

A Language-agnostic Approach to Exact Informative Tweets during Emergency Situations

Jacopo Longhini
Politecnico di Torino
Torino, Italy
jacopo.longhini@polito.it

Claudio Rossi
Istituto Superiore Mario Boella
Torino, Italy
rossi@ismb.it

Claudio Casetti
Politecnico di Torino
Torino, Italy
casetti@polito.it

Federico Angaramo
Istituto Superiore Mario Boella
Torino, Italy
angaramo@ismb.it

Abstract—In this paper, we propose a machine learning approach to automatically classify non-informative and informative contents shared on Twitter during disasters caused by natural hazards. In particular, we leverage on previously sampled and labeled datasets of messages posted on Twitter during or in the aftermath of natural disasters. Starting from results obtained in previous studies, we propose a language-agnostic model. We define a base feature set considering only Twitter-specific metadata of each tweet, using classification results from this set as a reference. We introduce an additional feature, called the Source Feature, which is computed considering the device or platform used to post a *tweet*, and we evaluate its contribution in improving the classifier accuracy.

Index Terms—Disaster relief; social media analysis; classification; machine learning; real-world traces.

I. INTRODUCTION

It is estimated that, due to climate change, extreme weather events will occur more frequently and last longer [1]. Extreme weather events are defined as the weather conditions lying in the most unusual ten percent, given the recorded weather history. As the world warms up, the resulting greater evaporation of the air lead to increase water vapor in the atmosphere, which produces more intense precipitations increasing the likelihood of floods. Scientists claim that in the next decades rainfall will be more intense in many parts of the world, leading to more impacting floods. Also hurricanes and typhoons will be more intense because they will draw energy from hotter oceans [2], while frequent heat waves and forest fire could become common [3].

During crises, such a the ones cause by natural hazards, getting timely access to information is a matter of life and death. Recent technological developments resulted in the ubiquity of advanced mobile hand-held devices, which let us receive and transmit information within few seconds to any distance. According to the International Telecommunication Union [4], there are more than 6.8 billion mobile phone subscriptions in the world, accounting for almost one mobile phone subscription per person, and those numbers are globally rising each year. At the same time, the use of online social networks allowing to collect and share unstructured and heterogeneous user-generated contents has raised exponentially, transforming users into content generators (human sensors). Twitter is a social network platform that has largely grown during the last few years and that is now counting more than 300 millions

active users per month. Twitter is a micro-blogging service, which means that users can write messages with length up to 140 characters¹: these messages are called *tweets*. A Twitter user has *followers*, i.e., users who follow him/her and *followee*, i.e., users that he/she follows. Relevant users, such as Governments or public personalities can be verified by Twitter. Other users can be referenced in a message by the character @ followed by their usernames. Within *tweets*, words followed by # are called *hashtags*, which are used to label content. Users can share messages that were already written and shared by other users doing the so called *retweet*. This allows for message propagation and virality within the network. Thanks to its conciseness and its extent all over the world, Twitter has become particularly important during emergencies, where real-time information can enhance situational awareness [15]. In fact, during recent crisis events, some information was available on Twitter before any other new channel. An example of how social networks are used in time of crisis is shown in Fig. 1, where tweets are related to a bombing in New York. People post information about the event (first explosion at 8:31 pm), share links to streaming channels, and declare that Twitter is more effective compared to TV channels and other medias.

There is however, a downside to the use of social media.

¹at the time of the experiments. In November 2017 Twitter increased the number of characters up to 280 [6].



Fig. 1: Tweets share soon after the New York and the New Jersey bombings in September 2016

As stated in [5], due to the growing of social network users, information generated during crises are getting increasingly large and very hard to analyze. An overflow of information and data can be as paralyzing as the absence of it, especially when it comes to mobilizing response both locally and internationally. Some user-generated content posted on social medias can be informative and relevant for disaster response, but collecting, cleaning and analysing these datasets often takes a disproportionate amount of time. The field of 'advanced computing' has developed two ways to manage such big data: human computing and machine computing. The former uses crowd-sourcing and micro-tasking platforms to distribute tasks that are easily completed by a 'crowd' of humans. In contrast, the latter automate such tasks using data mining and machine learning algorithms, which are faster, hence scalable to huge volume of data. Both approaches can be used to manage the big crisis data challenge and the problem of verifying user-generated content [5].

The focus of this work is to study Twitter contents during high impact events, and to propose a service able to automatically filter the huge amount of information that is exchanged via Twitter during disasters. Our goal is to give in near real-time useful information to decision makers in order to help them in having a better operational picture and ultimately in taking better decisions. Starting from already sampled and labeled datasets, we train a language-agnostic machine learning model to classify tweets and retain only informative contents, where the concept of informativeness is defined as everything that can be useful to improve the situational awareness for both citizens and authorities during an emergency event. We select a language-agnostic approach in order to provide a baseline tool that can be used at any location worldwide. Starting from results achieved in previous works, we propose a model that works on Twitter-specific features and we study the contribution of a novel feature, which is derived from the source application that generated the content, evaluating its contribution in improving the classification accuracy.

II. RELATED WORKS

Recently, the use of social media during emergencies, and how it can be exploited to enhance situational awareness, has received much attention. In the work done by Olteanu et al. [9], the authors present the result of a manual labeling campaign to describe what to expect from social media data across a variety of emergencies (natural disasters, terrorist attacks, explosions, etc.) in terms of volume, informative level, information type and source. In the emergency context, studies have also been done on the reliability of viral tweets and the retweet activity, as the one done by Mendoza et al. [10], which show how the propagation of rumors on Twitter differs from the one of real news during the 2010 earthquake in Chile. The work done by Jackoway et al. [11] extracts event-related tweets using information coming from constantly updated news articles, but does not look at the informativeness of a tweet and does not work on unexpected events (like earthquakes). Several works have been done concerning the classification of online data into

information classes or topics. The closest work to ours on the emergency context is the one by Caragea et al. [12], which analyzes text messages written during the Haiti earthquake and gathered by the Ushahidi platform² to classify each post into an information classes using several approaches in comparison to Bag of Words. A similar work [14] was done to address information credibility on Twitter on trending topics with a focus on how tweets are propagated (re-tweeted). Our work is different, because we propose a language-agnostic approach to automatically filter Twitter streams during disaster caused by natural hazards to retain only informative contents, where the concept of informativeness is defined as everything that can be useful to improve the situational awareness for both citizens and authorities about the emergency event. Our objective is to identify the baseline classification performance using only Twitter-specific features and to study the contribution of a novel feature.

III. LABELING THE DATASET

As explained above, we aim to use machine learning to classify tweets and filter out irrelevant information. We select a *supervised* learning approach, which is trained on a labeled set of data.

We use the labeled dataset [8], which is available for research purposes and that was collected during several large crisis events. The dataset contains *tweets* gathered by Olteanu et al. [9] that were retrospectively sampled on the 1% of the public Twitter data stream using the Twitter Streaming API, including crisis events that took place from 2012 to 2013. A dataset of each crisis was extracted by specific keywords, which were manually defined according to the event type, location and language of the affected geographical area. The keywords used include *hashtags* related to the specific event, especially binomial hashtags formed by different composition and/or abbreviation of two words, being often the 'event name' and the 'location name'. For example, a widely used hashtag during the 2011 New Zealand earthquake that was *#eqnz*, where 'eq' was use to reference the earthquake and 'nz' to specify its location [13]. The authors of the dataset looked for other relevant hashtags of each crisis to get all relevant tweets. From the collected tweets, those shorter than 3 words were removed, because in the majority of the cases they do not contain any useful information. In Table I there are the natural crisis events that we select. They are divided in three main category: *Earthquake*, *Floods* and *Fire*. We define the aggregation of all the aforementioned as *Natural events*. We select *Flood* and *Fire* because they are the natural events mostly affected by climate change, and we include *Earthquake* because it is among the most impacting natural hazards, both in terms of human and economic losses.

Due to the huge amount of collected data, it was impractical to ask the help of some experts to manually label the tweets, so a crowd-sourcing platform was used, i.e., CrowdFlower. Crowdworkers were asked to analyze each tweet and assign

²<https://www.ushahidi.com/>

TABLE I: Considered crisis events.

| Crisis name | Year | Type |
|-----------------------|------|------------|
| Bohol earthquake | 2013 | Earthquake |
| Costa Rica earthquake | 2012 | Earthquake |
| Guatemala earthquake | 2012 | Earthquake |
| Italy earthquakes | 2012 | Earthquake |
| Sardinia floods | 2013 | Floods |
| Alberta floods | 2013 | Floods |
| Philippines floods | 2012 | Floods |
| Colorado floods | 2013 | Floods |
| Queensland floods | 2013 | Floods |
| Manila floods | 2013 | Floods |
| Australia bushfire | 2013 | Fire |
| Colorado wildfires | 2012 | Fire |

to it different labels [9]. The one we are interested in this work is the *Informativeness*, which distinguishes the *tweets* that can add useful information and improve situational awareness from the ones that do not. The Crisislex dataset includes only the tweet ID, its text, and the assigned labels. Hence, we use Twitter’s Rest APIs [7] in order to retrieve the full information about each *tweet*, and the associated user profile.

Having the *tweet* ID, we fully retrieve the tweet object together with its metadata using *GET status/lookup*. Then, having the user screen name we call the *GET users/lookup* to gather the full user profile, which contains fields such as the number of followers, the number of followed accounts, the number of generated *tweets*, etc. We temporary store the retrieved data on the Microsoft Azure infrastructure using an *Azure SQL Database* for further processing. In total, we processed around 2M elements (*tweets* and *retweets*) and 1.2M unique user profiles. About 23.75% of tweet were lost, meaning that they were removed from the social media for some reason or that the user profile was removed.

Fig. 2 represents a simplified flow chart of our data acquisition process.

IV. BASELINE FEATURE DEFINITION

In this section we present some insights on the data that we collected and the baseline features we selected considering previous works and our background knowledge. Table II shows such baseline features.

The feature ΔT is the time difference between the time of the oldest tweet of the associated crisis and the considered

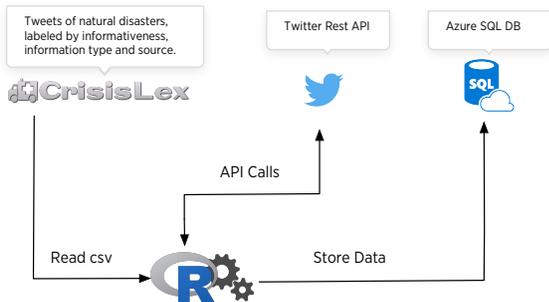


Fig. 2: Simplified flow chart of the data acquisition process.

TABLE II: Table showing which features compose the baseline feature set to train the machine learning model.

| Feature | Description |
|------------------|--|
| ΔT | Delta time since first <i>tweet</i> of the crisis |
| nFollowers | Number of followers (Fw) of the user |
| nFollowee | Number of followees (Fe) of the user |
| $\Delta regDate$ | user age in days with respect to its registration date |
| nTotTw | Total <i>tweets</i> posted by the user |
| vUser | Boolean: TRUE if User is Verified |
| geoTag | Boolean: TRUE if the <i>tweet</i> is geotagged |
| nHash | Number of # in the <i>tweet</i> |
| nLink | Number of URLs in the <i>tweet</i> |
| nAt | Number of @ sign in the <i>tweet</i> |

tweet. We assume to know exactly the starting date of the crisis, after which the data collection process from Twitter starts. This variable is quite important because we notice that informative content are usually generated within the first hours after the event, especially for unexpected, not predictable and instantaneous events such are earthquakes. In order to extend our approach to a monitoring use-case, where there is not any ongoing event, this feature has to be removed.

nFollowers and **nFollowee** are the numbers of followers and followee of the used who posted the *tweet*, respectively. Fig. 3a and Fig. 3b show the Informativeness distribution based on the number of follower and followee, respectively.

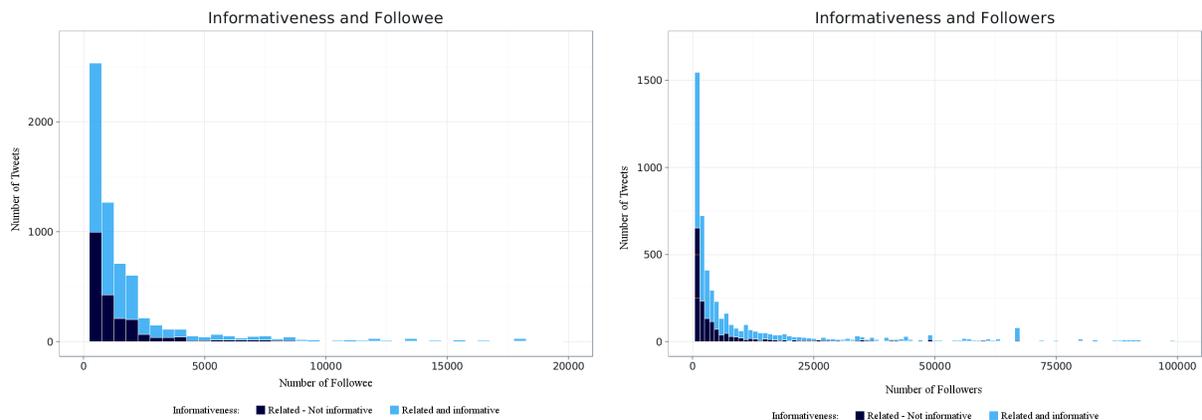
nTotTw is the total number of tweets posted by the user since the registration date. The rational is that the higher the number of tweets, the better the user reputation within the social media, hence the higher the probability that the user is not a spammer.

vUser is an attribute given by Twitter to tag special accounts, such as government, politics, media, journalists. If TRUE, it means that the account is authentic and verified [20]. As shown in Fig. 3c, the percentage of informative content is higher for verified users.

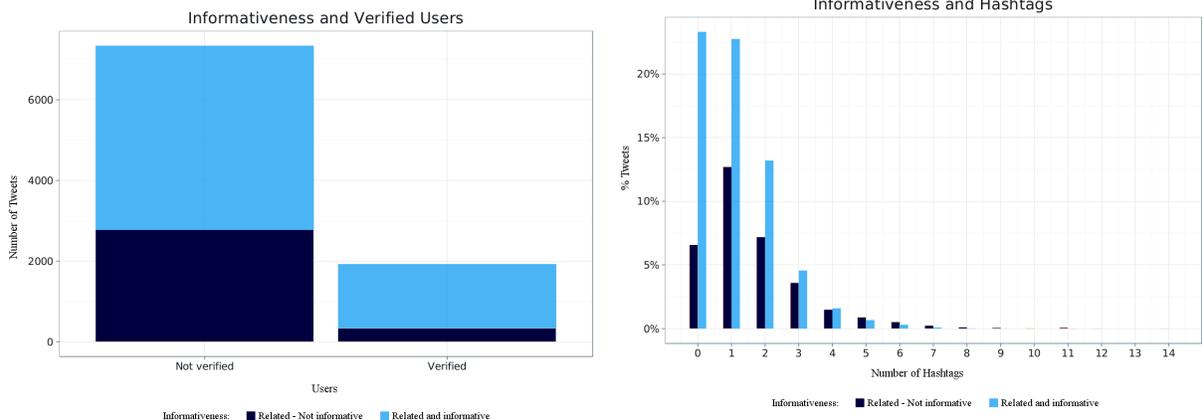
$\Delta regDate$ is the time difference between the *tweet* and the user registration date. It represent how long the user has been subscribed to the platform before creating the analyzed *tweet*. Usually, fake accounts are very young because after a while they get identified and removed from the platform. Hence, among old users there is an higher amount of informative *tweets*.

geoTag is a boolean variable indicating whether the user has chosen to share his location. The assumption here is that if a user agrees to be located is more reliable when providing information. However, we observe a very limited amount of geolocated (with latitude and longitude) *tweeta*. Only the 0,23% of the tweets were geolocalized in the entire dataset. This reduces the relevance of this feature.

The last variables computed and inserted in the baseline features set are: **nHash**, **nLink** and **nAt**, which are the number of hashtags, numbers of URLs, and the number of mentions included in the *tweet*, respectively. As Kwak et al. [16] argue, spammers tend to use more trending topic hashtags than others to appear in the results for any Twitter search and to gain more attention. We acknowledge this behavior, as shown in Fig. 3d.



(a) Relation between informativeness and number of Followee. (b) Relation between informativeness and number of Followers.



(c) Relation between informativeness and verified users. (d) Relation between informativeness and number of used hashtags.

Fig. 3: Relation between informativeness and other features.

Another common behavior of spammers is to put many URLs in the *tweets* in order to attract users to external pages. They usually use URLs shortening services to cope with the 140 character limit and to mask domain-level information. Also, an high number of mentions (@) can be a bad signal because they are often used to gather the attention of well-known people, i.e., users having an higher number of followers, and hope in their *retweet*. For brevity, we omit plot related to links, mentions, total tweets, registration data and ΔT as some of their behavior has been described in other studies [14].

V. CLASSIFICATION THROUGH SUPERVISED LEARNING

The task of predicting to which category an example belongs to is known as *classification*, while supervised learning is a machine learning task for inferring a function from labeled training data, where during the learning part the algorithm attempts to discover and model the relationship between the target feature (the informativeness in our case) and the other input features.

Then, the trained model is evaluated using a labeled *test dataset* and prediction errors can be measured in terms of misclassification [18].

We use a supervised approach over the Crisislex dataset, trying several state-of-art classifiers with default parameters (e.g., Support Vector Machines, decision trees, neural networks). The algorithm that achieve the best performance is the Random Forest (RF) algorithm [19]. All results presented are obtained using RF, which is trained and evaluated with 10 fold cross-validation.

In order to mitigate the effect of class unbalance, which is known to affect the performance of classifiers such as RF, we try different techniques, namely: *downsampling*, which randomly subsets the classes to match the least prevalent class; *upsampling*, which randomly samples the minority class with replacement to match the frequencies of the majority class; *SMOTE*, i.e., a hybrid method that downsamples the majority class and synthesizes new data in the minority class. We achieve the best classification results using the downsampling technique on the majority class. Hence, we present all the subsequent results using downsampling.

Binary classification classes are commonly called positive class and negative class. We consider the informativeness classes and define:

- Positive class: as the *Not-Informative* class that represent the *tweets* we want to discard;
- Negative class: as the *Informative* class that represent the *tweets* that contain useful information.

The relationships between the positive class and the negative class predictions can be represented as a confusion matrix that tabulates whether predictions fall into one of the four categories:

- *True positive* (TP): if the sample is positive and is classified as positive;
- *False positive* (FP): if the sample is negative and is classified as positive;
- *True negative* (TN): if the sample is negative and is classified as negative;
- *False negative* (FN): if the sample is positive and is classified as negative.

Misclassifications happen when a positive sample is classified as negative or viceversa.

As a statistical measure of the performance, we use three metrics. The *accuracy*, defined as the ratio between the True predictions and all predictions:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (1)$$

The *sensitivity* (also called the true positive rate) that measures the proportion of positives samples that are correctly identified:

$$\text{Sensitivity} = \frac{TP}{TP + FN} \quad (2)$$

The *specificity* (also called the true negative rate) that measures the proportion of negatives samples that are correctly identified:

$$\text{Specificity} = \frac{TN}{TN + FP} \quad (3)$$

We use the accuracy to define the overall performance of the classifier, the sensitivity to represents how many *tweets* are correctly identified as *Not-informative*, and the specificity to know how many *tweets* are correctly identified as *Informative*.

VI. SOURCING THE TWEETS

The Twitter REST APIs also provides for each *tweet* its source, i.e., the application that was used to post the *tweet*. Twitter has an ecosystem of services and applications made by third-party developers, who aim to improve Twitter usage by allowing users to send *tweets* from all kind of sources, ranging from smartphones to smart televisions.

Analyzing the collected data, we count 572 different sources.

We tried to make sense of this plethora of sources and understand whether knowing the source could provide additional information on its dependability. Thus, we group sources using different methods, as described below.

Binary Approach: A first approach to group the different sources is to distinguish between tweets sent from *mobile devices* and tweets produced by the *websites and other sources*.

To test our idea we added a boolean feature to each tweet for classifying them into *mobile* or *other* according to their source. We define as *mobile* all tweet containing in *statusSource* at least one among the following keywords: Android, BlackBerry, Instagram, iOS, iPad, iPhone, Windows Phone, Mobile.

The aforementioned keywords have been designed to aggregate all Twitter’s native application for mobile operating systems, the Twitter mobile website and third party clients developed for mobile platforms. We also included *Instagram*, which is another popular social media that allows users to upload photos exclusively through the mobile application and to share the content also on other social networks, including Twitter.

Discard Rate approach: The second approach is to calculate a Discard Rate for each source. The discard rate *DR* of a source *s* is calculated as described in Eq. (4), considering the number of not-informative tweets coming from that source over the total number of tweets of the source.

$$DR_s = \frac{TN_s}{TI_s + TN_s} \quad (4)$$

where:

- s* is the source (e.g., *Twitter for iPhone*)
- TN_s is the count of how many tweets are labeled as not-Informative given the source *s*
- TI_s is the count of how many tweets are labeled as Informative given the source *s*

In order to calculate the *DR* we use only *tweets* in the training set, avoiding biased results.

The main shortcoming of this approach is that we cannot evaluate a *tweet* (in the test set) coming from a new source (which is not present in any of the training set *tweets*), because we do not have data to compute its *DR*. A possible countermeasure would be to assign the average *DR* to those sources.

Clustering by discard rate: We equally divide the range of the *DR* into five different clusters, and create a new feature with the corresponding class in the Likert Quality scale [21]: ‘Very Good’, ‘Good’, ‘Acceptable’, ‘Poor’, ‘Very Poor’. Hence, a source having a *DR* of 0, will be assigned to the ‘Very Good’ class, while a source with a *DR* of 1 will be ‘Very Poor’. The aim is to have sources aggregated by the percentage of informative tweets over the total considering the elements in our training set. The scatterplot of the obtained clusters (considering, only here, all *tweets*) is shown in Fig. 4. Of course, also by grouping features according to their *DR* we have the shortcoming of new sources in the test set. Following the same approach, we can initially assign to new sources the average class, i.e., ‘Acceptable’, and then update it following a continuous learning approach.

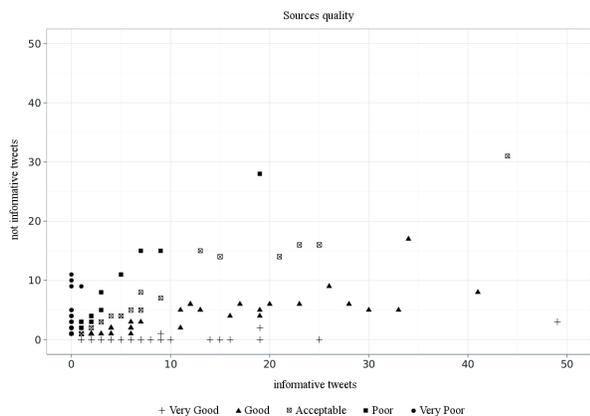


Fig. 4: Clusters obtained by grouping sources by *DR* according to the Likert scale.

Results: Surprisingly, we find that mobile sources produce more not-informative contents respect to “fixed” ones, with a DR of 48% versus 20%, respectively. This is probably due to the larger plethora of unspecific mobile applications allowing to produce *tweets*, and their popularity among citizens. Conversely, “fixed” sources are dominated by business applications mainly used by professional companies, who rely on platforms that automatically share their content to Twitter. This can be clearly observed by looking at Table III, which is computed using *all* tweets and ranks the top 14 sources by the total number of tweets, listing them in ascending order by DR. The first two more informative sources, *twitterfeed* and *dlvr.it*, are two applications that let items published in a RSS feed to be automatically shared in social media networks, including Twitter. Note that RSS feeds are very popular in websites of news and official announcements. The third ranked source is *Hootsuite*, which is an application that lets users to manage multiple social media in one place. The ranking continues with the Twitter desktop applications (*Twitter for Websites*, *TweetDeck*, *Twitter Web Client*) and *Facebook*, followed by the mobile versions of Twitter for the different mobile systems, i.e., BlackBerry, iOS, Android. Interestingly, BlackBerry users produce more informative content compared to Android and Apple ones. This may be due to the market share of Blackberry systems, which have their widest penetration in the business sector. Conversely, Instagram, a newest social media platform very popular among the youngest generations, have an high discard rate (53.11%). Other authoritative sources (like *BBC News*) are not listed because, despite their low DR, they have few tweets in our dataset.

Now, we compare the baseline feature-set and the proposed new feature based on the source. We define the following feature groups:

- *B* is the baseline feature set (Base);
- *BB* is Base and source feature with the Binary approach;
- *BD* is Base and source feature with the Discard Rate;
- *BC* is Base and source feature with the Clustered ap-

TABLE III: Ranking of the TOP 14 sources by number of total *tweets*, listed in ascending order by Discard Rate (DR).

| Source | Tweets | DR |
|---------------------------|--------|--------|
| twitterfeed | 1420 | 4.93% |
| dlvr.it | 476 | 6.72% |
| Hootsuite | 1084 | 13.56% |
| Twitter for Websites | 920 | 14.02% |
| TweetDeck | 1619 | 15.26% |
| Facebook | 540 | 19.81% |
| Twitter Web Client | 3985 | 35.71% |
| Twitter for BlackBerry | 548 | 40.69% |
| UberSocial for BlackBerry | 167 | 41.32% |
| Twitter for iPad | 416 | 42.55% |
| Tweetbot for iOS | 154 | 44.81% |
| Twitter for Android | 959 | 52.55% |
| Twitter for iPhone | 2234 | 52.69% |
| Instagram | 209 | 53.11% |

proach.

Hence, we added each of the 3 proposed feature to the base ones, and evaluated each case separately. Furthermore we evaluate each feature group aggregating similar events (Earthquake, Fire, Hazard and Flood). We can state that the tweets’ source increase the classification performance for all metrics: accuracy, specificity and sensitivity. In Figure 5a, Fig. 5b, Fig. 5c we present the results obtained, highlighting the different metrics.

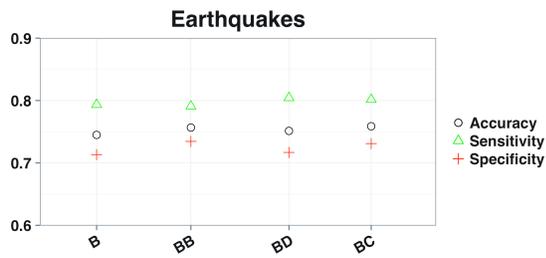
As shown in Table IV, the *Clustered* features achieves the better results, on average, improving the accuracy by 0.24%, 1.39% and 2.56% for floods, earthquakes and fire, respectively.

VII. CONCLUSIONS AND FUTURE WORKS

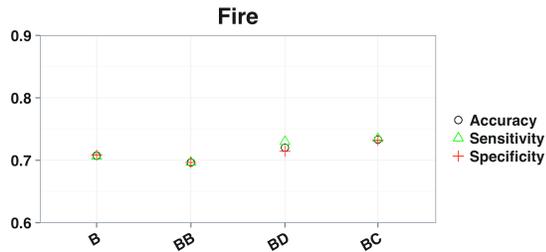
In this paper we designed a language-agnostic approach for the automatic classification of tweets during emergency in order to retain only informative content while excluding not informative ones. We tested the performance of our model using real labeled data from past emergencies, selecting relevant features from recent literature, and studied in details the effect of a new feature, called Source Feature, which improved the overall accuracy of the classifier from 0.24% (floods) up to 2.56% (fire).

TABLE IV: Classification Results

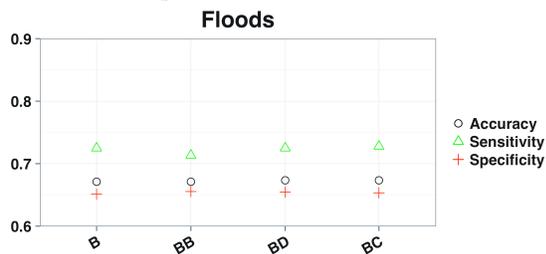
| Feature set | Earthquakes | Fire | Floods | Mean |
|-------------|---------------|---------------|---------------|---------------|
| Accuracy | | | | |
| B | 0.7449 | 0.7075 | 0.6710 | 0.7061 |
| BB | 0.7567 | 0.6965 | 0.6710 | 0.7064 |
| BD | 0.7513 | 0.7203 | 0.6734 | 0.7155 |
| BC | 0.7588 | 0.7331 | 0.6734 | 0.7179 |
| Sensitivity | | | | |
| B | 0.7935 | 0.7062 | 0.7249 | 0.7439 |
| BB | 0.7908 | 0.6967 | 0.7135 | 0.7389 |
| BD | 0.8043 | 0.7299 | 0.7249 | 0.7520 |
| BC | 0.8016 | 0.7346 | 0.7278 | 0.7525 |
| Specificity | | | | |
| B | 0.7133 | 0.7083 | 0.6512 | 0.6868 |
| BB | 0.7345 | 0.6964 | 0.6554 | 0.6898 |
| BD | 0.7168 | 0.7143 | 0.6543 | 0.6964 |
| BC | 0.7310 | 0.7321 | 0.6533 | 0.7004 |



(a) Results comparison for the Source Feature (Earthquakes).



(b) Results comparison for the Source Feature (Fire).



(c) Results comparison for the Source Feature (Floods)

Fig. 5: Classification results of the different feature groups by hazard type.

Future works will include the study of additional feature derived from the processing of images included in the *tweet* as well as the behavior in terms of *tweet* propagation. Furthermore, a multi-class classifier will be designed and evaluated in order to predict the information type according to emergency responder needs.

ACKNOWLEDGEMENT

This work was partially funded by the European Commission through the I-REACT project (H2020-DRS-1-2015), grant agreement n.700256.

REFERENCES

[1] "The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change" Climate Change, IPCC, 2013

[2] About Climate Change, The New York Times, http://www.nytimes.com/interactive/2015/11/28/science/what-is-climate-change.html?_r=0

[3] Climate change consequences, http://ec.europa.eu/clima/change/consequences/index_en.htm

[4] ICT Facts and Figures, 2013 <http://www.itu.int/en/ITU-D/Statistics/Documents/facts/ICTFactsFigures2013-e.pdf>

[5] World Disasters Report, International Red Cross and Red Crescent Movement, 2013, <http://www.ifrc.org/PageFiles/134658/WDR%202013%20complete.pdf>

[6] Aliza Rosen, Tweeting Made Easier, Twitter, https://blog.twitter.com/official/en_us/topics/product/2017/tweetingmadeeasier.html

[7] Twitter REST APIs, <https://dev.twitter.com/rest/public>

[8] Crisislex, <http://crisislex.org/>

[9] Olteanu, Alexandra and Vieweg, Sarah and Castillo, Carlos, "What to Expect When the Unexpected Happens: Social Media Communications Across Crises", Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing, Vancouver, BC, Canada, 2015, pp. 994-1009

[10] Mendoza, M. and Poblete, B. and Castillo, C., "Twitter under crisis: can we trust what we RT?", Proceedings of the First Workshop on Social Media Analytics, Washington D.C., District of Columbia, 2015, pp. 71-79

[11] Jackoway, A. and Samet, H. and Sankaranarayanan, J., "Identification of Live News Events using Twitter", Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Location-Based Social Networks, Chicago, Illinois, 2011, pp. 25-32

[12] Caragea, Cornelia and McNeese, Nathan and Jaiswal, Anuj and Traylor, Greg and Kim, Hyun-Woo and Mitra, Prasenjit and Wu, Dinghao and Tapia, Andrea H. and Giles, Lee and Jansen, Bernard J. and Yen, John, "Classifying Text Messages for the Haiti Earthquake", Proceedings of the 8th International ISCRAM Conference, Lisbon, Portugal, 2011, pp. 1-10

[13] Potts, Liza and Seitzinger, Joyce and Jones, Dave and Harrison, Angela, "Tweeting Disaster: Hashtag Constructions and Collisions", Proceedings of the 29th ACM International Conference on Design of Communication, Pisa, Italy, 2011, pp. 235-240

[14] Castillo, Carlos and Mendoza, Marcelo and Poblete, Barbara, "Information Credibility on Twitter", Proceedings of the 20th International Conference on World Wide Web (WWW), Hyderabad, India, 2011 pp. 675-684

[15] Vieweg, Sarah and Hughes, Amanda L. and Starbird, Kate and Palen, Leysia, "Microblogging During Two Natural Hazards Events: What Twitter May Contribute to Situational Awareness", Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Atlanta, Georgia, USA, 2010 pp. 1079-1088

[16] Kwak, Haewoon and Lee, Changhyun and Park, Hosung and Moon, Sue, "What is Twitter, a Social Network or a News Media?", Proceedings of the 19th International Conference on World Wide Web, Raleigh, North Carolina, USA 2010, pp. 591-600

[17] Brett Lantz, "Machine Learning with R, 2nd Edition", Packt, 2015

[18] Loh, Wei-Yin, "Classification and regression trees", Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, 2011, pp. 14-23

[19] Breiman, Leo, "Random Forests", Machine Learning, 2001, pp. 5-32

[20] Twitter: About verified accounts, <https://www.extension.iastate.edu/Documents/ANR/LikertScaleExamplesforSurveys.pdf>

[21] Likert Scale Examples for Surveys, Sorrel Brown, "Classification and regression trees", Iowa State University, <https://www.extension.iastate.edu/Documents/ANR/LikertScaleExamplesforSurveys.pdf>