Interactive Monitoring of Critical Situational Information on Social Media

Michael Aupetit
Qatar Computing Research Institute, HBKU
Doha, Qatar
maupetit@hbku.edu.qa

Muhammad Imran
Qatar Computing Research Institute, HBKU
Doha, Qatar
mimran@hbku.edu.qa

ABSTRACT

According to many existing studies, the data available on social media platforms such as Twitter at the onset of a crisis situation could be useful for disaster response and management. However, making sense of this huge data coming at high-rate is still a challenging task for crisis managers. In this work, we present an interactive social media monitoring tool that uses a supervised classification engine and natural language processing techniques to provide a detailed view of an on-going situation. The tool allows users to apply various filtering options using interactive timelines, critical entities, and other logical operators to get quick access to situational information. The evaluation of the tool conducted with crisis managers shows its significance for situational awareness and other crisis management related tasks.

Keywords

Social media, disaster management, information visualization

INTRODUCTION

The role of social networks such as Twitter and Facebook to rapidly monitor different events and activities happening around the world is increasingly acknowledged. During natural or man-made disasters, people use social media platforms to get latest updates about the disaster event, post situational information, report suspicious activities around them, ask for help, and also offer help (Starbird et al. 2010; Cameron et al. 2012; Hughes and Palen 2009). This online information, in the form of textual messages, imagery content, and videos, contains valuable information useful for humanitarian organizations to gain early insights from an on-going situation (Imran, Castillo, Diaz, et al. 2015). If processed timely and effectively, studies have shown the utility of this online content for disaster response and management (Nguyen et al. 2016; Rudra et al. 2016).

Despite its usefulness, consuming and making sense of large amounts of social media content is a challenging task due to a number of reasons. For instance, a big proportion of social media content can be considered as noisy data due to irrelevant and duplicate messages. Many approaches, based on Natural Language Processing (NLP), have been proposed in the past. Despite advances in NLP techniques, automatic classification of tweets to identify their usefulness is still challenging because tweets are short (i.e. 140 characters only), often misspelled, contain slangs, abbreviations, and informal language (Imran, Mitra, et al. 2016). Among others, one effective approach to process tweets is through supervised classification techniques. In this case, a user decides categories of interest to which he/she wants tweets to be assigned to. For instance, such categories could be *report of injured or dead people, urgent needs, missing person, critical infrastructure damage* etc.

However, even if an automatic classification system is in place, making sense of the large amounts of classified messages (several thousand) in each category of interest requires substantial manual efforts. Moreover, if the disaster event spans over days and weeks, this could create a filter-failure issue. To overcome these issues, in this work we propose an interactive monitoring tool to provide rapid access to the most *critical* and *important situational information* during an on-going disaster event. Based on our interactions and learning from many crisis responders, we observe that the most critical information needs of humanitarian organizations or affected people, have certain high-level themes (e.g. categories) and information within each category is centered around some certain entities. For instance, in the case of an earthquake or a typhoon, such critical entities could be the locations where critical

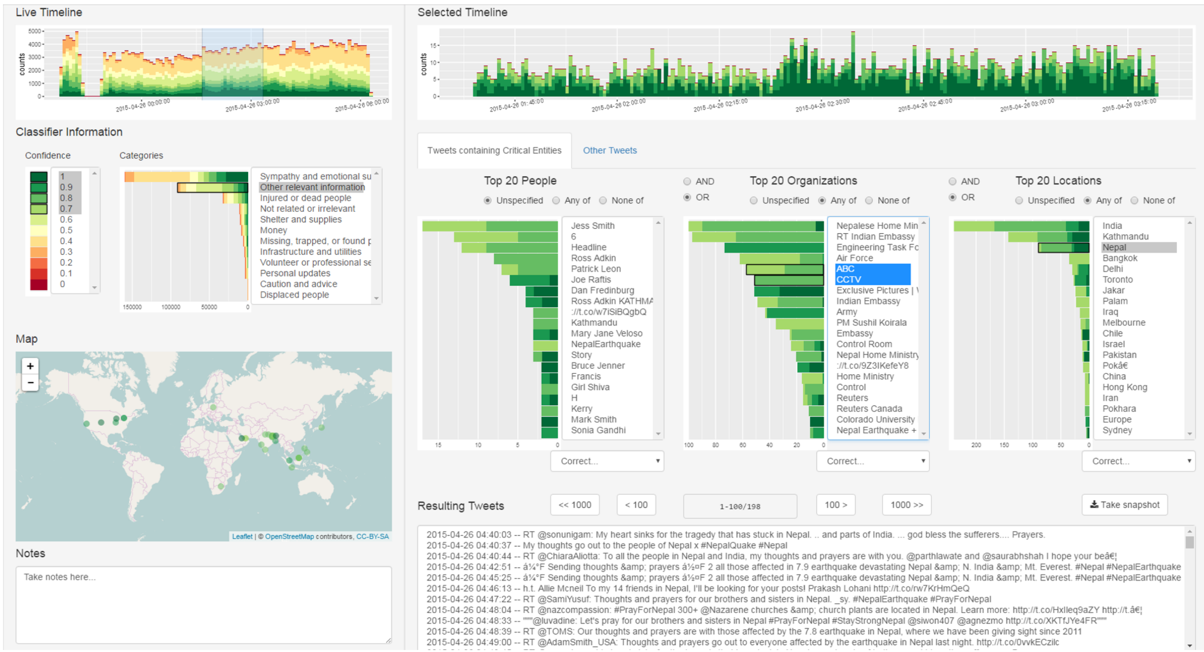

Figure 1. Interactive Monitoring tool full view showing various components like interactive timelines, top entities, map, classifier's options, and filter-view of tweets

infrastructure damage has happened or organizations that are providing aid (e.g. shelter, food, water) to affected people, or names of missing persons, a list of donation providers and so on.

To fulfill such diverse information needs, we employ supervised classification techniques to classify data into different categories of interest of a user. This step filters out irrelevant content. Classified data (i.e., only relevant messages) are then used to extract critical entities using state-of-the-art named-entity recognition techniques. To find most critical entities within each category, we combine our classifiers' classification confidence scores with the proliferation of entities observed over a time-window to fetch top-N most important results. Furthermore, the tool provides an interactive interface to dynamically select a category, a range of confidence scores, and a number of entities to present a filtered view of the messages. Moreover, the tool provides ways to take multiple snapshots of different filtered-views with notes for future reference, which is one of the requirements we learned from crisis managers. Figure 1 shows a full view of the interactive tool. Once the work on this tool is completed, we will make it online at `http://aidr.qcri.org`.

In the sequel, we first present main types of critical information learned from crisis managers. We then describe our interactive tool for rapid monitoring of emergencies. System evaluation details and discussion is presented after that. Related work followed by conclusions and future work directions are presented in the end.

UNDERSTANDING THE TYPES OF CRITICAL INFORMATION

Data, Accessibility, and Interaction with Crisis Managers

People use social media platforms for a variety of purposes. Other than being social and to be in contact with their family and friends, social media platforms, in particular microblogging sites such as Twitter, are ideal to get latest updates on an event. In Twitter, one can search the entire Twitter data repository using simple keywords. Moreover, many social media platforms provide automated ways to consume live information (i.e. real-time access to live posts). For instance, Twitter has a streaming API, which once subscribed using some keywords or hashtags, provides real-time access to posts which match the provided keywords or hashtags.

Such accessibility to live data creates numerous opportunities for humanitarian organizations as well as for affected or concerned people to consume live information during a disaster event. The value of such information becomes more important when no other means of information such as traditional News media, radio etc. are available, which is typically the case in the first few hours of a crisis event. However, not all information on social media is relevant. A big proportion is comprised of irrelevant and duplicate content.

To understand the types of information useful for disaster response and management, we studied the use of our online supervised classification tool (Imran, Castillo, Lucas, et al. 2014). The tool is operational since 2013 and has been used several times by a number of humanitarian organizations such as UN OCHA during a number of major disaster events including disasters such as the 2015 Nepal Earthquake, the 2014 Typhoon Hagupit, and the 2013 Super Typhoon Yolanda. During and after a disaster event, we engage with those humanitarian organizations to get their feedback about the tool.Most of these interactions with crisis managers happen over Skype calls, through emails, or through a post-disaster report that we usually request after a disaster ends.

Based on our interactions, in this section, we summarize our findings regarding different types of information needs of various stakeholders, including crisis managers and individuals who use our platform during disasters. These findings help us understand gaps between the existing features of our technologies and users' information needs. Moreover, it helps determine new requirements for the design of our interactive monitoring tool, which is the contribution of this paper.

High-level Situational Requirements

It has been observed that the information needs of humanitarian organizations vary from one event to another. Past studies attempt to break down these information needs into a set of categories. For instance, Vieweg 2012 presents 28 information categories, while Bruns et al. 2011 report about 13 categories, and Imran, Elbassuoni, et al. 2013 proposes hierarchical categorization consisting of 10 categories. However, we observe an overlap between these categorization schemes. In below, we present a common categorization scheme containing important categories, which we have observed in the majority of the cases during real disasters. Each category is shown with an example tweet from the 2015 Nepal earthquake disaster.

- **Infrastructure and utilities damage:** Reports of damage to critical infrastructure, utilities services, government organizations, communication pools, hospitals, potential shelter houses

  `RT @AshkaITsolution: Nepal's historic Dharahara Tower collapses in massive earthquake: The historic Dharahara tower, a landmark`

- **Missing, trapped, or found people:** Reports of missing, trapped or found people due to the disaster

  `RT @prissy_anne: HELP find my baby sister Ballantyne Forder, she is in Nepal, Kathmandu`

- **Shelter and supplies:** Reports of shelter and supplies requests in different areas and offers of shelter and supplies

  `RT @Gurmeetramrahim: 2day 170 Nepal quake victim families r provided wth tents; relief material - blankets, kitchen commodities, medicines`

- **Injured or dead people:** Reports of affected people due to injuries or dead people

  `More than 900 people dead in #Nepal quake, deadly avalanche on Everest -URL-`

- **Donation of money, food, services:** Requests of donations such as money, food, water, services, blood, and/or offers

  `RT @derasachasauda: Mission NepalDisasterReliefByMSG @Gurmeetramrahim Ji provided relief material; medical aid to earthquake hit Nepal`

Critical Information Centroids

In addition to the high-level situational categorization scheme described above, we observe from our interaction with our users that majority of critical information is centered around certain entities. For example, from disaster response and management point of view, if we closely observe each category mentioned above with the given examples, we clearly see the presence of one or more such entities. For instance, in the case of *Infrastructure and utilities damage* category, the damage is reported to an entity called "`Dharahara Tower`". Similarly, in the case of *Missing, trapped, or found people* category, the report of a missing person including the name of the person "`Ballantyne Forder`" is posted. And, in the case of *Donation of money, food, services* category, an organization is providing relief material and medical aid. These are important reports associated with some entities. However, as it can be seen in the case of the *Shelter and supplies* example, an entity is not always present in each message.

Therefore, rapid access to important entities and the information associated with them is an effective way to quickly filter through large amounts of data during an emergency event. Many crisis managers believe such a functionality

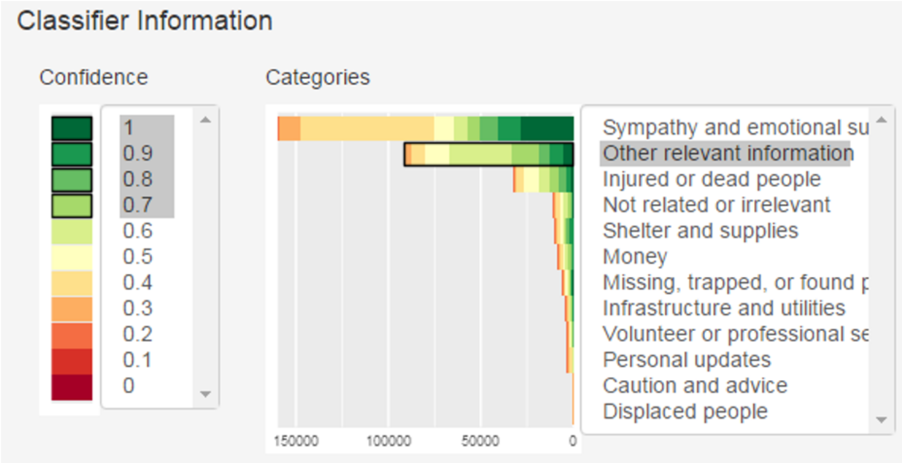

Figure 2. The *Classifier Information* component shows the classifier categories (right bar chart) ordered by decreasing amount of tweets assigned to them, and color-coded classifier confidence score distribution as color segments. The color scale is interactive to filter tweets based on a range of confidence scores, and the user can select the category to focus on. Selected elements appear with a thick black border.

would greatly help analyze crucial information quickly. Moreover, it will also help reduce the burden to go through the remaining situational messages (i.e. reports without entities).

Another requirement we learned from crisis managers is to understand where a piece of information is originated from. For example, this allows identifying the location of the author of a Twitter post. If important reports are put together on a map, collectively, it helps identify big clusters in terms of geographical locations of messages. Obviously, messages posted from a disaster zone would then be given high-priority to analyze and act accordingly. Among other requirements, the ability to interact and apply different filtering criteria, taking snap-shots of the filter-view of data, saving the snapshot with personal notes for future reference, are also considered as requirements for building an effective crisis monitoring tool.

INTERACTIVE MONITORING SYSTEM

To fulfill the requirements gathered from our users (reported in the previous section), we design an interactive tool to monitor Twitter data in real-time. As mentioned earlier, this tool operates on top of our previously built platform (Imran, Castillo, Lucas, et al. 2014), which is responsible for data collection and classification. Before going into the details of our visualization tool, we briefly discuss the data collection and classification parts first.

Data Collection and Building Classifiers

To collect and classify social media data at the onset of a crisis situation, we use our existing AIDR platform. In this paper, we only focus on Twitter data collection, classification, and visualization, so the next discussion will only be on Twitter. At the onset of a crisis, a user defines either a set of keywords/hashtags or a set of geographical bounding boxes or a combination of both. The selected search strategy is used by the system to continuously collect data from Twitter streaming API. A big proportion of social media data may consist of irrelevant information let alone useful for emergency management. To overcome this issue, the user, based on his/her information needs, defines a list of categories to which the incoming tweets will be automatically classified by the system. To enable automatic classification, the user provides a handful of training examples to the system. After one of more classifiers are up, all subsequent tweets are then automatically categorized. Given a tweet, a classifier assigns a category and a classifier confidence score to it. These are the inputs to our interactive tool we present next.

Interactive Monitoring Tool

Figure 1 shows the complete dashboard interface of the proposed tool including its ten important components such as live timeline, top-n entities, maps, classifier scores, tweets view etc. The dashboard mainly consists of two types of components: (i) primary filtering components, (ii) secondary filtering components. The purpose of a primary filtering component is to set high-level filters on large amounts of data to get a focused view. Tweets that pass through high-level filters are displayed on the secondary components. Next, we describe each component in detail. We give details about the graphical design rationale we followed, which are all grounded in the reference books (Munzner 2014) and (Ware 2004).

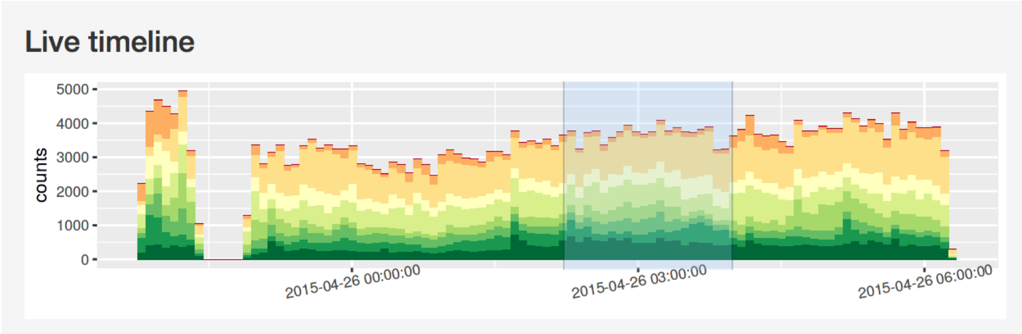

Figure 3. *Live Timeline* component shows beginning to end live timeline of an event. Each bar height codes for a number of tweets collected during a short time period (fixed width of the bar) at a specific date and time. Color segments show the distribution of the classifier confidence scores for this set of tweets.

Primary Filtering Components

There are two primary filtering components on the dashboard: the *Classifier Information* view and the *Live Timeline* view.

Classifier Information View: Confidence Scores and User-defined Categories

As already discussed above, the user defines a set of categories to which the tweets are automatically classified by the system. Each classified tweet gets a category and a classifier confidence/probability score ranging from 0.0 to 1 from the classifier. This score shows the classifier confidence in assigning a category to the tweet. Higher scores (near 1) are preferable and show tweets with desired information. Figure 2 shows the interface for the classifier part.

On the left side, the confidence scores are displayed with a diverging color scale to emphasize good (greenish) and bad (reddish) scores following the green-orange-red traffic-light widespread convention. For instance, the dark green color represents confidence score = 1 and the dark red color represents confidence score = 0. It is important to note that information shown on all other components are based on the classifier confidence scores and associated coloring scheme. A range of scores can also be selected e.g. as shown in Figure 2.

On the right side, we show the list of user-defined categories ordered by their prevalence (i.e. the number of tweets in a category). The length of the bars encodes the prevalence and the classifier confidence score distribution of each category is displayed as color segments for each bar. Greenish colors encode most reliable scores and bars are aligned closer to the vertical reference axis to ease frequency comparison. The bars are presented first on the left as they provide the pre-attentive perceptual information (color, position, and length) that drives the selection of a category. Bars and their category names are aligned on a central vertical axis to ease mutual matching. The user can select one category at a time (e.g. *Other relevant information* is selected in the Figure 2).

The *Classifier information* view is one of the primary filter views on the dashboard. Any selection on this component affects the data shown on all the secondary filtering components. For instance, if a user selects a new category e.g. *Shelter and supplies* then all secondary components will show data from the selected category.

Live Timeline View

To monitor an event at a high-level and to understand how it unfolds, we provide the *Live Timeline* view. Figure 3 shows its interface. This interface depicts the event-level timeline from the beginning of an event to its end (i.e. this could be a view over months of data). On the x-axis, the interface shows date-time and on the y-axis, it shows the frequency of tweets received in a unit time. Each bar contains multiple colors that represent the proportion of tweets classified with a certain confidence score using the same confidence score color-scale as the *Classifier Information* component. The greenish colors bearing the highest confidence levels primarily looked after by users, these segments are aligned with the horizontal reference axis to ease frequency comparison across the timeline.

The timeline is periodically refreshed automatically to continuously update data received since the last refresh. This provides a comprehensive view of how the event is unfolding, how the classifiers are behaving, and what are the distributions of confidence scores. For instance, we can see on Figure 3 that no data have been collected for about half an hour (gap on the left), and that many tweets have a medium or low confidence score. Furthermore, this is our second primary filtering component, that means a selection made on this component affects all secondary views. A user can make a selection of a narrow window-size, as shown with the shaded rectangle area on Figure 3. All the secondary view components will then only show data from the selected time window.

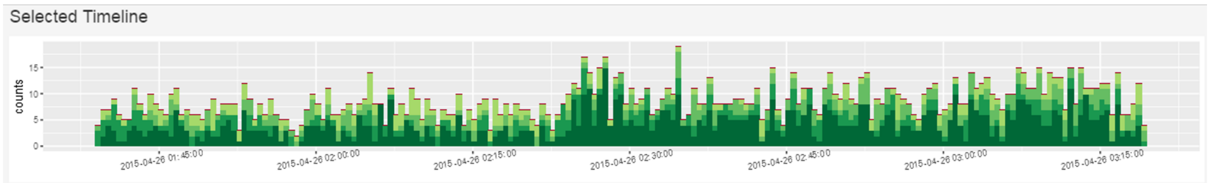

Figure 4. The Selected Timeline component shows a fine-grained view of a filtered timeline of messages

Secondary Filtering Components

As a result of a selection made at any of the primary filtering components, all the secondary components automatically update their view with the selected data (i.e. confidence scores, categories and time window). The purpose of the secondary filtering components is to apply data-driven filters on the selected subset of data obtained from the primary components. For instance, a data-driven filter could be to view all the tweets which talk about a specific location. Next, we describe all of our secondary view components.

Selected Timeline View

The *Selected Timeline* is similar to the *Live Timeline* showing time and tweets frequencies. However, instead of the whole event timeline, it only shows the data within the selected portion of the primary *Live Timeline* component. Hence, the user can have a detailed view of the event from time A to time B. Furthermore, the user can make a further selection on this component to get a more fine-grained focus, in this case, the timeline will automatically update to show only the selected area. As a result of the new selection on this timeline, all the secondary filtering components get updated with the latest information from tweets posted within the selected time window.

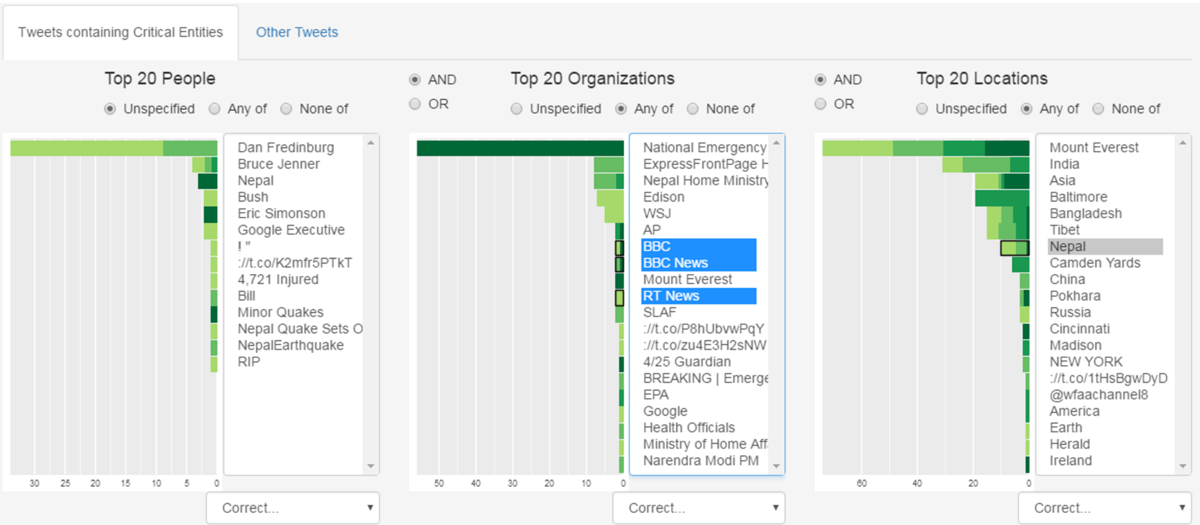

Figure 5. The critical entities component shows three types of named-entities (people, organizations, and locations) along with different filtering options.

Critical Entities View

One of the most important components of our dashboard is the *Critical Entities* view. Figure 5 shows the interface of this component. Based on our findings regarding critical information centroids (presented in the previous section), we extract important entities from the data. For this purpose, we use a state-of-the-art named-entity recognition tool from NLTK[1]. Given a text document (in our case a tweet), the named-entity recognizer extracts entities such as names of persons, organizations, and locations from the text. For each extracted entity, we compute its number of occurrences in the data. We select the top-20 people, organizations, and locations and display them on this component. Each entity is represented by a horizontal bar where each bar shows the distribution of classifier confidence scores along with its frequency. The design of these bar-charts-and-names components and the rationale behind them are identical to that of the one used for category selection in the *Classifier Information* view. Moreover,

[1]http://www.nltk.org/

in this component, we introduce the multi-selection of entities feature (e.g. see the selection of three organizations in the figure 5).

Twitter data are often noisy, e.g. tweets may contain misspellings, shortened words, slangs. Due to such issues, named-entity recognition techniques do not perform well. Often wrong entities are extracted and missed the correct ones. To deal with such issues, we provide a feedback mechanism to allow users to move wrongly categorized entities (e.g. a location name shows up in the organizations) to their correct places. Moreover, entries which are not actual named-entities can also be removed. For this purpose, each entity view has a drop-down menu (see Figure 6) with options to remove or move entities across.

Furthermore, this interface provides a number of filtering options. For instance, one can select one or more entities (e.g. from top people) and select the *Any of* filtering option. This will return a list of tweets in which at least one of the selected entities appears (i.e. logical OR). The *None of* filter will restrict the results to tweets in which none of the selected entities appear (i.e. logical NOT-OR). And the *Unspecified* option is the default to be used when the user does not care about presence or absence of entities of that type in the results. Moreover, one can make a filter across different types of entities. For example, if one wants to see tweets in which a person X and an organization Y appear, then he/she can select the *AND* option which applies between all types of entities at once (both radio button selectors are linked to show the same state). Similarly, the *OR* filtering options can also be applied to all three types of entities.

This component alone has the capability to show, for example, what is happening in some locations, names of missing people, or which organizations are taking part in some emergency relief operations and so on.

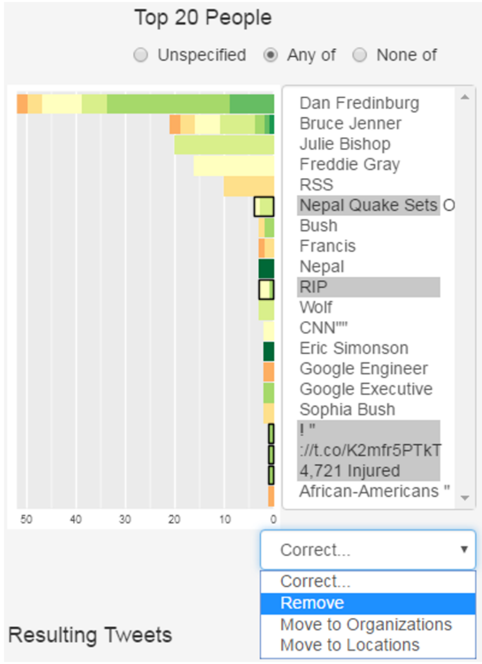

Figure 6. Drop-down menu with options to remove erroneous entities, or to move them to their correct place.

Map View

One of the most demanded components in our requirement list is the map view component (shown in the Figure 7). Geographically-tagged information is always more informative than a non-geo-tagged information. However, on Twitter only 3% to 5% of the tweets are geo-tagged and the rest remains without any geo information. The map view component simply shows tweets that are geotagged that are selected by the primary filtering components.Each dot on the map is a message (tweet), which can be seen by clicking on it. Each dot on the map has a color representing classifier confidence score and transparency is used to see through multiple dots in crowded areas. The map implements pan and zoom with the top left $+/-$ buttons and the mouse wheel for rapid focus on areas of interest. With this feature, one can immediately get an idea about geo-locations that are producing more relevant data for a given information need.

Resulting Tweets View

The *Resulting Tweets* view is the component where we show the actual tweets as a result of the filter selection made on primary as well as secondary view components (Figure 8). Any filter applied to any other components affects the content of the tweets view. It always shows the updated filtered list of tweets. In this view, we show date and text content of one tweet per row. The tweets are ordered from top to bottom in increasing date and time. The set of tweets is segmented in pages to allow quick navigation through hundreds or thousands of tweets quickly. For this purpose, navigational buttons at the top of the component are used. The total number of tweets found and the range displayed are shown at the top center of the component.

Notes View and Snapshot

Below the map, an editable text area allows the user to take notes on her analysis and findings (see bottom left of figure 1. This feature is useful to take notes on the analysis performed by a user. Moreover, a user can copy-paste tweets by manually selecting them from the *Resulting Tweets* view. When pressed, the *Take snapshot* button on the top right corner of the component allows downloading the data into a text file containing all the resulting

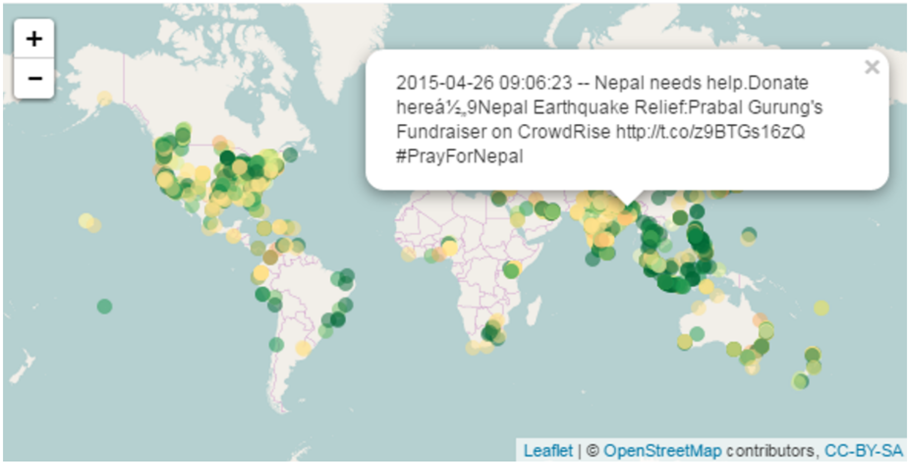

Figure 7. The Map view component show geo-tagged messages

tweets together with a textual description of the boolean filters applied that the user defined through her interactive selections, and the textual content of the *Notes* view. For instance, using this feature a user can create different filters and can save multiple snapshots of the data.

Other Tweets View

Below the *Selected Timeline* secondary filtering component, the user can select either the *Tweets containing Critical Entities* tab or the *Other Tweets* tab. The former shows the *Critical Entities* and resulting tweets containing at least one named entity despite its type or value. The latter shows all the tweets selected from the primary filtering components for which no named entity has been identified automatically. This view is important to keep track of important information without critical entities.

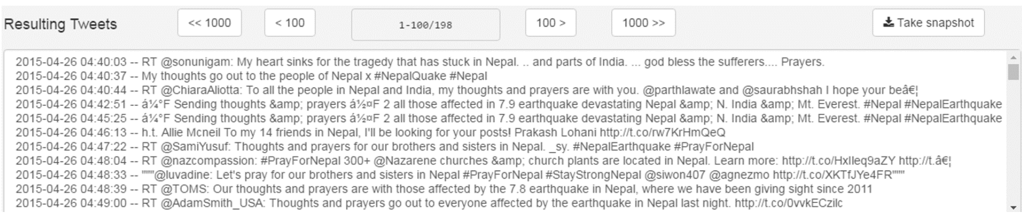

Figure 8. The Resulting Tweets view component shows date, time and text of the tweets resulting from the overall filtering process

USER STUDIES AND EVALUATION

To evaluate the usefulness of our interactive monitoring tool, we interviewed three experienced people A, B, and C who work in the field of crisis response and management. Participants A and B are crisis managers and they work at an emergency management organization in New York, USA. Participant C also closely works with crisis managers and has almost five years of experience building technologies for crisis response and management. For each participant, we first give a general introduction and a live demonstration of the tool showing historical data (the live timeline was static during our demo) from the Nepal Earthquake. Specifically, we use the following demo script.

Demo script

Each interview started with an introduction and motivation behind building the interactive dashboard. Moreover, necessary details regarding our AIDR technology, which is used to classify social media messages, was provided. Next, a demo of the tool including detailed description and functionalities of each component, their interactions, and various filtering options such as entities refinements features, and filtered results analysis are shown. The live demo includes a detailed overview of each component in which selection of different options (e.g. selections on the timeline, changing confidence score ranges etc.). Each demo took 30 minutes on average. We were flexible to change the demo script on participants request. For example, if a participant asked to perform a different filtering option than our planned one, we did so.

Each session has been recorded for a detailed analysis of the discussion that took place during the interviews. After the detailed demo sessions, we asked participants if they had any questions or doubts regarding any component or feature. Each participant asked a few questions. Finally, when participants were satisfied with their understanding of the tool, we asked the following question to each of them.

Evaluation questions:

(i) Do you think the tool can help you quickly analyze large amounts of tweets during an event as compared to your earlier approach?
(ii) Do you think the tool can provide a high-level overview of the event as it unfolds using the live timeline and selected timeline?
(iii) Do you think the critical entities we show are good to understand the situation?
(iv) Any suggestions regarding other such entities that we should consider adding?
(v) Do you think the tool can provide a fine-grained view of the critical information (using its critical entities view) from tweets?
(vi) Do you think the primary and secondary filtering component and options are useful and intuitive?
(vii) Is the flow from primary to secondary filtering useful and intuitive?
(viii) Do you think the filtering options on entities components are useful and intuitive?
(ix) Do you think the feature to remove or move a wrong entity is useful?
(x) Do you think the map is a useful component in this context?
(xi) Do you think the notes component and the snapshot features would be useful for you?
(xii) Overall, what is your opinion regarding this dashboard, is it useful for emergency situational awareness?
(xiii) Any other general comments, suggestion, or feedback?

Summary of participants feedback

One general observation we learned is that all the participants look very excited and interested to see the tool. Some initial comments from the participants were "very impressive" (A), "terrific" (B), "great" (C). Participant (A) reminded that a better situational awareness is to understand a situation from different angles, so the map and the timeline components were regarded as very important as they let users catch both temporal and spatial information of an unfolding crisis, and the *Critical Entities* view allows to focus on specific problems concerning population.

Besides the standard pan and zoom of the *Map* and *Timeline* views, the possibility to interactively filter based on categories and on named entities was considered very helpful. (A) liked a lot the interactive color scale to filter tweets by confidence scores of the classifier categories. The *Critical Entities* view attracted most of the interest as it allows the user to design complex filtering rules by simple clicks. (A) and (B) closely examined the top Organizations entities, and told that being able to filter based on these entities would be very useful. For instance, both participants described cases in which it is important to distinguish situations described by official organizations like NGOs or mass media, and also information coming from other sources like people closely involved in the field but with no official role.

The map shows only geo-tagged information and can be used to distinguish people on the crisis site from people in remote locations. However, all the participants wanted to geo-locate all the tweets using the geographical information present in their content. (B) also wondered if a location entity could be automatically breakdown into subdivisions, like a city into its neighborhoods to filter even more precisely. (A) and (B) evoked practical cases in New York where they wish they could use this tool to monitor possible crisis like a hurricane coming to Long Island, based on geographical filtering combined with this named-entities-based filter.

At last all participants questioned the issue of the reliability of the Named Entities and found very useful to be able to clean the set of entities manually when the system makes an error.

Participant (C) is the most technical person among all three. He admired the ability to make corrections of the system through our critical entities view and being able to tune the system while analyzing live data. He also suggested us to map the non-geo-tagged tweets on the map view component. This could be possible by first disambiguating geo-locations and then finding geographical coordinates for tweets mapping. We take this as our future work. Moreover, he suggested us to include a help button with each component so that first-time user can easily understand the system.

Overall, all the participants agreed that the tool helped them analyze large amounts of tweets in timely manner. They found all the components useful and intuitive to use. Overall, all of them seem to like the concept of primary and secondary filtering options that we offer through our components. The notes and snapshot features were considered very useful. Overall, all agreed that such a tool would definitely help crisis managers get early insights from the social media data and ultimately will advance crisis response and management efforts.

RELATED WORK

Dashboards visualize different types of visual components usually containing maps, timelines, and charts in various ways for a broad range of application domains. Some of these tools use social media data for analysis and provide advanced filters. For example, in the crisis response domain, Abel et al. 2012 presents Twitcident, a tool for filtering, searching and analyzing information about incidents that people publish on Twitter. MacEachren et al. 2011 provide a map-based, interactive web application that enables information searching and sense-making using tweets indexing. They display the information based on place and time.

Regarding the design of filters, some works have been proposed to support users to set up a filter based on graphical representations. Kaleidoquery (Murray et al. 2000) is a visual query language to support large data exploration. This language maps SQL statements to graphical elements, providing a visual data navigation technique to select or filter out a subset of the data using a flow-diagram type of graphical representation. Decision trees can also be used to represent filtering of a data stream (Elzen and Wijk 2011; Gangavarapu et al. 2016). In these approaches, a tree represents a filter that is applied to the data passing from the root to the leaves. While decision trees are useful to visualize the data filtered at each level they spend unnecessary space to represent the tree structure at the expense of the details of the filters encoded at each node. VizFilt is another approach designed for the analysis of data streams in the cyber security domain (Aupetit et al. 2016). It uses a bar chart representing a filter where bars both encode the items of the filter and the amount of data filtered by these items in a very compact way. The problem is that these bars are not visually linked to the categorical values they encode (IP addresses, source and destination ports, countries) other than using a discrete color scale. In our case linking such bars to the critical entities (people, organizations or location) help select most important entities easily. So, we propose a similar approach but with grouped bars instead of stacked bars.

A dashboard of the data stream from social media called NStreamAware has been proposed in (Fischer and Keim 2014) to gain situational awareness in the cyber-security domain, but it is applicable to other domains as well. It focuses on extracting relevant features and visualizing time slices from a stream of network data. Word clouds and node-link diagrams are the main graphical metaphors used to visualize network information of interest. However, this tool does not provide quantitative indicators like bar graphs to support the user to include entities of interest in a filter or to let her know the confidence score of the automatically clustered entities. It does not allow for correction of the erroneous classifications either.

A haze monitoring system[2] centered on a map component and information from social media (Pulse 2016). Timelines and filtering components based on critical entities are displayed in a dedicated "Analysis of user generated content" panel. However this information is static, they are not linked to the map or the timeline. No confidence score of the keywords and categories like "Haze Health" or "Haze Impact" is visualized and it does not allow interactive filtering of messages related to specific keywords nor to see the messages themselves to get the context where these keywords were used.

CONCLUSION

Information available on social media platform during an emergency situation proved to be useful for crisis response and management. However, analyzing large amounts of social media data pose serious challenges to crisis managers, especially under time-critical situations. Even if a supervised classification approach is employed, analyzing thousands of machine-classified messages is a non-trivial task. In this paper, we presented an interactive tool to effectively monitor social media big crisis data in a timely manner. The tool provides a number of useful components to gain early insights by applying various filtering options. The interactive timelines provide both coarse-grained and fine-grained view of the event, and the critical entities filter allows detecting relevant entities that are being reported in a given time-window. Evaluation performed with experienced crisis managers demonstrated the significance of the tool.

Future work: Based on the feedback obtained through users interviews, we aim to improve the interface by including the map in the filtering process and to display both geo-tagged as well as geo-coded tweets for which geo-location can be inferred from the content of a Twitter message. We aim to cluster related tweets and entities to ease the selection and filtering of tweets. Moreover, we aim to improve the performance of the classifiers by employing feedback mechanism from users. This will not only improve the named-entity recognition system, but also the automatic classifiers that categorize tweets into different humanitarian categories.

[2]http://hazegazer.org/home

REFERENCES

Abel, F., Hauff, C., Houben, G.-J., Stronkman, R., and Tao, K. (2012). "Twitcident: fighting fire with information from social web streams". In: *Proceedings of the 21st International Conference on World Wide Web*. ACM, pp. 305–308.

Aupetit, M., Zhauniarovich, Y., Vasiliadis, G., Dacier, M., and Boshmaf, Y. (2016). "Visualization of actionable knowledge to mitigate DRDoS attacks." In: *VizSEC*. IEEE Computer Society, pp. 1–8.

Bruns, A., Burgess, J., Crawford, K., and Shaw, F. (2011). *Crisis Communication on Twitter in the 2011 South East Queensland Floods: Report Addressing the Use of Social Media Services During National Disasters, with a Focus on the Queensland Police Force's Use of Twitter*.

Cameron, M. A., Power, R., Robinson, B., and Yin, J. (2012). "Emergency situation awareness from twitter for crisis management". In: *Proceedings of the 21st International Conference on World Wide Web*. ACM, pp. 695–698.

Elzen, S. van den and Wijk, J. J. van (2011). "BaobabView: Interactive construction and analysis of decision trees." In: *IEEE VAST*. IEEE Computer Society, pp. 151–160.

Fischer, F. and Keim, D. A. (2014). "NStreamAware: Real-time Visual Analytics for Data Streams to Enhance Situational Awareness". In: *Proceedings of the Eleventh Workshop on Visualization for Cyber Security*. VizSec '14. Paris, France: ACM, pp. 65–72.

Gangavarapu, K., Babji, V., Meibner, T., Su, A. I., and Good, B. M. (2016). "Branch: an interactive, web-based tool for testing hypotheses and developing predictive models." In: *Bioinformatics* 32.13, pp. 2072–2074.

Hughes, A. L. and Palen, L. (2009). "Twitter adoption and use in mass convergence and emergency events". In: *International Journal of Emergency Management* 6.3-4, pp. 248–260.

Imran, M., Castillo, C., Diaz, F., and Vieweg, S. (2015). "Processing social media messages in mass emergency: A survey". In: *ACM Computing Surveys (CSUR)* 47.4, p. 67.

Imran, M., Castillo, C., Lucas, J., Meier, P., and Vieweg, S. (2014). "AIDR: Artificial intelligence for disaster response". In: *Proceedings of the 23rd International Conference on World Wide Web*. ACM, pp. 159–162.

Imran, M., Elbassuoni, S. M., Castillo, C., Diaz, F., and Meier, P. (2013). "Extracting information nuggets from disaster-related messages in social media". In: *Proc. of ISCRAM, Baden-Baden, Germany*.

Imran, M., Mitra, P., and Castillo, C. (2016). "Twitter as a Lifeline: Human-annotated Twitter Corpora for NLP of Crisis-related Messages". In: *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*. Portoroz, Slovenia: European Language Resources Association (ELRA).

MacEachren, A. M., Robinson, A. C., Jaiswal, A., Pezanowski, S., Savelyev, A., Blanford, J., and Mitra, P. (2011). "Geo-twitter analytics: Applications in crisis management". In: *25th International Cartographic Conference*, pp. 3–8.

Munzner, T. (2014). *Visualization Analysis and Design.* A.K. Peters visualization series. A K Peters, pp. I–XXIII, 1–404.

Murray, N., Paton, N. W., Goble, C. A., and Bryce, J. (2000). "Kaleidoquery-A Flow-based Visual Language and its Evaluation." In: *J. Vis. Lang. Comput.* 11.2, pp. 151–189.

Nguyen, D. T., Joty, S., Imran, M., Sajjad, H., and Mitra, P. (2016). "Applications of Online Deep Learning for Crisis Response Using Social Media Information". In: *arXiv preprint arXiv:1610.01030*.

Pulse, U. G. (2016). "Haze Crisis Analysis and Visualization Tool". In: *Big Data Tools Series* 2.

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

Starbird, K., Palen, L., Hughes, A. L., and Vieweg, S. (2010). "Chatter on the red: what hazards threat reveals about the social life of microblogged information". In: *2010 ACM conference on Computer supported cooperative work*, pp. 241–250.

Vieweg, S. E. (2012). "Situational awareness in mass emergency: A behavioral and linguistic analysis of microblogged communications". In: *http://works.bepress.com/vieweg/15/*.

Ware, C. (2004). *Information Visualization: Perception for Design*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.