

Image4Act: Online Social Media Image Processing for Disaster Response

Firoj Alam, Muhammad Imran, Ferda Ofli
Qatar Computing Research Institute, HBKU, Doha, Qatar
{fialam, mimran, fofli}@hbku.edu.qa

Abstract—We present an end-to-end social media image processing system called *Image4Act*. The system aims at collecting, denoising, and classifying imagery content posted on social media platforms to help humanitarian organizations in gaining situational awareness and launching relief operations. It combines human computation and machine learning techniques to process high-volume social media imagery content in real time during natural and human-made disasters. To cope with the noisy nature of the social media imagery data, we use a deep neural network and perceptual hashing techniques to filter out irrelevant and duplicate images. Furthermore, we present a specific use case to assess the severity of infrastructure damage incurred by a disaster. The evaluations of the system on existing disaster datasets as well as a real-world deployment during a recent cyclone prove the effectiveness of the system.

I. INTRODUCTION

The extensive use of social media platforms such as Twitter, Facebook, Instagram at the time of mass emergency situations due to natural or human-made disasters has created a number of opportunities for information seekers to gain timely access to valuable insights. During such events, bystanders and affected people post situational updates including reports of injured or dead people, infrastructure damage, requests for urgent needs such as food, water, shelter, donation offers and so on. This online data on social networks arrive in a variety of forms such as textual messages, images, and videos [1]–[4]. Rapid access to these critical and situation-sensitive updates through social media networks is useful for a number of real-world applications and can also help to fulfill various information needs [5]–[7].

Among other applications, rapid crisis response and management is the focus of this work. Formal humanitarian organizations, law enforcement agencies, and other volunteer groups look for timely information to gain situational awareness and to plan relief operations [8]–[10]. Research studies show the importance of social media data for an effective crisis response [1], [3]. Moreover, a number of techniques

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ASONAM '17, July 31 - August 03, 2017, Sydney, Australia

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-4993-2/17/07?/\$15.00

<http://dx.doi.org/10.1145/3110025.3110164>



Fig. 1: Relevant (first two columns) and irrelevant images (third column) collected during different disasters on Twitter.

based on artificial intelligence and machine learning have been developed to process this data [2]. However, the majority of these studies focus primarily on analyzing the textual content, ignoring the rich information provided by the visual content. In this work, we present a system to address this limitation by introducing a real-time social media image processing pipeline to help humanitarian organizations enhance disaster response and management operations [11].

Analyzing high-volume and high-velocity social media imagery data in real time is a challenging task. Moreover, a large portion of social media images contains redundant or irrelevant content, and hence, results in a low signal-to-noise ratio. Figure 1 shows examples of useful relevant images in the first two columns, and irrelevant images (e.g., cartoons, banners, celebrities) in the third column. Therefore, prior to in-depth analysis of the visual social media content to extract actionable information for humanitarian organizations, the raw social media imagery data needs to be filtered from this noise. In this paper, we present an online image processing pipeline that comprises of de-duplication and relevancy filtering modules to collect and filter social media images in real-time during a crisis event. The system combines human computation and

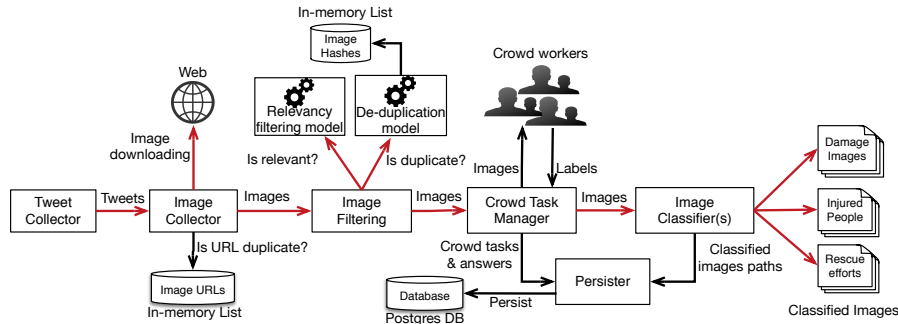


Fig. 2: *Image4Act* system architecture for online social media image processing pipeline.

machine learning techniques to process high-velocity Twitter stream. We employ Stand-By-Task-Force (SBTF)¹ volunteers at the time of emergencies to help in tagging images, which are then used to train machine learning models for specific humanitarian use cases, e.g., damage assessment. Removing duplicate and irrelevant images is crucial before employing humans for image annotation to save crowdsourcing budget both in terms of time and money. That is, human volunteers should not waste their time in tagging irrelevant or duplicate images repeatedly. To detect duplicate and irrelevant images, we use deep neural networks and perceptual hashing techniques. Furthermore, the denoised stream of images are analyzed to assess the severity of infrastructure damage shown in the images. The system evaluation performed on existing disaster datasets and a deployment during a recent real-world crisis event (Cyclone Debbie that hit Queensland, Australia in March 2017) demonstrates its effectiveness. The *Image4Act* system is integrated with the AIDR system [12], which can be accessed at <http://aidr.qcri.org/>.

II. IMAGE PROCESSING PIPELINE

Figure 2 depicts the modules of the proposed *Image4Act* system. The modules communicate (i.e., data flow) with each other using Redis channels². Moreover, each module has a set of RESTful APIs to enable external interactions (e.g., UI interactions) and to set parameter values, if required. Red arrows in the figure represent live streams carrying data items, whereas black arrows show non-streaming communications. The system is implemented using the Java Enterprise Edition (J2EE) programming language.

A. Tweet Collector

The *Tweet Collector* module is responsible for collecting live tweets from the Twitter streaming API³. To create a collection for a specific event (e.g., an earthquake), the user specifies keywords, hashtags, geographical bounding boxes, and/or Twitter users. In the case of geographical bounding box option, only geo-tagged tweets are collected, however, one can use both the keywords and bounding box options to get tweets matching either one of the two criteria.

¹<http://www.standbytaskforce.org/>

²<https://redis.io/>

³<https://dev.twitter.com/streaming/overview>

B. Image Collector

Tweets collected by the *Tweet Collector* are ingested by the *Image Collector* module to extract image URLs from the collected tweets. Each URL is checked against an in-memory key-value pair database, which stores unique URLs, for de-duplication purposes. Images with unique URLs are queued to be downloaded from the Web and published to a Redis channel. All subscribers of this channel immediately receive images as they are downloaded.

C. Image Filtering

Modeling *relevancy* is a challenging problem, as the context of relevancy varies across disaster events, humanitarian organizations, and even within a long-running event (e.g., wars, conflicts). On the contrary, what is deemed *irrelevant* seems consistent across disasters and to many humanitarian organizations. That is, images showing cartoons, celebrities, advertisements, banners are all examples of irrelevant content, hence not useful for disaster response and management. The *Image Filtering* module employs deep neural networks and perceptual hashing techniques to determine whether a newly arrived image (i) is relevant for a given disaster response context and (ii) is not a duplicate of previously collected images. Specifically, we use Convolutional Neural Networks (CNN), VGG-16 [13] architecture, in particular, to train a relevancy model using the DeepLearning4J library⁴. For *duplicate* detection, we use the Perceptual Hashing (pHash) technique [14], [15] to compute a pHash value for each image, which is then kept in an in-memory database. A newly arrived image pHash is compared against the stored hashes using the Hamming distance to detect duplicates and near-duplicates.

D. Crowd Task Manager

The *Crowd Task Manager* module is responsible for assigning image tagging tasks to SBTF volunteers. An end-user creates a task, which we also call classifier (more details regarding the classifiers follow in the next section), that consists of a set of classes (e.g., severe damage, mild damage, no damage). The *Crowd Task Manager* shows an image and the list of classes to a human labeler. The labeler selects an

⁴<https://deeplearning4j.org/>

appropriate label for the image, which is then considered as a training example.

E. Image Classifiers

The system allows end-users to define one or more classifiers. A classifier may consist of two (binary) or more classes (multi-class). Human-labeled images obtained from the *Crowd Task Manager* are used to train these user-defined image classifiers. Since several studies in computer vision literature, e.g., [16], [17], have already shown that the features learned by CNNs on generic large-scale visual recognition tasks (i.e., millions of parameters, trained on millions of images from the ImageNet dataset [18]) are proven to be transferable and effective when used in other specific tasks, particularly when training data are limited. For example, in our case, it is very early stages of data collection. We adopt a transfer learning approach where we use the VGG-16 network [13] pre-trained on ImageNet data set as an initialization for fine-tuning the same network on our own training dataset. We also adapt the last layer of the network to comply with the number of classes specified by the user instead of the original 1,000-class classification. Hence, this transfer learning approach allows us to transfer the features and the parameters of the network from the broad domain (i.e., large-scale image classification) to the specific one (i.e., relevant-image classification).

F. Persister

The *Persister* module is responsible for all database-specific operations such as insertion of images’ meta-data, storage, and retrieval of models’ predictions. Moreover, it also persists machine-tagged images into the file system.

III. DATASET AND SYSTEM EVALUATION

Datasets: Images posted on Twitter during four natural disasters were used for the evaluation of the proposed *Image4Act* system. Table I shows details of the events and label distribution obtained using human-volunteers. The original crowdsourcing task was to determine whether an image shows any *severe damage*, *mild damage*, or *no damage*. To train the relevancy filter (binary classifier), 3,518 images were randomly selected from the *severe* and *mild* categories and considered as *relevant*. Images in the *none* category were labeled by the ImageNet 1,000-class VGG-16 model [13], and with human-supervision, top 12 most-frequent, irrelevant categories were selected (e.g., website, suit, lab coat, menu, etc.). In total 3,518 images were obtained from the *none* category, all deemed as *irrelevant*. We used 60:20:20 split as our training, validation, and test sets, respectively.

Relevancy Evaluation: Given we have limited labeled data, we followed a transfer learning approach to initialize the VGG-16 network [13] using the weights of the original ImageNet [18] model. We then fine-tuned the same network using our labeled data after adapting the last (i.e., softmax) layer of the network for the binary classification task. We achieved an AUC = 0.98.

TABLE I: Datasets – **NE**: Nepal Earthquake, **EE**: Ecuador Earthquake, **TR**: Typhoon Ruby, **HM**: Hurricane Matthew.

| Class | NE | EE | TR | HM | Total |
|--------------|---------------|--------------|--------------|------------|---------------|
| Severe | 8,927 | 955 | 88 | 110 | 10,080 |
| Mild | 2,257 | 89 | 338 | 94 | 2,778 |
| None | 14,239 | 946 | 6,387 | 132 | 21,704 |
| Total | 25,423 | 1,990 | 6,813 | 336 | 34,562 |

De-duplication Evaluation: The performance of the de-duplication filter depends on the Hamming distance threshold value. To determine this value, we randomly sampled 1,100 images and computed their pHashes. Image pairs with a distance between 0 to 20 were manually examined and learned that $d \leq 10$ is the optimal distance value at which the system maintains an accuracy of 0.98.

Damage severity assessment: To obtain the training data for the damage assessment task, we applied the de-duplication and relevancy model on the raw data and obtain a cleaned set from which a random sample of 6k images was selected (severe=1,765; mild=483; none=3,751). Following the classifier details described in Section II-E, we trained a damage assessment model using 60:20:20 split as training, validation, and test sets, respectively. We achieved an AUC = 0.72 on average, which is reasonably in the acceptable range.

Deployment during a Real-World Disaster: The system was deployed during the recent Cyclone Debbie that made landfall in Queensland Australia on March 28th, 2017. Out of 76k tweets that were collected from 28/03/2017 to 03/04/2017 in real-time using the hashtag #CycloneDebbie, 7k tweets were with images. We used the already trained relevancy model to predict whether an incoming image is relevant or not. For the evaluation, 500 machine-classified images were sampled and examined by the authors of this paper. The precision scores were 0.67 and 0.92 for the irrelevant and duplicate cases, respectively. Figure 3 shows examples of relevant and irrelevant images automatically classified by the system during the ongoing disaster. It can be clearly seen that the system was successfully able to detect both relevant and irrelevant images.

IV. RELATED WORK

There has been a very few works on image processing for disaster response. Most of the work in the literature are from the remote-sensing domain. In [19], Liu et al. present the utility of images for disaster response, in which they collected images using FORMOSAT-2 satellite during South Asia tsunami in 2004. Studies of satellite image processing for crisis response also include [20]–[22] and [23]. Similar studies in remote sensing research area show the evidences of damage level assessment from aerial [5], [6] and satellite [24], [25] images collected from disaster-hit regions.

The importance of social media images for disaster management has been recently highlighted in [10]. The authors analyzed tweets and messages from Flickr and Instagram for the flood event in Saxony (2013). They found that the existence of images within on-topic messages are more relevant to

REFERENCES

- [1] K. Starbird, L. Palen, A. L. Hughes, and S. Vieweg, "Chatter on the red: what hazards threat reveals about the social life of microblogged information," in *CSCW*, 2010, pp. 241–250.
- [2] M. Imran, C. Castillo, F. Diaz, and S. Vieweg, "Processing social media messages in mass emergency: A survey," *ACM Comput Surv (CSUR)*, vol. 47, no. 4, p. 67, 2015.
- [3] C. Castillo, *Big Crisis Data*. Cambridge University Press, 2016.
- [4] F. Offi, P. Meier, M. Imran, C. Castillo, D. Tuia, N. Rey, J. Briant, P. Millet, F. Reinhard, M. Parkan *et al.*, "Combining human computing and machine learning to make sense of big (aerial) data for disaster response," *Big data*, vol. 4, no. 1, pp. 47–59, 2016.
- [5] M. Turker and B. T. San, "Detection of collapsed buildings caused by the 1999 izmit, turkey earthquake through digital analysis of post-event aerial photographs," *Int J Remote Sens*, pp. 4701–4714, 2004.
- [6] J. Fernandez Galarreta, N. Kerle, and M. Gerke, "UAV-based urban structural damage assessment using object-based image analysis and semantic reasoning," *Natural Hazards and Earth System Science*, vol. 15, no. 6, pp. 1087–1101, 2015.
- [7] S. Daly and J. A. Thom, "Mining and classifying image posts on social media to analyse fires," in *ISCRAM*, May 2016, pp. 1–14.
- [8] M. Imran, S. M. Elbassuoni, C. Castillo, F. Diaz, and P. Meier, "Extracting information nuggets from disaster-related messages in social media," in *ISCRAM*, 2013.
- [9] K. Rudra, S. Banerjee, N. Ganguly, P. Goyal, M. Imran, and P. Mitra, "Summarizing situational tweets in crisis scenario," in *ACM Conference on Hypertext and Social Media*, 2016, pp. 137–147.
- [10] R. Peters and P. d. A. Joao, "Investigating images as indicators for relevant social media messages in disaster management," in *ISCRAM*, 2015.
- [11] D. T. Nguyen, F. Alam, F. Offi, and M. Imran, "Automatic image filtering on social networks using deep learning and perceptual hashing during crises," in *ISCRAM*, May 2017.
- [12] M. Imran, C. Castillo, J. Lucas, P. Meier, and S. Vieweg, "AIDR: Artificial intelligence for disaster response," in *Proceedings of WWW Conference*. ACM, 2014, pp. 159–162.
- [13] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [14] Y. Lei, Y. Wang, and J. Huang, "Robust image hash in radon transform domain for authentication," *Signal Processing: Image Communication*, vol. 26, no. 6, pp. 280–288, 2011.
- [15] C. Zauner, "Implementation and benchmarking of perceptual image hash functions," in *PhD Dissertation*. Fraunhofer Institute for Secure Information Technology, 2010.
- [16] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" in *NIPS*, 2014, pp. 3320–3328.
- [17] G. Ozbulak, Y. Aytar, and H. K. Ekenel, "How transferable are cnn-based features for age and gender classification?" in *International Conference of the Biometrics Special Interest Group*, Sept 2016, pp. 1–6.
- [18] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *Int J Comput Vision*, vol. 115, no. 3, pp. 211–252, 2015.
- [19] C.-C. Liu, J.-G. Liu, C.-W. Lin, A.-M. Wu, S.-H. Liu, and C.-L. Shieh, "Image processing of formosat-2 data for monitoring the south asia tsunami," *Int J Remote Sens*, vol. 28, no. 13-14, pp. 3093–3111, 2007.
- [20] S. Voigt, T. Kemper, T. Riedlinger, R. Kiefl, K. Scholte, and H. Mehl, "Satellite image analysis for disaster and crisis-management support," *IEEE T Geosci Remote*, vol. 45, no. 6, pp. 1520–1528, 2007.
- [21] S. M. Adams and C. J. Friedland, "A survey of unmanned aerial vehicle (uav) usage for imagery collection in disaster research and management," Louisiana State University, Tech. Rep., 2011.
- [22] T. Lillesand, R. W. Kiefer, and J. Chipman, *Remote sensing and image interpretation*. John Wiley & Sons, 2014.
- [23] I. Nourbakhsh, R. Sargent, A. Wright, K. Cramer, B. McClendon, and M. Jones, "Mapping disaster zones," *Nature*, pp. 787–788, 2006.
- [24] M. Pesaresi, A. Gerhardinger, and F. Haag, "Rapid damage assessment of built-up structures using vhr satellite data in tsunami-affected areas," *Int J Remote Sens*, vol. 28, no. 13-14, pp. 3013–3036, 2007.
- [25] T. Feng, Z. Hong, Q. Fu, S. Ma, X. Jie, H. Wu, C. Jiang, and X. Tong, "Application and prospect of a high-resolution remote sensing and geo-information system in estimating earthquake casualties," *Natural Hazards and Earth System Sciences*, vol. 14, no. 8, pp. 2165–2178, 2014.

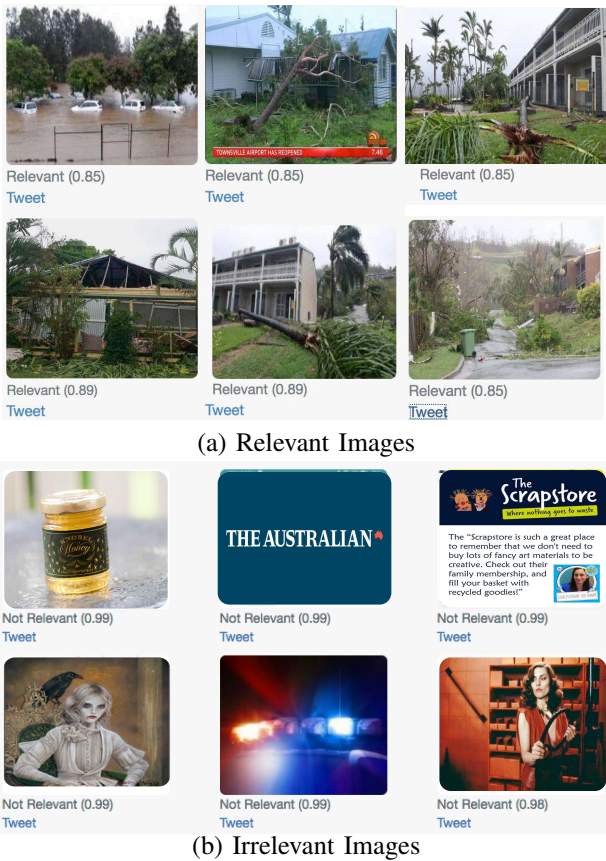


Fig. 3: Sample images processed by *Image4Act* during Cyclone Debbie with the classifier confidence scores.

the disaster event, and the image content can also provide important information related to the event. In another study, Daly et al. [7] focused on classifying images extracted from social media data, i.e, Flickr, and analyzed whether a fire event occurred at a particular time and place [7]. Their study also analyzed spatio-temporal meta-data associated with the images and suggested that geotags are useful to locate the fire affected area.

Our work on social media image processing is different compared to the previous studies. For instance, (i) We provide an end-to-end system, which includes collecting images from social media, filtering irrelevant images, removing duplicates, and assessing damage level, (ii) for the image filtering and damage assessment we use deep learning techniques, which are current state-of-art and demonstrate efficient performance.

V. CONCLUSIONS

We presented *Image4Act*, a system that can ingest and process imagery content on Twitter in real-time to help humanitarian organizations to understand the severity of a crisis for better decision-making. The system comprised of two crucial image filtering modules to filter out the noisy content, and to help crisis managers build more fined-grained classifiers, e.g., damage assessment from images, with the help of the crowd workers. We evaluated the system both offline and online during a real-world disaster to demonstrate its effectiveness.