2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)

CitizenHelper-Adaptive: Expert-augmented Streaming Analytics System for Emergency Services and Humanitarian Organizations

Rahul Pandey and Hemant Purohit
Department of Information Sciences and Technology
George Mason University
Fairfax, United States
{rpandey4 & hpurohit}@gmu.edu

***Abstract*—There is an increasing amount of information posted on Web, especially on social media during real world events. Likewise, there is a vast amount of information and opinions posted about humanitarian issues on social media. Mining such data can provide timely knowledge to inform disaster resource allocation for who needs what and where as well as policies for humanitarian causes. However, information overload is a key challenge in leveraging this big data resource for organizations. We present an interactive user-feedback based streaming analytics system '*CitizenHelper-Adaptive*' to mine social media, news, and other public Web data streams for emergency services and humanitarian organizations. The system aims to collect, organize, and visualize the vast amounts of data across various user and content-based information attributes using the adaptive machine learning models, such as intent classification models to continuously identify requests for help or offers of help during disasters. This demonstration shows the first application of transfer-active learning methods for time-critical events, when there is an availability of abundant labeled data from past events but a scarcity of the sufficient labeled data for the ongoing event. The proposed system provides a user interface to solicit expert feedback on the predicted instances from pretrained models and actively learns to improve the models for efficient information processing and organization. Finally, the system regularly updates the predicted information categories in the visualization dashboard. We will demo CitizenHelper-*Adaptive* system for case studies in both mass emergency events and humanitarian related topics such as gender violence using datasets of more than 50 million Twitter messages and news streams collected between 2016 and 2018.**

Keywords—adaptive systems, online streams, social media mining, transfer active learning, humanitarian technology

I. Introduction

Information networks on Web and social media have become the most popular way to share timely information during the time-critical emergency events (Castillo, 2016) or or sharing opinions in any real world events with varied intents (Purohit and Pandey, 2018). With the click of a button, a user can share crucial intent like requesting for help or offering help during mass emergencies or disasters. Yet, the emergency response services face a lot of difficulties finding these intent-specific messages from the noisy data streams by naïve keyword search approaches. Mining behaviors from social media streams as well as intelligence from open Web sources including news streams during the real-world events present an unprecedented opportunity to extract relevant knowledge for improving decision making in the disaster management organizations. This can further enhance the processes of complex humanitarian organizations by enriching information (Meier, 2015).

Existing social and web analytics systems for large-scale humanitarian events and natural hazards are limited and highly dependent on the single data source based analyses, such as UN Global Pulse's monitoring system (United Nations Global Pulse, 2014) and CrisisTracker (Rogstadius et al., 2013). Single data source may not provide diverse perspectives about narratives of events (e.g., Twitter versus blogs versus local news versus global news). Similarly, current social media mining systems for disasters (e.g., ESA (Cameron et al.), AIDR (Imran et al., 2014)) and general events (e.g., BlogTracker (Agarwal et al., 2009), TweetTracker (Kumar et al., 2011), Truthy (McKelvey and Menczer, 2013), Twitris (Sheth et al., 2017)) do not provide features for adaptive social media mining that are essential for studying the highly dynamic nature of the events. Existing systems do not have a generalizable approach and features for soliciting and incorporating feedbacks by experts on old data and essentially adapting to the feedbacks for analyzing the new data in near realtime.

Hence, we propose *CitizenHelper-Adaptive*, an adaptive system which will not only mine the historic event datasets but also provide a medium for effective adaptation and retraining, via active learning with the expert feedbacks using an online visual interface. Figure 1 shows the architecture of our system and Figure 2 shows an illustration of event analysis using the proposed visual interface.

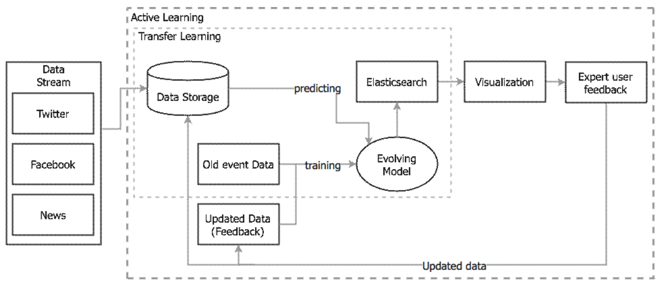

Figure 1. System architecture for user-feedback based transfer-active learning for the streaming analytics pipeline.

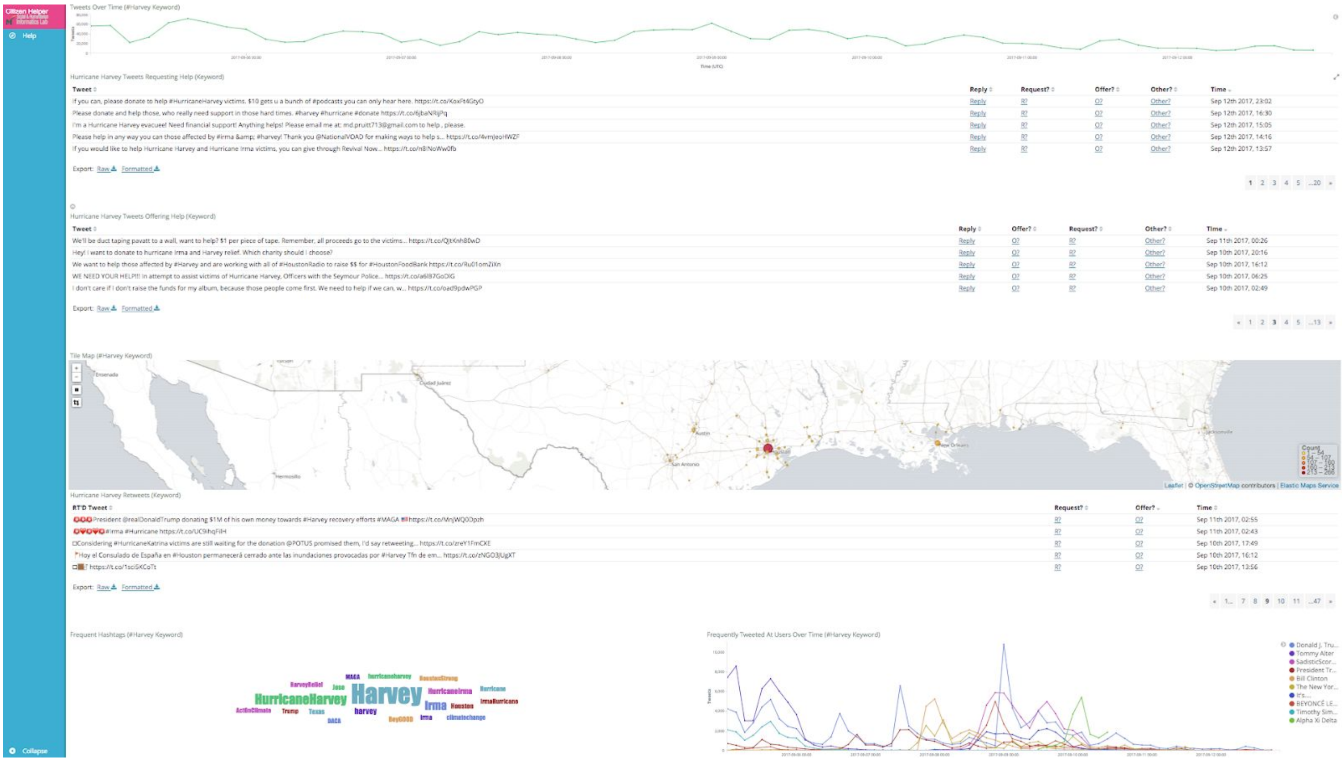

Figure 2. Illustration of analysis using *CitizenHelper-Adaptive* system widgets for *Hurricane Harvey*, a disaster event data stream, which provides insights on participating demographics and user intent classified messages for requesting and offering help. Furthermore, analytical widgets for tag-clouds, geospatial information distribution, and trending users provide additional analytics to understand the evolving context of an event.

The rest of the paper is organised as follows. We first give an overview of our proposed system pipeline, including the detailed description of our transfer-active learning approach. Then the analysis of the proposed tool with the examples of events from both emergency and humanitarian domain. The conclusion summarize our plan for the demo.

II. CitizenHepler-Adaptive System Pipeline

This section describes the pipeline for various features of the proposed *CitizenHelper-Adaptive* system, which is an extension of our prior system (Karuna et al., 2017) for adaptive social media mining capabilities.

A. Data Stream Collection

We use Apache *Kafka*, which is an open source distributed computing platform. The advantage of using *Kafka* is mainly flexibility and scaling, in addition to a streaming data buffer that is valuable for slow downstream processors. *CitizenHelper-Adaptive* supports real-time data collection using social media APIs including but not limited to Twitter. Additionally, it is capable of news collections (including GDELT (Leetaru and Schrodt 2013)) and collecting blog streams as well as static web knowledge bases like Wikipedia, GeoNames, and OpenGov data. We have used Apache *Spark:* an open source processor for data stream, for extracting metadata. These metadata are used to generate features which is then predicted and optimized with the help of transfer and active learning methods. We store the processed data in *Elasticsearch* database that seamlessly support a frontend visualization dashboard, i.e. *Kibana* for streaming analytics.

B. Transfer Learning

The metadata provided by the analytics processors and web services during the process of data collection include the higher-level content and behavioral metadata inference, such as intents like requesting or offering help. For creating such analytical predictors, we used crowdsourcing platform to leverage and label the training data of previous disaster events, which is then used for training models of higher level content analysis such as intent mining methods. We used traditional *Bag of Words* features to train a linear logistic regression model, which is used as a pretrained model for our real-time system. For our incoming data stream, we create the same features and classify them with our pretrained model into multiple classes based on the training sets. For instance, considering the disaster data context, we have three classes: request help or offer help or other. The results are then stored in *Elasticsearch* database and visualized in *Kibana*.

C. Active Learning

When designing the frontend for our streaming analytics, we provide an option for the end users (e.g., public information officer in emergency services) to give feedback for each message whether it is correctly predicted. For instance, considering the disaster data, the prediction for a message on the interface can fall under a request or offer help intent category. Based on the expert feedback provided, we store the current features with their corrected label as the additional training instances of the new event and retrain the pretrained model. Finally, the new model is used to predict the ongoing stream of data with effective predictions.

Hurricane Harvey Tweets Offering Help (Keyword)

Tweet	Reply	Offer?	Request?	Other?
i'm crying so hard, i want to go out and help https://t.co/vxlTy1K3bf	Reply	O?	R?	Other?
Well said. Today, I want to thank first responders who put themselves in danger to help others. #September11 https://t.co/wnKGVUvX52	Reply	O?	R?	Other?
We'll be duct taping pavatt to a wall, want to help? $1 per piece of tape. Remember, all proceeds go to the victims... https://t.co/QJtKnh80wD	Reply	O?	R?	Other?
We want to help those affected by #Harvey and are working with all of #HoustonRadio to raise $$ for #HoustonFoodBank https://t.co/Ru01omZiXn	Reply	O?	R?	Other?
We need to multi-task, save lives and end horrible suffering. So many need help. #Irma #Harvey https://t.co/drx7sfSoVb	Reply	O?	R?	Other?

Figure 3. Case Study #1: [Before] Messages with offering intent as a result of transfer learning on the disaster dataset of hurricane Harvey.

Hurricane Harvey Tweets Offering Help (Keyword)

Tweet	Reply	Offer?	Request?	Other?
i'm crying so hard, i want to go out and help https://t.co/vxlTy1K3bf	Reply	O?	R?	Other?
We'll be duct taping pavatt to a wall, want to help? $1 per piece of tape. Remember, all proceeds go to the victims... https://t.co/QJtKnh80wD	Reply	O?	R?	Other?
We want to help those affected by #Harvey and are working with all of #HoustonRadio to raise $$ for #HoustonFoodBank https://t.co/Ru01omZiXn	Reply	O?	R?	Other?
We need to multi-task, save lives and end horrible suffering. So many need help. #Irma #Harvey https://t.co/drx7sfSoVb	Reply	O?	R?	Other?
Hey! I want to donate to hurricane Irma and Harvey relief. Which charity should I choose?	Reply	O?	R?	Other?

Figure 4. Case Study #1: [After] Messages with offering intent as a result of the transfer and active learning on the disaster dataset of hurricane Harvey.

Since the training instances for our work are mostly human-generated (crowdsourced), the amount is low for optimal training. Hence, our system not only helps in providing a way with updating the current labels but also retraining the model with high quality training instances, and ultimately, results in a robust model.

D. Visualization

We have used *Kibana* for visualizing the data stream. *Kibana* is integrated with *Elasticsearch* database, thus it is easily compatible and provides real-time streaming analytics. Visualization interface allows us to select a particular geography or time range, which can be absolute or relative. Different widgets show metadata of the predicted classes using the transfer-active learners. The visualization consists of 6 analytical widgets as shown in Figure 2. The first widget is a volume-trend graph of incoming messages over time. The next two widgets display all the messages classified with the request help intent and offer help intent respectively. These two widgets also have an option to provide feedback for any of the messages as if they belong to correct classes or not. For an emergency event, it would be request, offer or other classes. The fourth widget shows the word cloud of the most frequent hashtags popular during that time range of event-data stream. The fifth widget shows the multi-line chart demonstrating the frequency of most active users over the time during the selected time range. Finally, the last section includes a geospatial map, showing the frequency-based intensity of the most active location of the incoming data streams around the world. These widgets are refreshed every *t minutes* (t=5) for updating the event analysis visualization with the new data or new predicted labels from the updated predictors.

III. Event Analysis Using *CitizenHelper-Adaptive*

We have tracked and analyzed several disasters and humanitarian related events since August 2016 using the proposed system. We propose two case studies to demonstrate the working of our transfer and active learning based system. *Case Study 1* comprises of the disaster related data of the natural hazard event of *Hurricane Harvey* and we have predicted the request or offer help intent of the incoming tweet messages. We trained the model from the crowdsourced labeled data of *Hurricane Sandy* and classified each incoming tweet message of *Harvey* for a given time span. We provide the snapshots of the "before" and "after" effect of our transfer-active learning feedback approach to retrain and re-predict the incorrect data. In Figure 3, the "before" phase messages for offering help, the second *message* was incorrect in the context of offering help intent. So, a feedback was given by an author for this message with *other intent* category label. After providing such human feedback, the model was re-trained including the newly labeled/corrected messages, followed by re-predicting the whole message set again. After periodic refresh of the visual interface in Kibana, we got new predicted messages for offer help intent as shown in Figure 4 for "after" phase. We can see that now all the messages show offer intent with the first four being the same as Figure 3, but the 5th message with offer intent of a potential donor, who inquires for a charitable organization to donate.

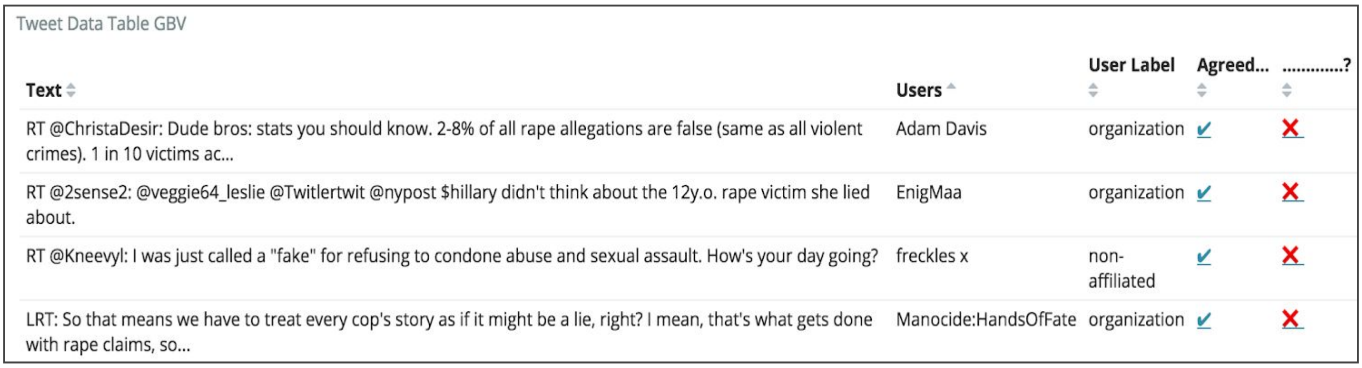

Tweet Data Table GBV

Text	Users	User Label	Agreed...	............?
RT @ChristaDesir: Dude bros: stats you should know. 2-8% of all rape allegations are false (same as all violent crimes). 1 in 10 victims ac...	Adam Davis	organization	✔	✘
RT @2sense2: @veggie64_leslie @Twitlertwit @nypost $hillary didn't think about the 12y.o. rape victim she lied about.	EnigMaa	organization	✔	✘
RT @Kneevyl: I was just called a "fake" for refusing to condone abuse and sexual assault. How's your day going?	freckles x	non-affiliated	✔	✘
LRT: So that means we have to treat every cop's story as if it might be a lie, right? I mean, that's what gets done with rape claims, so...	Manocide:HandsOfFate	organization	✔	✘

Figure 5. Case Study #2: Widget design for soliciting feedback on the metadata of classified user type label, for the analytics to assist humanitarian agencies on the events related to gender violence.

Similarly, for the context of humanitarian-related events, we visualized the analytical results based on the user-type prediction, which relies on the user metadata with a tweet message in an online stream (see Figure 5). We collected gender-violence related *Twitter* data using keyword-based crawling in consultation with social scientists. We crowdsourced the labeled data for user types for creating a user classifier. We created a classifier model to predict the type of users as *organization (official organization), affiliated (individual showing an organizational affiliation), non-affiliated (individual showing no organizational affiliation), and none (bots, fake profile, etc.)*. We visualized the results with our system (Figure 5) and provided feedback for each user classification based on the associated user metadata with their tweets, whether it has been correctly classified and thus, evaluated the iterative approach to learning.

Demo Plan And Conclusion

We will demo the *CitizenHelper-Adaptive* system for multiple events. The system can analyze both humanitarian and disaster event data, with an ability to provide online feedback to transfer and active learning predictors for improving the automated machine classification. The above-discussed case studies analyzed two different datasets with different analytical objectives of classification, first being the mass emergency event data for the understanding of request/offer help intent, and second being the humanitarian event data for the understanding of the nature of user type or information source. We discussed the case studies to show the effect of the active learning system. For the live demo, we will consider different intent mining and user classification approaches across millions of data instances collected from different humanitarian and emergency events and in the background, we will show the analytical effect of transfer-active learning approach in real-time.

Acknowledgement

Authors thank U.S. National Science Foundation (NSF) for partially supporting this research on intent mining through grant award IIS-1657379.

References

[1] S. Kumar, G. Barbier, M. A. Abbasi, and H. Liu, "TweetTracker: An Analysis Tool for Humanitarian and Disaster Relief", in *Proceedings of International Conference on Weblogs and Social Media*, pp. 661–662, 2011.

[2] K. Leetaru and P. A. Schrodt, "Gdelt: Global data on events, location, and tone, 1979–2012", *ISA Annual Convention*, vol. 2, no. 4, pp 1–49, 2013.

[3] P. Meier, *Digital humanitarians: how big data is changing the face of humanitarian response*, CRC Press, 2015.

[4] J. Rogstadius, M. Vukovic, C. Teixeira, V. Kostakos, E. Karapanos, and J. A. Laredo, "Crisistracker: Crowdsourced social media curation for disaster awareness", *IBM Journal of Research and Development*, vol. 57, no. 5, pp. 4–1, 2013.

[5] C. Castillo, *Big crisis data: Social media in disasters and time-critical situations*, Cambridge University Press, 2016.

[6] A. Sheth, H. Purohit, G. A. Smith, J. Brunn, A. Jadhav, P. Kapanipathi, C. Lu, and W. Wang, "Twitris: A system for collective social intelligence", in *Encyclopedia of Social Network Analysis and Mining*, Reda Alhajj, Jon Rokne: Springer, 2017, DOI: 10.1007/978-1-4614-7163-9_345-1.

[7] M. Imran, C. Castillo, J. Lucas, P. Meier, and S. Vieweg, "AIDR: Artificial intelligence for disaster response", in *Proceedings of ACM International Conference on World Wide Web*, pp. 159–162, 2014.

[8] N. Agarwal, S. Kumar, H. Liu, and M. Woodward, "BlogTrackers: A Tool for Sociologists to Track and Analyze Blogosphere", in *Proceedings of International AAAI Conference on Web and Social Media*, 2009.

[9] K. R. McKelvey and F. Menczer, "Truthy: Enabling the study of online social networks", in *Proceedings of ACM International Conference on Computer Supported Cooperative Work Companion*, pp. 23–26, 2013.

[10] M. A. Cameron, R. Power, B. Robinson, and J. Yin, "Emergency situation awareness from twitter for crisis management", in *Proceedings of ACM International Conference on World Wide Web*, pp. 695–698, 2012.

[11] P. Karuna, M. Rana, and H. Purohit, "CitizenHelper: A Streaming Analytics System to Mine Citizen and Web Data for Humanitarian Organizations", in *Proceedings of International AAAI Conference on Web and Social Media*, pp. 729-730, 2017.

[12] H. Purohit, and R. Pandey, "Intent Mining for the Good, Bad & Ugly Use of Social Web: Concepts, Methods, and Challenges", in *Emerging Research Challenges and Opportunities in Computational Social Network Analysis and Mining*, N. Agarwal, N. Dokoohaki, S. Tokdemir: Springer, 2018 (in press), DOI: 0.1007/978-3-319-94105-9. [Online]. Available: http://ist.gmu.edu/~hpurohit/informatics-lab/papers/snam-intent.pdf

[13] Tweets About Global Development Topics.[Online]. Available: http://post2015.unglobalpulse.net/