DOI: 10.1111/mice.12576

ORIGINAL ARTICLE

COMPUTER-AIDED CIVIL AND INFRASTRUCTURE ENGINEERING WILEY

Semiautomated social media analytics for sensing societal impacts due to community disruptions during disasters

Cheng Zhang[1] | Wenlin Yao[2] | Yang Yang[2] | Ruihong Huang[2] | Ali Mostafavi[1]

[1]Zachry Department of Civil Engineering, Texas A&M University, College Station, TX, USA

[2]Department of Computer Science & Engineering, Texas A&M University, College Station, TX, USA

Correspondence
Ali Mostafavi, Zachry Department of Civil Engineering, Texas A&M University, 400 Bizzell St, College Station, TX 77843, USA.
Email: amostafavi@civil.tamu.edu

Funding information
National Science Foundation, Grant/Award Numbers: IIS-1759537, CMMI-1846069; Amazon Web Services (AWS) Machine Learning Award

Abstract

Understanding the societal impacts caused by community disruptions (e.g., power outages and road closures), particularly during the response stage, with timeliness and sufficient detail is an underexplored, yet important, consideration. It is critical for effective decision-making and coordination in disaster response and relief activities as well as post-disaster virtual reconnaissance activities. This study proposes a semiautomated social media analytics approach for social sensing of Disaster Impacts and Societal Considerations (SocialDISC). This approach addresses two limitations of existing social media analytics approaches: lacking adaptability to the need of different analyzers or different disasters and missing the information related to subjective feelings, emotions, and opinions of the people. SocialDISC labels and clusters social media posts in each disruption category to facilitate scanning by analyzers. Analyzers, in this paper, are persons who acquire social impact information from social media data (e.g., infrastructure management personnel, volunteers, researchers from academia, and some residents impacted by the disaster). Furthermore, SocialDISC enables analyzers to quickly parse topics and emotion signals of each subevent to assess the societal impacts caused by disruption events. To demonstrate the performance of SocialDISC, the authors proposed a case study based on Hurricane Harvey, one of the costliest disasters in U.S. history, and analyzed the disruptions and corresponding societal impacts in different aspects. The analysis result shows that Houstonians suffered greatly from flooded houses, lack of access to food and water, and power outages. SocialDISC can foster an understanding of the relationship between disruptions of infrastructures and societal impacts, expectations of the public when facing disasters, and infrastructure interdependency and cascading failures. SocialDISC's provision of timely information about the societal impacts of people may help disaster response decision-making.

1 | INTRODUCTION

Disaster-induced community disruptions include physical infrastructure and social disruptions, such as hazards, lack of supplies, business interruption losses, and evacuation orders (Batouli & Mostafavi, 2018; Nejat & Damnjanovic, 2012; Ng, Park, & Waller, 2010). The disruption events can cause significant impacts on residents' quality of life and well-being (i.e., societal impacts; Gardoni & Murphy, 2010), such as loss of life, injury, economic hardship, business interruption, and economic decline (Guo & Li, 2016; Lindell & Prater, 2004; Morss, Cuite, Demuth, Hallman, & Shwom, 2018; Toya & Skidmore, 2007; Uchida, Takahashi, & Kawahara, 2014; Ward & Shively, 2017). Sensing social impacts due to community disruptions when disaster strikes is an underexplored facet of the disaster response (NIST, 2016; Felts, Leh, & McElvaney, 2016; NIST, 2015; Othman, Beydoun, & Sugumaran, 2014). It is during the response stage of a disaster when community disruptions occur and cause the most severe impacts. Improving the understanding of societal impacts is pivotal for infrastructure management during disaster responses for three reasons. First, a majority of community impacts are due to physical infrastructure disruptions, such as road closures, power outages, and drainage overflows. Understanding the societal impacts caused by infrastructure disruptions would provide insights for enhanced resource allocation and prioritization of disaster response processes (NIST, 2016; Felts et al., 2016; NIST, 2015). Second, social disruptions, such as casualty, people in danger, and evacuation orders, are other significant causes of community disruptions and are highly related to physical infrastructure systems. Third, societal impacts due to infrastructure and social disruptions can jointly influence the long-term recovery and trajectories of long-lasting impacts for urban infrastructure and social systems (Elliott & Pais, 2006; Gardoni & Murphy, 2010; Murphy & Gardoni, 2007; Nomura et al., 2016). Hence, timely and detailed assessment of community disruption events during the disaster response stage can provide important insights for infrastructure managers and operators to proactively preventing short- and long-term societal impacts.

There is a dearth of studies on sensing societal impacts during disaster response stages in the literature (NIST, 2016; Felts et al., 2016; NIST, 2015). Existing studies on disaster societal impact rely mainly on statistical data (number of death, affected population; Gardoni & Murphy, 2010; Lindell & Prater, 2004; Murphy & Gardoni, 2007; Toya & Skidmore, 2007; Ward & Shively, 2017) and surveys focusing on mid- to long-term (≥3 months) disaster impacts (Elliott & Pais, 2006; Guo & Li, 2016; Morss et al., 2018; Nomura et al., 2016; Uchida et al., 2014). Such data lack the level of detail and timeliness to reflect the societal impacts during the response stage, due mainly to the absence of appropriate techniques to capture and analyze the time-sensitive societal impact information in the disaster response stage. Residents' memory may fade after the disruption passes, especially for short- and medium-term disruptions, which limits the effectiveness of using standard survey instruments after the disaster. On the other hand, designing, deploying, and analyzing surveys is a lengthy process that makes it nearly impossible to evaluate societal impacts in a timely enough manner to inform disaster response and relief efforts. This limitation can be addressed using social media analytics, which refers to the use of user-generated social media data for analyzing a particular event or phenomenon.

Publicly available social media contents contain rich information about the descriptions of and people's reactions to disruption events, which could support a timely assessment of societal impacts (Dutt, Basu, Ghosh, & Ghosh, 2019; Fan & Mostafavi, 2019; Yang et al., 2019). However, a practical social media analytics approach for analyzing social impacts is still lacking (Hiltz, Kushma, & Plotnick, 2014; Reuter, Ludwig, Kaufhold, & Spielhofer, 2016). According to a survey with 761 emergency service staff in Europe (Reuter et al., 2016), more than 60% of participants thought that "general situational updates" and "information about the public mood" is "useful" or "very useful." However, only 23% stated that they "often" or "sometimes" searched social media sites to gain situational awareness. Furthermore, 68% thought that the software tools to access social media data needed to be improved. Social media analytics for disaster response are many in existing studies (Laylavi, Rajabifard, & Kalantari, 2017; Reuter & Kaufhold, 2018; Atefeh & Khreich, 2015; Yuan, Liu, & Wu, 2019). Current social media analytics tools, however, have limited capability to support the analyzers to process both quantitative and qualitative data analysis promptly, including filtering posts according to different needs, generating statistical data, and quickly scanning social media posts (Hiltz et al., 2014; Zhang, Fan, Yao, Hu, & Mostafavi, 2019). Analyzers in this paper refer to, for instance, infrastructure managers, disaster responders, and academic researchers.

The ineffectiveness of existing social sensing approaches for disaster response is due to two limitations. First, domain and event adaptation capability, which refers to adapting one approach developed according to one disaster to a different disaster, is missing for most social media analytics approaches (Zhang et al., 2019). This adaptation is crucial for practical societal impact analysis approaches because different analyzers (or organizations) have different analysis objectives for different disasters. Second, the existing social media analytics approaches focus more on retrieving facts (e.g., number of casualties; X. Chen, Elmes, Ye, & Chang, 2016; Rudra et al., 2016; Xu et al., 2016, 2017) and miss the information related to subjective feelings, emotions, and opinions of the people. Such information reflects the experience and hardship of residents facing disruptions, and, hence, could indicate the nature and extent of societal impacts.

This study proposes a social media analytics system named social sensing of Disaster Impacts and Societal Considerations (SocialDISC) to address the two abovementioned limitations. SocialDISC can adaptively identify and aggregate social media clusters as subevents discussing different categories of impacts according to taxonomies that are defined by the users. For each subevent, SocialDISC further distills the content into salient information to reduce the number of

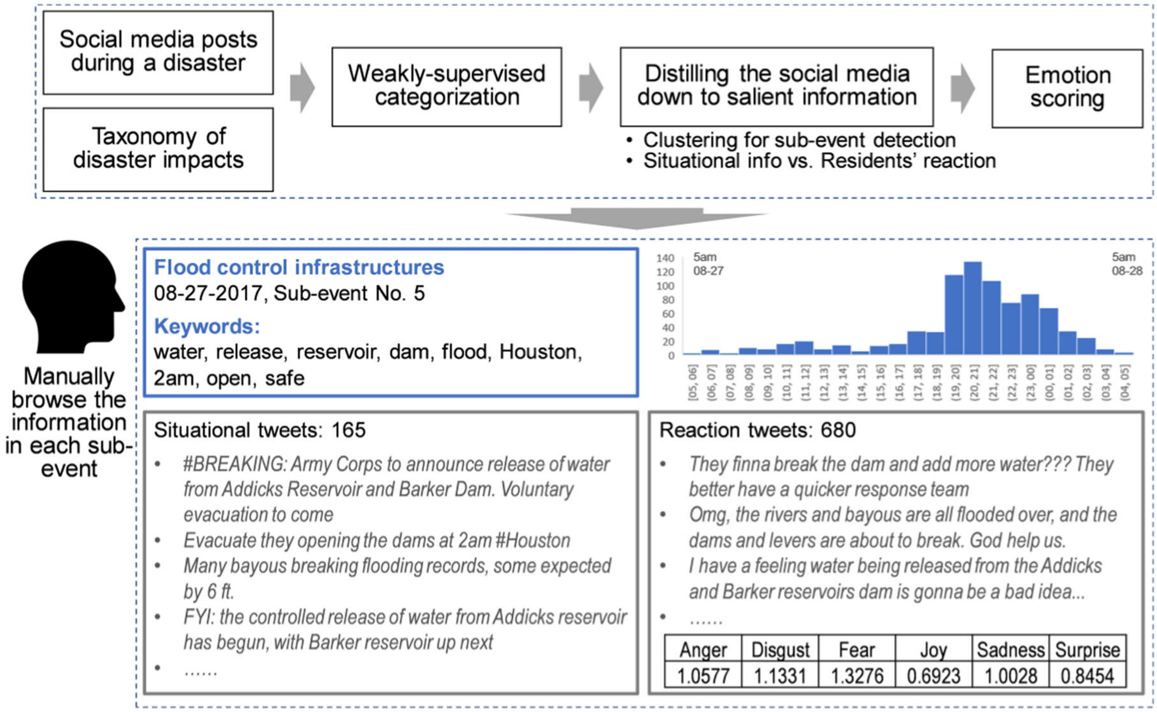

FIGURE 1 Overview of SocialDISC (social sensing of Disaster Impacts and Societal Considerations) framework and end-user product

posts for analyzers to scan. SocialDISC also aggregates the emotional signals in social media posts. The acquired emotion signal enables the analyzers to observe the patterns of emotion change in different topics for the unfolding of the disaster and its impacts. Associating the evolution of emotional signals with disruptions can provide unique insights regarding the societal impacts in SocialDISC. This study validates the performance of SocialDISC in terms of comprehensiveness, level of detail, and efficiency through a case study analyzing the societal impacts of Houstonian people when responding to one of the costliest disasters, Hurricane Harvey.

The rest of this paper is organized as follows: Section 2 describes the SocialDISC methodology. Section 3 presents the results of this approach and its application in the context of Hurricane Harvey. Section 4 concludes this paper by discussing the findings, contributions, and future works.

2 | ARCHITECTURE AND METHODOLOGY

Social media data often consist of fragmented sentences written in a casual language; retrieving salient information from this unconventional manner of expression is a challenge. As mentioned earlier, an analyzer needs to process both quantitative and qualitative analysis for social media analytics. Quantitative analysis requires the use of navigational metadata, which includes emotion tags, keywords, and metrics describing the societal impacts (Hiltz et al., 2014). For example, U.S. emergency managers believe that twittering filtering techniques by needs, subevents, or customized categorization is a popular possible technological enhancement (Hiltz et al., 2014). The process of generating such metadata should be adaptive to different social media datasets from different disaster contexts. On the other hand, to perform qualitative analysis, analyzers should be able to browse the vast amount of social media posts effectively as useful information may hide in the content of social media posts and evade capture by the metadata. Therefore, a practical approach for retrieving and aggregating information about disaster disruption and corresponding societal impacts should have the following features: (a) filtering out irrelevant information, (b) sorting information according to subevent, need, or customized categorization, (c) reflecting the societal reaction to disaster impacts, (d) providing concise results so that people can read easily, and (e) processing data in a timely manner.

Figure 1 depicts an overview of the SocialDISC approach proposed in this study. The upper part of Figure 1 visualizes the framework of the proposed method of social media analytics for retrieving disruptions of critical infrastructures and societal impacts. SocialDISC labels disaster-related social media posts according to predefined taxonomy (Table 1) indicating different types of community disruptions and filters out irrelevant posts (Yao, Zhang, Saravanan, Huang, & Mostafavi, 2020). Within each category in the taxonomy, a content distilling process enables the analyzers to determine the topic of each subevent efficiently by scanning the situational information posts and the keywords with high frequency. In the

TABLE 1 The taxonomy of community disruption events

Category name (ID)	Description	Initial keywords
Built-environment hazards (HAZ)	The damage or risks that may cause injury or death related to the built environment, such as fire, explosion, contamination, electric shock, collapse, and so on.	Fire, explosion, collapse, poison, electrocute
Business, work, school (BWS)	Disaster's impact on businesses, works, and schools, for example, business closed/open, school closed/open, and so on.	Office, school, close, open, work
Casualty (CAS)	Disaster-caused deaths and injuries.	Die, drown, injure, hurt
Flood control infrastructures (FCI)	The impact on or damage to the reservoir, bayou, canal, dam, and so on.	Reservoir, bayou, canal, dam, levee
Help and rescue (RES)	People provide, receive, or seek face-to-face help in disastrous environments, including indirect help such as donating money, supply, and providing services.	Rescue, boat, help, red-cross, donate, guard
Housing (HOU)	Reporting damages of a house, apartment, home, and so on.	House, home, room, apartment
Preventative measure (PRE)	The processes that people do to avoid hurricane damage, including evacuation, sheltering, and so on.	Evacuate, shelter, refugee
Transportation (TRA)	The impact on the traffic, bus services, or the closure of a road, airport, highway, and so on.	Plane, flight, airport, highway, freeway, road, avenue
Utilities and supplies (UTI)	Impacts on, gas, water, power, communication facility, food, grocery stores, and so on.	Power, electricity, gas, store, food, supply

content distilling process, a network-based clustering algorithm first clusters social media posts in each category into clusters according to content similarity, enabling users to read similar social media posts more efficiently. Within each information cluster, a pretrained classifier (Yao et al., 2020) then separates social media posts into situational information (posts describing the situation using a formal and objective language) or residents' reaction (posts expressing personal feelings, comments, and complaints about the situation using a casual and subjective language). Finally, SocialDISC will assess emotion scores of the societal reaction, a quantitative indicator of the societal impact caused by the disruption event for each subevent. The lower part of Figure 1 shows the conceptual user interface of SocialDISC with data processing outcomes, which enables the analyzer to quickly scan social media posts based on subevents in different categories.

2.1 | Taxonomy of tweets for analyzing disaster impacts

We developed a standard event taxonomy describing community disruption event categories and employed it in analyzing Twitter content during Hurricane Harvey (see Subsection 3.1 for details about the dataset). This taxonomy consists of nine categories of keywords related to human activities and built-environment disruptions, inspired by previous studies (Huang & Xiao, 2015; Sutton et al., 2015). Table 1 summarizes the event taxonomy and corresponding keywords. SocialDISC classifies tweets according to the categories of this taxonomy in the next step.

2.2 | Labeling of social media posts according to the taxonomy

The first step of retrieving societal impact information is to identify social media posts that belong to specific community disruption categories, as is described in the taxonomy. Identifying posts related to different types of disruptions and impacts is a fine-grained classification task, which is much more challenging to achieve than the existing classification approaches, which classify posts as to disaster-relatedness. The mainstream of fine-grained approaches is based on supervised learning from hundreds of thousands of labeled posts (Burel, Saif, & Alani, 2017; Nguyen, Yang, Li, Cao, & Jin, 2018; Zhang et al., 2019). Due to the uniqueness of each disaster setting and community, achieving a comprehensive training dataset is not feasible. Also, acquiring and properly labeling such training data for each is resource-intensive, especially during the disaster response stage. Furthermore, domain adaptation of supervised classification is challenging overall, which exacerbates the limitation of supervised learning caused by the paucity of training data. To increase the adaptivity and timeliness, SocialDISC uses a weakly supervised approach to label social media posts (Yao et al., 2020).

This labeling approach mainly consists of three main phases (Figure 2). In phase one, we applied automatic clustering of social media posts (i.e., tweets in this study) containing preliminary keywords for each event category (based on the taxonomy in Table 1) and then conducted manual word sense disambiguation (WSD) on the clusters to generate training instances for the classifier. In phase two, we trained a

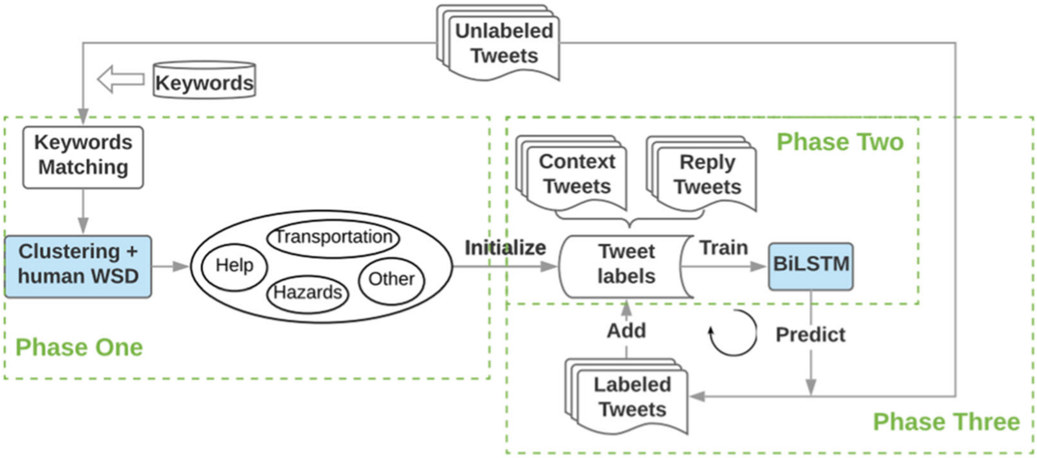

FIGURE 2 Overview of the weakly supervised labeling of social media posts (Yao et al., 2020)

multichannel Bidirectional Long Short-Term Memory (BiLSTM) classifier (an artificial recurrent neural network architecture used in the field of deep learning; Graves & Schmidhuber, 2005) using tweets together with their context tweets and reply tweets. We chose the BiLSTM model because it is proved to be the state-of-the-art in classifying natural language data and capturing the semantics of sentences (T. Chen, Xu, He, & Wang, 2017; Peters et al., 2018; Zhou et al., 2016). In phase three, we iteratively retrain the multichannel classifier to improve its performance of event recognition. To see more technical details about the weakly supervised labeling algorithm, specifically for the presentation of the WSD algorithm and the BiLSTM classifier, please refer to another work from the authors (Yao et al., 2020).

2.2.1 | Phase one: Clustering-assisted manual WSD

For each event category, we first retrieved posts containing event keywords and then applied a clustering algorithm to form post clusters. The clustering algorithm is based on the Speaker–Listener Label Propagation Algorithm (SLPA; Xie, Szymanski, & Liu, 2011), which is a clustering algorithm that allows a node to be clustered into multiple overlapping clusters. SLPA is proven to be one of the best algorithms for detecting overlapping communities (Xie, Kelley, & Szymanski, 2013). Given a set of tweets, we constructed an undirected graph $G(V; E)$, where V represents all tweets and E represents weighted edges between nodes. The weight of an edge between two posts is calculated based on the content similarity of the two tweets. The similarity score between post u and v is the number of common important words/(length of $u \times$ length of v). Empirically, we found this similarity measure—considering only important words—performs better than the straightforward cosine similarity measure using all words.

The word importance measure—an important function of the SLPA clustering algorithm—is derived from an approach for learning universal sentence representations (Conneau, Kiela, Schwenk, Barrault, & Bordes, 2017). Specifically, for a given post with T words $\{w_t\}$, we compute T hidden vectors with D dimensions for each word in the post ($\{h_t \in \boldsymbol{R}^D\}$, where $t = 1,2,\ldots,T$) by applying the sentence encoder in the universal sentence representations (Conneau et al., 2017). A word w_t is more important if the elements in the vector h_t are larger than the counterpart of other words in the post. Therefore, for a specific word w_t, we define the importance score of a word I_{w_t} is the number of dimensions the element of which in h_t is the maximum among all the T words in the post divided by D. We determine a word as important when its importance scores above the average of that of all words in the dataset. In the following tweet, for example, the algorithm will label the bold words as important: "**Houston woke** up to **catastrophic flooding**, with five people **dead**."

To facilitate manual WSD after the clustering, we ranked tweet clusters based on the number of tweets. A domain expert then manually judges whether each cluster (from largest to smallest) captures the pertinent meaning of an event keyword based on five randomly sampled posts (visualized in Figure 3). The domain expert stops scrutiny once 20 pertinent clusters are identified for each event category. Note that if a cluster uses a word sense that is irrelevant to any event category, it assigned the label "Other" to it, and we later used the posts in such clusters as negative training instances.

2.2.2 | Phase two: Multichannel tweet classification

To incorporate the nature that social media contents are highly correlated with the context in terms of time and the connected user network, we applied a feedforward neural network to directly map the enriched representation of a

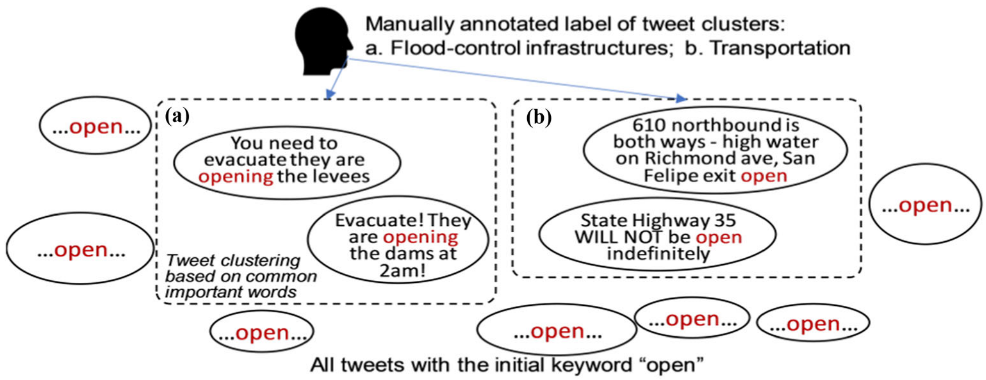

FIGURE 3 Visualization of the clustering-assisted manual word sense disambiguation. Example keyword: open

targeted post to ten classes in taxonomy (nine event categories + other). Here, the enriched representation is the concatenation of the target post, a weighted average of context posts, and the simple average of replies; we retrieved at most five recent posts immediately before the target post as context posts. Meanwhile, we observed that topic relatedness between the target post and context posts decreases exponentially with time. To reflect such observation, we applied the function $w_i = 0.8^{mi}$ to assign weights for each context post, where mi is the time interval (in minutes) between ith context post (repost or reply) and the target post. To incorporate additional evidence for classifying the target post, we applied three separate BiLSTM encoders to obtain sentence embedding for the target posts, context posts, and reply posts.

In training, we use negative training instances to compete with positive training instances (tweets that have event labels for a particular category c) for improving precision. The negative training instances of a category c consist of posts labeled as $c' \neq c$ as well as randomly sampled posts labeled as "other." The number of posts labeled as "other" included in the negative training instance equals to the number of posts labeled as all nine categories to reflect the fact that there are generally more posts reporting no event.

2.2.3 | Phase three: Improve coverage with bootstrapping

In phase one, by exploiting SLPA clustering with quick human-involved WSD, we collected hundreds of thousands of labeled posts; however, each event category may include some cases not matched by keywords. Therefore, we apply bootstrapping to iteratively retrain the context-and-reply-enhanced classifier obtained from phase two for further exploring semantic space. Specifically, we applied the initial multichannel classifier to the unlabeled posts to identify new posts in each event category. Newly labeled posts, together with their context posts and replies, are used to retrain the model. The bootstrapping process repeats until no more new posts can be discovered.

2.2.4 | Validation of the adaptive and weakly supervised labeling algorithm

For evaluation purposes, we asked two human annotators to exhaustively annotate all 11,782 tweets posted by Twitter users whose profile location is "Houston, TX" from August 28, 1:00 to 2:00 p.m., for Hurricane Harvey, and then evaluated our system and baselines on annotated data. The detailed data collection description is in Subsection 3.1. To compare the performance of the proposed weakly supervised approach, we compared the result of the proposed algorithm with the following baselines. First, we trained a standard BiLSTM sentence classifier in 10-fold cross-validation using human-annotated tweets, which is a solid baseline. Specifically, we use 70% annotated tweets as training data and 30% as testing data because, based on our experiments, further increase of the percentage of training instance will not significantly increase the performance. The third baseline system is keywords matching, which simply labels each tweet into an event category if it contains any keywords in that category.

Table 2 shows the experimental results for validating our weakly supervised event detection approach. We present an F1 score for each event category and macro average Precision/Recall/F1 score across all event categories in the last column. The final macro average F1 is 65.5%, which is 3.6% higher than the supervised classifier when we used 90% and 10% of annotated data in training and testing, respectively. Overall, the three-phase approach can significantly improve the precision of keyword matching by 19% and recall by 9%.

2.3 | Distilling the social media down to salient information

Distilling the content down to salient information is necessary and challenging. The distilling process refers to detecting the subevent in each category and identifying the rare posts delivering salient situational information in each subevent so that the analyzers can browse efficiently. During disasters, it is impossible for response personnel to read numerous social media posts line by line, although useful information hides in

TABLE 2 Experimental results for validating our weakly supervised event detection approach. Bold values show the highest score among the proposed and baseline methods for each column

	PRE	RES	CAS	HOU	UTI	TRA	FCI	BWS	HAZ	Overall		
	F1									Precision	Recall	F1
Keywords	73.9	56.6	26.2	36.4	54.3	38	54.4	55.5	43.1	51.1	52.5	51.8
Supervised learning	80.8	**72.2**	**48**	45.3	56.3	**67.9**	65.9	**71.2**	45.3	**73.2**	53.6	61.9
Proposed method	**83.9**	67.8	36.7	**45.7**	**66.1**	61.3	**74.8**	69.1	**57.7**	70.1	**61.6**	**65.5**

the noisy, repetitive, and casual posts (X. Chen et al., 2016; Rudra et al., 2016; Xu et al., 2016; Xi et al., 2017). The distilling process, therefore, greatly reduces the labor for an analyzer to browse the social media posts for understanding the community disruptions.

SocialDISC uses a clustering algorithm to detect the subevents in each category after we labeled each post according to the taxonomy. We observe that tweets adopting one common stance of an event keyword often share content words and can be easily grouped, which supports the weakly supervised labeling method introduced above. Such observation is consistent with previous research on unsupervised WSD (Navigli & Lapata, 2010). Therefore, the SLPA algorithm (introduced in Subsection 2.2) can cluster tweets in each category into clusters, with each tweet cluster consisting of tweets with content similarity becoming a subevent in that category. As is shown in Figure 1, this method yields outputs related to the top keywords based on the word frequency in each subevent to help the analyzer identify the topics of each subevent quickly.

The next step of the information distilling process is identifying the situational posts. We observed that two types of posts exist in each subevent (i.e., cluster in each category): situational posts and nonsituational posts. Situational posts contain information indicating an awareness of the scope of the disaster as well as specific details about the situation (such as numbers). Situational posts are often written in formal language with fewer personal pronouns, exclamations, and intensifiers compared with nonsituational tweets (Rudra, Ghosh, Ganguly, Goyal, & Ghosh, 2015). In each subevent, the situational posts are relatively rare (account for 20% of all social media posts according to the case study) and with content similarity, which can be easily browsed by the analyzers. In this study, we use the pretrained Support Vector Machine (SVM) classifier to identify situational posts and nonsituational posts (Rudra et al., 2015). This classifier uses the following features for classification: the proportion of subjective words, count of personal pronouns, count of numerals, count of exclamations, count of question marks, count of modal verbs, count of *wh*-words, count of intensifiers, count of nonsituational words.

2.4 | Emotion scoring

Existing studies usually highlight the value of situational posts (discussed in the previous subsection). Nonsituational tweets, however, contain significant information about emotion signals than can inform societal impacts. The information related to subjective feelings, comments, and opinions of people, which contain rich information reflecting the impacts caused by community disruptions, is often ignored. The emotion signals in social media posts reflect the experience and hardship of residents facing disruptions, and, hence, could indicate the nature and extent of societal impacts. Therefore, this study considers nonsituational posts as residents' reaction posts, and SocialDISC detects the emotion signal within the residents' reaction posts through an emotion scoring approach.

TABLE 3 Modifications made to the original lexicon for disaster scenarios

Word	Anger	Disgust	Fear	Joy
Food				Removed
Storm	Removed			
Volunteer			Removed	Removed
Child				Removed
Found				Removed
F**k	Add	Add		

To analyze emotion signals in the residents' reaction posts, we processed emotion analysis using emotional lexicon collected and curated by the National Research Council of Canada (NRC; Mohammad & Turney, 2013), which is well-accepted and widely employed in the literature (Vosoughi, Roy, & Aral, 2018). This lexicon contains a comprehensive list of 141,820 English words and their associations with eight emotions: anger, fear, anticipation, trust, surprise, sadness, joy, and disgust. In our study, we focused only on the six basic emotions—anger, fear, surprise, sadness, joy, and disgust—proposed by Ekman (1992) because this framework focuses more on negative emotions, as expected in the case of community disruptions. Only joy was considered as a positive emotion to capture emotion signals related to the restoration of disruptions and return to normal situations. Besides, we modified the lexicon according to the disaster scenario. For example, we removed the emotion label "Joy" of the word "food" in the original lexicon labels because many people complained about the lack of food during Harvey. In addition, we added a swear word to the lexicon that was frequently used in social media posts to reflect the emotion of people properly. Table 3 summarizes the modification made to the original lexicon (Mohammad & Turney, 2013).

To illustrate the process of calculating the emotion score for each tweet, we used anger as an example without losing the generality. We first calculate the rate of words associated with anger in tweet i (r_{anger}).

$$r_{\text{anger},\,i} = \frac{\text{Number of words associated with anger in tweet } i}{\text{Number of words in the tweet}}.$$

As we care more about comparing the changing pattern of each emotion over time and the difference in emotions among different categories of community disruption events, we calculated the relative emotion score for each tweet. Still using anger as an example, we calculated the anger score of the post i as follows:

$$\text{Angerscore}_i = \frac{r_{\text{anger},\,i}}{\text{Average of } r_{\text{anger}} \text{ of all tweets related to Harvey}}.$$

Using the same approach, we can replace anger with the other five emotions and calculate their emotion score.

2.5 | Manual scanning of subevents

After processing all the collected tweets using SocialDISC, an analyzer can scan the information of each subevent detected in each category. The lower half of Figure 1 visualizes an example interaction paradigm between an analyzer and a detected subevent. Based on the experience of the authors, it is an efficient and effective practice for an analyzer to browse the top 10 largest subevents in each large category, which covers about 50–90% tweets of that category, to acquire situational awareness information. The small subevents were ignored during the scanning process in the analysis due to their lack of representativeness. The analyzers examined the keywords and the situational posts to understand the content of each subevent. The analyzer can also look at the emotion scores of the residents' reaction tweets to understand the severity of societal impacts caused by the community disruptions in each category.

3 | CASE STUDY: ANALYZING THE SOCIETAL IMPACT OF HURRICANE HARVEY IN THE RESPONSE STAGE

This section shows a case study of analyzing the societal impact of Hurricane Harvey in the response stage. The purpose of this case study is twofold: (a) to verify whether SocialDISC enables analyzers to examine social media posts in a major disaster in a timely manner; and (b) to show the analysis result to highlight the practical and theoretical implications of SocialDISC. Specifically, this section aims at testing the performance of SocialDISC by exploring three aspects: (a) validating the comprehensiveness and timeliness of event detection using SocialDISC (Subsection 3.2); (b) verifying its performance of assessing social impacts by summarizing the societal impacts information acquired from social media posts during a major disaster (Subsection 3.3); and (c) verifying its efficiency by analyzing the time and labor cost for an analyzer to browse the tweets posted by Houstonians during Harvey for (Subsection 3.4).

3.1 | Data collection

To demonstrate the application of the SocialDISC framework, we used a tweet dataset from Hurricane Harvey as the testbed. Hurricane Harvey, which struck the Houston metropolitan area and Southeast Texas from August 17, 2017 to September 2, 2017, ranks as the second-costliest hurricane ($125 billion in damage) on record for the United States (National Hurricane Center, 2018). We chose August 25–30 for the time most impacted by Hurricane Harvey. In this study, we choose 5:00 a.m. instead of midnight as the boundary between two consecutive days (i.e., August 27 refers to the period between 5:00 a.m., August 27 and 5:00 a.m., August 28); this time was chosen because the fewest tweets were posted during the early morning. To retrieve tweets in affected areas through Twitter PowerTrack API, we consider two constraints in Twitter crawling: tweets with geo-location within a bounding box of Houston, or tweets by authors with profile location in affected areas. In total, we collected 4,714,105 tweets, including original posts, retweets, and replies. As we aim to recognize tweet messages reporting events for infrastructure management during disasters, among three tweet types (i.e., original post, retweet, and reply) in our data collection, we only consider 1,121,363 original posts as target tweets to be classified across all the experiments. Table 4 shows the number of tweets collected on each day during Harvey landing.

3.2 | Validating the event detection performance of SocialDISC

To validate the credibility of the identified societal impact information through SocialDISC, we compared the detected events with the events listed on Wikipedia. Specifically, we first asked one of the authors to play the role of an analyzer who needs to acquire situational information during Hurricane Harvey. The analyzer used SocialDISC to browse the important subevents (i.e., top 10 largest tweet clusters) in each category for each day from August 25 to 30, 2017 and then summarized the topic of each subevent. Next, we extracted all the events shown in the section of the "Houston metropolitan area flooding" on the Wikipedia page of Hurricane Harvey

TABLE 4 Number of tweets collected on each day during Harvey landing

	8/25	8/26	8/27	8/28	8/29	8/30
Number of tweets	518,250	636,684	901,972	986,364	926,359	744,476
Original posts	160,453	158,808	220,010	212,938	198,718	170,436

TABLE 5 The comparison between event detection based on SocialDISC and the events listed on Wikipedia page. Bold dates indicate early detections while the underlined date indicates a late detection

Events on Wikipedia	Category	Posted date by media	Date detected on SocialDISC	Daily ranking	Example tweet
More than 800 Houston area flights were canceled	TRA	August 27	August 27	3	All incoming and outgoing Houston flights canceled through Wednesday. All schools and classes canceled through…
Several tornadoes were spawned in the area	HOU	August 26	August 26	2	Hurricane Harvey: Tornado Damages Home In Cypress
A tornado damaged or destroyed the roofs of dozens of homes in Sienna Plantation	HOU	August 26	August 26	4	Tornado damages 50 homes in Sienna Plantation #HurricaneHarvey
Six from the same family who died when their van was swept off a flooded bridge	CAS	August 29	<u>August 30</u>	2	#UPDATE: Family of six presumed dead after van sinks in Harvey floods
A police officer drowned while trying to escape rising waters	CAS	August 29	August 29	1	BREAKING: Houston police officer dies in #Harvey flood emergency
A mandatory evacuation was issued for all of Bay City on August 27	PRE	August 28	**August 27**	1	Bay City under mandatory evacuation by 1:00 pm tomorrow	10-ft flood expected
Evacuations took place in Conroe on August 28 following release of water from the Lake Conroe dam	PRE	August 29	**August 27**	1	New details on Conroe and Missouri City mandatory evacuations
A levee along Columbia Lakes in Brazoria County was breached, prompting officials to urgently request for everyone in the area to evacuate	FCI	August 29	August 29	2	HAPPENING NOW: Brazos River breaches levee at Columbia Lakes in Brazoria County. Residents urged to evacuate immediately! #ABC13 #Hounews
The U.S. Army Corps of Engineers began controlled water releases from Addicks and Barker Reservoirs in the Buffalo Bayou watershed	FCI	August 28	**August 27**	2	The Army Corp of Engineers will manage this release. Addicks Reservoir will release first. Barker will release a day later.
Many people began evacuating the reservoir release area, fearing a levee breach	FCI	August 29	**August 28**	1	MANDATORY EVACUATIONS for Inverness Forest Subdivision bc of imminent levee breach! Cypress Ck 11 ft over bank.
Addicks Reservoir reached capacity on the morning of August 29 and began spilling out	FCI	August 29	**August 28**	2	Look at the map! Barker Reservoir at capacity so they are slowly releasing water. Up to 5 miles west of reservoir expected to get flooded.

and listed them in Table 5. Finally, we examined whether each important event listed on Wikipedia is successfully detected by an analyzer using. The rationale for this validation process is that a practical social sensing tool should be able to help people extract the important events during the unfolding of a disaster. Also, the authors assume that the events mentioned on Wikipedia pages are important ones in the timeline of Hurricane Harvey.

Table 5 shows the comparison between event detection based on SocialDISC and the events listed on the Wikipedia page. SocialDISC captured all the events listed on the Wikipedia page. This result indicates that SocialDISC is a reliable tool to capture situational information during disasters. Second, SocialDISC can identify many events before documentation by news media, the example of which includes but is not limited to “A mandatory evacuation was issued for

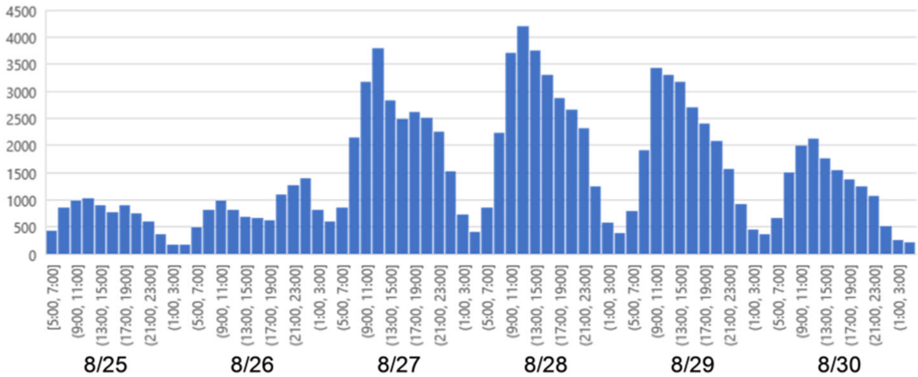

FIGURE 4 The histogram of the number of tweets posted in every 2 hr from 5:00 a.m., August 25 to 5:00 a.m., August 31

all of Bay City on August 27"; "many people began evacuating the reservoir release area, fearing a levee breach"; and "Addicks Reservoir reached capacity on the morning of August 29 and began spilling out." On August 27, for example, people were heavily discussing the release of Addicks and Barker reservoirs as soon as this information was released by the Army, while this information is posted by media the second day. This early detection feature of SocialDISC makes it a practical tool during disasters for acquiring situational awareness regarding the impacts of infrastructure disruptions. The SocialDISC users can acquire the information related to significant events without digging into the small subevents in each category. Overall, the comparison between the events detected using SocialDISC and the ones documented on Wikipedia page validates that SocialDISC can help the analyzers to acquire disaster situational information of important events in a comprehensive and time-efficient manner.

3.3 | Societal impact analysis results for important categories

This section will summarize the societal impacts of Houstonians during Hurricane Harvey to verify the performance of SocialDISC in providing useful insights about the societal impacts of the residents due to disruption events. We start by reviewing the number of tweets related to the taxonomy categories discussed in Subsection 2.1. Figure 4 shows the histogram of the number of hurricane-related tweets posted in nine categories every 2 hr from 5:00 a.m., August 25 to 5:00 a.m., August 31. This figure shows that the number of hurricane-related tweets increased dramatically on August 27 and decreased gradually after August 29.

Figure 5 depicts the societal impact analysis result of Hurricane Harvey for each disruption event category. This figure uses red, green, purple, yellow, blue, and orange with different transparency to visualize the emotion scores of anger, disgust, fear, joy, sadness, and surprise for each category in each day. To be specific, the emotion score of each cell refers to the average emotion score of all tweets in a certain category in a certain day. Figure 5a shows that the scores of the negative emotions (i.e., anger, disgust, fear, and sadness) were decreasing gradually during the landing time of Harvey. The following subsections will discuss the detailed analysis result of social impacts in each category. Figure 5b–j shows the amount and the emotion score of the tweets in each day for each category, which shows that different patterns exist in different categories. Considering both the number of tweets and the emotion score, we found that the Houstonians suffered greatly and broadly from the impacts on *Utility and supplies*, *Housing*, and *Preventative measure*. On the other hand, the tweets in the categories of *Casualty* and *Built-environment hazards* contained extreme emotion but with a tiny size. This phenomenon means that the tragedies related to life loss and built-environment hazards were not the primary concern of the general Houstonians during Harvey. Considering the constraints on the length of the paper, Subsections 3.3.1 to 3.3.3 will introduce the findings on societal impacts due to Harvey in three categories in detail: utilities and supplies, housing, and preventative measure. The appendix shows the societal impacts due to Harvey in other categories (see supporting information).

3.3.1 | Utilities and supplies

Overall, the tweets in the category of utilities and supplies contained three subtopics: food and water, electricity, and gas. The *Utility and supplies* section of Figure 5b showed strong negative emotions (i.e., anger, disgust, fear, and sadness) August 25–26, and the negative emotion gradually decreased during the next 4 days of Harvey. This decrease is attributed to the fact that the primary source of negative emotion, which was "stores running out of food and water" and "power outage," dominant in the first 2–3 days, gradually diminished.

For the subcategory related to "purchasing food and water," the complaints on Twitter focused on the following topics: "food sold out," "long line for grocery stores," "grocery store

(a) All categories

Date	Anger	Disgust	Fear	Joy	Sadness	Surprise	Number
8/25	1.16	1.05	1.20	0.93	1.01	0.93	8152
8/26	1.36	1.25	1.19	0.90	1.09	0.94	10431
8/27	1.06	1.07	1.07	1.07	1.08	1.06	28721
8/28	1.03	1.03	1.09	0.98	1.09	1.00	28239
8/29	0.85	0.89	0.85	0.98	0.89	0.98	23307
8/30	0.77	0.80	0.70	1.05	0.78	0.99	14427

(b) Utilities & supplies

Date	Anger	Disgust	Fear	Joy	Sadness	Surprise	Number
8/25	1.49	1.26	1.14	0.90	1.13	1.02	3785
8/26	1.59	1.53	1.20	1.04	1.37	1.12	1582
8/27	1.16	1.24	0.80	1.07	1.05	1.07	2659
8/28	1.04	1.08	0.70	0.94	1.07	0.86	2964
8/29	0.93	1.10	0.68	0.87	0.88	0.88	3639
8/30	0.96	1.05	0.64	0.84	0.90	0.81	2457

(c) Housing

Date	Anger	Disgust	Fear	Joy	Sadness	Surprise	Number
8/25	0.99	1.79	0.22	1.08	0.72	1.86	23
8/26	1.66	1.45	1.29	0.75	1.18	0.88	3235
8/27	1.41	1.35	1.03	0.98	1.52	1.02	6410
8/28	1.59	1.52	1.11	1.09	1.86	1.11	4903
8/29	1.32	1.56	0.90	1.20	1.43	1.33	2570
8/30	1.13	1.58	0.78	1.45	1.34	1.61	1369

(d) Preventative measures

Date	Anger	Disgust	Fear	Joy	Sadness	Surprise	Number
8/25	1.20	1.11	2.88	0.87	1.26	0.95	843
8/26	0.95	0.87	2.60	0.85	1.13	0.89	466
8/27	0.90	0.91	2.22	0.74	1.06	0.79	3370
8/28	0.84	0.90	2.05	0.72	0.97	0.80	4259
8/29	0.57	0.55	1.09	0.75	0.71	0.58	3016
8/30	0.66	0.50	0.99	0.83	0.72	0.69	1440

(e) Help & rescue

Date	Anger	Disgust	Fear	Joy	Sadness	Surprise	Number
8/25	1.12	1.16	0.58	0.64	0.65	0.48	44
8/26	1.03	0.99	1.49	1.33	1.14	1.66	917
8/27	0.83	0.83	1.01	1.32	0.89	1.39	9062
8/28	0.88	0.89	0.97	1.06	0.87	1.14	10720
8/29	0.76	0.61	0.84	1.06	0.79	0.97	9478
8/30	0.64	0.48	0.78	1.01	0.67	0.86	5539

(f) Business, work, & school

Date	Anger	Disgust	Fear	Joy	Sadness	Surprise	Number
8/25	0.78	0.80	0.81	1.00	0.78	0.78	3137
8/26	1.02	1.02	0.69	1.06	0.73	0.71	1792
8/27	0.90	1.00	0.57	1.10	0.82	0.78	4473
8/28	0.72	0.75	0.62	0.94	0.68	0.78	4477
8/29	0.70	0.85	0.51	0.94	0.64	0.88	4027
8/30	0.61	0.78	0.32	1.11	0.55	1.07	2956

(g) Flood control infrastructures

Date	Anger	Disgust	Fear	Joy	Sadness	Surprise	Number
8/25	0.27	0.71	0.96	0.17	0.56	0.55	43
8/26	0.57	0.47	0.73	0.55	0.62	0.61	773
8/27	1.01	1.04	1.05	0.62	0.89	0.86	2294
8/28	0.83	0.84	1.05	0.70	0.79	0.84	1216
8/29	0.69	0.90	1.09	0.55	1.08	0.86	955
8/30	0.58	0.46	0.67	0.97	0.65	0.77	493

(h) Transportation

Date	Anger	Disgust	Fear	Joy	Sadness	Surprise	Number
8/25	0.94	0.74	1.19	1.09	1.41	1.55	264
8/26	0.89	0.75	0.75	0.70	0.77	0.68	1772
8/27	0.82	0.84	0.65	0.80	0.70	0.62	2103
8/28	0.77	0.56	0.39	0.55	0.51	0.31	922
8/29	0.53	0.84	0.74	0.68	0.78	0.71	1287
8/30	0.56	0.75	0.46	0.75	0.54	0.53	648

(i) Casualty

Date	Anger	Disgust	Fear	Joy	Sadness	Surprise	Number
8/25	1.73	2.94	1.60	0.00	1.40	0.52	37
8/26	3.50	3.99	3.47	0.52	4.75	2.01	146
8/27	2.79	4.05	2.47	0.43	3.87	2.25	131
8/28	1.94	2.06	2.68	0.73	3.55	1.06	336
8/29	1.83	2.16	2.67	0.98	2.76	2.03	469
8/30	3.41	3.47	2.64	1.86	2.65	3.92	143

(j) Built-environment hazards

Date	Anger	Disgust	Fear	Joy	Sadness	Surprise	Number
8/25	1.17	1.20	4.15	0.47	1.00	0.71	171
8/26	0.84	0.90	2.92	0.17	0.44	0.67	239
8/27	0.64	0.95	3.89	0.93	1.67	1.40	138
8/28	1.37	0.91	4.33	0.34	0.90	1.69	220
8/29	0.57	4.03	4.28	0.57	2.79	1.67	203
8/30	1.41	1.74	3.58	0.21	2.24	2.63	271

Emotion legend: Anger Disgust Fear Joy Sadness Surprise

FIGURE 5 Overview of the societal impact analysis result of Hurricane Harvey

and restaurant closed," and "eating junk food." On August 25, "purchasing the supplies for the hurricane" was the major topic on Twitter, which was also the source of negative emotions such as anger and disgust. Among the top 10 largest subevents (covering 73% of utility and supply tweets posted on August 25), four subevents are about "purchasing food and water in grocery stores and supermarkets." These subevents had an average anger score of 1.18. Residents were complaining about "empty shelves" and "food and water being sold out," which registered anger and disgust emotions. (e.g., "This is the 4th grocery store I've been to this morning, and there's literally no water left"). On August 26, the subevent of "grocery store is packed/empty" disappears from the top 10 largest subevents. Instead, the subevent of "grocery store closed" emerged with an anger score of 1.19 and a disgust score of 1.13. Besides, the primary contributor to negative emotion on this day was the subevent of "eating junk food due to hurricane" (anger: 1.97, disgust: 2.01), and "preparing hurricane supplies" (anger: 1.61, disgust: 1.60) due to the swear words. August 27–28, the leading subtopics on "food and supplies" was related to "gathering and distributing food," such as "help on food," "food place open/close," and "running low on food." These tweets contain less negative emotion compared with the tweets on August 25 and 26. On August 29–30, negative emotions in food-related tweets further decreased because "grocery stores and food places open" started to emerge as a major subevent. The main contributor to anger and disgust during August 27–30 was the subevent of "eating junk food due to the hurricane," which only occupies about 6–8% of the tweets in the top 10 largest subevents.

"Power outage" was another vital subcategory in *Utilities and supplies*. In terms of magnitude, the number of tweets on August 25–27 was large (about 50% of the tweets in the top 10 subevents in *Utilities and supplies*). On August 28, tweets about power outages decreased significantly (about 40% of the number on August 25 or 26), and people started to talk about "Power is back." Finally, on August 29–30, the number of tweets about power outage became very small (about 1/10 of the number on August 25 or 26). In terms of emotion score, the subevents directly mentioning "lose power" always had very high scores of anger, disgust, and sad (about 2–3) because the keyword of "lose" indicates the three negative emotions according to the emotion lexicon (Mohammad & Turney, 2013). The subevent of "lose power" existed in the top 10 largest subevents on August 25–29. For other power-related subevents without the word "lose," the score of negative emotions was also high (about 2–3) on August 25 and 26, and then diminishes dramatically after August 27.

Besides food/water and power outage, "gas station" was another smaller topic in the category of *Utility and supplies*. This topic occupied one of the top 10 largest subevents in four of six analyzed days. People were complaining about "gas stations being out of gas" on August 25 with anger (1.52) and disgust (1.29). Then, people were complaining about "the difficulty to find an open gas station" on August 27–29 with a relatively mild emotional status.

3.3.2 | Housing

Overall, tweets in the category of *Housing* have a high score of anger, disgust, and sadness. As shown in Figure 5c, the emotion score for "anger" and "fear" reached their maximum on August 26, while the score for "sad" and "disgust" reached their maximum on August 28. This inconsistency indicates the shift in the emotions of the citizens: on August 26, people were worried about their house being flooded, and the worries started to become a reality on August 27, which caused sadness and disgust emotions.

Housing-related tweets barely exist on August 25. Tweets on housing increased significantly on August 26, reached a maximum on August 28, and then gradually decreased afterward. Tweets about "flooded house" occupied the majority of this category during August 26–30 and with a relatively high score of negative emotions (Figure 5c). Specifically, on August 26, the subevent of "tornado threat to houses" emerged (anger: 1.77, disgust: 1.34, fear: 2.82). During August 27–28, more than half of the tweets in this category were about flooded houses/apartments. The second-largest subevent during this time is "losing home/house" (49 tweets on August 27 and 1,237 tweets on August 28), which contributed significantly to anger (2.07), disgust (1.30), and sadness (2.84). Finally, the subevent of clean the house emerged on August 30, showing the start of recovery efforts from flood impacts on housing.

3.3.3 | Preventative measure

Two major topics exist in the category of Preventative measure: evacuation and sheltering. Overall, the keyword "evacuate" caused a high score of fear for related subevents. According to Figure 5d, the negative emotion in the tweets in this category went down gradually since August 25, while the number of tweets reached its peak on August 27. This result shows that the Houstonians were less panicky and getting emotionally ready for evacuation and sheltering when facing the landfall of Harvey August 27–28.

The number of tweets related to both evacuation and sheltering was small compared with later days on August 25–26. Discussion on evacuation activities was dominating the tweets in this category (e.g., "mandatory/voluntary evacuation," "reluctant to leave," "family member's reaction to evacuation order"), and they often contained strong fear emotion (fear: about 3.0). Discussions about sheltering also existed, such as "sheltering for the tornado" and "animal shelters." Several topics caused strong fear emotions, such as: "officials have divergence on whether people should evacuate" (anger: 1.77, fear: 2.71); "homeless people cannot find shelters" (anger: 2.59, disgust: 3.52, fear: 2.33, sadness: 2.80); "mayor of Rockport just told the residence of Rockport if they didn't evacuate to please mark their name and SSN on their arms" (fear: 3.00), and "evacuate from home" (anger: 1.4, fear: 3.17, sadness: 2.42).

On August 27–28, the number of tweets in this category increased to more than 500% of the number during August 25–26. During these days, mandatory and voluntary were evacuations issued across the Houston area while the mayor of Houston defended the decision not to evacuate, generating heated discussion on Twitter. One opinion was that the government should have announced the evacuation order earlier (e.g., "they're trying to evacuate everyone when it's already too late"), but more people supported Houston mayor's decision (e.g., "Too true. We evacuated for Rita and it was a bad idea. Logistical nightmare. I think everyone learned their lesson"). Tweets related to the discussion on evacuation orders during August 27–28 had a relatively high score of fear (August 27: 2.53, August 28: 3.00). On the other hand, discussions on the shelter were also abundant. Tweets sharing information on shelter openings formed several major subevents. It is worth noting that people promoted the trending event that "the owner of Gallery Furniture provided shelter to people whose homes were flooded." Contrarily, people condemned the behavior of a church pastor, Joel Osteen, who "didn't open his church as a shelter during Harvey."

During August 29–30, the number of tweets on preventative measure decreased, which was mainly due to the reduction in the number of tweets about evacuation. The major topics in this category were the "mandatory evacuation orders," "the complaints about evacuating from home," "sharing shelter information," and "the need for help, resource, and domination of the shelters."

3.4 | Data processing and manual scanning time of using SocialDISC

Prompt processing and delivery of SocialDISC to reveal the societal impacts of community disruptions is key to the usefulness of that information for disaster response decision-making and later for accelerating post-disaster virtual reconnaissance activities. This subsection discusses the time spent

TABLE 6 Data processing and manual scanning time for one user to use SocialDISC to analyze 1-day tweets in Houston during Harvey

Activity	Steps	Time spent	Minutes (low)	Minutes (high)
SocialDISC (social sensing of Disaster Impacts and Societal Considerations) semiautomated data processing	Extracting important words from each Twitter message	40 min	40	40
	Generating semantic clusters	20 min	20	20
	Annotation of the word-sense disambiguation	2–3 hr	120	180
	Bootstrapping process	7–9 iterations, 10 min per iteration	70	90
	Total	About 4–5.5 hr	250	330
Manually scanning generated subevents	Scanning large categories	Five categories, 30 min per category	150	150
	Scanning small categories	Four categories, 10–20 min per category	40	80
	Total	About 3–4 hr	190	230
Total		About 7–9.5 hr	440 min or 7.3 hr	560 min or 9.3 hr

on data processing and manual scanning time of using SocialDISC for analyzing 1-day Harvey tweets in Houston.

The most time-consuming phase of SocialDISC semiautomated data processing is the adaptive and weakly supervised labeling; time spent on other phases of data processing is negligible. For this phase, all experiments are executed on a Linux operating system computer (Ubuntu 14.04) with the PyTorch deep-learning framework (https://pytorch.org/), with CUDA 9.0 dependencies for graphics processing unit acceleration. The computer contains an Intel Xeon E5450 3.0-GHz processor, 1 TB of hard disk storage, 64 GB of RAM, and an NVIDIA GeForce GTX 1080 GPU. The processing of tweets begins with extracting important words from each Twitter message requiring, on average, 40-min processing by the sentence encoder for encoding 1-day Harvey data. The SLPA clustering algorithm takes about 20 min to generate semantic clusters on 1-day Twitter messages. Annotation of the WSD by the domain expert required about 2–3 hr of manual work, depending on the number of tweets posted on the day and the diversity of the topics. The bootstrapping process normally stopped after seven to nine iterations, and each iteration took around 10 min. In summary, processing 1-day tweets (more than 200K tweets, excluding retweets and replies) representing Harvey landfall and time in Houston took about 4–5.5 hr.

Manually scanning the subevents by the analyzer is another time-consuming activity. It takes the analyzer (who, in this case, is not a native English speaker) about 3 min to scan one large tweet subevents consisting of 100+ tweets. Finishing the scanning of 1-day tweets in one large category (i.e., *Preventative measure, Utility and supplies, Housing, Help and rescue, and Business, work, and school*) takes about half an hour. On the other hand, scanning the tweets in small categories is much quicker, which costs about 10–20 min. This is because both the number of subevents in each category and the number of tweets in each subevent are fewer. According to the actual process of scanning the Harvey data, scanning of major subevents in all nine categories takes about 3–4 hr for August 27–29, 2017, when the impact of Harvey reached its maximum.

In summary, assessing the societal impact of the residents in a large city (like Houston) from 1-day tweets when the residents suffer the most takes up to 7–9.5 hr of processing and manual scanning time. Table 6 summarizes the analysis time. This result shows the ability of SocialDISC to generate timely information regarding the societal impacts of community disruptions. Furthermore, reducing the period of the analysis (i.e., cutting the analysis time to half a day instead of one whole day) can significantly reduce the data processing time in case one needs the results urgently. Adding more annotators and analyzers will also increase data processing productivity because the run-time analysis is based on a single annotator and analyzer in this study. This time spent can be further reduced if the historical annotations are recycled to avoid repetitive annotation on similar contents, which could be a future research focus.

4 | DISCUSSION AND CONCLUDING REMARKS

This paper proposes SocialDISC, a social media analytics framework, to give infrastructure management personnel access to social impacts and situational awareness of the population during a disaster. SocialDISC quickly quantitatively and qualitatively analyzes thousands of tweets posted during disasters to assess the societal impacts on affected residents. To demonstrate the performance of SocialDISC, the authors proposed a case study based on Hurricane Harvey, which is one of the costliest disasters in U.S. history and analyzed disruptions and corresponding societal impacts in different aspects. The analysis result showed that Houstonians suffered from the lack of food, lack of access to medical care, power outages, and flooded houses. SocialDISC can potentially help people understand the relationship between disruptions of infrastructures and societal impacts, expectations of the public when facing disasters, and infrastructure interdependency and cascading failures. Also, SocialDISC provides timely information about the societal impacts, which could be a factor in the decision-making in managing lifeline infrastructures during disaster response and relief activities.

4.1 | Insights gained from the observed societal impacts

The proposed social media analytics framework helps the users acquire information regarding the societal impacts caused by community disruptions. In the case of Hurricane Harvey, considering the magnitude and the emotional score of the tweets in each category, disruptions in housing and utilities/supplies caused the most significant impacts on people's lives. On the other hand, extreme events about *Casualty* and *Built-environment hazard* raise relatively fewer discussions in terms of the number of tweets in these two categories. One possible explanation of such a phenomenon is that the impact of Harvey is so broad and severe that the residents tended to express their concerns about their own situation or to disseminate information about disaster response, which occupies the bandwidth of expressing compassion on tragedies happening to strangers. Besides, tweets about *Help and rescue* is the most among the nine categories and the emotion signal of which is the most positive as well. This phenomenon shows the strong will power and social cohesion of Houstonians when facing Harvey.

The analysis results provide insights into relationships between community disruptions, emotion signals, and societal impacts and expectations. The signal of food and water shortage appeared on August 25, and the shortage and inaccessibility to food deepened in the following days. Such an unsatisfied need for food shows a gap between the public's expectation for food access and the capability of providing food to the retail outlets. If used in conjunction with other data, such as storage and supply chain records of grocery stores and supermarkets, this societal impact information can provide a reference for evaluating mitigation and preparedness plans to serve the public's expectation for future disasters better.

The analysis result shows a clear timeline about how Hurricane Harvey impacts the different aspects of Houston residents by comparing the amount and emotion score of tweets in different categories. On August 25 and 26, the dominant community disruption events are in the "Utility and supplies" category, which are food and water out of stock in supermarkets and grocery stores. During this period, the impact on other aspects, such as housing and rescuing, was minimum. On August 27 and 28, the impact of Hurricane Harvey reached the maximum: people evacuated from flooded areas, sought water rescue, and even lost their homes. The impacts started to fade during August 29 and 30. Grocery stores, businesses, and schools started to re-open, and people expressed their joy for being able to work again. Such a timeline indicates the interaction between dynamic disaster situations and residents' active and adaptive reactions, which enriches the understanding of the patterns of community resilience under evolving disaster impacts.

The analysis results also provide a rich context for understanding the interdependency between infrastructure systems and the cascading effects of infrastructure-related failures. This information is essential in virtual post-disaster reconnaissance efforts. Figure 6 shows the identified interdependency relationships between different infrastructure systems captured directly from SocialDISC. Analysis results provide records and details about some interdependency relationships identified in existing literature, such as business closures mean unavailability of food and water, or empty storage tanks in gas stations influence transportation. Also, the results show unique and severe cases of interdependency relationships and cascading effects, that is, the water release of Addicks and Barker reservoirs led to flooded roads that significantly hindered the emergency relief operations. Finally, sufficient details in the analysis result also unveil the underexplored interdependency patterns in the existing literature. For example, road closures and power outages exacerbate the food shortage challenge because people confined to their homes without diversions or tasks tended to eat all hurricane food storage. To sum up, societal impact information retrieved using SocialDISC from social media posts can develop the disaster-specific models and assessment tools to describe infrastructure interdependencies that influence community resilience.

4.2 | Contribution of the analytics framework

The major outcome of this study is an analytics framework that enables the timely acquisition of societal impact information related to disaster-induced physical infrastructure and social disruptions extracted from social media data. Compared with existing methods, SocialDISC has two significant advantages. First, analyzers can adaptively modify the social media classification taxonomy according to different disaster characteristics. Based on their domain knowledge, disaster responders, infrastructure managers, and academic researchers can provide the topics related to various

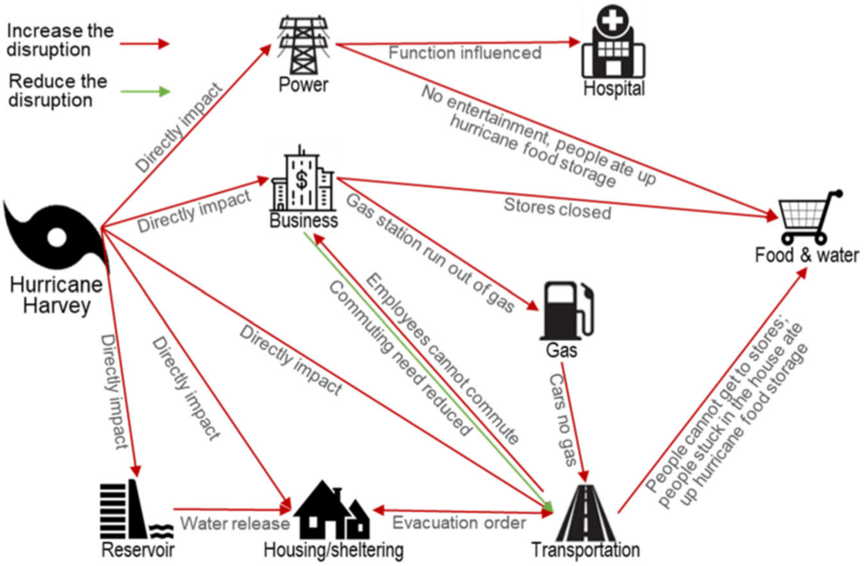

FIGURE 6 Identified interdependency relationships between infrastructure systems directly obtained from SocialDISC (social sensing of Disaster Impacts and Societal Considerations)

infrastructure and social disruptions as well as corresponding keywords to initiate the analytics framework in different disaster settings. Second, analyzers can employ this timely analytics framework to expedite the analysis of social media data for the understanding of disruptions, emotion signals, and societal impacts. Such adaptivity and timeliness are due to pipelining a series of weakly supervised or reliable pretrained machine learning approaches.

From a scientific perspective, SocialDISC addresses the limitation of existing social sensing approaches by developing a new computational model for analyzing the societal impacts of infrastructure and social disruptions based on social media data. Specifically, the semiautomated data analytics approach created and tested in this study integrates the advantage of AI-based data analytics capacity and the domain knowledge of the disaster response and infrastructure management professionals. Such integration offers a new capability for timely and fine-grained classification of events related to different disaster impacts to better examine societal impacts. As social media become an important infrastructure for communities in coping with disasters, fundamental understanding of assessing societal impacts due to infrastructure and social disruptions may hold the key for urban systems to become more intelligent and resilient in responding to disruptions. Therefore, the proposed SocialDISC system sheds light on data-enriched infrastructure management processes and systems for disaster resilience that solicit the power of big data generated by the crowd.

From a practical perspective, the outcomes of this research may provide new tools and insights to decision makers, emergency managers, and public officials regarding ways to improve intelligence in infrastructure management and emergency response. SocialDISC provides a quick generation of information about the disruption events and the reactions and impacts on the residents. Abnormal reactions to a particular type of disruption, for instance, grocery stores running out of food and water, from the public perspective can provide a timely signal for the execution of corresponding response plans. Media agents can ameliorate the effects on the population by posting or broadcasting specific preparedness or response-related information based on the topics that trigger devastating negative emotions.

4.3 | Limitation and future work

One limitation of SocialDISC is the vacancy of the geo-parsing function. Very few social media posts contain geo-coordination information (Zhang et al., 2019). Geo-parsing, which refers to retrieving geographic coordination mentioned in each social media post, will extend the dimension from the time domain to space domain (Hoang & Mothe, 2018; Khodabandeh Shahraki, Fatemi, & Malazi, 2019). However, existing geo-parsing techniques have limitations in terms of the level of detail and level of accuracy for disaster situational information retrieval tasks. They usually refer to Wikipedia or DBpedia to acquire city-level location lists but are less effective in acquiring street-level locations. Named entity recognition techniques (Finkel, Grenager, & Manning, 2005), which identifies the location names within social media posts, could be the solution if combined with a detailed GIS database about the names of all the names of the entities in the built environment (e.g., names of all the roads, schools, and hospitals). The on-going work of the authors focuses on adding geo-parsing information to SocialDISC for retrieving the geo-location of mentioned infrastructures and facilities in social media posts.

A potential future direction is exploring the viability and reliability of assessing the well-being of the residents using social media data. Existing studies have supported the argument that emotion is an important indicator of people's well-being (Bourke, Douglas, & Porter, 2010; Finucane, Dima, Ferreira, & Halvorsen, 2012; Martin & Dahlen, 2005).

SocialDISC sheds light on using social media posts to understand what disruption events impacted the well-being of the residents and to what degree. However, many gaps exist between state-of-the-art social media analytics techniques (e.g., SocialDISC) and well-being assessment in disasters. For example, whether social media users are a representative sample of the residents in the disaster-impact areas? Can emotion signals in social media posts comprehensively reflect different dimensions (e.g., long-term or short-term, physical, or mental) of the well-being of people? Do the underserved communities (e.g., the minorities, the low-income communities, and the elderly) have enough exposure to the social media so that their well-being status is appropriately reflected? What are the natural limitations of social media data in terms of assessing people's well-being during disasters? Future works will focus on fulfilling the proposed research gaps to build the paradigm of efferent and effective sensing of people's societal impacts and well-being in disasters using social media.

Another important future direction is assessing and improving the practical use of SocialDISC via participatory approaches. SocalDISC can generate timely and detailed societal impact information regarding various community disruptions. Future studies could hold participatory panels or workshops to validate the practical use case of SocialDISC with infrastructure managers and operators and emergency responders and solicit feedback on improving its human-computer interaction interfaces. The participatory approaches will also explore the necessity and practicality to integrate the assessment of societal impact using SocialDISC with existing disaster response and infrastructure management processes under different disaster types and scenarios.

Finally, future studies can test the performance of SocialDISC in different disaster scenarios, especially for different types of disasters. It would be valuable to identify the common patterns and differences between the societal impacts caused by different disasters at different locations, such as earthquakes, floods, and wildfires. For example, when integrated with other information sources (Panakkat & Adeli, 2007, 2008; Rafiei & Adeli, 2017), timely analysis of the societal impacts based on social media data can help predict the damage and disruption caused by different natural hazards, which is critical for the response, warning, and emergent relief.

ACKNOWLEDGMENTS

This material is based in part upon work supported by the National Science Foundation under Grant Number IIS-1759537 and the Amazon Web Services (AWS) Machine Learning Award. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation and Amazon Web Services (AWS).

REFERENCES

Atefeh, F., & Khreich, W. (2015). A survey of techniques for event detection in twitter. *Computational Intelligence*, *31*(1), 133–164. https://doi.org/10.1111/coin.12017.

Batouli, M., & Mostafavi, A. (2018). Multiagent simulation for complex adaptive modeling of road infrastructure resilience to sea-level rise. *Computer-Aided Civil and Infrastructure Engineering*, *33*(5), 393–410. https://doi.org/10.1111/mice.12348.

Bourke, C., Douglas, K., & Porter, R. (2010). Processing of facial emotion expression in major depression: A review. *Australian & New Zealand Journal of Psychiatry*, *44*(8), 681–696. https://doi.org/10.3109/00048674.2010.496359.

Burel, G., Saif, H., & Alani, H. (2017, October). Semantic wide and deep learning for detecting crisis-information categories on social media. In *International semantic web conference* (pp. 138–155). Cham: Springer. https://doi.org/10.1007/978-3-319-68288-4_9.

Chen, T., Xu, R., He, Y., & Wang, X. (2017). Improving sentiment analysis via sentence type classification using BiLSTM-CRF and CNN. *Expert Systems with Applications*, *72*(April), 221–230. https://doi.org/10.1016/j.eswa.2016.10.065.

Chen, X., Elmes, G., Ye, X., & Chang, J. (2016). Implementing a real-time twitter-based system for resource dispatch in disaster management. *GeoJournal*, *81*(6), 863–873. https://doi.org/10.1007/s10708-016-9745-8.

Conneau, A., Kiela, D., Schwenk, H., Barrault, L., & Bordes, A. (2017). *Supervised learning of universal sentence representations from natural language inference data*. Retrieved from http://arxiv.org/abs/1705.02364

Dutt, R., Basu, M., Ghosh, K., & Ghosh, S. (2019). Utilizing microblogs for assisting post-disaster relief operations via matching resource needs and availabilities. *Information Processing & Management*, *56*(5), 1680–1697. https://doi.org/10.1016/j.ipm.2019.05.010.

Ekman, P. (1992). An argument for basic emotions. *Cognition and Emotion*, *6*(3–4), 169–200. https://doi.org/10.1080/02699939208411068.

Elliott, J. R., & Pais, J. (2006). Race, class, and Hurricane Katrina: Social differences in human responses to disaster. *Social Science Research*, *35*, 295–321. https://doi.org/10.1016/j.ssresearch.2006.02.003.

Fan, C., & Mostafavi, A. (2019). A graph-based method for social sensing of infrastructure disruptions in disasters. *Computer-Aided Civil and Infrastructure Engineering*, *34*(12), 1055–1070. https://doi.org/10.1111/mice.12457.

Felts, R., Leh, M., & McElvaney, T. (2016). *Public safety analytics R&D roadmap*. (NIST Technical Note 1917). U.S. Department of Commerce, National Institute of Standards and Technology. https://doi.org/10.6028/NIST.TN.1917.

Finkel, J. R., Grenager, T., & Manning, C. (2005). Incorporating non-local information into information extraction systems by Gibbs sampling. *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics—ACL '05* (pp. 363–370), Morristown, NJ: Association for Computational Linguistics. https://doi.org/10.3115/1219840.1219885.

Finucane, A. M., Dima, A., Ferreira, N., & Halvorsen, M. (2012). Basic emotion profiles in healthy, chronic pain, depressed and PTSD individuals. *Clinical Psychology & Psychotherapy*, *19*(1), 14–24. https://doi.org/10.1002/cpp.733.

Gardoni, P., & Murphy, C. (2010). Gauging the societal impacts of natural disasters using a capability approach. *Disasters*, *34*(3), 619–636. https://doi.org/10.1111/j.0361-3666.2010.01160.x.

Graves, A., & Schmidhuber, J. (2005). Framewise phoneme classification with bidirectional LSTM and other neural network architectures. *Neural Networks*, *18*(5–6), 602–610. https://doi.org/10.1016/j.neunet.2005.06.042.

Guo, Y., & Li, Y. (2016). Getting ready for mega disasters : The role of past experience in changing disaster consciousness. *Disaster Prevention and Management: An International Journal*, *25*(4), 492–505. https://doi.org/10.1108/DPM-01-2016-0008.

Hiltz, S. R., Kushma, J. A., & Plotnick, L. (2014). Use of social media by U.S. public sector emergency managers: Barriers and wish lists. *Proceedings of the 11th International ISCRAM Conference* (pp. 602–611), University Park, PA. https://doi.org/10.13140/2.1.3122.4005.

Hoang, T. B. N., & Mothe, J. (2018). Location extraction from tweets. *Information Processing & Management*, *54*(2), 129–144. https://doi.org/10.1016/j.ipm.2017.11.001.

Huang, Q., & Xiao, Y. (2015). Geographic situational awareness: Mining tweets for disaster preparedness, emergency response, impact, and recovery. *ISPRS International Journal of Geo-Information*, *4*(3), 1549–1568. https://doi.org/10.3390/ijgi4031549.

Khodabandeh Shahraki, Z., Fatemi, A., & Malazi, H. T. (2019). Evidential fine-grained event localization using twitter. *Information Processing & Management*, *56*(6), 102045. https://doi.org/10.1016/j.ipm.2019.05.006.

Laylavi, F., Rajabifard, A., & Kalantari, M. (2017). Event relatedness assessment of twitter messages for emergency response. *Information Processing & Management*, *53*(1), 266–280. https://doi.org/10.1016/j.ipm.2016.09.002.

Lindell, M. K., & Prater, C. S. (2004). Assessing community impacts of natural disasters. *Natural Hazards Review*, *4*(4), 176–185.

Martin, R. C., & Dahlen, E. R. (2005). Cognitive emotion regulation in the prediction of depression, anxiety, stress, and anger. *Personality and Individual Differences*, *39*(7), 1249–1260. https://doi.org/10.1016/j.paid.2005.06.004.

Mohammad, S. M., & Turney, P. D. (2013). Crowdsourcing a word-emotion association lexicon. *Computational Intelligence*, *29*(3), 436–465. https://doi.org/10.1111/j.1467-8640.2012.00460.x.

Morss, R. E., Cuite, C. L., Demuth, J. L., Hallman, W. K., & Shwom, R. L. (2018). Is storm surge scary ? The influence of hazard, impact, and fear-based messages and individual differences on responses to hurricane risks in the USA. *International Journal of Disaster Risk Reduction*, *30*(September), 44–58. https://doi.org/10.1016/j.ijdrr.2018.01.023.

Murphy, C., & Gardoni, P. (2007). Determining public policy and resource allocation priorities for mitigating natural hazards: A capabilities-based approach. *Science and Engineering Ethics*, *13*(4), 489–504. https://doi.org/10.1007/s11948-007-9019-4.

National Hurricane Center. (2018). *Costliest U.S. tropical cyclones tables updated*. Miami, FL: Author.

Navigli, R., & Lapata, M. (2010). An experimental study of graph connectivity for unsupervised word sense disambiguation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *32*(4), 678–692. https://doi.org/10.1109/TPAMI.2009.36.

Nejat, A., & Damnjanovic, I. (2012). Agent-based modeling of behavioral housing recovery following disasters. *Computer-Aided Civil and Infrastructure Engineering*, *27*(10), 748–763. https://doi.org/10.1111/j.1467-8667.2012.00787.x.

Ng, M. W., Park, J., & Waller, S. T. (2010). A hybrid bilevel model for the optimal shelter assignment in emergency evacuations. *Computer-Aided Civil and Infrastructure Engineering*, *25*(8), 547–556. https://doi.org/10.1111/j.1467-8667.2010.00669.x.

Nguyen, L., Yang, Z., Li, J., Cao, G., & Jin, F. (2018). *Forecasting people's needs in hurricane events from social network*. Retrieved from http://arxiv.org/abs/1811.04577

NIST. (2015). *Community resilience planning guide for buildings and infrastructure systems: Volume II*. Gaithersburg, MD: Author. https://doi.org/10.6028/NIST.SP.1190v1.

NIST. (2016). Critical assessment of lifeline system performance: Understanding societal needs in disaster recovery (NIST GCR 16-917-39). Prepared for U.S. Department of Commerce National Institute of Standards and Technology (NIST) Engineering Laboratory by Applied Technology Council. https://doi.org/10.6028/NIST.GCR.16-917-39.

Nomura, S., Parsons, A. J. Q., Hirabayashi, M., Kinoshita, R., Liao, Y., & Hodgson, S. (2016). Social determinants of mid- to long-term disaster impacts on health: A systematic review. *International Journal of Disaster Risk Reduction*, *16*, 53–67. https://doi.org/10.1016/j.ijdrr.2016.01.013.

Othman, S. H., Beydoun, G., & Sugumaran, V. (2014). Development and validation of a disaster management metamodel (DMM). *Information Processing & Management*, *50*(2), 235–271. https://doi.org/10.1016/j.ipm.2013.11.001.

Panakkat, A., & Adeli, H. (2007). Neural network models for earthquake magnitude prediction using multiple seismicity indicators. *International Journal of Neural Systems*, *17*(01), 13–33. https://doi.org/10.1142/S0129065707000890.

Panakkat, A., & Adeli, H. (2008). Recent efforts in earthquake prediction (1990–2007). *Natural Hazards Review*, *9*(2), 70–80. https://doi.org/10.1061/(ASCE)1527-6988(2008)9:2(70).

Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., & Zettlemoyer, L. (2018). *Deep contextualized word representations*. Retrieved from http://arxiv.org/abs/1802.05365

Rafiei, M. H., & Adeli, H. (2017). NEEWS: A novel earthquake early warning model using neural dynamic classification and neural dynamic optimization. *Soil Dynamics and Earthquake Engineering*, *100*(September), 417–427. https://doi.org/10.1016/j.soildyn.2017.05.013.

Reuter, C., & Kaufhold, M-A. (2018). Fifteen years of social media in emergencies: A retrospective review and future directions for crisis informatics. *Journal of Contingencies and Crisis Management*, *26*(1), 41–57. https://doi.org/10.1111/1468-5973.12196.

Reuter, C., Ludwig, T., Kaufhold, M-A., & Spielhofer, T. (2016). Emergency services′ attitudes towards social media: A quantitative and qualitative survey across Europe. *International Journal of Human-Computer Studies*, *95*(November), 96–111. https://doi.org/10.1016/j.ijhcs.2016.03.005.

Rudra, K., Banerjee, S., Ganguly, N., Goyal, P., Imran, M., & Mitra, P. (2016). Summarizing situational tweets in crisis scenario. *Proceedings of the 27th ACM Conference on Hypertext and Social Media—HT '16*, 129 (pp. 137–147), New York, NY: ACM Press. https://doi.org/10.1145/2914586.2914600.

Rudra, K., Ghosh, S., Ganguly, N., Goyal, P., & Ghosh, S. (2015). Extracting situational information from microblogs during disaster events: A classification-summarization approach. *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*, CIKM '15 (pp. 583–592). New York, NY: ACM. https://doi.org/10.1145/2806416.2806485.

Sutton, J., Ben Gibson, C., Phillips, N. E., Spiro, E. S., League, C., Johnson, B., … Butts, C. T. (2015). A cross-hazard analysis of terse message retransmission on twitter. *Proceedings of the National Academy of Sciences*, *112*(48), 14793–14798. https://doi.org/10.1073/pnas.1508916112.

Toya, H., & Skidmore, M. (2007). Economic development and the impacts of natural disasters. *Economics Letters*, *94*, 20–25. https://doi.org/10.1016/j.econlet.2006.06.020.

Uchida, Y., Takahashi, Y., & Kawahara, K. (2014). Changes in hedonic and eudaimonic well-being after a severe nationwide disaster: The case of the Great East Japan Earthquake.*Journal of Happiness Studies*, *15*, 207–221. https://doi.org/10.1007/s10902-013-9463-6.

Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, *359*(6380), 1146–1151. https://doi.org/10.1126/science.aap9559.

Ward, P. S., & Shively, G. E. (2017). Disaster risk, social vulnerability, and economic development. *Disasters*, *41*(2), 324–351. https://doi.org/10.1111/disa.12199.

Xie, J., Kelley, S., & Szymanski, B. K. (2013). Overlapping community detection in networks. *ACM Computing Surveys*, *45*(4), 1–35. https://doi.org/10.1145/2501654.2501657.

Xie, J., Szymanski, B. K., & Liu, X. (2011, December). SLPA: Uncovering overlapping communities in social networks via a speaker-listener interaction dynamic process. *In 2011 IEEE 11th International Conference on Data Mining Workshops* (pp. 344–349), IEEE.

Xu, Z., Liu, Y., Yen, N., Mei, L., Luo, X., Wei, X., & Hu, C. (2016). Crowdsourcing based description of urban emergency events using social media big data. *IEEE Transactions on Cloud Computing*, *7161*(c), 1–1. https://doi.org/10.1109/TCC.2016.2517638.

Xu, Z., Liu, Y., Zhang, H., Luo, X., Mei, L., & Hu, C. (2017). Building the multi-modal storytelling of urban emergency events based on crowdsensing of social media analytics. *Mobile Networks and Applications*, *22*(2), 218–227. https://doi.org/10.1007/s11036-016-0789-2.

Yang, Y., Zhang, C., Fan, C., Yao, W., Huang, R., & Mostafavi, A. (2019). Exploring the emergence of influential users on social media during natural disasters. *International Journal of Disaster Risk Reduction*, *38*(August), 101204. https://doi.org/10.1016/j.ijdrr.2019.101204.

Yao, W., Zhang, C., Saravanan, S., Huang, R., & Mostafavi, A. (2020). Weakly-supervised fine-grained event recognition on social media texts for disaster management. *The Thirty-Fourth AAAI Conference on Artificial Intelligence*, New York, NY.

Yuan, K., Liu, G., & Wu, J. (2019). Whose posts to read: Finding social sensors for effective information acquisition. *Information Processing & Management*, *56*(4), 1204–1219. https://doi.org/10.1016/j.ipm.2019.01.009.

Zhang, C., Fan, C., Yao, W., Hu, X., & Mostafavi, A. (2019). Social media for intelligent public information and warning in disasters: An interdisciplinary review. *International Journal of Information Management*, *49*(December), 190–207. https://doi.org/10.1016/j.ijinfomgt.2019.04.004.

Zhou, P., Qi, Z., Zheng, S., Xu, J., Bao, H., & Xu, B. (2016). *Text classification improved by integrating bidirectional LSTM with two-dimensional max pooling*. Retrieved from http://arxiv.org/abs/1611.06639

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

How to cite this article: Zhang C, Yao W, Yang Y, Huang R, Mostafavi A. Semiautomated social media analytics for sensing societal impacts due to community disruptions during disasters. *Comput Aided Civ Inf*. 2020;1–18. https://doi.org/10.1111/mice.12576