Resonance+: Augmenting Collective Attention to Find Information on Public Cognition and Perception of Risk

Di Wang*
University of Utah, School of Computing
dwang0127@gmail.com

Marina Kogan
University of Utah, School of Computing
kogan@cs.utah.edu

ABSTRACT

Microblogging platforms have been increasingly used by the public and crisis managers in crisis. The increasing volume of data has made such platforms more difficult for officials to find on-the-ground information and understand the public's perception of the evolving risks. The crisis informatics literature has proposed various technological solutions to find relevant information from social media. However, the cognitive processes of the affected population and their subsequent responses, such as perceptions, emotional and behavioral responses, are still under-examined at scale. Yet, such information is important for gauging public perception of risks, an important task for PIOs and emergency managers. In this work, we leverage the noise-cutting power of collective attention and take cues from the Protective Action Decision Model, to propose a method that estimates shifts in collective attention with a special focus on the cognitive processes of those affected and their subsequent responses.

Keywords

Crisis Informatics, Social Media Data, Word Embedding, Cognitive Process, Protective Action Decision Model

INTRODUCTION

Social media platforms are increasingly used by the general public and crisis management agencies during crises, providing public information officers (PIOs) and crisis managers with new opportunities to gauge the severity and impact of the crisis and disseminate consistent crisis communication to the affected communities (Martín et al. 2017; Kogan et al. 2015). However, according to prior research, the majority of PIOs and crisis managers predominantly distribute information on social media platforms (Houston et al. 2015; Lin et al. 2016; Lachlan et al. 2014; Plotnick and Hiltz 2016; Reuter and Kaufhold 2018; San et al. 2012; Wukich and Mergel 2015). While interested in utilizing the information gathered from the general public, they often cannot leverage information gathered from social media to identify operational needs or gauge public's perception of incident response and recovery efforts (Knox 2022) due in part to the overabundance of information.

To help PIOs navigate through the data deluge on social media and better utilize social media content produced by the public, researchers in the crisis informatics and information retrieval have come up with important ontologies and baselines that are useful for filtering a social media stream down to information supporting situational awareness at scale or for routing that information to the appropriate stakeholder (Zahra et al. 2020; Rudra et al. 2017; Mazloom et al. 2018; McCreadie, C. L. Buntain, et al. 2019; McCreadie, C. Buntain, et al. 2020). Such information is useful for identifying unmet operational needs, which is the first task of PIOs in crisis. However, the other social media-related task of PIOs is gauging the public's perception of incident response and recovery efforts (*NIMS Basic Guidance for PIOs* 2020). This task is still under-addressed in crisis informatics. For such tasks, another type of information about people's cognition around risks is needed. This includes people's cognitive processes and their outcomes during crises, such as risk perceptions, and emotional and behavioral responses to those perceived risks. However, such information has only been studied qualitatively (Demuth et al. 2018; Stowe et al. 2018) and such case studies could not practically assist PIOs assessing the situation in future events. To address such information need at scale, we take cues from the theoretical cognitive model — Protective Action Decision Model (PADM) (Lindell

*corresponding author

and Perry 2012), which conceptualizes individual response to environmental hazards, including the cognitive role of information exposure and attention in forming risk perception and subsequent response. Thus in our work in progress, we operationalize the exposure and attention facets of PADM to better support the second information need of PIOs.

Though PIOs and crisis managers seek to gauge public risk perceptions with the use of information and communication technology (ICT), they are also hesitant to trust and act on information reported by regular users from social media platforms due to the inability to validate message content within a very short time frame (A. Crowe 2010; G. D. Haddow and K. S. Haddow 2014; Mehta et al. 2017; Tapia and Moore 2014; A. L. Hughes and Tapia 2015). According to the risk communication guidelines for public officials (*NIMS Basic Guidance for PIOs* 2020), PIOs and crisis managers are oftentimes advised to instead use social media platforms to attend to official, credible sources such as national agencies, publicized individuals, news outlets and more recently well-trained digital volunteers (Tapia and Moore 2014; A. L. Hughes and Tapia 2015). To ensure the messages produced by regular users could be trusted and used by PIOs, one approach by emergency management organizations is to validate social media content through intelligence gathering (Mehta et al. 2017). This presumes that the collective intelligence of the crowd is more effective at detecting important signals than the individual users, making them more trustworthy. This also suggests that the people's attention and decisions are based on aggregate patterns. Thus social media users should be treated as agents and observed from an aggregate, bird's-eye perspective in the information verification process by emergency management organizations or PIOs (Mehta et al. 2017). To achieve this task, the attention of the affected population needs to be studied collectively.

Collective attention refers to the degree to which public interest and awareness are focused on individual events, entities, or topics (Sasahara et al. 2013). Due to the complex dynamics of the emergence and dissipation of collective attention, many previous studies used simple proxies such as volume of posting and sharing activity (Mitra et al. 2017; Leavitt and Clark 2014), popularity of hashtags (Lehmann et al. 2011) or more nuanced information such as context descriptors (place names) (Stewart et al. 2020) in the message to measure collective attention on social media. Though such research provides insightful summaries that aid in understanding collective behavior, the results do not encompass the information about the cognitive processes. Instead, existing proxies generally target pre-defined ontologies in the literature (McCreadie, C. Buntain, et al. 2020; McCreadie, C. L. Buntain, et al. 2019) such as information supporting situational awareness. In addition, they also tend to focus on "overview" of the crisis or contentious topics such as politics, sports, or societal issues that do not benefit the crisis managers in monitoring public risk perceptions during crises. Thus, in this work in progress, we study the shifts of collective attention of the affected populations and pay special attention to their cognitive processes to better support the second task of PIOs and crisis managers.

BACKGROUND

Crisis Informatics

Crisis informatics is an interdisciplinary field of study that focuses on the use of Information and Communication Technologies (ICTs) and their adaptations in relation to mass emergencies (Hagar and Haythornthwaite 2005; Palen, Vieweg, et al. 2009). Many research strands in crisis informatics revolve around the issues of social computing in disaster—specifically, how people navigate and leverage the affordances of various sociotechnical systems in order to share information, collectively problem-solve, and self- organize in crisis (Starbird and Palen 2011). Similarly to early disaster sociology (Fritz and Mathewson 1957; Dynes 1970; Kendra and Wachtendorf 2003), crisis informatics literature emphasizes the active participation of the affected populations in the disaster-related information sharing and collective problem solving—now through the use and appropriation of ICTs. Moreover, crisis informatics research highlights the importance and value of the information produced and shared through this active participation and collective sensemaking by the affected users (A. Hughes et al. 2008; Liu et al. 2008; Palen, Vieweg, et al. 2009; Sutton et al. 2008). Therefore, we build on this historical emphasis of crisis informatics and focus on the salient social media content reflecting the psychological and behavioral processes of those affected.

The active participation and collective sensemaking of the affected populations, emphasized by crisis informatics research, contribute to the ever-increasing volumes of information being shared online in crisis. This is partly due to the long-documented phenomenon of convergence where people, resources, and information coalesce around physical sites of the disaster and the digital conversation around it (Fritz and Mathewson 1957; Kendra and Wachtendorf 2003; A. Hughes et al. 2008). This data deluge is increasingly difficult to navigate, especially considering the highly temporal and dynamic nature of crises. This is especially the case for PIOs and emergency management agencies, who aim to keep abreast of the public perceptions of the crisis, but struggle to do so due to the data deluge (Knox 2022). Thus, we seek to support PIOs' social media needs by following the cues of the collective attention to cut through the data deluge.

Finding Information Useful for Emergency Responders

According to recent studies focused on PIOs and crisis managers, PIOs are mainly seeking two types of information on social media platforms: information supporting situational awareness and information representing people's cognition and perception of risks during crises (*NIMS Basic Guidance for PIOs* 2020). To identify such information, previous work generally turns to machine-learning information classification tasks based on human annotations (Imran et al. 2016; Olteanu, Castillo, et al. 2014; Alam, Ofli, et al. 2018; McCreadie, C. L. Buntain, et al. 2019; Olteanu, Vieweg, et al. 2015). These classification tasks have evolved from classification of relevance (Olteanu, Castillo, et al. 2014; Stowe et al. 2018; CrowdFlower 2015) and disaster type (Stowe et al. 2018; Imran et al. 2016; Olteanu, Castillo, et al. 2014; Alam, Ofli, et al. 2018) to more nuanced information types such as information pertaining to donations, volunteering, advice, or eyewitness messages (McCreadie, C. L. Buntain, et al. 2019; Zahra et al. 2020). However, the existing classification tasks do not focus on understanding the cognitive processes of those locally affected, which would be useful for PIOs to understand public perception and responses. In our work, we build a scalable pipeline for finding information potentially useful to PIOs and other officials, by relying on collective attention to locate persistent messages (echoed by others) and augment it with localized focus on people's cognitive processes.

Researchers in crisis informatics have long aimed at surfacing the voices of the affected populations (J. Anderson et al. 2016). They have historically sought to highlight affected users' cognitive processes, including their threat perceptions, emotional and behavioral responses, in order to understand the dynamic ways in which people interact with risk information, think and feel about risks. However, these studies (Demuth et al. 2018; Stowe et al. 2018) have relied on qualitative analyses that can be prohibitive for the crisis managers. Thus, information about the cognitive processes of the affected population still needs to be located quantitatively, at scale, to help the public information officers and crisis managers navigate through the nuanced information produced by the locals. We seek to fill this gap with our work.

Protective Action Decision Model

We bridge this gap by taking cues from theoretical models in the risk communication and hazards literature that represent the individual sensemaking processes during crisis such as risk perceptions, emotional and behavioral responses. One such model is the Protective Action Decision Model (PADM) by Lindell and Perry (Lindell and Perry 2012). The model integrates the processing of information derived from social and environmental cues with messages that information sources transmit through communication channels to those at risk. The information flow in the PADM characterizes the way people make decisions about adopting actions to protect against environmental hazards. The processing of information begins with social/environmental cues or official risk communication messages that initiate a series of pre-decisional processes. The pre-decision processes are automatic processes of exposure, attention, and comprehension that take place outside of conscious processing for the information receivers. These subconscious processes can form the basis for decisions about how to respond to disaster situations, followed by protective action decision making and behavioral responses. Based on PADM, previous research has shown that social media data could be leveraged to deeply understand people's risk assessments (Demuth et al. 2018; J. Anderson et al. 2016), including risk information, interpretation and response from Twitter narratives. However, the pre-decision processes in PADM are still under-addressed, because people generally do not explicitly post such information on social media. Yet, with the availability of digital trace data, this process can be operationalized with contextual information (Palen and K. M. Anderson 2016), where the exposure and attention could be inferred from people's behavioral traces without explicitly reporting them. Thus, we scope the appropriate social media data to include information produced by those locally-affected, as well as the contextual information they are potentially exposed to estimated by the proxy of hashtags. We rely on hashtags because they represent ad hoc publics or communities that organize activity (Bruns and Burgess 2011) and collective attention on Twitter (Lehmann et al. 2011; Sasahara et al. 2013; Mitra et al. 2017; Leavitt and Clark 2014; Stewart et al. 2020), likely reflecting the kind of information to which a user would be exposed. We also operationalize the information exposure in our method in order to capture people's cognitive processes.

Trust and Collective Attention

Though social media enables PIOs and crisis managers to monitor online conversations through multi-way communications, this is problematic in some cases when the trust relationships between PIOs and individual users are not established. According to the literature (Mehta et al. 2017), emergency management organizations are hesitant to trust and act on information reported by regular users from social media platforms. This is due to the inability to validate message content within a very short time frame (A. S. Crowe 2012; G. D. Haddow and K. S. Haddow 2014; Mehta et al. 2017; Tapia and Moore 2014; A. L. Hughes and Tapia 2015). Thus they have utilized three broad approaches to validate social media content during disasters (Mehta et al. 2017). The three models are intelligence gathering, quasi-journalistic verification and crowdsourcing. The intelligence gathering model seeks to discover unknown activities or track known events by aggregating patterns of social media activity among large numbers of users instead of considering individual posts in-depth. Compared with the latter two approaches that require journalistic practices or the involvement of crowdsource workers and are more resource-intensive, the intelligence gathering model is the fastest and largely automated approach to extracting information from social media feeds for real-time applications. Thus we base our work on this model, focusing on people's collective attention and opinions over time.

Previous studies have analyzed collective attention on social media platforms to understand its dynamic fluctuations (Wu and Huberman 2007; Sasahara et al. 2013; Ratkiewicz et al. 2010; Stewart et al. 2020). Collective attention refers to the degree to which public interest and awareness is focused on individual events, entities, or topics (Sasahara et al. 2013). Collective attention often shifts in response to major events, such as sports matches, natural disasters, and political events (Stewart et al. 2020), exhibiting burst-like dynamics. Previous studies demonstrate that the attention strengthens to signify popularity shifts when particular real-world events occur (Sasahara et al. 2013; Ratkiewicz et al. 2010). In addition, the attention may also get counterbalanced when the novelty of the content fades with time (Wu and Huberman 2007), which could be caused by either the habituation or competition from other attention-garnering content. Thus, in tracking collective attention, it is important to study both the novelty of the content, as well as its persistence. In our work, we conceptualize the social media collective attention through the lens of resonance — combination of novelty and persistence — of individual content (Barron et al. 2017). Further, we augment resonance with aspects from PADM to better leverage it as a noise-cutting measure, specifically suited for the localized social media information needs of PIOs and other crisis managers.

METHODS

Data Collection and Scoping

As part of operationalization of the pre-decision exposure process in PADM, we scoped the data to focus on the geo-vulnerable Twitter users and their global information exposure. The complete workflow for data collection and processing is shown in Table 1. Our data collection consists of two parts: 1) identification of local users based on cross-filtering of keyword and geo-location based data collections and 2) data expansion to capture the hurricane-related tweets to which the local users are potentially exposed. We collected two datasets for Hurricane Harvey—a major Hurricane from the 2017 Atlantic hurricane season—using GNIP historical PowerTrack API: one based on the hurricane-related keywords and location names of the affected area (Keyword Data), and the other based on the geolocation information using a bounding box of the affected region (Geo-Vulnerable Dataset) shown in Figure 1. The datasets cover Aug. 17th and Sep. 9th, which spans a week before Harvey's landfall in Houston through a week after its dissipation. We cross-filtered the Keyword Dataset with the Geo-Vulnerable Dataset to identify the Harvey-related tweets geo-located in the affected area— Geo-Vulnerable Keyword Dataset. We consider users who produced these tweets to be the Geo-Vulnerable User group (29,325 users)—twitterers with one or more Harvey-related tweets within the affected region. For further data cleaning, we relied on the crisis-specific lexicon [CrisisLex] (Olteanu, Castillo, et al. 2014) to remove 20 users with the majority of tweets not containing any crisis-related lexicon terms. We also removed 14 automated accounts producing multiple tweets per second.

We then expanded the Geo-Vulnerable Keyword Dataset beyond those locally-affected to include the Harvey-related information that the locals are likely exposed to. Since prior literature has demonstrated that the use of Twitter hashtags has enabled the rapid formation of ad hoc issue publics (Bruns and Burgess 2011), capturing people's response to emerging issues and acute events, we use hashtags as a proxy for communities and thus the information to which people are exposed. Thus, to get the global exposure of the Geo-Vulnerable Users, we expanded the data based on the hashtag communities of the local users. The Harvey-related tweets produced by the locals contained a total of 16,287 hashtags. Thus, we expanded the Geo-Vulnerable Keyword Dataset by including all tweets from the Keyword Dataset containing at least one of these hashtags. This produced a dataset of 1,778,233 tweets, which includes the locals' posts, as well as their global Harvey-related information exposure — Geo-Vulnerable Expanded Dataset.

For this work, the data collected via geolocation information is used to crossfilter with the data collected via keywords for the data scoping. For other work when the Geo-Vulnerable Dataset is not available, researchers can use the geolocation information in the keyword data, applying to it a bounding box of the region affected by the disaster to identify the affected population. Then our method could be followed to identify the global exposure of the affected users. For Hurricane Harvey, we found the difference between the Geo-Vulnerable Expanded Datasets using these two methods is 0.8%.

Figure 1. Bounding Box for Data Collection

Table 1. Data Construction and Processing

DATASET NAME	DATA COLLECTION & CLEANING PROCESS	USER VOLUME	TWEET VOLUME
Keyword	Data collection with predefined list of keywords	2,772,713	9,807,408
Geo-Vulnerable	Data collection within bounding box of the affected region	79,903	759,195
Geo-Vulnerable with Keyword	Removing tweets without predefined Harvey related keywords	36,927	216,303
	Removing tweets with <3 tokens after preprocessing	29,359	139,872
	Removing 20 high frequency users with >80% tweets without any crisis related lexicon (CrisisLex)	29,339	136,709
	Removing 14 bot-like users	29,325	132,528
Geo-Vulnerable Expanded	Data expansion with hashtags by the locals	544,984	1,778,233

Resonance as a Metric of Collective Attention

In the following section we detail two methodological approaches to measuring collective attention: a simple resonance metric that is based solely on the novelty and persistence of the content and the augmented version of resonance that accounts for PADM-inspired exposure and local persistence— to more directly match the information needs of the PIOs and crisis managers.

Resonance: Content Novelty and Persistence

Collective attention bursts with the emergence of new topics(Lehmann et al. 2011; Sasahara et al. 2013). Eventually, the novelty of the content may fade due to habituation or competition of other information(Wu and Huberman 2007). Thus we build on existing work around the dynamics of collective attention and propose a metric to compute the asymmetry between content's novelty and its fading with time. Prior work used topic modeling and information theoretic approaches to track the formation and dissipation of topics (Barron et al. 2017) in historical manuscripts. Instead, we trained a word2vec model to estimate vectors representing each message, which works better for short documents and diverse online conversations (Kenter and Rijke 2015).

We applied standard NLP preprocessing and trained a word2vec model (Kenter and Rijke 2015; Mikolov, Sutskever, et al. 2013; Mikolov, Chen, et al. 2013) to represent the embedding vectors for each tweet. For each tweet, we calculated average cosine similarity of its embedding vector to those of tweets within the preceding time window of length w—its novelty(N_w). Cosine similarity is a measure of similarity between two vectors based on the cosine of the angle between them, often used to measure document similarity in text analysis. For Twitter, the information diffusion process generally has a short lifespan (Wu and Huberman 2007), with content rising and falling in popularity within hours. This is even more salient in the fast-changing context of disaster, thus we used 30 minutes as the time window, where the median number of tweets per 30 minutes is w = 1191 tweets. Similarly, the persistence of the current tweet is computed by averaging its cosine similarity with the tweets in the succeeding window of length w. Resonance is computed as the sum of novelty and persistence of a tweet: tweets with high resonance correspond to novel ideas that also persist over time.

Resonance Augmented with Exposure and Local Persistence

To ensure that the resonance metric fulfills PIOs' needs and highlights information representing people's cognitive processes and responses, we propose an augmented resonance metric by operationalizing the pre-decision processes in the theoretical Protective Action Decision Model (PADM). Specifically, we modify the resonance metric by transforming its components into in-conversation novelty and local persistence. According to the PADM, the input information on social media will not lead to the initiation of risk perception or appropriate protective actions unless people are exposed to, heed, and accurately interpret it. Thus, the novelty is adjusted to prioritize information produced by users with high exposure to the past information—users well-versed in similar prior conversations. Thus, we call it in-conversation novelty; it is a more focused metric than general novelty, which is easily diverted by the many parallel conversations on social media. The content's persistence is also adjusted to account for only the collective attention of those locally-affected, producing local persistence. We then combined the two into the augmented metric, **resonance+**, to represent the locals' attention with specific focus on cognitive processes.

We estimate the exposure of a user to prior conversations based on hashtags. We rely on the notion of hashtags as representing ad hoc issue publics (Bruns and Burgess 2011), essentially organizing Twitter discourse into conversational threads. For example, if a user posts with #houstonflood, they are more likely to have been exposed to other information posted under this hashtag. Since various hashtags may occupy a similar semantic space, we propose to treat such similar hashtags as related conversations. To accomplish this, we trained another word2vec model using only hashtags to estimate the similarities between them using cosine similarity. The exposure vector is then computed using embedding vectors representing the hashtag(s) to account for the similarity between hashtag(s) in a tweet and those in the preceding tweets (in window sized w). The exposure vector is then used to augment the novelty in order to estimate the user's information exposure to the past content. Specifically, the dot product of novelty vector N_w and exposure vector W_w is computed as the in-conversation novelty:

$$Novelty_{in\_conv} = (N_w . * W_w), \tag{1}$$

where $N_w = (n_1, n_2, ..., n_w)$, $W_w = (w_1, w_2, ..., w_w)$, and $w = 1191$.

Thus, high in-conversation novelty captures the shifts of attention (high N_w) within the same overall conversation — operationalized through related, similar hashtags. At the same time, the metric also captures the fact that Twitterer has higher exposure to prior related content (high W_w). In this way, we ensure that the metric is more focused on people's cognitive processes, since information generally does not initiate risk perceptions unless people are exposed to, heed, and accurately interpret it (Lindell and Perry 2012).

To thoughtfully augment the persistence of ideas, we compute a message's persistence among only the locals, to measure the dissipation of the collective attention of the locally-affected population over time. Specifically, we use mean cosine-similarity by comparing with tweets produced by the locals in the succeeding time window (median window size is 74 tweets).

Then the augmented resonance metric—-resonance+— is computed by adding the local persistence to the in-conversation novelty. The high resonance+ corresponds to the shift of local collective attention that has not faded with time, contributed by users who are more exposed to past information. As the aim of this work is to better support PIOs' information needs around public cognition and perception of the event, in our evaluation and results we only focus on the augmented resonance+ (we tested resonance as well, and it did not produce information around public's risk perceptions potentially useful to the PIOs).

Evaluation

To evaluate our method, we compare the tweets with high resonance scores to the highly-retweeted messages. Number of retweets is the most common proxy for engagement on Twitter and is often used to measure influence and importance of content. Several studies have specifically used retweets as a proxy for collective attention (Lehmann et al. 2011; Sasahara et al. 2013): one has to pay attention to content before propagating it to their social media network, especially as it is seen as an implicit recommendation for the content (Kogan et al. 2015). Thus, we found retweets to be a suitable comparison for our metric of collective attention. However, the resonance metric has been augmented to account for the attention of the locals. Therefore, we used the retweet count contributed only by the locals, providing another measure of local collective attention during the hyperlocal time of the disaster (Kogan et al.

2015). We compared the messages recommended by resonance and by the local-retweet metric quantitatively. We evaluated the results based on information source and information type of interest, namely, the cognitive processes and their outcomes.

Quantitative Evaluation using LIWC

To quantitatively compare content of high resonance tweets and the posts highly-retweeted by the affected population, we used categories in LIWC 2015 (Linguistic Inquiry and Word Count)(Pennebaker, Booth, et al. 2007; Pennebaker, Boyd, et al. 2015), which is widely used to study the emotional, cognitive, and structural components present in individuals' verbal and written speech samples. Based on our task, we focused on LIWC categories that represent people's cognitive processes, perceptual processes, emotions, as well as categories related to behaviors.

For the **cognitive processes**, we used several measures that quantify levels of multiple aspects of cognitive processing to analyze how people think, digest information, problem-solve, and make decisions. The categories we used for the cognitive processes are shown in Table 2 (Pennebaker, Booth, et al. 2007; Pennebaker, Boyd, et al. 2015; Receptiviti Inc. 2022).

Table 2. Cognitive Processes in LIWC

Category in LIWC	Summary
cogproc(Cognitive Process)	A measure of how people pay attention to and process social / environmental signals.
causation	A measure for people's engagement in causal thinking.
certainty	A measure of how a person uses words reflecting certainty or with the intention of claiming that something is true.
differentiation	The degree to which words are used to distinguish between entities people, or ideas.
discrepancies	The degree to which a person compares or articulates the difference between current states and alternative states.
insight	The degree to which a person is focused on understanding or gaining insight.
tentative	The degree to which a person signals uncertainty or uses non-definitive or hedging language.
comparison	The degree to which language is used for comparison.

For **perceptions**, we used the category **Perceptual Process**, which includes multiple sensory and perceptual dimensions associated with the people's senses. Here we also included subcategories of **See**, **Hear** and **Feel** to understand individual's focus of attention and their sensory experiences.

For **emotion**, we used the affective process in LIWC, which includes **Positive emotion** and **Negative emotion** (**Anxiety**, **Anger** and **Sadness**). In addition, we also included the category **Swear**, which is also often used to express negative emotions.

For **behavioral responses**, we included subcategories in Drives, which includes words related to motivational and behavioral drives. Specifically, we included the subcategory **Risk**, which describes references to aspects of danger, concerns, and things to avoid. We also included the category **Social Process**, which includes all non-first-person-singular personal pronouns as well as verbs that suggest human interactions. Additionally, we included several categories from personal concerns, including **Work**, **Leisure**, **Home**, **Money** and **Death**. Those categories would reflect people's activities as well as the potential impacts of the disaster. Because of the high-tempo, high-stakes nature of the crisis, we also included **Time**, **Space** and **Motion** categories from Relativity, all of which are highly related to people's physical movement.

FINDINGS

Since our work mainly focuses on the cognitive processes and the subsequent responses to better support PIOs' information needs around public cognition and perception of the event, we only report the results for the augmented resonance (resonance+). However, we also tested regular resonance metric, which performed similarly to local retweets in terms of recommending information reflecting people's cognitive processes, perceptions, emotional and behavioral responses. In this section, we first compare the posts recommended by high resonance+ with the highly retweeted messages among the locals in terms of information sources. Then we utilize LIWC to compare the content in terms of the information types of interest.

Information Source Distribution

The guidance from National Incident Management System (NIMS) suggests that PIOs need more information from the affected community (*NIMS Basic Guidance for PIOs* 2020). Thus, in this section we explore the information sources—the authors of tweets—highlighted by each metric.

We find that the top contributors for the top 1% highly-retweeted posts in the geo-vulnerable region were highly influential official users such as abc13houston (16%), NWSHouston (12%), JeffLindner1 (3.7%), HCSOTexas (3.4%), HoustonOEM (2.9%), HellerWeather (2.8%), ReadyHarris (2.7%), NWSSanAntonio (2.7%), SylvesterTurner (2.6%), HoustonISD (2.4%). These represent news organizations, weather and warning accounts, and local and state government officials. The authors of the top highly-retweeted tweets have a median number of followers of 25,266. On the other hand, there were largely no high-influence accounts contributing to tweets of high resonance, and the median number of followers was 796. Though organizations like FEMA have suggested that they are increasingly interested in capturing the bottom-up on-the-ground voices of those affected, including their social media activities, the results indicate that during a high-tempo crisis situation it is still hard for regular geo-vulnerable users without excessive social influence to have their voices heard (and propagated) by the general public. It also further emphasizes the need for a metric that highlights the voices of the regular affected individuals. Such information would provide a different perspective from both the traditional top-down risk communication and influencer-centered focus of social media platforms.

Information Type Comparison with LIWC

We started our analysis with simple resonance as a measure of the collective attention and then further augmented it with PADM-based information exposure and local attention. Results show that the augmented resonance metric (resonance+, with exposure and local attention) is better reflecting information about people's cognitive processes, perceptions, emotional and behavioral responses. For each category of interest, we compare the top 1% of high resonance+ tweets with that of highly-retweeted posts by the affected population in terms of LIWC categories they represent. For each LIWC category, the LIWC score is computed as the percentage of occurrences of the dictionary words in the document, ranging from 0 to 100. Then a two-proportion Z-test is used to compare the LIWC scores (divided by 100) for the tweets recommended by the two methods ($\alpha = 0.05$).

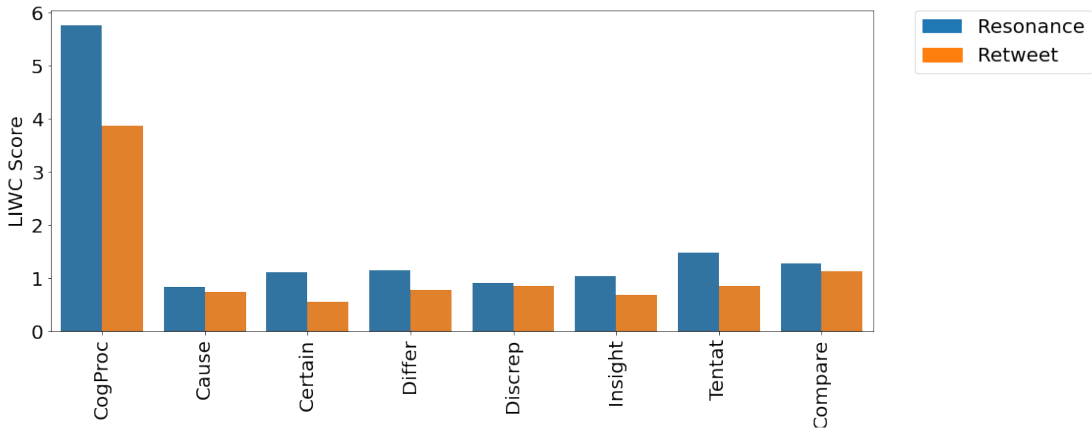

Figure 2. LIWC Score for Cognitive Processes

All the LIWC categories we chose to represent cognitive processes show a statistically significant difference between resonance+ and local retweets (p-value$<$0.05). Figure 2 shows the average LIWC score per tweet for each LIWC category of interest in the cognitive processes category. The results indicate that tweets recommended using resonance+ contain more words in the cognitive processes category than messages highly retweeted by the affected population. Two dimensions with higher LIWC scores in high resonance tweets are **Certainty** and **Tentativeness** (p-value$\ll$0.001 for both). The higher LIWC score in **Certainty** category suggests that some users have a high confidence in their decisions and actions and are aware of the preparations or protective actions they should take. The following example shows that the user is aware of the situation and is certain about the actions they should take.

`Aug 24, 2017,` **`Twitterer1:`**
`School has been canceled, have our batteries, candles, water & all the canned goods I could buy. We are ready for #Harvey. @***`

On the other hand, the higher score in **Tentativeness** suggests that some users are more cautious and unsure about their decision-making during the disaster. The following is an example showing that the user is not sure about the decision before Harvey's landfall.

`Aug 25, 2017,` **`Twitterer2:`**
`I'm debating on staying here or loading the kids up and heading to MS before #Harvey rolls in.....`

The cognitive processes captured by these dimensions will be useful for PIOs to understand people's cognition of risks and decision-making.

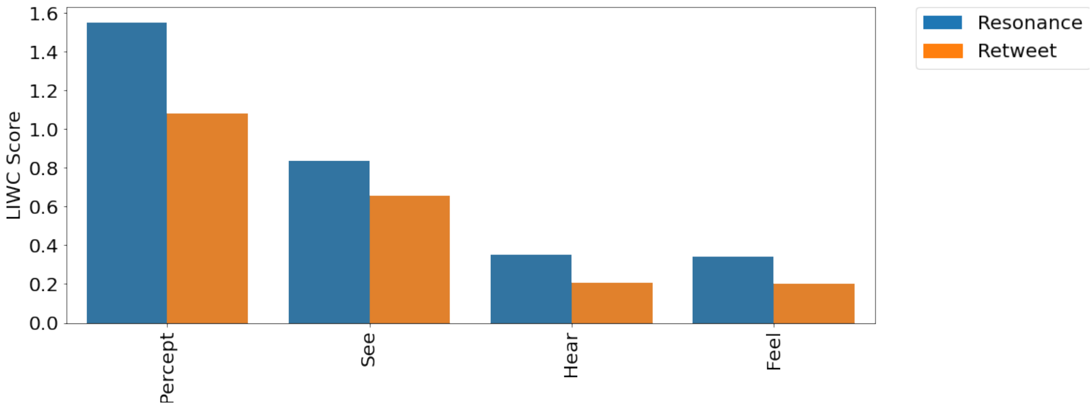

Figure 3. LIWC Score for Perceptions

All the LIWC categories we chose to represent perceptions are statistically significantly different (p-value<0.05) for resonance+ and local retweets (Figure 3). The messages related to perceptions captured via high resonance are mostly experiences of the affected individuals and how they perceive social and environmental cues or the evolving risk, whereas the highly retweeted messages are generally opinions or suggestions from official accounts. The results show that messages recommended by resonance+ provide more details about how the affected population experiences and makes sense of the situation, and thus may be useful to PIOs and crisis managers.

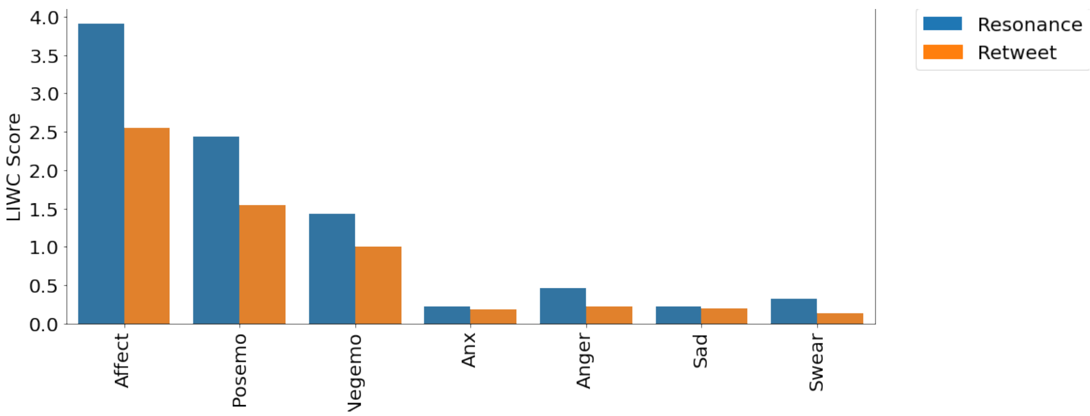

Figure 4. LIWC Score for Emotions

All the LIWC categories related to emotions are statistically significantly (p-value<0.05) except for **Anxiety** and **Sadness**(Figure 4). That means there is no evidence that the resonance metric outperforms the metric of retweets for these two emotions. Yet, results suggest that the resonance+ is more capable of capturing messages conveying other emotions. Specifically, the emotions expressed in high resonance tweets are mostly related to the experiences of the affected individuals such as disaster preparations, compared with prayers, worries or gratitude in the highly retweeted messages. The following tweet is an example of emotion captured by resonance+. However, such emotions are less frequent in highly retweeted tweets.

`Aug 25, 2017,` **`Twitterer3:`**
`People stocking up their houses with water and food.  #HurricaneHarvey is scaring me <Facepalming> <url>`

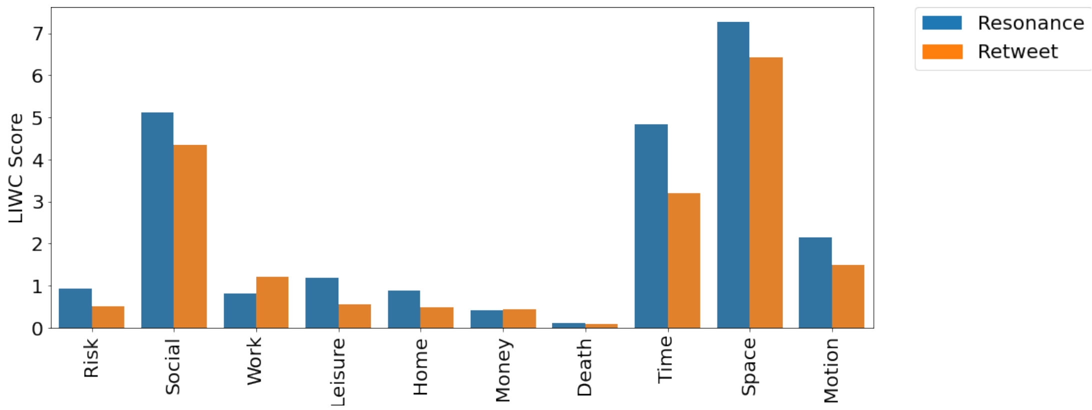

Figure 5. LIWC Score for Behavioral Responses

All the LIWC categories we chose to represent behavioral responses show a statistically significant difference between resonance+ and highly-retweeted posts among the locals (p-value<0.05), except for **Money** (Figure 5). However, though both high resonance and high retweets are similar in the **Money** category, they are used quite differently. For high resonance messages, such posts are related to expenses during disaster preparations. The following tweet is an example of personal expenses during disaster preparations from high resonance tweets. This is unlikely to be retweeted by other people. However, such message reveals the user's decision-making and response during the disaster preparation stage.

`Aug 25, 2017,` **`Twitterer4`**`:`
`Wife finished shopping in prep for #Harvey Spent $500 on lord only knows what but I hope we are good for the next 5 days`

For highly retweeted messages, such messages are generally related to donations. Though the latter (donation) represent call-for-action messages that help disaster recovery, the former ones are more related to behavioral responses of the affected population facing the disaster. The high resonance tweets have higher LIWC scores in **Risk**, **Leisure**, **Home**, **Motion**, **Space** and **Time** (p-value≪0.001) categories, indicating that high resonance tweets are more likely to be related to risk, coping behavior, people's movement in space (such as evacuation) or temporally-sensitive information. On the other hand, the highly retweeted messages have higher scores in the category **Work** (p-value≪0.001). This is due to the fact that many of these messages are official risk communication messages, including volunteering or recovery efforts. Though such information may be of equal importance for the first information need of PIOs , it is not related to subsequent behavioral responses of people's cognitive processes, which fulfills the second information need of PIOs — the focus of this work.

DISCUSSION AND FUTURE WORK

The increasing public engagement on the microblogging platforms has made it easier than ever to access the information from the affected populations. However, finding important, pivotal messages in such a high-volume environment proves to be a difficult task. According to a study of crisis After Action Reports (AARs) and focus groups with emergency managers (Knox 2022), challenges of social media use by the responders still remain, despite the increasingly available sociotechnical systems for disaster management (Starbird and Palen 2011; Lachlan et al. 2014; Plotnick and Hiltz 2016; Hilts et al. 2015). Challenges include information overload, which may increase the stress level for emergency managers and affect their decision making (Misra et al. 2020), locating information aiding in situational awareness (Vieweg 2012; Vieweg et al. 2010), and early and continuous monitoring of social media platforms to understand public attention and combat negative sentiments (Yeo et al. 2020).

Resonance+ leverages both the noise-cutting power of collective attention of the affected population and the cognitive role of information exposure and attention in forming risk perception and subsequent response from the theoretical PADM model. By recommending more information about people's cognitive processes, perceptions and subsequent responses, such a method could greatly benefit the PIOs and crisis managers in terms of gauging public's perceptions of the event and the official response. Existing tools that PIOs use for social listening (Hootsuite, Sprout Social, Brandwatch, TweetDeck or Keyhole) fall into two major categories: customer service tools and public perception tools (*NIMS Basic Guidance for PIOs* 2020). They allow to monitor social media accounts or specific keywords, hashtags, and mentions, with some providing real-time data visualization and sentiment analysis.

However, many of these require lots of manual decision-makings and are easily biased by the volume of activity. Resonance+ does not correlate with the volume of activity and prioritizes the messages that represent the shift of attention for people's cognition and perceptions. It could be integrated into the existing infrastructure the PIOs use and replace/complement the social listening tools mentioned above, where the recommended messages will provide more insights in terms of cognition and perception of risks or risk communications. This will lead to more effective risk communication or decision-making.

Resonance+ ranks messages based on the collective attention, and the results are aggregated and anonymized during the analysis to protect the privacy of individual users, consistent with guidelines on ethical social media research (Fiesler and Proferes 2018).

We evaluated this work in progress using LIWC — a lexicon-based tool that may suffer from sparsity. We will conduct two other evaluations for information representing people's cognitive and perceptual processes in future work. As the base-rates for word usage in each category is different, normalized LIWC Score will be computed to account for the differences in the messages recommended by resonance+ (and baseline) from the general word usage by the affected population. Secondly, we will leverage existing coding schemes around risk perceptions and response to the environmental hazards (Demuth et al. 2018) to qualitatively evaluate the results.

To make resonance+ even more useful for PIOs and crisis managers, the next step includes leveraging important ontologies in the crisis informatics literature (Olteanu, Castillo, et al. 2014; Alam, Ofli, et al. 2018; McCreadie, C. Buntain, et al. 2020; McCreadie, C. L. Buntain, et al. 2019; Alam, Sajjad, et al. 2021; Imran et al. 2016) to demonstrate its ability to recommend information supporting situational awareness, the first information need of PIOs. Preliminary results suggest that resonance+ performs similarly in categories such as affected individuals, displaced and evacuations, infrastructure and utilities damage, injured or dead people, missing and found people (Alam, Sajjad, et al. 2021). It under-prioritizes certain 'call-for-action' information types such as donations and volunteering, consistent with our LIWC-based observations. More effort could be made for the task of information supporting situational awareness in the future, in addition to the risk perceptions prioritized in this work. Finally, resonance+ needs to be evaluated against other baselines. And to ensure its applicability for real-time events, it also needs to be tested in real-time scenarios to show its usability for PIOs and crisis managers for future events.

REFERENCES

Alam, F., Ofli, F., and Imran, M. (June 2018). "CrisisMMD: Multimodal Twitter Datasets from Natural Disasters". In: *Proceedings of the 12th International AAAI Conference on Web and Social Media (ICWSM)*. USA.

Alam, F., Sajjad, H., Imran, M., and Ofli, F. (May 2021). "CrisisBench: Benchmarking Crisis-related Social Media Datasets for Humanitarian Information Processing". In: *Proceedings of the International AAAI Conference on Web and Social Media* 15.1, pp. 923–932.

Anderson, J., Kogan, M., Bica, M., Palen, L., Anderson, K., Stowe, K., Morss, R., Demuth, J., Lazrus, H., Wilhelmi, O., et al. (May 2016). "Far Far Away in Far Rockaway Far Far Away in Far Rockaway: Responses to Risks and Impacts during Hurricane Sandy through First-Person Social Media Narratives". In.

Barron, A., Huang, J., Spang, R., and DeDeo, S. (Oct. 2017). "Individuals, Institutions, and Innovation in the Debates of the French Revolution". In: *Proceedings of the National Academy of Sciences* 115.

Bruns, A. and Burgess, J. (2011). "The use of Twitter hashtags in the formation of ad hoc publics". In.

CrowdFlower (2015). *Disasters on Social Media*.

Crowe, A. (Jan. 2010). "The Elephant in the JIC: The Fundamental Flaw of Emergency Public Information within the NIMS Framework". In: *Journal of Homeland Security and Emergency Management* 7.

Crowe, A. S. (2012). "Disasters 2.0: The Application of Social Media Systems for Modern Emergency Management". In.

Demuth, J. L., Morss, R., Palen, L., Anderson, K. M., Anderson, J., Kogan, M., Stowe, K., Bica, M., Lazrus, H., Wilhelmi, O. V., et al. (2018). ""Sometimes da #beachlife ain't always da wave": Understanding People's Evolving Hurricane Risk Communication, Risk Assessments, and Responses Using Twitter Narratives". In: *Weather, Climate, and Society*.

Dynes, R. R. (1970). *Organized behavior in disaster*. Heath Lexington Books.

Fiesler, C. and Proferes, N. (2018). ""Participant" perceptions of Twitter research ethics". In: *Social Media+ Society* 4.1, p. 2056305118763366.

Fritz, C. E. and Mathewson, J. H. (1957). *Convergence behavior in disasters : a problem in social control.* Washington : National Academy of Sciences-National Research Councel.

Haddow, G. D. and Haddow, K. S. (2014). “Disaster Communications in a Changing Media World, Second Edition”. In.

Hagar, C. and Haythornthwaite, C. (June 2005). “Crisis, Farming amp; Community”. In: *The Journal of Community Informatics* 1.3.

Hilts, A. S., Mack, S., Eidson, M., Nguyen, T. Q., and Birkhead, G. S. (2015). “New York State Public Health System Response to Hurricane Sandy: An Analysis of Emergency Reports”. In: *Disaster Medicine and Public Health Preparedness* 10, pp. 308–313.

Houston, J. B., Hawthorne, J., Perreault, M. F., Park, E. H., Hode, M. G., Halliwell, M., McGowen, S. E. T., Davis, R., Vaid, S., McElderry, J. A., et al. (2015). “Social media and disasters: a functional framework for social media use in disaster planning, response, and research.” In: *Disasters* 39 1, pp. 1–22.

Hughes, A., Palen, L., Sutton, J., Liu, S., and Vieweg, S. (Apr. 2008). “Site-seeing in disaster: An examination of on-line social convergence”. In.

Hughes, A. L. and Tapia, A. H. (2015). In: *Journal of Homeland Security and Emergency Management* 12.3, pp. 679–706.

Imran, M., Mitra, P., and Castillo, C. (2016). *Twitter as a Lifeline: Human-annotated Twitter Corpora for NLP of Crisis-related Messages*. arXiv: 1605.05894 [cs.CL].

Kendra, J. and Wachtendorf, T. (Oct. 2003). “Reconsidering convergence and converger legitimacy in response to the World Trade Center disaster”. In: *Research in Social Problems and Public Policy* 11, pp. 97–122.

Kenter, T. and Rijke, M. de (2015). “Short Text Similarity with Word Embeddings”. In: CIKM ’15. Melbourne, Australia: Association for Computing Machinery, pp. 1411–1420.

Knox, C. C. (2022). “Local emergency management’s use of social media during disasters: case study of Hurricane Irma”. In: *Disasters*.

Kogan, M., Palen, L., and Anderson, K. M. (2015). “Think Local, Retweet Global: Retweeting by the Geographically-Vulnerable During Hurricane Sandy”. In: *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. CSCW ’15. Vancouver, BC, Canada: ACM, pp. 981–993.

Lachlan, K. A., Spence, P. R., Lin, X., and Greco, M. del (2014). “Screaming into the Wind: Examining the Volume and Content of Tweets Associated with Hurricane Sandy”. In: *Communication Studies* 65, pp. 500–518.

Leavitt, A. and Clark, J. (Apr. 2014). “Upvoting hurricane Sandy: event-based news production processes on a social news site”. In: *Conference on Human Factors in Computing Systems - Proceedings*.

Lehmann, J., Gonçalves, B., Ramasco, J. J., and Cattuto, C. (Nov. 2011). “Dynamical Classes of Collective Attention in Twitter”. In: *WWW’12 - Proceedings of the 21st Annual Conference on World Wide Web*.

Lin, X., Spence, P. R., Sellnow, T. L., and Lachlan, K. A. (Dec. 2016). “Crisis Communication, Learning and Responding”. In: *Comput. Hum. Behav.* 65.C, pp. 601–605.

Lindell, M. and Perry, R. (Jan. 2012). “The protective action decision model: theoretical modifications and additional evidence. In: Risk analysis, vol 32(4)”. In: *Risk Anal : Off Publ Soc Risk Anal* 32, pp. 616–632.

Liu, S., Palen, L., Sutton, J., Hughes, A., and Vieweg, S. (Jan. 2008). “In search of the bigger picture: The emergent role of on-line photo sharing in times of disaster”. In.

Martín, Y., Li, Z., and Cutter, S. (July 2017). “Leveraging Twitter to gauge evacuation compliance: Spatiotemporal analysis of Hurricane Matthew”. In: *PLoS ONE* 12.

Mazloom, R., Li, H., Caragea, D., Imran, M., and Caragea, C. (2018). “Classification of Twitter Disaster Data Using a Hybrid Feature-Instance Adaptation Approach”. In: *ISCRAM*.

McCreadie, R., Buntain, C., and Soboroff, I. (2020). *Incident Streams 2019: Actionable Insights and How to Find Them*.

McCreadie, R., Buntain, C. L., and Soboroff, I. (2019). “TREC Incident Streams: Finding Actionable Information on Social Media”. In: *ISCRAM*.

Mehta, A. M., Bruns, A., and Newton, J. (2017). “Trust, but verify: social media models for disaster management”. In: *Disasters* 41.3, pp. 549–565. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/disa.12218.

Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013). *Efficient Estimation of Word Representations in Vector Space*. arXiv: 1301.3781 [cs.CL].

Mikolov, T., Sutskever, I., Chen, K., Corrado, G., and Dean, J. (2013). *Distributed Representations of Words and Phrases and their Compositionality*.

Misra, S., Roberts, P., and Rhodes, M. (2020). "Information overload, stress, and emergency managerial thinking". In: *International Journal of Disaster Risk Reduction* 51, p. 101762.

Mitra, T., Wright, G., and Gilbert, E. (Dec. 2017). "Credibility and the Dynamics of Collective Attention". In: *Proc. ACM Hum.-Comput. Interact.* 1.CSCW, 80:1–80:17.

NIMS Basic Guidance for PIOs (Dec. 2020).

Olteanu, A., Castillo, C., Diaz, F., and Vieweg, S. (2014). "CrisisLex: A Lexicon for Collecting and Filtering Microblogged Communications in Crises". In: *Proceedings of the Eighth International Conference on Weblogs and Social Media, ICWSM 2014, Ann Arbor, Michigan, USA, June 1-4, 2014*. Ed. by E. Adar, P. Resnick, M. D. Choudhury, B. Hogan, and A. H. Oh. The AAAI Press.

Olteanu, A., Vieweg, S., and Castillo, C. (2015). "What to Expect When the Unexpected Happens: Social Media Communications Across Crises". In: *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work Social Computing*. CSCW '15. Vancouver, BC, Canada: Association for Computing Machinery, pp. 994–1009.

Palen, L. and Anderson, K. M. (2016). "Crisis informatics—New data for extraordinary times". In: *Science* 353.6296, pp. 224–225. eprint: https://science.sciencemag.org/content/353/6296/224.full.pdf.

Palen, L., Vieweg, S., Liu, S., and Hughes, A. (Oct. 2009). "Crisis in a Networked World". In: *Social Science Computer Review - SOC SCI COMPUT REV* 27, pp. 467–480.

Pennebaker, J. W., Booth, R. J., and Francis, M. E. (2007). "Linguistic Inquiry and Word Count (LIWC2007)". In.

Pennebaker, J. W., Boyd, R. L., Jordan, K., and Blackburn, K. G. (2015). "The Development and Psychometric Properties of LIWC2015". In.

Plotnick, L. and Hiltz, S. R. (2016). "Barriers to Use of Social Media by Emergency Managers". In: *Journal of Homeland Security and Emergency Management* 13, pp. 247–277.

Ratkiewicz, J., Fortunato, S., Flammini, A., Menczer, F., and Vespignani, A. (Oct. 2010). "Characterizing and Modeling the Dynamics of Online Popularity". In: *Phys. Rev. Lett.* 105 (15), p. 158701.

Receptiviti Inc. (2022). *receptiviti: Text Analysis Through the Receptiviti API*.

Reuter, C. and Kaufhold, M.-A. (2018). "Fifteen Years of Social Media in Emergencies: A Retrospective Review and Future Directions for Crisis Informatics". In: *POL: Other Change Management Strategy (Topic)*.

Rudra, K., Sharma, A., Ganguly, N., and Imran, M. (2017). "Classifying Information from Microblogs during Epidemics". In: *Proceedings of the 2017 International Conference on Digital Health*. DH '17. London, United Kingdom: Association for Computing Machinery, pp. 104–108.

San, Y., Clarence, S., Iii, W., Thorkildsen, Z., Giovachino, M., and Director, M. (2012). *Social Media in the Emergency Management Field 2012 Survey Results*.

Sasahara, K., Hirata, Y., Toyoda, M., Kitsuregawa, M., and Aihara, K. (Apr. 2013). "Quantifying Collective Attention from Tweet Stream". In: *PLOS ONE* 8.4, pp. 1–10.

Starbird, K. and Palen, L. (2011). ""Voluntweeters": Self-Organizing by Digital Volunteers in Times of Crisis". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '11. Vancouver, BC, Canada: Association for Computing Machinery, pp. 1071–1080.

Stewart, I., Yang, D., and Eisenstein, J. (2020). "Characterizing Collective Attention via Descriptor Context: A Case Study of Public Discussions of Crisis Events". In: *ICWSM*.

Stowe, K., Palmer, M., Anderson, J., Kogan, M., Palen, L., Anderson, K. M., Morss, R., Demuth, J., and Lazrus, H. (Aug. 2018). "Developing and Evaluating Annotation Procedures for Twitter Data during Hazard Events". In: *Proceedings of the Joint Workshop on Linguistic Annotation, Multiword Expressions and Constructions (LAW-MWE-CxG-2018)*. Santa Fe, New Mexico, USA: Association for Computational Linguistics, pp. 133–143.

Sutton, J., Palen, L., and Shklovski, I. (Apr. 2008). "Backchannels on the Front Lines: Emergent Uses of Social Media in the 2007 Southern California Wildfires". In: *Proceedings of the 5th International ISCRAM Conference*.

Tapia, A. H. and Moore, K. (Dec. 2014). "Good Enough is Good Enough: Overcoming Disaster Response Organizations' Slow Social Media Data Adoption". In: *Comput. Supported Coop. Work* 23.4–6, pp. 483–512.

Vieweg, S. (2012). "Situational Awareness in Mass Emergency: A Behavioral and Linguistic Analysis of Microblogged Communications". In.

Vieweg, S., Hughes, A. L., Starbird, K., and Palen, L. (2010). "Microblogging During Two Natural Hazards Events: What Twitter May Contribute to Situational Awareness". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '10. Atlanta, Georgia, USA: ACM, pp. 1079–1088.

Wu, F. and Huberman, B. A. (2007). "Novelty and collective attention". In: *Proceedings of the National Academy of Sciences* 104, pp. 17599–17601.

Wukich, C. and Mergel, I. (2015). "Closing the Citizen-Government Communication Gap: Content, Audience, and Network Analysis of Government Tweets". In: *Journal of Homeland Security and Emergency Management* 12, pp. 707–735.

Yeo, J., Knox, C., and Hu, Q. (Dec. 2020). "Disaster Recovery Communication in the Digital Era: Social Media and the 2016 Southern Louisiana Flood". In: *Risk Analysis*.

Zahra, K., Imran, M., and Ostermann, F. O. (2020). "Automatic identification of eyewitness messages on twitter during disasters". In: *Information Processing Management* 57.1, p. 102107.