# Enabling Rapid Disaster Response Using Artificial Intelligence and Social Media

Muhammad Imran, Firoj Alam, Ferda Ofli, Micheal Aupetit
Qatar Computing Research Institute,
Hamad Bin Khalifa University,
Doha, Qatar
{mimran, falam, fofli, maupetit}@hbku.edu.qa

## ABSTRACT

Disasters and emergencies often bring unanticipated situations for the members of public and formal response organizations. During such times, access to rapidly changing information plays an important role to understand the situation as it unfolds. However, information scarcity, especially in the early hours, hinders this task and ultimately delays response operations. Social media platforms such as Twitter and Facebook are gaining attention of both formal response organizations and affected individuals. People use social media to share a variety of information online including requests for help, and other urgent needs such as food, water, and shelter. This online information can be useful for humanitarian organizations if processed timely. This work presents a number of Artificial Intelligence (AI) technologies developed at the Qatar Computing Research Institute (QCRI) to aid disaster response and management. These technologies include machine learning platforms to automatically process textual messages (e.g., tweets) into humanitarian categories such as reports of needs, injured or dead people, and to process imagery data. The platforms work in real-time to collect and analyze data on social media to help humanitarian organization gain situational awareness and extract actionable information. These technologies were deployed during the Hurricane Maria and in this paper we present our findings and insights from this real-world natural disaster.

**Keywords:** Disaster response, Social media, Artificial intelligence

## INTRODUCTION

According to The Economist[1], meteorological disasters caused by extreme weather including cyclones, blizzards, heat waves, hurricanes, droughts are increasing in the 21st century. While more people are being suffered by these disasters, the number of fatalities is actually decreasing as response to these disasters improves. However, the devastation in terms of human-lives and economic damage caused by these disasters is still huge. For instance, the recent Hurricane Harvey[2] is being estimated as the second-costliest natural disaster in the U.S. history after Hurricane Katrina[3] in 2005. Disasters and emergencies bring uncertain situations in which formal responders and members of the affected communities struggle to find useful information, in particularly first-hand and trustable information to improve decision-making and to launch relief operations. For instance, the United Nations Office for the Coordination of Humanitarian Affairs (UN OCHA) seeks to gain situational awareness in the first 48 to 72 hours after a sudden-onset disaster to understand urgent needs of affected population (e.g., food, water, shelter, medical assistance), disaster severity, and local government capacity to respond, among other factors (Whipkey & Verity, 2015). Gathering such information at the sudden-onset of a disaster, especially in the first 24 hours is nearly impossible, as it requires human experts and crisis severity assessors to visit the disaster areas, interview the affected people, and observe their urgent needs.

---

[1] https://www.economist.com/blogs/graphicdetail/2017/08/daily-chart-19

[2] https://en.wikipedia.org/wiki/Hurricane_Harvey

[3] https://en.wikipedia.org/wiki/Hurricane_Katrina

Access to rapid information, as the crisis situation unfolds, is a challenging tasks particularly in the early hours of a disaster event when no other information sources are available (Imran, Castillo, Diaz, & Vieweg, 2015). To bridge this information scarcity gap, humanitarian relief organizations have recently started acknowledging the role of social media platforms such as Twitter and Facebook as vital information sources for disaster response and management (Whipkey & Verity, 2015). The widespread use of social media sites during disasters and emergencies provides convenient ways to share and consume public information as fast and easy as never before (Hughes & Palen, 2009; Vieweg, Hughes, Starbird, & Palen, 2010). People use web and mobile technologies to share different types of information including textual and multimedia content such as images and videos (Dat Tien Nguyen, Alam, Ofli, & Imran, 2017). If accessed and processed properly, this information could be useful for a variety of disaster response operations (Imran, et al., 2015).

Despite the fact that information available on social media is timely and potentially useful, making-sense of it is challenging due to a number of reasons. Among others, the time-critical analysis of social media streams requires processing millions of messages arriving at high-velocity, real-time parsing of brief and informal content, handling information overload, determining information credibility, and prioritizing useful information for stakeholders (Boersma, Fonio, Imran, & Meier, 2017; Castillo, 2016; Imran, et al., 2015).

This work presents a number of Artificial Intelligence technologies developed by the Qatar Computing Research Institute (QCRI) towards addressing the challenges in the use of social media for disaster response and to help relief organizations. On these lines, one of QCRI's flagship AI technologies is "Artificial Intelligence for Digital Response (AIDR)" (Imran, Castillo, Lucas, Meier, & Sarah, 2014). AIDR is a system conceived and developed at QCRI to harness information from real-time tweets that emerge from an area struck by a natural disaster to help coordinate relief activities. The AIDR system combines human and machine intelligence to categorize crisis-related messages during the sudden-onset of natural or man-made disasters. The United Nations OCHA now routinely uses AIDR for the coordination of humanitarian affairs as well as by many other emergency departments in the world. The system has been used during several major disasters such as the 2015 Nepal Earthquake, the 2014 typhoon Hagupit, the 2013 typhoon Hayan by a number of humanitarian organizations including UN OCHA and UNICEF.

In addition to textual messages, multimedia content such as images and videos on social media can be valuable for response organizations in a number of ways (Daly & Thom, 2016). For example, images can be used to determine damage severity done by the disaster (Ofli, et al., 2016). For this purpose, the AIDR system, which was originally designed to process textual content in real-time, has been extended to process imagery data posted on social media (Alam, Imran, & Ofli, 2017). In order to make sense of imagery content on social media and to use it to aid relief operations, a number of challenges have to be addressed. Among others, these challenges include identifying and removing irrelevant and duplicate images, classifying images into different humanitarian categories such as injured people scenes, damage scenes, shelter location scenes and so on. As a first step, currently we focus on the task of disaster damage severity assessment. That is, we aim to use social media images to measure the extent of damage to critical infrastructure such as buildings, bridges, and roads. Specifically, we have developed Deep Neural Networks based models to assess the severity of damage shown in an image in three levels: SEVERE damage, MILD damage, and NO damage (Dat T Nguyen, Ofli, Imran, & Mitra, 2017).

Although, employing AI techniques to process social media potentially noisy data helps reduce noise, however, often even after an intelligent filtering of noisy content, the remaining data, which is presumably useful, can still consist of tens of thousands of messages. Manual analysis of these thousands of messages remains a challenging endeavor for decision-makers to understand it, as it creates an information overload issue during an ongoing disaster—i.e., new data arrives before human-assessors finished analyzing the old data. To overcome this information overload issue, we have developed an interactive dashboard that helps crisis responders quickly sift through thousands of messages (Aupetit & Imran, 2017). The dashboard extracts critical entities (mentions of people, places, and organizations) from the data and allows users to apply multiple complex filtering criteria to find useful data nuggets.

To determine the effectiveness of all these technologies, we deployed all these systems during the 2017 Hurricane Maria. We use AIDR machine learning classifiers to automatically classify Twitter messages into humanitarian categories and we use images posted during the hurricane to assess the severity of damage. Furthermore, we use our interactive dashboard technology to find useful information during the event. This work presents the results of our analysis performed on the data collected from Twitter during Hurricane Maria.

The rest of the paper is organized as follows. The next section describes the artificial intelligence technologies for text and image processing. We present details of the dataset in the "Data collection and Description" section. Experimental results and discussion are presented in the next section. Finally we conclude and present future directions in the last section.

## ARTIFICIAL INTELLIGENCE TECHNOLOGIES FOR DISASTER RESPONSE

For rapid crisis response, real-time insights are important for emergency responders (Imran, Castillo, Lucas, Meier, & Rogstadius, 2014). To identify actionable and tactical information pieces from a growing stack of social media data and to inform decision-making processes as early as possible, messages need to be processed as they arrive. Given the high volume of social media messages, we need to triage them by categorizing them into different actionable bins such as food, supplies, financial, logistics, etc. Manual analysis of these messages coming at high rates is not possible. Even employing a large workforce could somehow be amassed to keep up with the high velocity of social media data, having humans do something that can be automated is inefficient and a misallocation of scarce resources. The problem gets even worse during prolonged crises, such as civil wars, or during the long recovery phase following major natural disasters.

Research on building automatic computational methods to overcome information overload issues has demonstrated significant advantages in many application domains (Dat Tien Nguyen, Al-Mannai, et al., 2017; Dat Tien Nguyen, Joty, Imran, Sajjad, & Mitra, 2016; Rudra, et al., 2016). For instance, machine learning techniques help recommend readers news articles (Hwang, Park, & Kim, 2015), natural language processing techniques to automate question-answering tasks both general-purpose (Sasikumar, 2014) as well as domain-specific e.g., helping health-experts to answer HIV-related questions via SMS (Imran, Meier, Castillo, Lesa, & Herranz, 2016). Despite advances in the field artificial intelligence for standard web data sources, in which the techniques mainly deal with structured and errors-free data, processing social media web data is still challenging. Often messages on social media are brief, full of slangs, and contain mistakes such as misspellings and shortened forms.

### Automatic Textual Content Processing

To overcome the above-mentioned issues and shortcomings of the exiting data processing techniques, we have developed an AI-based social media processing platform called "Artificial Intelligence for Digital Response (AIDR)". AIDR combines human-intelligence and machine learning techniques to process data. Specifically, AIDR uses supervised machine learning techniques to classify each message/tweet as it is posted on social media in real-time. AIDR is specifically designed to scale up the abilities of human workers by intelligently removing noise form the data e.g., in the form of duplicate and irrelevant messages. At the onset of a disaster event, humanitarian experts or members of affected communities start AIDR data collection. The data collection in AIDR can be configured either using event-specific keywords or hashtags; or by using geographical bounding box. One can define multiple bounding boxes to cover different disaster areas. In that case, geo-tagged tweets originating from the areas under surveillance are collected. However, we have observed that only 1% to 3% of the tweets are geo-tagged. For this reason, AIDR also provides a combination of keywords and geographical bounding box data collection strategy in which case a user defines both keywords and bounding boxes to collect tweets either matching with one of the keywords or geo-tagged within one of the bounding boxes.

After setting the data collection, the user enables machine classification process. This process involves defining a set of user-defined categories of interest to process text messages. For instance, many humanitarian organizations are interested in learning about injured or dead people, infrastructure damage, urgent needs of affected people in terms of food, water, shelter, utilities, money etc. Depending on information needs, the user defines categories in AIDR. Since AIDR uses supervised machine learning techniques, the next stage is to provide a handful of training examples for machine training. For this purpose, the user tags a few hundreds of tweets i.e., assigning appropriate category to tweets. AIDR learns new models using the training examples. New models are learned as more human-tagged tweets become available. AIDR uses Random Forest, a well-know decision trees based learning scheme, to learn models. As for the features, uni-grams and bi-grams are used. A well-know feature selection technique called "Information Gain" is employed to select top 1,000 features.

**Automatic Multimedia Content Processing**

In addition to textual messages, images available on social media could be useful for a number of humanitarian tasks. One of the core tasks, among others, that we are working on is automatic damage assessment from images. Disaster damage severity detection is an important situational awareness task through which humanitarian organizations estimate the scale of the disaster and accordingly plan response and reconstruction efforts. Critical infrastructure that is severely damaged gets high priority while repair or reconstruction. For this purpose, AIDR has been extended to process images on social media to determine the damage severity.

For image processing, AIDR employs state of the art computer vision techniques. The image processing pipeline mostly works as the text processing pipeline works as explained earlier. For image processing, we also employ supervised machine learning techniques that means humans are in the loop to provide some sort of supervision to the machine. The supervision in this case is human-tagged images. The data collection in this case is exactly same as explained earlier. The user specifies a set of keywords or hashtags or defines geographical bounding boxes for tweets collection. Each collected tweet is then checked to find if it contains any image link. In case of an image associated with the tweet, the image is downloaded from the Web. We have observed a large proportion of images posted during disaster events consist of duplicate content. Other than exact duplication, cropped, resized, recolored are other examples of duplicate in which the main scene stays similar to the original image. Furthermore, another big proportion of images on social media is irrelevant showing cartoons, banners, celebrities, advertisements etc.

To tackle the image duplicate and relevancy issues, AIDR have employed state of the art techniques. For the de-duplication of images, we use Perceptual Hashing technique. To determine, an image relevancy, we use Convolutional Neural Networks (CNN). Specifically, we use VGG-16 trained on a large number of labeled images from the ImageNet repository (Jia Deng, et al., 2009). We fine-tuned the same network using our labeled data after adapting the last layer.

After removing noise (i.e., duplicate and irrelevant) from the collected images, next we perform damage assessment. We use three categories to represent damage or not damage as follows: we use SEVERE damage category to represent non-useable buildings, non-drivable road or bridges, we use MILD damage category to represent partially destroyed infrastructure, and NO damage to represent little to no damage. Given these categories, we employed human workers to tag images. Human-tagged images are then used to train CNN classifiers. Specifically, we adapt a transfer learning approach in which pre-trained models using the ImageNet data are fine-tuned using our own labeled data. For this purpose, we sampled 6,000 images from four disaster events namely the 2015 Nepal earthquake, the 2016 Ecuador earthquake, Typhoon Ruby (2014), and Hurricane Matthew (2016). The human-tagged images distribution is as follows: SEVERE=1,765; MILD=483; NO=3,751.

**Critical Information Extraction**

As presented earlier this section that automatic approaches like AIDR help eliminate noise form the data. However, even after removing noisy content, the data may still be too big to be manually processed. Moreover, if the disaster event spans over weeks or months, this could create a filter failure issue. To overcome these issues, we have developed an interactive monitoring tool to help crisis responders and decision makers provide rapid access to the most critical and important situational information during an on-going disaster event. For this purpose, we begin by extracting critical entities mentioned in tweets. These entities represent people, places, or organization names. We use Stanford Named Entity Recognition tool (Finkel, Grenager, & Manning, 2005) for Named-Entity extraction. Furthermore, we combine AIDR's classifiers' output (i.e., classified tweets), use-defined categories, timelines, and map components to help users quickly sift through thousands of tweets. Figure 3 shows the interactive dashboard and its different visual components including timelines (top), machine classifiers (left), top critical entities (people, organizations, locations), and tweets view (bottom).


**Figure 1: Interactive dashboard to quickly find critical and useful information during disasters**

## HURRICANE MARIA: DATA COLLECTION AND DESCRIPTION

We collected Twitter data during the 2017 Category 5 Hurricane Maria that made landfall around 20-September-2017 and caused catastrophic damage in Puerto Rico, Dominica, and Dominica Republic. Around 129 people were killed by the hurricane in total, from which 55 are killed in Puerto Rico and 57 in Dominica. The hurricane caused more than 51 billion dollars worth of damage[4].

To collect messages during the hurricane, we used the AIDR platform. AIDR uses the Twitter streaming API and offers different data collection strategies such as collection by keywords or hashtags, by geographical bounding boxes, by following specific Twitter users. We used the keywords-based data collection approach and formed a query consisting of event-specific keywords including *"Hurricane Maria", "Maria storm", "Maria cyclone", "Tropical storm Maria"* and so on. As a result, in total, we collected around 2 million tweets from 20-September-2017 to 06-October-2017. Figure 1 shows the distribution of tweets over these days. From the collected tweets, we found ~80k tweets with images. Figure 3 shows a few images showing the destruction of the Hurricane Maria.

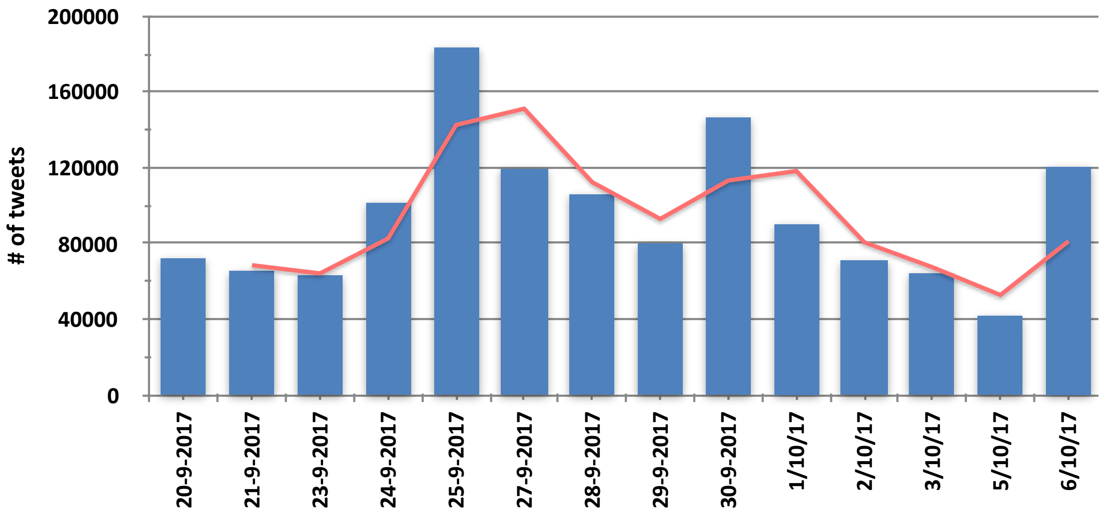


**Figure 2: Hurricane Maria number of tweets per day. The red trend line shows the moving average**

**Figure 3: Sample images collected during Hurricane Maria**

---

[4] https://en.wikipedia.org/wiki/Hurricane_Maria

## EXPERIMENTAL RESULTS & DISCUSSION

We employed AIDR's machine learning classifiers to understand the types of information present in the collected tweets. We trained a text classifier using ~30,000 human-labeled tweets, which were originally collected from a number of past disaster events. Each tweet has been tagged by at least three human workers using the humanitarian categories shown in Table 1.

To train the classifier, we take 80% of the human-tagged data. For this purpose, we use Random Forest learning scheme. The classifier evaluation is performed using a hold out test set consisting of 20% of the labeled data. Table 1 shows the evaluation results in terms of Precision, Recall, and F1. Overall, the classifier achieved an AUC score=0.64, which is a reasonably acceptable performance. However, individually some classes have better scores such as Injured or dead people F1=0.84, Affected individual F1=0.71, whereas some other classes have lower scores such as Relevant information F1=0.50. Largely, a classifier's performance depends on the size of training data available for each class. The dataset in this case is imbalanced, which is one potential reason of low performance for some classes.

**Table 1: Text classifier evaluation results in terms of Precision, Recall, and F1**

| Category name | Precision | Recall | F1-score |
|---|---|---|---|
| Affected individual | 0.71 | 0.71 | 0.71 |
| Caution and advice | 0.61 | 0.64 | 0.62 |
| Donation and volunteering | 0.64 | 0.77 | 0.7 |
| Infrastructure and utilities damage | 0.65 | 0.65 | 0.65 |
| Injured or dead people | 0.82 | 0.87 | 0.84 |
| Personal updates | 0.66 | 0.62 | 0.64 |
| Other relevant information | 0.66 | 0.4 | 0.5 |
| Sympathy and support | 0.54 | 0.76 | 0.63 |
| Not related or irrelevant | 0.66 | 0.72 | 0.69 |

We applied the trained classifier on the whole Hurricane Maria's dataset. Figure 2 shows the prediction results in terms of distribution of tweets into different classes. Consistent with our past observation while analyzing social media data during disasters and emergencies that the amount of irrelevant messages is always higher than the useful messages, in this case we also observed a similar pattern. The "Not related or irrelevant" category contains around 28% of the whole data in the Hurricane Maria case. The "Other relevant information" category emerged as the second largest category, which is also not surprising. The "Other relevant information" category contains messages that do not belong to any other informative class but these messages are important for humanitarian purposes. Research studies have found that this category contains from about 10% to 30% of the messages on social media during disasters. For more information about why AIDR uses this category and why this category could be useful to learn about several small-scale unanticipated events, we refer to these papers (Imran & Castillo, 2015; Imran, Chawla, & Castillo, 2016).

In the remaining categories, the "Donation and volunteering" and "Sympathy and support" contain substantial amounts of messages, specifically 26% and 14%, respectively. Among other relatively small categories, the "Infrastructure and utilities damage" category contain around 3% of the full data.

Next, we wanted to look into these individual categories and find our what kind of information these automatically categorized messages contain. However, we noticed that even after the automatic classification of tweets and filtering-out irrelevant ones, the remaining messages are still in thousands. Manual analysis of these thousands of messages (e.g., ~27k donation-related messages) is a hectic and time-consuming task for crisis responders, especially during an on-going disaster event. To overcome this information overload issue and to help crisis responders sift through enormous data as quick as possible, we have developed an interactive dashboard technology that provides convenient ways to find critical information from big live data. We deployed the dashboard to analyze the Hurricane Maria's data. In below, we provide a few most critical tweets found during data analysis through the dashboard.

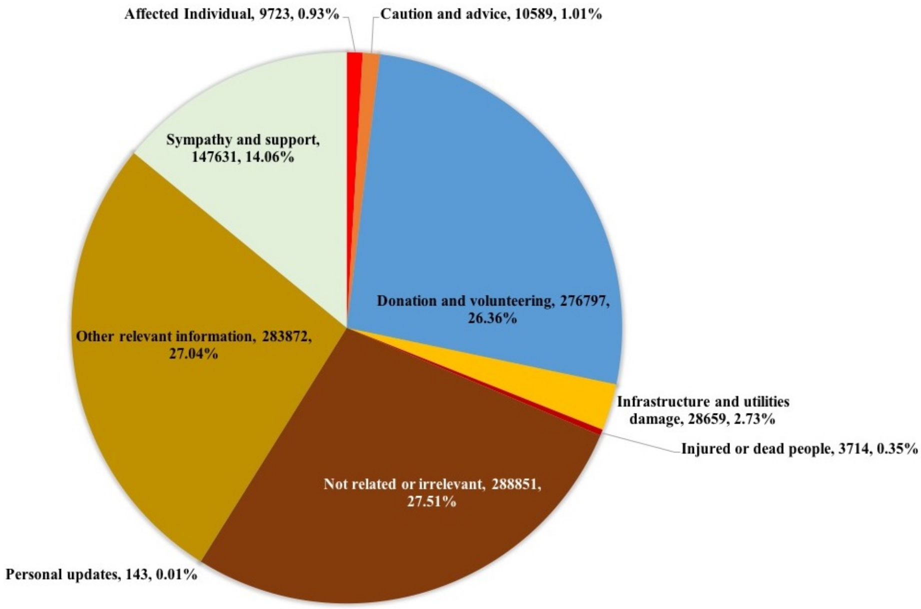


**Figure 4: Tweets distribution in different categories by the text classifier**

**Donations and volunteering requests or offers:**

- *RT @DorothyKidd1: Food, water, power, gasoline desperately needed in Puerto Rico https://t.co/z3tVTc7pAU*
- *Puerto Rico needs our help. Bring diapers, baby food, batteries, first aid supplies, feminine hygiene products to:… https://t.co/KWi6uCWMxz*
- *Our neighbors need our help! Click on the link on our page and donate what you can. Thank you! https://t.co/zyITNsMUHy*

**Infrastructure and Utility Damage reports:**

- *Bavaro, Dominican Republic A gift shop on Cofrecito beach is damaged after being hit by #HurricaneMaria Photograph:… https://t.co/51CmpqAtG2*
- *@Reuters Destroyed homes are seen from a Marine Corps Osprey surveying damage from #HurricaneMaria in St. Croix,… https://t.co/T7Xtc8vslP*
- *RT @youthmappers: There are 70k people downstream of this dam damaged by H.Maria https://t.co/cyki5UMZyq*
- *Hurricane Maria leaves Arecibo radio telescope damaged and dark https://t.co/b1G7EW9BEI via @theregister*

- *Hurricane Maria Damages Dominica's Main Hospital, Leaves 'War Zone' Conditions* *https://t.co/JuCd4XIkW8*
- *RT @NWSPittsburgh: Maria did a lot of damage to our Doppler Radar in San Juan, PR. Here is what a radar should look like.…*
- *San Juan airport remains crippled by Hurricane Maria damage* *https://t.co/167dRw8OkX*
- *Maria's Devastation Damaged Hospitals Running Out of Supplies in U.S. Virgin Islands @weatherchannel https://t.co/BWfZmCRU1y*

**Missing or found people reports:**

- *RT @Jasamsdestiny: Amy is missing – plz help find her. #Miami #southmiamiheights #palmettoestates #miamidadecounty #MiamiBeach…*
- *#hurricanemaria QUESTELL FAMILY MISSING 5 DAYS IN PONCE. PLEASE HELP FIND THEM*

During the Hurricane Maria, AIDR's image processing pipeline was also activated to identify images that show infrastructure damage. In total, we found around 80k tweets containing images. However, ~75% of these images were found duplicate. The remaining 25% (~20k) of the images were automatically classified by the AIDR's damage assessment classifier. As stated above, we use three classes to represent the extent of damage in an image namely: SEVERE damage, MILD damage, and NO damage. From the predictions, we observe that almost 78% of the images we collected on Twitter do not show any damage, mainly because these images depict hurricane paths, people, and other info-graphics rather than built structures. In Figure 5, the first row shows some examples of such images. From the predictions, 10.9% of images show MILD damage and 11.2% of the images show SEVERE damage. Figure 5 shows a few images automatically classified by our image classifier into the three damage classes.

**AIDR Damage Severity Assessment for Hurricane Maria Images from Twitter**

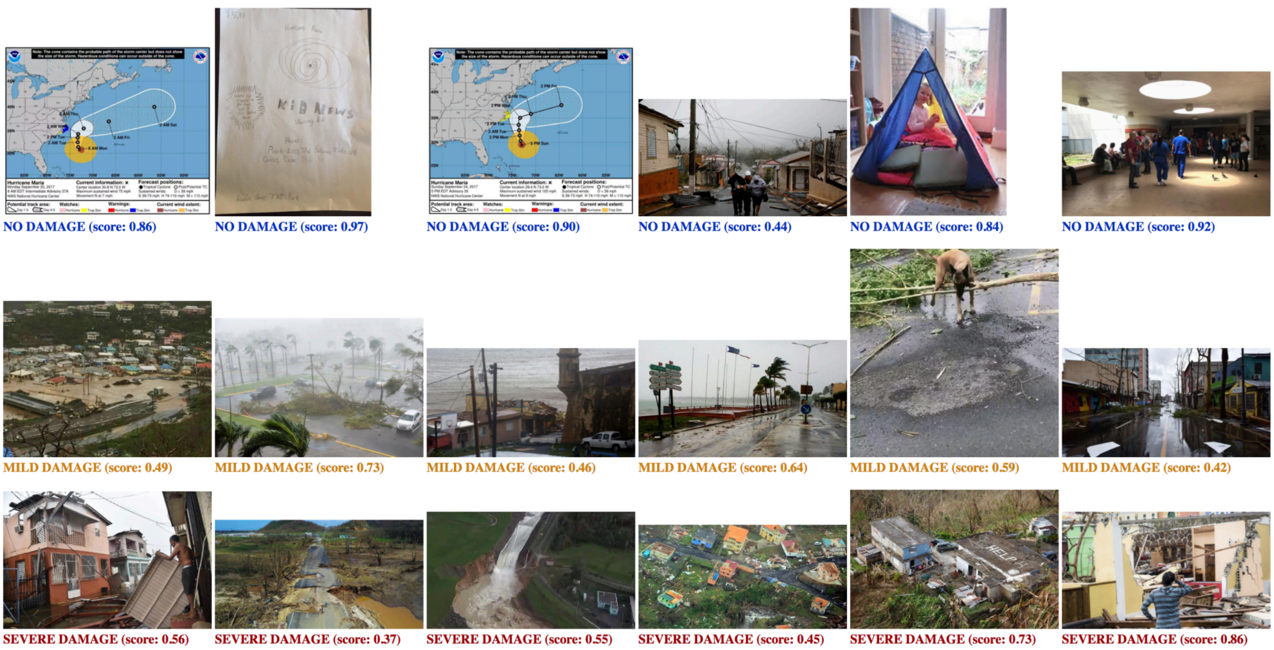


**Figure 5: Damage assessment results using images collected during the Hurricane Maria**

By looking at the actual image content, we can extract more information about the devastation caused by the disaster than relying solely on the textual content provided by the users. The image processing pipeline could be useful for a number of humanitarian use cases. For instance, with the automatic damage assessment, instead of trying to look at all the images, humanitarian organizations and emergency responders can simply take a look at the images containing MILD or SEVERE damage to quickly understand sense of the level of destruction incurred by the disaster.

## CONCLUSIONS

Humanitarian organizations rely on credible and timely information to act against natural disasters and emergencies. Social media has been emerged as an alternative information source that contains information that directly comes from affected people. This information consists variety of useful reports such as reports of urgent needs of affected people and other situational information. However, real-time processing of social media data is a challenging task. In this work, we have presented a number of technologies that employ state of the art artificial intelligence techniques to collect and process social media information in real-time to aid humanitarian organizations.

## REFERENCES


Alam, F., Imran, M., & Ofli, F. (2017). Image4Act: Online Social Media Image Processing for Disaster Response. In *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*.

Aupetit, M., & Imran, M. (2017). Interactive Monitoring of Critical Situational Information on Social Media. In *International Conference on Information Systems for Crisis Response And Management (ISCRAM)*.

Boersma, K., Fonio, C., Imran, M., & Meier, P. (2017). *Big Data, Surveillance and Crisis Management*. Routledge.

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

Daly, S., & Thom, J. A. (2016). Mining and Classifying Image Posts on Social Media to Analyse Fires. *13th International Conference on Information Systems for Crisis Response and Management*, (May), 1–14.

Finkel, J., Grenager, T., & Manning, C. (2005). Incorporating non-local information into information extraction systems by gibbs sampling. *... of the 43rd Annual Meeting on ...*, (June), 363–370. https://doi.org/10.3115/1219840.1219885

Hughes, A. L., & Palen, L. (2009). Twitter adoption and use in mass convergence and emergency events. *International Journal of Emergency Management*, *6*(3/4), 248. https://doi.org/10.1504/IJEM.2009.031564

Hwang, W.-S., Park, J., & Kim, S.-W. (2015). A method for recommending the latest news articles via MinHash and LSH. In *Proceedings of the 9th International Conference on Ubiquitous Information Management and Communication - IMCOM '15* (pp. 1–6). https://doi.org/10.1145/2701126.2701205

Imran, M., & Castillo, C. (2015). Towards a Data-driven Approach to Identify Crisis-Related

Topics in Social Media Streams. In *Proceedings of the 24th International Conference on World Wide Web Companion* (pp. 1205–1210).

Imran, M., Castillo, C., Diaz, F., & Vieweg, S. (2015). Processing social media messages in mass emergency: A survey. *ACM Computing Surveys (CSUR)*, *47*(4), 67.

Imran, M., Castillo, C., Lucas, J., Meier, P., & Rogstadius, J. (2014). Coordinating Human and Machine Intelligence to Classify Microblog Communications in Crises. In *11th International Conference on Information Systems for Crisis Response and Management*.

Imran, M., Castillo, C., Lucas, J., Meier, P., & Sarah, V. (2014). AIDR: Artificial Intelligence for Disaster Response. In *In Proceedings of the 23rd International World Wide Web Conference (WWW '14)*.

Imran, M., Chawla, S., & Castillo, C. (2016). A Robust Framework for Classifying Evolving Document Streams in an Expert-Machine-Crowd Setting. In *In proceedings of the 16th International Conference on Data Mining (ICDM)*.

Imran, M., Meier, P., Castillo, C., Lesa, A., & Herranz, M. G. (2016). Enabling Digital Health by Automatic Classification of Short Messages. In *In Proc. of the ACM Digital Health 2016 Conference*.

Jia Deng, Wei Dong, Socher, R., Li-Jia Li, Kai Li, & Li Fei-Fei. (2009). ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition* (pp. 248–255). https://doi.org/10.1109/CVPRW.2009.5206848

Nguyen, D. T., Al-Mannai, K., Joty, S. R., Sajjad, H., Imran, M., & Mitra, P. (2017). Robust Classification of Crisis-Related Data on Social Networks Using Convolutional Neural Networks. In *ICWSM* (pp. 632–635).

Nguyen, D. T., Alam, F., Ofli, F., & Imran, M. (2017). Automatic Image Filtering on Social Networks Using Deep Learning and Perceptual Hashing During Crises. *Info. Sys for Crisis Response and Management (ISCRAM).*

Nguyen, D. T., Joty, S., Imran, M., Sajjad, H., & Mitra, P. (2016). Applications of Online Deep Learning for Crisis Response Using Social Media Information. *Social Web for Disaster Management (SWDM 2016) - Co-Located with CIKM 2016 - arXiv Preprint arXiv:1610.01030*.

Nguyen, D. T., Ofli, F., Imran, M., & Mitra, P. (2017). Damage Assessment from Social Media Imagery Data During Disasters. In *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*.

Ofli, F., Meier, P., Imran, M., Castillo, C., Tuia, D., Rey, N., … others. (2016). Combining Human Computing and Machine Learning to Make Sense of Big (Aerial) Data for Disaster Response. *Big Data Journal*, *4*(1), 47–59.

Rudra, K., Banerjee, S., Ganguly, N., Goyal, P., Imran, M., & Mitra, P. (2016). Summarizing situational tweets in crisis scenario. In *Proceedings of the 27th ACM Conference on Hypertext and Social Media* (pp. 137–147).

Sasikumar, U. (2014). A Survey of Natural Language Question Answering System. *International Journal of Computer Applications*, *108*(15), 975–8887.

Vieweg, S., Hughes, A. L., Starbird, K., & Palen, L. (2010). Microblogging During Two Natural Hazards Events : What Twitter May Contribute to Situational Awareness. *CHI 2010: Crisis Informatics*, 10. https://doi.org/10.1145/1753326.1753486

Whipkey, K., & Verity, A. (2015). Guidance for incorporating Big Data into humanitarian operations. *Digital Humanitarian Network Document*.