti earthquake, perhaps as many as 70% of the 4636 messages contained errors, such as requests to locate victims by people who knew the victims to be dead [34].

Photo sharing during disaster has also seen some degree of academic study, though more work is needed. In 2008, Liu et al. looked at how Flickr had been used in response to seven different disasters [35]. They observed that individuals were posting photos of damaged areas for a variety of different reasons, united by an over-arching theme of documenting the crisis. The different photographs can generally be categorized as depicting a particular event, capturing on-line social convergence (e.g. screen shots of Facebook posts), listing the missing, and showing personal belongings (taken for inventory purposes).

Flickr can be understood as fulfilling some of the same needs as text-based services: individuals post representative images of disaster sharing information

about what they understand the situation to be. It's used to help communities organize and share information. It's also used to serve practical, individual needs, such as inventorying possessions. The authors tie this back to the medium itself: a photograph is a richer data source than a 140 character message. Twitter isn't an efficient tool for cataloging possessions.

Regardless of precise intent, when placed in disaster situations individuals broadcast and examine data using social media. Academic research has tried to categorize these usages, often noting that actionable information is hard to find, that information can be false, and that primary source information from local users can be rare. Additionally, platform-specific practices can potentially be subverted for additional information: we can characterize true and false statements seen during a disaster based on the number of debunking statements noticed in response. We can infer that photographs taken in disasters of peoples' possessions are being used to inventory property.

Just as social media are being leveraged by individuals during disasters, so too are they being used by relief organizations, both government affiliated and independent. While the specific purposes behind the uses may be different and complementary, the uses of the platforms are similar. We discuss these in the next section.

2.2 Organizations

First responder organizations, which include government agencies, police, firemen, medical and public health organizations, military responders and not-for-profits play critical roles in disaster response. These groups generally have access to data and analytical tools that are not available to the public, as well as the resources needed to rescue and aid individuals who are in distress. While people may distrust the accounts of unknown strangers reporting rumors, these organizations have established brands that often temper or heighten critical attitudes towards their own postings. First responder organizations are increasingly turning to social media to identify actionable needs and orient the response, gauge the scope of impact of the event, provide information to the public, track and mitigate firestorms and counter false information, and to try to identify potential secondary disasters before they occur [36].

Where social media has provided individuals with new spaces in which to mingle and interact, it has provided organizations with new spaces in which to research ongoing disasters and communicate with both victims and the general public. St. Denis et al. explored this phenomenon by looking at how a Virtual Operations Support Team (VOST) dealt with the 2011 Shadow Lake Fire [37]. The concept of the VOST was developed by emergency manager Jeff Phillips as a way for an organization to coordinate its responses on and to social media coverage of a disaster, and has been propagated by other emergency managers [38–40]. According to Phillips, the VOST should “integrate[e] ‘trusted agents’ into [emergency management] operations by creating a virtual team whose focus

is to establish and monitor social media communication, manage communication channels with the public, and handle matters that can be executed remotely through digital means such as the management of donations or volunteers.” This list of the roles taken on by the VOST encapsulates the ways that organizations active in the relief sphere must address social media when disasters strike.

After the start of the 2011 Shadow Lake Fire, the Portland branch of the National Incident Management Organization (NIMO) recruited an all-volunteer VOST in order to help coordinate their online responses. In their postmortem interviews, St. Denis and her fellow researchers found that the groups settled into a routine: NIMO would draft a press release each evening, consult with the VOST about the need for updates or amendments in the morning, and then release the statement with according changes. Meanwhile, the VOST took charge of updating a blog, Facebook page, and Twitter account set up by NIMO to provide updates on the fire. The VOST provided feedback to Facebook users who posted to group’s wall, relaying information back to NIMO and did its best to maintain some sense of community among those coming to the page. Still, according to Eriksen, the VOST’s most important accomplishment was locating a blogger who was concerned about fire trucks getting routed over unsuitable back roads. This blogger possessed niche, critical knowledge that couldn’t have been found without the VOST, and NIMO was able to contact the blogger directly to get more information from him.

During the 2011 London riots, local police authorities used Twitter as a way to communicate with citizens. The bulk of the posted tweets, as analyzed by Panagiotopolous et al., encouraged people to participate in “cleanup” activities after the riots had ended, commented on how well communities were doing in coming together after the riots had ended, and described the situation on the ground; they also posted requests for information, albeit much more rarely. The researchers suggest that the authorities’ posting about clean up actions may have played a role in getting the public out to help clean, though definitively proving this is beyond the paper’s scope [41].

When Sarcevic et al. examined the practices of individuals affiliated with medical organizations using Twitter in Haiti, they observed a widespread phenomenon they termed “beaconing”: the broadcasting of requests for information or material to Twitter because of uncertainty about how to obtain them [42]. While this practice is also used by individuals as a general facet of information-seeking, its application by individuals affiliated with organizations tied to the crisis is important because it suggests that the organization itself has a need that it can’t address internally or through established contacts.

The Haiti earthquake itself was a critical proving ground for social media’s use in disaster. In its immediate wake Haiti saw an influx of aid workers and organizations, many of which planned to use sophisticated technological solutions in order to help provide disaster relief. Meanwhile, other volunteer organizations operating remotely helped to collate social media data arriving from victims, analyze it, and get the results to other organizations on Haiti.

In the previous section, we briefly mentioned the Harvard Humanitarian Initiative's (HHI's) study of responses to the Haiti earthquake. This study focused on the responses of and relationship between official relief organizations (such as those operated by the UN) and these "volunteer and technical communities" (V&TCs), an umbrella term covering both volunteer, not-for-profit, and for-profit groups active in the disaster [34]. V&TCs outside of Haiti played a key role parsing and analyzing local distress. That said, the critical problem on the ground in Haiti was lack of access to the Internet and the aid organizations' lack of a single, unifying platform. Difficulty integrating the output of different programs made it hard to combine results and merge workflows. Unlike with the VOST, the interface between the V&TCs and humanitarian organizations was less well defined.

We mentioned Mission 4636 in the previous section as an example of how individuals used social media during disaster. It bears revisiting here from an organizational standpoint, as the project fits the HHI's description of a V&TC, and played a key role in addressing the disaster. It was also essentially a one-time effort; while it can be replicated, these particular volunteers have separated.

Ralph Munro, one of the lead organizers of Mission 4636, has emphasized the front-facing aspect of 4636. He stresses that the project's success was due to its being a largely Haitian initiative. Without a robust group of Haitian expatriates, neither the back-end translation nor the SMS shortcode would have been useful. Indeed, he suggests that the primary role of social media other than SMS during the crisis was as a recruiting and advertising platform. Volunteers working for 4636 claimed to have been posting alerts about the project to Facebook so often that they were being threatened with bans for acting like spammers [43]. If Mission 4636 wished to reorganize, the leaders would need to turn to social media to again advertise the service and recruit volunteers for support.

The critical accomplishments of Mission 4636 were getting promoted to the Haitian community for local use, organizing and motivating volunteers, and providing a consistent pipeline of data to other V&TCs and relief organizations. It was specifically intended to connect victims and relief organizations; they provided little in the way of feedback to those outside of the relief loop. While victims were aware of Mission 4636 through the existence of the short code, other relief groups operating off the ground were effectively individual to Haitians, interacting only with responders and the public. As such, we want to briefly highlight Haiti Ushahidi, one of the V&TCs that received data from Mission 4636 and had no specific on-the-ground presence.

Ushahidi is an online platform to which individuals can post reports about distress in disaster situations. These reports can then be coded to fit particular categories and get pinned to locations on a map. Developed by Ory Okolloh for use during the Kenyan election crisis of 2009, the platform has since been deployed in other crisis situations [44]. Haiti Ushahidi was a particular instance of the Ushahidi platform set up by students at Tufts University. In marked contrast with Mission 4636, the Haiti Ushahidi project was less well known to average

Haitians. The maps produced by the project, however, were used by groups on the ground, and the project is often mentioned in close connection with Mission 4636 despite functioning independently [34]. The Haiti Ushahidi project presented a better public face than did Mission 4636 despite processing significantly less information. Further, by choosing to release a large subset of the disaster messages they were working with to the public, they helped put a face on the disaster in a way that the directed channel of SMS generally does not.

During the 2011 East Japan earthquake, a group of computer scientists and engineers formed a new, one-off aid group called ANPI_NLP to help get relevant information from tweets [45]. The researchers sought to parse tweets to find references to individuals who had gone missing or been found and then updating records in Google Person Finder, a missing persons database.

While a one-time effort like Mission 4636, in general ANPI_NLP is a V&TC effort in the Ushahidi Haiti mold. The researchers didn't present themselves in a way that would be perceived by the Japanese populace, and the results that they produced were stored in a database maintained by another V&TC (Google) which then dealt with relief organizations. Where ANPI_NLP differed from Ushahidi Haiti was in using up-to-date natural language processing to speed up the task of extracting information from tweets. The researchers rapidly created a pipeline for morphologically analyzing tweets and that both extracted named entities and locations, and classified the nature of the information expressed. The researchers had to perform some manual coding to create gold standard data and to vet results, but in general this was an automated process. They also point out the existence of problems similar to those described by Munro: translation is difficult, and human resources are critical. To the members of ANPI_NLP, the solution lies in better automated systems, and in tools that can more rapidly adapt to training data.

Our discussion of how organizations used social media is framed by the understanding that at some level they use social media the same way individuals do: they search for information, and contribute in order to participate in the conversation as fits their mission. Relief organizations must deal with the larger challenge of managing their presence in particular social media spaces and must understand the information that is coming to them via the different interfaces. One way to deal with this is through a dedicated group such as a VOST. Further, a host of small organizations are appearing to help work with social media data in particular crises, leveraging local knowledge and deploying new technologies. In its report, the HHI both noted the importance of this small organization in the Haitian Earthquake's aftermath while also airing the concerns of relief workers that in Haiti it that few of these tools had an established, dependable reputation.

It is impossible to review all of the tools that exist to help relief organizations and analysts mine meaning from social media data. In this next section, we approach this challenge by provide a useful framework for considering them, as well as descriptions of different tools that fit into the framework.

3 A Data Analytics Framework and Associated Tools

A variety of different tools have been and are being developed and deployed to help people and institutions work with social media. In this section, our primary focus is on those that are useful to analysts trying to mine social media data, particularly from Twitter, in disaster response situations. They range from libraries for programming languages to sophisticated GUI-based tools for responders who need quick assessments of information to platforms for recruiting other workers to help with tasks. Technically savvy responders and analysts chain the output produced by multiple tools together in order to create meaningful results.

In this section we describe a mix of these tools, broken down into a rough framework corresponding to different data mining tasks. Some of our distinctions would not exist in a conventional data mining text, but speak to our particular focus on social media in disaster. More specifically, in this section we discuss tools that support data collection; that support workflow management by way of third-party tool interoperability and enabling data retrieval; that support narrative construction from fragments of social media data; that support data processing for quantitative analysis and disaster response; that support pinning social media data to maps based on geolocation data; and that support quantitative text analysis for use with machine learning. We also cover an additional, slightly different category: those used to broadcast on social media and to manage collection and publication of data. These won't be relevant to analysts or investigators looking at specific data published on social media services, but can be important for developing a holistic understanding of how different platforms are being used. Such tools are of particular interest to organizations doing their best to manage all aspects of their social media presence.

3.1 Data Collection

The central problem for researchers wanting to take a quantitative, data mining approach to analyzing social media data is that it can be hard to obtain, store, or trade. In Twitter's case, few canonical data sets are available for study due to the company's restrictions on data storage. The corpus of tweets made available for the 2011 TREC Twitter competition¹ is a useful exception, but is limited in scope.

While archives of data are useful, analysts and relief workers also need methods for gleaning facts from Twitter in real time, that limit the amount of effort that they have to put into monitoring social media.

If an analyst is skilled at programming, the basic way of approaching social media data is to obtain it using a website's API. Twitter², Flickr³, and many other social media platforms invite developers to access some portion of the website's

¹ <http://trec.nist.gov/data/tweets/>

² <https://dev.twitter.com>

³ <https://secure.flickr.com/services/api/>

data programmatically. In the case of Twitter, roughly 1% of tweets are made available via the API. These limits depend on the site and API. In the case of Twitter, if a researcher wants to access a larger percentage of the Twitter stream than is available from the API they must deal with a data warehouse such as Spinn3r⁴ or GNIP⁵, which provide access blog data, the full Twitter stream, and a variety of other social media data. Limitations on data consumption via API are dependent on each site's Terms of Use.

If an analyst doesn't wish to work directly with the API they can turn to third party tools that will obtain the data for them and possibly provide some analysis. For example, TweetTracker, developed at Arizona State University, allows users to filter the stream of tweets in real-time based on keyword and location [46]. These tweets are then archived and stored for future analysis. The ORA network analysis tool⁶ supports importing ego network data from both individual Facebook accounts and email boxes [47–49]. Ushahidi (the company behind the platform of the same name) has worked on its own tool, SwiftRiver⁷, which uses crowd-based validation of data. As different RSS and Twitter streams are passed into the platform, users can remotely coordinate to annotate particular items regarding their accuracy or inaccuracy. Tools such as Social Radar, CRAFT, and SORASCS (discussed in more detail in the next section) provide platforms in which multiple tools can interoperate to create flexible disaster response systems and scalable data storage systems that support social media collection and analysis. Within these platforms, third-party tools can be used as components of larger workflows; a data collection tool such as TweetTracker can be paired with different analytical tools such as ORA to provide richer insights into data.

Yahoo!'s Pipes platform⁸ is another option in this area, albeit a middle ground between pure coding and pure GUI solutions. It allows users to tie together a mixture of data from different RSS feeds, conditioned on different events occurring. Different pipes can be configured via an API or a graphical user interface. In a similar vein, CMU's Rapid Ethnographic Assessment (REA) system allows users to pull in data from Facebook, RSS feeds and Lexis-Nexis.

3.2 *Workflow Management*

Researchers wanting to take a "big data" approach to dealing with social media are faced with a plethora of challenges. As described above, social media data can be difficult to store, obtain, or trade. Additionally, the quantity of data makes it difficult to intuit critical patterns and characteristics when exploring. There is also no inherent guarantee of accuracy regarding the data's provenance. A fourth problem is that the data providers often do not maintain archives of the messages,

⁴ <http://spinn3r.com>

⁵ <http://gnip.com>

⁶ <http://www.casos.cs.cmu.edu/projects/ora/>

⁷ <http://www.ushahidi.com/products/swiftriver-platform>

⁸ <http://pipes.yahoo.com>

so if all messages back to a particular date are needed, a database needs to be built and maintained with the relevant data and all associated meta-data. No one tool exists to address all these challenges. As we will see in subsequent parts of this section, many different tools are emerging to handle pieces of these tasks. Correspondingly, new tools are emerging to manage workflows between these more focused tools and the larger process of cleaning and analyzing social media data.

Social Radar, CRAFT, and SORASCS⁹ are three tools that address this problem. Each is a web-based system that supports disaster response by helping analysts and responders chain together third-party tools for sequential data analyses. All three tools work by collecting social media data from a data warehouse or via a particular third-party tool that access a social media platform's API. The collected data is then archived and can be sent to different integrated tools (or sequences of tools) for further processing. These tools often address text-mining, network analysis, sentiment analysis, geo-spatial analysis, and visualization. While some are used interactively, others process data in a silent and opaque manner, converting them from one form to another.

Many of the tools incorporated in Social Radar, developed by MITRE, are aimed at detecting sentiment in Twitter [50, 51]. It provides a web interface for looking at trends in Twitter over time such as total sentiment (derived from the presence of particular sentiment charged terms), heavily retweeted users, and the prevalence of particular keywords.

CRAFT, developed by General Dynamics, is similar to these other workflow management tools but also supports an associated environment for general mashups. Files can be linked to Google Drive, and the platform supports a "playback" mode that allows disaster response training exercises to be run with archived social media data collected during prior disasters.

SORASCS, developed at the CASOS Center at Carnegie Mellon University, supports workflow management and sharing [52, 53]. Unlike CRAFT and Social Radar, which require outside tools to be integrated before deployment, SORASCS is an open architecture to which analysts can independently attach their own tools. It allows analysts to preserve, share, and modify particular workflows by saving them to files. SORASCS's open design would make it eligible to serve as a coordinating under-structure behind CRAFT or Social Radar. While the latter tools have stronger user interfaces from a crisis responder's perspective, they provide no facilities to preserve particular workflows for future use. Unlike CRAFT and Social Radar, SORASCS does not necessarily convert all data into a common database; the user is responsible for supplying a database component themselves. In a sense, SORASCS is at a different level of application hierarchy than CRAFT and Social Radar. It could serve as middleware using either platform as a front end. This could provide some benefits to analysts because Social Radar

⁹ <http://www.casos.cs.cmu.edu/projects/project.php?ID=20&Name=SORASCS>

and CRAFT put the third-party tools in an open unstructured environment and don't support the development of automated and streamlined workflows as does SORASCS.

3.3 *Narrative Construction*

Social media data, composed of textual and other artifacts produced by millions of individuals, can be construed as a digital history of some aspects of the modern world. To parse the history of a particular disaster –or any other event- requires tools for composing narratives.

Appropriately aggregated data can naturally lend itself to this end. Indeed, data mining's focus on using big data demands that analysts use a combination of aggregation and culling for story-telling. Tools such as TweetTracker, ORA, and Social Radar can be used to plot the use of particular keywords and topics over time. As these terms fall into and out of use, they tell the story of what issues matter to particular users. ORA, as a network visualization tool, can be used to display the changing relationships between sets of entities graphically. In the case of Twitter data, this may refer to the relationships between individuals, individuals and the topics or keywords they have mentioned, and the topics and keywords themselves. These relationships can be rendered as a static snapshot or as a series of networks evolving over time. Newspapers have also turned to sophisticated visualization programming libraries in order to tell stories. The New York Times, for example, uses the D3.js JavaScript library¹⁰ to create graphics for data-driven news stories [54–57].

It can also important to understand the course of an event through a collection of specific tweets or other social media postings, each of which provides a fragment of the story. Andy Carvin of NPR has made heavy use of the Storify¹¹ platform to collate individual social media postings to document news events [58]. Blogging tools such as Blogger¹², WordPress¹³, and Tumblr¹⁴ can be used solely for reposting entries, thus providing some measure of the service provided by Storify. Timeline publishing services such as Dipity¹⁵ provide another alternative for describing the chronology of a particular event. By focusing on the specific rather than the aggregate these methods differ from conventional data mining approaches. However, given the emotional appeal of individual stories over general descriptions, analysts may want to direct some of their efforts towards finding those individual stories within the larger collection of data that can best serve as representatives of the whole.

¹⁰ <http://d3js.org>

¹¹ <http://storify.com>

¹² <http://www.blogger.com>

¹³ <http://www.wordpress.com>

¹⁴ <http://www.tumblr.com>

¹⁵ <http://www.dipity.com>

3.4 Data Processing for Relevance

Collected social media data must be processed to determine its meaning. There are a variety of ways in which data can be processed, and relief workers must focus on those that can best cater to a particular set of needs: predicting if a disaster is going to occur, assessing the scope of an ongoing disaster, identifying the key entities and actors involved in a disaster, and a variety of case-specific needs. Many of the tools we have already been discussed have been used by relief workers to address different parts of these challenges. Several critical methods for helping resolve this challenge are by leveraging keywords, annotating the data using crowdsourcing, using sentiment dictionaries to code text, and leveraging network analysis to identify key entities

3.4.1 Keyword-Based Labeling

Searching social media for particular disaster-related keywords is a simple but often effective technique for tracking disaster information. Because people often post news relating to disasters before it is reported in the mass media, a keyword search on a social network can provide early news about a disaster. Twitter determines its “trending topics” by processing large numbers of tweets to determine when keywords move into and out of currency [59]. When using a tool like TweetTracker to find disaster news, the underlying calls to the API are often simply looking for words mentioning certain keywords. Similarly, data warehouses like GNIP will often provide a separate listing of keywords that they have determined to be relevant in the requested tweets. While crude, individuals in distress who engaging in beaconing behaviors on Twitter to seek aid aren’t trying to be deceitful and so will likely use the obvious and expected keywords. That said, keyword based searches have limits: individuals can make typos and spelling mistakes, and the particular keywords relevant to a disaster can evolve and change. It is a static approach to a dynamic situation.

For social media that isn’t text based, an analyst can attempt to initially reduce the quantity of data by using any sort of qualitative textual label assigned to the particular object – the tags assigned to a Flickr image, for example. The assumption is that even if the choice of a particular tag or keyword will cause us to miss a few images, because the vast majority will be retained the amount of useful structure lost will be insignificant. This assumption needs additional empirical study. Martin et al found that tags are acceptable for the general flow but miss local information [60]. In crisis response, such local information may be critical.

3.4.2 Crowdsourcing-Based Labeling

While keyword-based coding can be useful for culling data down to general matches, the reduced data must often still be codified for relevance, actionability, and accuracy. This can be partially accomplished by automated processing of the data using trained machine learning algorithms, as in the ANPI_NLP project, but

is often handled manually. A human workforce with domain expertise can be used to provide sophisticated labeling to disaster data.

We've discussed the role played by Ushahidi in Haiti, but the platform bears revisiting here. Individual Ushahidi deployments can be used to categorize disaster reports and then post them to a map. This system provides a basic architecture for splitting the coding task across a group of individuals in order to streamline the completion of particular tasks. Analysts can also label messages post-facto, making Ushahidi a useful system for individuals seeking to place particular messages onto a map. The QuickNets platform¹⁶, built using Ushahidi's source code as a base, further subdivides the crowdsourcing process in order to make coding tasks easier for individuals to complete.

When a crowdsourcing workforce for coding data must be raised quickly, the fastest method is to use a dedicated crowdsourcing platform. Amazon Mechanical Turk¹⁷ is the archetypal example of an online labor market but there are many alternatives. As Mission 43636's popularity increased during the Haiti earthquake's aftermath, it switched from its informal organization system over to using CrowdFlower¹⁸ and Samasource¹⁹ to managing their many volunteer workers who spoke Kreyol and could translate the text.

Volunteers will often feel motivated to contribute time and energy to addressing disasters and working with disaster data, particularly for very large disasters. Dedicated communities of "Crisis Mappers" have formed around the idea of collecting geospatial data from afflicted regions and annotating it with relevant information²⁰. Similarly, sparked.com²¹ has focused on recruiting volunteers interested in contributing to meaningful causes. The best annotators for data may not be those obtained from a crowdsourcing marketplace but rather from within these and other communities of skilled volunteers with a specific investment in helping to resolve disasters.

3.4.3 Sentiment-Based Labeling

By measuring sentiment first responders can gauge the attitudes of populations to the ongoing disaster response and determine how they should adapt their activities. The field is very broad, and its state as of 2008 is described in detail by Pang and Lee [61]. The TweetTracker-ORA combination, Social Radar, Ushahidi, Google Crisis Maps, and ESRI ArcGIS are all being adapted to better incorporate methods for dealing with sentiment data. (We discuss the latter two programs further in the Geolocation section.)

¹⁶ <http://www.quick-nets.org>

¹⁷ <http://www.mturk.com>

¹⁸ <http://crowdfower.com>

¹⁹ <http://samasource.org>

²⁰ <http://crisismappers.net>

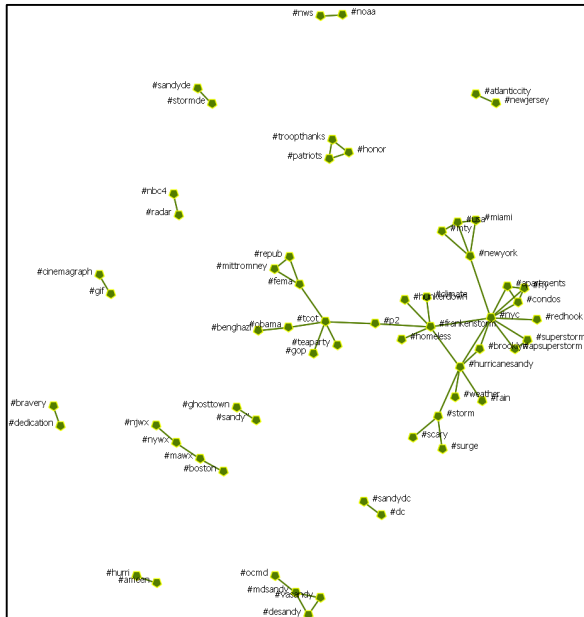
²¹ <http://sparked.com>

While training an algorithm to code social media data with sentiment information tool may be beyond the scope of most analysts working during a disaster, a dictionary labeling of terms with defined sentiment analysis is not. While simple and potentially prone to error, the method is rapid and lends itself to the brevity of social media. Several different dictionaries exist, but notably ones include: Linguistic Inquiry and Word Count (LIWC), developed by a team of judges evaluating large lists of words²²; the dictionaries of affective meanings collected by Heise when surveying different populations²³; and SentiWordNet²⁴, developed by Baccianella, Esuli, and Sebastiani by choosing a collecting of seed words on WordNet and carrying out random walks across the network.

3.4.4 Network Analysis for Relevance

As noted at the start of this section, responders must be able to identify the key entities and players in responding to the disaster. While the coding methods described above can help users filter social media data, network methods can help analysts and responders look at the structure of social media data in order to infer relevant structural information in the communications themselves.

Fig. 1 A network diagram of hashtags used in New York City related to Hurricane Sandy on October 29, 2012. Hashtags are linked if they co-occurred at least five times. (Reproduced from <http://www.pfeffer.at/sandy/>)



Network data connects victims and responders to both locations and the needs they mentioned. These expressions can be used to identify critical actors and places that must be reached by responders. For ease of use, the ORA network

²² <http://www.liwc.net>
²³ <http://www.indiana.edu/~socpsy/ACT/data.html>
²⁴ <http://sentiwordnet.isti.cnr.it>

analysis tool bakes a variety of useful social network metrics into reports to provide overall assessments of different situations. One of these reports has been designed to work with data from TweetTracker. It transforms the data to extract networks of retweets, hashtag co-occurrences, users and content, user and locations, and popular keyword distributions. It then processes these networks to identify influential Twitter users, core topics, and changing regions of concern. A similar technology has been built with ORA using REA for analyzing Lexis-Nexis data²⁵. This technology has been use with respect to natural and man-initiated crises [62]. Its simplicity lends itself to first response. Figure 1 shows a network created from Twitter data using ORA.

For analysts who wish to go deeper and possibly conduct richer data mining on network structures, ORA supports the extraction of a variety of different social network metrics. Other GUI-based tools such as Gephi²⁶ and Cytoscape²⁷ also provide methods for analysts to approach the data, but a variety exist. In contrast, if the analyst wants to take a programming approach and develop their own network metrics they may want to work with the statnet package²⁸ for R or the NetworkX library²⁹ for Python.

3.5 Geolocation

The classic tool used for geo-spatial analysis in the crisis mapping area is ESRI ArcGIS³⁰. ArcGIS is widely used by large number of response units including many police departments and military units. It supports pinning a variety of latitude/longitude data to maps, as well as visualizing changes in its distribution over time. In addition, ArcGIS supports a full complement of spatial analytics, and a layered visualization scheme. ArcGIS can import and export shapefiles, demarcations of geographic shapes, and KML, the XML-based markup language developed for use with Google Earth³¹. An increasing number of crisis-mapping tools, particularly those used by the large first responders, are exporting data in KML to support interoperability. Open source GIS tools are appearing that contain many of the features inherent in ArcGIS.

However, since the advent of Google Maps³² eight years ago, an increasing number of crisis response tools are making use of it as an alternative. Since then, the quantity of data and tools available for working with geospatial data has only

²⁵ Illustrative results generating using ORA with Sandy data can be seen at

<http://www.pfeffer.at/sandy/>

²⁶ <http://gephi.org>

²⁷ <http://www.cytoscape.org>

²⁸ <https://statnet.csde.washington.edu/trac>

²⁹ <http://networkx.lanl.gov>

³⁰ <http://www.esri.com/software/arcgis>

³¹ <http://www.google.com/earth/index.html>

³² <http://maps.google.com>

increased. According to the HHI's report, the V&TC community active in the Haiti earthquake particularly shone in its use of geospatial data. This is due to the dedicated work of the crisis mapping community and the willing participation of organizations with access to satellite imagery in crisis situations. In Haiti, a partnership between Google and GeoEye provided high-resolution images of the disaster area from above. With the right data, communities could annotate maps and workers on the ground could plan their activities.

Even when corporate entities do not provide such useful material, the community is able to rely on open platforms like the mapping site OpenStreetMap³³, which has become a staple of the crisis community. All of the mapping data on OpenStreetMap has been contributed by volunteers; individuals upload GPS data to the site, and then annotate and edit it to keep it current. To deal with situations where internet access is limited or where users don't have access to GPS equipment, Michael Magurski released first the Walking Papers³⁴ and then Field Papers³⁵ tools. These allow users to download, print, annotate, and then upload the annotations to OpenStreetMap.

Google Maps has a growing presence in the crisis mapping community as well, and Google has itself devoted resources to creating maps specifically of crisis situations. They've provided crisis maps for specific incidents such as Superstorm Sandy that have been annotated with a variety of user data culled from the web [63]. Google also maintains a real-time crisis map³⁶ that uses similar culling of data to provide updates about potential and on-going crisis situations.

The TweetTracker tool developed at ASU visualizes extracted tweets on maps and lets users set spatial bounding boxes for selecting tweets by placing squares on maps. (See Figure 2 for an example of an exported map.) ORA also supports visualizing networks and other data on maps. It can import and export shape files and KML. In addition ORA allows users to cluster entities based on their particular region and then use that clustering as an element of a social network analysis.

While it would be incorrect to consider the challenge of properly representing data that has been connected with physical locations a solved problem, at this point there are a variety of tools that allow users to place information with specific latitudes and longitudes on a map. The research challenges are no longer about rendering these points in an informative manner. They are about developing new algorithms for deriving data from geographical clusters, and analyzing and forecasting the geographic distributions of social media postings in specific disasters.

³³ <http://www.openstreetmap.org>

³⁴ <http://walking-papers.org>

³⁵ <http://fieldpapers.org>

³⁶ <http://google.org/crisismap/>



Fig. 2 A Google Map of locations from which tweets were received during a disaster simulation exercise. Tweet data was obtained using TweetTracker. See [64] for more information on the exercise.

3.6 Text Analysis

Covering the complete realm of automated text analysis is beyond this paper's scope. The field is immense and growing. Part of this expansion has incorporated the development of a variety of tools to make it easier to break down text and treat as quantitative data.

In general, toolkits in this area rely on the analyst being moderately familiar with programming languages. A GUI-based tool will be markedly easier for an analyst to work with if they lack the time or ability to code, or if they fail to thoroughly familiarize themselves with the language before a disaster strikes. They may also be difficult for first responders to integrate into a workflow, depending on the other tools they are using.

For Java, LingPipe³⁷ provides a useful library for tokenizing sentences, calculating sentiment, and stemming words, among other features. MALLET³⁸, developed by the University of Massachusetts at Amherst, provides some similar functions but focuses on using text for machine learning. If an analyst would rather not work with code directly, they can use the packages as precompiled binaries. Weka³⁹, a toolkit for running machine learning experiments developed at the University of Waikato, while distributed as an application with a GUI, can also be used as a Java library. While not specifically for text, like all machine learning packages it can be trained on textual features that have been quantified.

For Python, the Natural Language Toolkit⁴⁰ (NLTK) provides some of the same functionality as LingPipe. NLTK supports tokenization, stemming, text tagging and other standard natural language processing techniques. While incorporating some learning algorithms, analysts may want to investigate dedicated solutions like mlpy.⁴¹

³⁷ <http://alias-i.com/lingpipe/index.html>

³⁸ <http://mallet.cs.umass.edu>

³⁹ <http://www.cs.waikato.ac.nz/ml/weka/>

⁴⁰ <http://nltk.org>

⁴¹ <http://mlpy.sourceforge.net>

A large variety of machine learning models for working with text have been implemented as packages for the statistical language R. The `tm` package⁴² bundles together standard natural language processing features for working with unstructured text. Once parsed, other packages oriented specifically towards data mining can be used with the text.

GUI-based tools for working with text data also exist, and may be easier for first responders to integrate into their workflows than a coding solution. One good example is `AutoMap`⁴³, a tool developed at the CASOS Center at Carnegie Mellon University that supports both GUI-based cleaning and an XML-based scripting language [65]. Like NLTK and other tools mentioned, `AutoMap` provides a number of methods for cleaning text documents like stemming words to their base forms, deleting stop words, and calculating the frequency of different multi-word sequences. `AutoMap`'s scripting GUI makes it relatively easy to improvise and modify cleaning processes on the fly. The program has also been significantly integrated with `ORA`, allowing analysts to use network metrics to identify prominent co-occurrences of particular words or entities mentioned in documents. These networks of texts can also be visualized and –if referencing geospatial data– can be pinned to maps. This approach was used by a team of Arizona State University and Carnegie Mellon University researchers with data from Superstorm Sandy to compare the difference in content between Twitter and the news media.

One difficulty of working with text data posted to Twitter and other microblogs is that it often doesn't fit the conventions expected in ordinary text. When `ANPI_NLP` developed their named entity recognizer, for example, they had to first train a morphological analyzer to correctly split a tweet into names. Analysts generally expect to have to train their own parsers when working with microblog syntax. While not a general purpose named entity recognizer, Gimpel et al. have developed a tokenizer and part-of-speech tagger for Twitter⁴⁴ that has since been improved by Owaputi et al. [66, 67]. The POS tagger correctly classifies emoticons and the roles of various acronyms (“lol”, “srsly”). While not critical for disaster on its own, in combination with the methods used by `ANPI_NLP` this could improve the speed and accuracy of other algorithms.

Translation of messages posted to social media in other countries remains a pressing problem, as we have discussed when describing the SMS messages translated by Mission 4636. This problem was also seen during the Egyptian Revolution and in the Yushu earthquake in China. While crowdsourcing markets are a proven solution for this problem, machine translation can also be used for potentially faster results. Google, for example, provides access to an API for automatic translation.⁴⁵ These will be less effective than native speakers of a particular language, but if it isn't possible to reasonably mobilize (or afford to mobilize) such a platform, machine translation is one possible alternative.

⁴² <http://tm.r-forge.r-project.org>

⁴³ <http://www.casos.cs.cmu.edu/projects/automap/>

⁴⁴ <http://www.ark.cs.cmu.edu/TweetNLP/>

⁴⁵ <http://developers.google.com/translate/>

3.7 *Broadcasting*

Broadcasting tools largely fall outside of the practical use case for analysts. They are, however, relevant for first responders attempting to leverage social media, so we mention them here briefly. One example of a broadcasting tool is HootSuite⁴⁶, which allows users to manage profiles on multiple social networks, time the broadcasting of particular tweets, and perform some analytics similar to those mentioned in our discussion of tools that can be used for data retrieval. TweetDeck⁴⁷, an application provided by Twitter, provides a few similar functions but only for Twitter: users can use the software to control multiple Twitter accounts, subdivide followers into different groups, and schedule particular tweets to be posted at certain times.

Regardless of these relatively sophisticated tools, first responders will often interact with followers through the main interfaces of whatever particular social media service they are using. If Twitter, it may simply be their organization's account from the web, or the smartphone application of an organization member on the ground.

4 **Research Directions**

A common need felt by both people and organizations who turn to social media in disaster is knowing what is happening on the ground as rapidly as possible. Solving this problem has become the thrust of many ongoing research projects in the field. That being said, it is important to recognize that there are two very different audiences to whom this chapter is speaking: first-responders and disaster researchers. Each group needs different tools to pursue their own ends. First responders need easy to use simple tools with pre-defined workflows, specialized interfaces, dashboards, and maps. The time constraints of disasters prevent them from turning to powerful but less intuitive or rapid tools such as programming languages. In contrast, disaster researchers need to be able to use and create new methods, new types of visualizations, with workflows that they develop as part of the research. In this case, real-time performance is less important than the ability to perform sophisticated analyses. A particular type of research, translational research is needed in the disaster response area that supports the movement of those findings and tools discovered or invented by disaster researchers that are the most valuable to first responder from the laboratory into the field [68].

We now discuss two families of approaches to this challenge. We will begin with attempts to leverage machine learning and crowdsourcing to automatically classify individuals based on whether they provide useful information. We will then move on to discussing several different methods for visualizing social media data to provide immediate, intuitive feedback.

⁴⁶ <http://hootsuite.com>

⁴⁷ <http://tweetdeck.com>

In our section on individuals, we discussed the challenge of locating tweets that contributed to situational awareness and brought up the work of several researchers who have developed different categorizations for twitter messages. As mentioned earlier, Vieweg has developed a hierarchy of three overarching categories and 35 specific categories for situationally aware tweets [20]. She also experimented with the possibility of using VerbNet to automatically categorize tweets according to her model.

VerbNet is a lexicon of English verbs. It is a collection of verbs linked together based on a variety of different features including word senses, syntactic frames, and thematic roles, similar to both WordNet [69, 70]. Because tweets are generally only one or two sentences long, the verb can often be used as a critical identifier of a tweet's meaning. Vieweg identified nine VerbNet classes that were routinely present in her collection of situationally aware tweets. Testing on a large sample of both situationally aware and ordinary tweets, she found that 32.6% of a random sample of 4000 coded tweets contained SA data. While not perfect, systems incorporating these VerbNet codes is one step towards correctly validating data without human intervention.

Verma, along with Vieweg and several other researchers, tested the possibility of training a machine learning classifier to identify situationally aware tweets in a variety of disasters [71]. Working with the same Twitter data used by Vieweg, the researchers trained a Maximum Entropy classifier to reach between 84.1% and 88.8% accuracy on each data set. Prior to training, the researchers generated not only unigram and bigram features but also predicted subjectivity/objectivity, formal/informal register, and personal/impersonal tone as predicted by several other classifiers. The data were also coded with parts of speech tags, with a primary focus on identifying adjective use.

Similarly, Starbird et al. have experimented with using Support Vector Machines to try and identify the small number of individuals tweeting locally [72]. Using tweets broadcast during Occupy Wall Street, the researchers trained their classifier on a set of profile features such as times retweeted, number of followers, and whether stated profile location changed over time. Their final classifier still only correctly classifies 67.9% of those tweeting locally. While useful, there is still significant room for improvement.

Given the effectiveness of using crowdsourcing to classify disaster data, there is a strong argument to be made for feeding volunteers that has been classified with some level of error and expecting them to filter out the bad from the good. Another possibility is to integrate the volunteer crowd with the algorithm itself, having the users correct and retrain the algorithm on the fly. Settles has implemented an example of one such system, Dualist [73]. Users interact with the program by both coding documents with correct labels and by correcting labels assigned by the classifier. The importance of the accomplishment in this case is not simply the integration of a user into a conventional machine learning classifier but also the interface for the classification. This is not just a problem of algorithm design but also of constructing a useful interface.

The research projects we have discussed so far have focused on trying to find the useful tweets within the broader pool of data. Some researchers have taken an alternate approach, opting to find general information from the general mass of tweets. For example, Sakaki et al. have used Twitter data to detect earthquake epicenters [74]. Using the small number of tweets that have location data for references to earthquakes, they combine both support vector machines and particle filters to account for the uncertainty of the reported physical locations and then calculate the likely epicenter. Their system is effective but contingent on having a large number of tweets tagged with particular locations.

Similarly, the Google Flu Trends project⁴⁸ uses search queries made to Google to identify outbreaks of influenza [75]. Flu Trends is a specialized version of Google Trends in general, which tries to identify trending searches on Google just as Twitter tries to identify trending topics discussed by its users. The tool's success depends on both the large number of searches and also a lack of bias in the search data.

Going beyond microblog text, Fontugne et al. have investigated Flickr's potential for disaster detection [76]. The researchers have developed a prototype system that tracks uploaded photographs, highlighting particular labels that are being uploaded by multiple users at once. Their method captured large bursts of activity in Miyagi prefecture in Japan after the Tohoku earthquake. While the system shows potency as an alarm system, the researchers also point out that only 7% of the photographs taken within 24 hours of the earthquake were uploaded within that 24 hours. This is a dramatically different usage pattern from Twitter, and one that should impact proposed research to leverage Flickr data.

Visualizations of social media data is another ongoing challenge for helping users comprehend the sea of social media information. While crowdsourcing and machine learning can help us prepare data, it is often a visualization that helps individuals understand what the data is saying.

Word clouds have become a staple of modern visualization, as websites such as Wordle⁴⁹ have made them easy to create from any readily available text. Researchers have also looked into optimizing the patterns of words in word clouds to make them easier to interpret [77]. One notable example of their practical use is the Eddi system developed by Bernstein et al. [78]. Eddi assigns a set of topic labels to particular tweets by treating them as web search queries and then identifying prominent terms in the resultant searches. These topic labels are displayed as tag clouds that can be used to identify prominent subject of discussion. Note that Eddi's primary achievement is its insightful method of finding categorizations for tweets. However, the system relies on simple tag cloud systems as a key component of its visualization scheme.

ORA also incorporates a word cloud visualization. When fed longitudinal data, it allows the user to render a sequence of word clouds as networks that can be monitored changing over time. This is then supplemented with the ability to track

⁴⁸ <http://www.google.org/flutrends/>

⁴⁹ <http://www.wordle.net/>

the criticality of topics (e.g., Hashtags) and actors (e.g. Tweeters) in the different clouds, tracking how different topics have come into or dropped out of prominence over the course of an event.

Kas et al. have had success using tree maps to display tweets prominently associated with particular topics [79]. The researchers calculate the co-occurrences of all words in tweets collected on particular topics, filter words based on how often they co-occur, and then calculate popularity within particular topics. The most prominent topic keywords are then placed in a tree map, sized based on the square roots of their overall frequencies. The researchers carried out a small user study comparing the effectiveness of using word clouds and tree maps to display the ranked words from Twitter. They found that in general tree maps were significantly more useful; test subjects both better identified data presented in the tree maps relative to that presented in the word clouds but also significantly preferred using the tree map visualization.

Word clouds and tree maps are both relatively established forms of visualization. Both methods are constrained by only displaying a static view of the world. Social media, however, is often in flux. To understand a particular sequence of events it can be useful to get back to the originator of a particular comment, tweet, or image in order to understand how it has come to have significance. Shahaf et al. have developed a new, alternative visualization, the metro map, that addresses this problem for longer documents but has potential for being adapted to the Twitter space [80]. The metro map visualization links together sequences of documents based on shared features. Documents are represented as “stations”, like a traditional metro map, arranged roughly chronologically. The documents are tied together by directed “tracks” derived from the amount of overlap in coherence, coverage and connectivity in the actual text of the documents. Coherence is measured based on the overlapping content of articles, coverage as the number of topics mentioned across the collection of documents, and connectivity as the number of connections that exist.

The visualizations we have discussed have all focused on social media as a general source of data. We cannot point to particular examples of visualizations of social media data that are disaster specific. For example, there is no visualization scheme based on Vieweg’s categories for social media messages posted in disaster. This is a notable gap, and one that research needs to speak to. Visualizations that cater to a specific end can be much more effective than a general tool. For example, Kamvar and Harris’s “We Feel Fine”, a set of visualizations of individual emotions on Twitter, has caused users to engage in introspection and personal probing [81]. This is partly due to the text, which consists of personal statements, and partly due to the way in which the text has been represented. Visualizations designed to highlight the features of disaster could provoke similarly reach responses from users while also speaking to relief workers and analyst’s needs to understand the situation on the ground.

5 Conclusion

In this chapter, we have reviewed how social media is used in disaster by individuals, first responders, and disaster researchers. We have also introduced a variety of software tools that can be used by analysts to work with social media, the utility of which will vary depending on whether the analyst is a first responder or a disaster researcher. We have concluded with a discussion of several different directions in which some of the research on social media usage in disaster is currently heading.

For individuals impacted by the event, we have sought to highlight that in crises people turn to technology in order to find information and to find each other. Some malicious individuals will turn to social media to spread havoc. Social media platforms like Twitter become avenues for people to both seek information and express distress when they aren't certain where else to go. It's also a venue for publicizing disasters, for becoming involved in the large pool of social interactions surrounding a particular disaster, and for propagating false information related to a disaster.

For first responders, a critical concern is that a small amount of locally actionable information is being lost within a large pool of irrelevant noise. Locating this information remains a key challenge. Researchers have been and continue to develop schemes for categorizing the different types of messages sent in disasters. They have also looked at how users of social media respond to the propagation of falsehoods, and at how groups of organized individuals can be mobilized to crowdsource the categorization of distress messages.

When organizations impacted by the event turn to social media, they do so for similar reasons to individuals: to find new information about ongoing disasters, to communicate with individuals looking to them as authorities, and to stay in contact with followers. First responders and relief organizations use social media in these ways as well. In addition, they will use teams of individuals to monitor social media to find the critical pieces of information posted by niche users that they can use in planning disaster responses. They also post their own updates and information, providing authority in what is often a sea of rumors, and stopping firestorms of false information. Volunteer-based communities come together to analyze social media data, and a slew of new tools have been developed by organizations to help them turn this new data into actionable information.

While some tools for handling social media data have been developed specifically in the context of helping to resolve disasters, many others have been developed for the broader market. For example, the same company that developed the Ushahidi platform for collating disaster information also created the SwiftRiver tool for collating different streams of social media information. While SwiftRiver has definite application during disaster, it can also be used in broader contexts to track the development of any sort of chain of events. Maintaining the distinction between disaster-focused tools and those that are more generally applicable can be counter-productive. Rather, we propose considering tools as

being situated in one of seven categories: data collection, workflow management, narrative construction, data processing, geolocation, text analysis, and broadcasting.

Current research speaks to these issues by trying to speed up our ability to comprehend what is being said on social media. This often takes the form of attempting to fit automated classifiers to data sets, as with Vieweg's fitting of VerbNet terms to tweets sent in distress. Given the large pools of volunteers interested in working with crisis data as well as the many markets for crowdsourced labor, other research has looked at the possibility of combining machine learning algorithms with human vetting, either by using machine learning to reduce the size of the data such that it can be handled by humans or by using humans to interactively train the machine learning algorithm. Researchers have also approached this problem from the standpoint of visualization. An insightful visual representation can rapidly summarize a large quantity of social media data. While word clouds and tree maps have been demonstrated to be useful, and metro maps provide an avenue for moving forward, the field remains open for new ideas in visualization. There is also a need for visualizations of disaster data that emphasizes the disaster aspect in tandem with that of social media. Researchers should look to these studies of communication in disaster and craft new visualizations that specifically highlight those interactions.

These two research threads are not intended to be an exhaustive catalog of the future. Rather, they are two trends that have fallen out of some of the ways in which social media has been used by individuals and organizations. No matter how the field progresses, how social media is being used should remain its guiding star. Only by understanding the stresses on individuals and organizations during disaster can research help them improve.

Acknowledgments. The authors would like to thank Dr. Huan Liu of Arizona State University for his great help in bringing this chapter to fruition.

References

1. Pfeffer, J., Carley, K.M.: Social Networks, Social Media, Social Change. In: Nicholson, D.M., Schmorow, D.D. (eds.) *Adv. Des. Cross-Cult. Act. Part II*, pp. 273–282. CRC Press (2012)
2. Pfeffer, J., Zorbach, T., Carley, K.M.: Understanding online firestorms: Negative word of mouth dynamics in social media networks. *J. Mark. Commun.* (2013)
3. Moloney, A.: Haiti must act to address housing crisis - Oxfam. Thompson Reuters Found. (2013)
4. Drabek, T.E.: *Human System Responses to Disaster: An Inventory of Sociological Findings*. Springer, New York (1986)
5. Dynes, R.R.: *Organized Behavior in Disaster*. Heath (1970)
6. Shklovski, I., Palen, L., Sutton, J.: Finding community through information and communication technology in disaster response. In: *Proc. 2008 ACM Conf. Comput. Support. Coop. Work*, pp. 127–136. ACM, San Diego (2008)

7. Shklovski, I., Burke, M., Kiesler, S., Kraut, R.: Technology Adoption and Use in the Aftermath of Hurricane Katrina in New Orleans. *Am. Behav. Sci.* 53, 1228–1246 (2010), doi:10.1177/0002764209356252
8. Arcenaux, N., Weiss, A.S.: Seems stupid until you try it: press coverage of Twitter, 2006-9. *New Media Soc.* 12, 1262–1279 (2010), doi:10.1177/1461444809360773
9. Sullivan, D.: Tracking Hurricane Sandy News Through Twitter. *Mark. Land.* (2012)
10. Carr, D.: How Hurricane Sandy Slapped the Sarcasm Out of Twitter, *New York Media Decod.* (2012)
11. Laird, S.: Sandy Sparks 20 Million Tweets. *Mashable* (2012)
12. Munro, R., Manning, C.D.: Short message communications: users, topics, and in-language processing. In: *Proc. 2nd ACM Symp. Comput. Dev.*, pp. 1–10. ACM, Atlanta (2012)
13. Kwak, H., Lee, C., Park, H., Moon, S.: What is Twitter, a social network or a news media? In: *Proc. 19th Int. Conf. World Wide Web*, pp. 591–600. ACM, Raleigh (2010)
14. Java, A., Song, X., Finin, T., Tseng, B.: Why we twitter: understanding microblogging usage and communities. In: *Proc. 9th Webkdd 1st Sna-Kdd 2007 Work. Web Min. Soc. Netw. Anal.*, pp. 56–65. ACM, San Jose (2007)
15. Naaman, M., Boase, J., Lai, C.-H.: Is it really about me?: message content in social awareness streams. In: *Proc. 2010 ACM Conf. Comput. Support. Coop. Work, Cscw*, pp. 189–192. ACM, Savannah (2010)
16. Bakshy, E., Hofman, J.M., Mason, W.A., Watts, D.J.: Everyone’s an influencer: quantifying influence on twitter. In: *Proc. Fourth ACM Int. Conf. Web Search Data Min.*, pp. 65–74. ACM, Hong Kong (2011)
17. Starbird, K., Palen, L., Hughes, A.L., Vieweg, S.E.: Chatter on the red: what hazards threat reveals about the social life of microblogged information. In: *Proc. 2010 ACM Conf. Comput. Support. Coop. Work*, pp. 241–250. ACM, Savannah (2010)
18. Sinnappan, S., Farrell, C., Stewart, E.: Priceless Tweets! A Study on Twitter Messages Posted During Crisis: Black Saturday. In: *Proc. 2010 Australas. Conf. Inf. Syst. Acis* (2010)
19. Qu, Y., Huang, C., Zhang, P., Zhang, J.: Microblogging after a major disaster in China: a case study of the 2010 Yushu earthquake. In: *Proc. Acm 2011 Conf. Comput. Support. Coop. Work, Cscsw*, pp. 25–34. ACM, Hangzhou (2011)
20. Vieweg, S.E.: *Situational Awareness in Mass Emergency: A Behavioral and Linguistic Analysis of Microblogged Communications.* University of Colorado at Boulder (2012)
21. Sutton, J.: *Twittering Tennessee: Distributed Networks and Collaboration Following a Technological Disaster.* In: *Proc. 7th Int. Conf. Inf. Syst. Crisis Response Manag.* (2010)
22. Starbird, K., Palen, L. (How) will the revolution be retweeted?: information diffusion and the 2011 Egyptian uprising. In: *Proc. ACM 2012 Conf. Comput. Support. Coop. Work, Cscw*, pp. 7–16. ACM, Seattle (2012)
23. Centola, D., Macy, M.: Complex Contagions and the Weakness of Long Ties. *Am. J. Sociol.* 113, 702–734 (2007)
24. NPR Staff, “Distant Witness”: Social Media’s “Journalism Revolution.” *Talk Naton* (2013)
25. @TwitterMedia NPR’s Andy Carvin Uses Twitter to Debunk A Hoax. #OnlyOnTwitter
26. Kaczynski, A.: How One Well-Connected Pseudonymous Twitter Spread Fake News About Hurricane Sandy. *Buzzfeed Polit* (2012)

27. Stuef, J.: The Man Behind @ComfortablySmug, Hurricane Sandy's Worst Twitter Villain. BuzzFeed Fwd. (2012)
28. Mendoza, M., Poblete, B., Castillo, C.: Twitter under crisis: can we trust what we RT? In: Proc. First Work. Soc. Media Anal. Soma, pp. 71–79. ACM, Washington D.C. (2010)
29. Madrigal, A.C.: It Wasn't Sunil Tripathi: The Anatomy of a Misinformation Disaster. The Atlantic (2013)
30. Weinstein, A.: Everybody Named the Wrong Boston Suspects Last Night and Promptly Forgot. Gawker (2013)
31. Martin, E.: Reflections on the Recent Boston Crisis. Reddit Blog (2013)
32. Keller, J.: How Boston Police Won the Twitter Wars During the Marathon Bomber Hunt. Bloom. Bussinessweek (2013)
33. Mission 4636, Collaborating organizations and History. Mission 4636 (2010)
34. Harvard Humanitarian Initiative, Disaster Relief 2.0: The future of Information Sharing in Humanitarian Emergencies. Harvard Humanitarian Initiative, UN Office for the Coordination of Humanitarian Affairs, United Nations Foundation (2011)
35. Liu, S.B., Palen, L., Sutton, J., et al.: In search of the bigger picture: The emergent role of online photo sharing in times of disaster. In: Proc. 5th Int. Conf. Inf. Syst. Crisis Response Manag. (2008)
36. Cohen, S.E.: Sandy Marked a Shift for Social Media Use in Disasters. Emerg. Manag. (2013)
37. St. Denis, L.A., Hughes, A.L., Palen, L.: Trial by Fire: The Deployment of Trusted Digital Volunteers in the 2011 Shadow Lake Fire. In: Proc. 9th Int. Conf. Inf. Syst. Crisis Response Manag. (2012)
38. Reuter, S.: What is a Virtual Operations Support Team? Idisaster 20 (2012)
39. Stephens, K.: Understanding VOSTs (Virtual Operations Support Teams) Hint: It's All About Trust. West. Mass Smem (2012)
40. VOSG.us, About. Virtual Oper. Support Group (2011)
41. Panagiotopoulos, P., Ziaee Bigdeli, A., Sams, S.: "5 Days in August" – How London Local Authorities Used Twitter During the 2011 Riots. In: Scholl, H.J., Janssen, M., Wimmer, M.A., Moe, C.E., Flak, L.S. (eds.) EGOV 2012. LNCS, vol. 7443, pp. 102–113. Springer, Heidelberg (2012)
42. Sarcevic, A., Palen, L., White, J., et al.: "Beacons of hope" in decentralized coordination: learning from on-the-ground medical twitterers during the 2010 Haiti earthquake. In: Proc. ACM 2012 Conf. Comput. Support. Coop. Work, Cscw, pp. 47–56. ACM, Seattle (2012)
43. Munro, R.: Crowdsourcing and the crisis-affected community: Lessons learned and looking forward from Mission 4636. Inf. Retr. 16, 210–266 (2013), doi:10.1007/s10791-012-9203-2
44. Okolloh, O.: Ushahidi, or "testimony": Web 2.0 tools for crowdsourcing crisis information. Particip. Learn. Action 59, 65–70 (2009)
45. Neubig, G., Yuichiroh, M., Masato, H., Koji, M.: Safety Information Mining — What can NLP do in a disaster—. In: Proc. 5th Int. Jt. Conf. Nat. Lang. Process. Asian Federation of Natural Language Processing, Chiang Mai, Thailand, pp. 965–973 (2011)
46. Kumar, S., Barbier, G., Abbasi, M.A., Liu, H.: TweetTracker: An Analysis Tool for Humanitarian and Disaster Relief. In: Proc. 2011 Int. Aaai Conf. Weblogs Soc. Media, pp. 661–662. AAI, Barcelona (2011)

47. Carley, K.M., Reminga, J., Storricks, J., Columbus, D.: *ORA User's Guide*, Carnegie Mellon University, School of Computer Science, Institute for Software Research, Pittsburgh, Pennsylvania (2013)
48. Carley, K.M., Columbus, D.: *Basic Lessons in ORA and AutoMap 2011*. Carnegie Mellon University, Pittsburgh (2011)
49. Carley, K.M., Pfeffer, J.: *Dynamic Network Analysis (DNA) and ORA*. *Adv. Des. Cross-Cult. Act. Part* (2012)
50. Costa, B., Boiney, J.: *Social Radar*. MITRE, McLean, Virginia, USA (2012)
51. Mathieu, J., Fulk, M., Lorber, M., et al.: *Social Radar Workflows, Dashboards, and Environments*. MITRE, Bedford (2012)
52. Schmerl, B., Garlan, D., Dwivedi, V., et al.: *SORASCS: a case study in SOA-based platform design for socio-cultural analysis*. In: *Proc. 33rd Int. Conf. Softw. Eng.*, pp. 643–652. ACM, Waikiki (2011)
53. Garlan, D., Schmerl, B., Dwivedi, V., et al.: *Specifying Workflows in SORASCS to Automate and Share Common HSCB Processes*. In: *Proc. Hscb Focus. Integrating Soc. Sci. Theory Anal. Methods Oper. Use* (2011), doi:10.1.1.190.7086
54. Bostock, M., Ogievetsky, V., Heer, J.: *D3 Data-Driven Documents*. *IEEE Trans. Vis. Comput. Graph.* 17, 2301–2309 (2011), doi:10.1109/TVCG.2011.185
55. Bostock, M., Carter, S.: *Wind Speeds Along Hurricane Sandy's Path - Interactive Feature*, New York (2012)
56. Bostock, M., Ericson, M., Leonhardt, D., Marsh, B.: *Across U.S. Companies, Tax Rates Vary Greatly*, New York (2013)
57. Bostock, M., Bradsher, K.: *China Still Dominates, but Some Manufacturers Look Else-where*, New York (2013)
58. Carvin, A.: *Andy Carvin's Social Stories*. *Andy Carvins Soc. Stories*
59. Lin, J., Snow, R., Morgan, W.: *Smoothing techniques for adaptive online language models: topic tracking in tweet streams*. In: *Proc. 17th Acm Sigkdd Int. Conf. Knowl. Discov. Data Min.*, pp. 422–429. ACM, San Diego (2011)
60. Martin, M.K., Pfeffer, J., Carley, K.M.: *Network text analysis of conceptual overlap in interviews, newspaper articles and keywords*. *Soc. Netw. Anal. Min.* (forthcoming)
61. Pang, B., Lee, L.: *Opinion Mining and Sentiment Analysis*. Now Publishers (2008)
62. Carley, K.M., Pfeffer, J., Morstatter, F., et al.: *Near Real Time Assessment of Social Media Using Geo-Temporal Network Analytics*. In: *Proc. 2013 Ieeeacm Int. Conf. Adv. Soc. Networks Anal. Min.* (2013)
63. Schroeder, S.: *Google Launches Crisis Map for Hurricane Sandy*. Mashable (2012)
64. Abbasi, M.-A., Kumar, S., Filho, J.A.A., Liu, H.: *Lessons Learned in Using Social Media for Disaster Relief - ASU Crisis Response Game* (2012)
65. Carley, K.M., Columbus, D., Landwehr, P.: *AutoMap User's Guide*, Carnegie-Mellon University, School of Computer Science, Institute for Software Research, Pittsburgh, Pennsylvania (2013)
66. Gimpel, K., Schneider, N., O'Connor, B., et al.: *Part-of-Speech Tagging for Twitter: Annotation, Features, and Experiments*. In: *Proc. 49th Annu. Meet. Assoc. Comput. Linguist. Hum. Lang. Technol.* (2011)
67. Owoputi, O., O'Connor, B., Dyer, C., et al.: *Part-of-Speech Tagging for Twitter: Word Clusters and Other Advances*. Carnegie Mellon University, Machine Learning Department, Pittsburgh, Pennsylvania, USA (2012)
68. Woolf, S.H.: *The Meaning of Translational Research and Why It Matters*. *J. Am. Med. Assoc.* 299, 211–213 (2008), doi:10.1001/jama.2007.26

69. Kipper Schuler, K.: *VerbNet: A broad-coverage, comprehensive verb lexicon*. University of Pennsylvania (2005)
70. Fellbaum, C.: *WordNet: an electronic lexical database*. The MIT Press (1998)
71. Verma, S., Vieweg, S.E., Corvey, W.J., et al.: *Natural Language Processing to the Rescue? Extracting “Situational Awareness” Tweets During Mass Emergency*. In: *Proc. 2011 Int. Aaai Conf. Weblogs Soc. Media* (2011)
72. Starbird, K., Muzny, G., Palen, L.: *Learning from the Crowd: Collaborative Filtering Techniques for Identifying On-the-Ground Twitterers during Mass Disruptions*. In: *Proc. 9th Int. Conf. Syst. Inf. Syst. Crisis Response Manag. Iscram* (2012)
73. Settles, B.: *Closing the loop: fast, interactive semi-supervised annotation with queries on features and instances*. In: *Proc. Conf. Empir. Methods Nat. Lang.*, pp. 1467–1478. Association for Computational Linguistics, Edinburgh (2011)
74. Sakaki, T., Okazaki, M., Matsuo, Y.: *Earthquake shakes Twitter users: real-time event detection by social sensors*. *Proc. 19th Int. Conf. World Wide Web*, pp. 851–860. ACM, Raleigh (2010)
75. Carneiro, H.A., Mylonakis, E.: *Google Trends: A Web-Based Tool for Real-Time Surveillance of Disease Outbreaks*. *Clin. Infect Dis.* 49, 1557–1564 (2009), doi:10.1086/630200
76. Fontugne, R., Cho, K., Won, Y., Fukuda, K.: *Disasters seen through Flickr cameras*. In: *Proc. Spec. Work. Internet Disasters*, pp. 1–10. ACM, Tokyo (2011)
77. Rivadeneira, A.W., Gruen, D.M., Muller, M.J., Millen, D.R.: *Getting our head in the clouds: toward evaluation studies of tagclouds*. In: *Proc. Sigchi Conf. Hum. Factors Comput. Syst.*, pp. 995–998. ACM, San Jose (2007)
78. Bernstein, M.S., Suh, B., Hong, L., et al.: *Eddi: interactive topic-based browsing of social status streams*. In: *Proc. 23rd Annu. ACM Symp. User Interface Softw. Technol.*, pp. 303–312. ACM, New York (2010)
79. Kas, M., Suh, B.: *Visual Summarization for Topical Clusters in Twitter Streams* (forthcoming, 2013)
80. Shahaf, D., Guestrin, C., Horvitz, E.: *Trains of thought: generating information maps*. In: *Proc. 21st Int. Conf. World Wide Web*, pp. 899–908. ACM, Lyon (2012)
81. Kamvar, S.D., Harris, J.: *We feel fine and searching the emotional web*. In: *Proc. Fourth ACM Int. Conf. Web Search Data Min.*, pp. 117–126. ACM, Hong Kong (2011)