Flood Relevance Estimation from Visual and Textual Content in Social Media Streams

Anastasia Mourtzidou, Stelios Andreadis, Ilias Gialampoukidis, Anastasios Karakostas, Stefanos Vrochidis, Ioannis Kompatsiaris

Information Technologies Institute - Centre for Research and Technology
Thessaloniki, Greece

{mourtzid, andreadisst, heliasgj, akarakos, stefanos, ikom}@iti.com

ABSTRACT

Disaster monitoring based on social media posts has raised a lot of interest in the domain of computer science the last decade, mainly due to the wide area of applications in public safety and security and due to the pervasiveness not solely on daily communication but also in life-threatening situations. Social media can be used as a valuable source for producing early warnings of eminent disasters. This paper presents a framework to analyse social media multimodal content, in order to decide if the content is relevant to flooding. This is very important since it enhances the crisis situational awareness and supports various crisis management procedures such as preparedness. Evaluation on a benchmark dataset shows very good performance in both text and image classification modules.

CCS CONCEPTS

• **Information systems** → **Online analytical processing**; *Data stream mining*;

KEYWORDS

social media, disaster monitoring, text classifier, visual classifier

ACM Reference Format:

Anastasia Mourtzidou, Stelios Andreadis, Ilias Gialampoukidis, Anastasios Karakostas, Stefanos Vrochidis, Ioannis Kompatsiaris. 2018. Flood Relevance Estimation from Visual and Textual Content in Social Media Streams. In *The 2018 Web Conference Companion, April 23–27, 2018, Lyon, France*. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3184558.3191620>

1 INTRODUCTION

The increasing popularity of social media has resulted in massive volumes of publicly available, user-generated multimodal content. Social media are not simply changing the way people communicate in their day-to-day lives, but also during disasters that endanger public health. Consequently, social media comprise an important source of information, which reports any major event including natural disasters [9]. This fact, coupled with various severe natural disasters that have taken place around the world such as the Haiti's 2010 earthquake, the 2010 Yushu earthquake, the 2010 Pakistan floods, the 2011 Tōhoku earthquake and tsunami, and the April 2015

Nepal earthquake, led to raising the interest of disaster monitoring based on social media in the domain of computer science.

It is observed that social media platforms, such as Twitter, are a rich source of information about real-world events, particularly during mass emergencies, from the citizens' point of view. The abundant nature of these data renders them as one of the most valuable sources to extract and deduct early warnings or identify an ongoing disaster [5]. Social media can support both early warnings mechanism and decision support systems since they offer real-time citizen observations, mainly textual and visual and they have been established as one of the most important communication channels.

In this work, we propose a framework for a social media monitoring tool that crawls, represents, and analyses content found in social media in order to decide whether the content is related or not to a natural disaster, using a combination of Deep Convolutional Neural Networks on visual content and Random Forests on textual features. The focus is on detecting flood events by using Twitter posts, mainly due to its real-time streaming nature, but it can be easily applied to other events and social media as well given that the framework presented is generic. The contribution of the work lies in the use of visual data, along with the textual, for determining whether the content is relevant to the disaster. The use of visual data also contributes to developing a language agnostic framework for an early warning mechanism.

2 RELATED WORK

There are several initiatives including projects and applications that exploit social media, such as Flickr tags [20], in order to create awareness about emergency situations or any other health related issues such as environmental conditions. First, within the hackAIR project [12], a platform has been developed for gathering and fusing environmental data and specifically Particulate Matter measurements from official sources and social media communities such as publicly available images shared through Instagram. In [21], the authors describe a framework that distinguishes between informational and conversational tweets shared during any major event and especially natural disasters. The framework uses a Naïve Bayes classifier for tweet classification and proposes the use of nine tweet-based features including emoticons, URLs, and instructional keywords. The framework was tested during hurricane Sandy and the results demonstrate that the nine features combined with the bag-of-words features (BoW) achieve over 85% accuracy. Another work with similar focus is that of [14] that aims at removing conversational data intermixed with informational data during natural disasters. The authors use the Geoparsing process that converts free text description to geographical coordinates in order to identify

This paper is published under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '18 Companion, April 23–27, 2018, Lyon, France

© 2018 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC BY 4.0 License.

ACM ISBN 978-1-4503-5640-4/18/04.

<https://doi.org/10.1145/3184558.3191620>

the relevant tweets. Eventually, in order to assess the severity of the natural disaster, sentiment analysis is performed. Furthermore, in [24], the authors analyse tweets generated during natural disasters, and apply burst detection for identifying early indicators of unexpected event, as well as classification and online clustering methods for filtering and summarising disaster-related information. The features used are the tweets' text itself and other tweet-related information such as mention and hashtag count, as also done in [1] for the estimation of the informativeness of a tweet. Another work towards this direction is that of [10], where the authors present a social media crisis mapping platform for natural disasters by using statistical analysis techniques for generating real-time crisis maps. They use locations from gazetteer, street map and volunteered geographic information sources for areas at risk of disaster which allows them to work at a building and street level resolution. Recently, Win et al. [22] introduced a tweet monitoring system that classifies messages in real-time by using LibLinear classifier and by considering linguistic features, sentiment lexicon based features and especially disaster lexicon based features. The performance was evaluated on four publicly available annotated datasets and showed that it outperformed the classifiers based on neural word embeddings and standard BoW models. Moreover, Fujitsu Laboratories developed technologies for disaster prevention and mitigation, by considering the knowledge of specialists. In [15], the authors describe an enhanced estimation technique involving social networking services, to quickly identify the location of a disaster.

Contrary to the above approaches, we follow a two-stage approach where relevance is assessed progressively. The classifier is a combination of Support Vector Machine (SVM) classification on visual features which are extracted using Deep Convolutional Neural Networks (DCNN) and Random Forests on textual features. In the sequel, we present briefly several approaches of visual and textual classification.

2.1 Image classification

Image classification involves the use of visual concept detection algorithms based on visual low-level features and classifiers for deciding whether an image shows evidence of flood.

Regarding low-level feature extraction, the most recent trend that seems to outperform all other previously developed methods is the DCNN-based features. DCNN-based features derive from the raw image pixels using Deep Convolutional Neural Networks (DCNNs), which consist of many layers of feature extractors and they can be used either as standalone classifiers, i.e., unlabeled images are passed through a pre-trained DCNN that performs the final class label prediction directly, or as generators of image features, i.e., the output of a hidden layer of the pre-trained DCNN is used as a global image representation [17], [8]. The latter type of features is referred to as DCNN-based. Several DCNN software libraries are available, such as Caffe [6], and different DCNN architectures have been proposed, such as GoogLeNet [19].

Classification involves the construction of models by using the low-level visual features, and then the application of these models for image labelling. Common classifiers that are used for learning the associations between the image representations and concept labels are the SVM and Logistic Regression (LR) [8].

2.2 Text classification

Text classification is a typical task in supervised machine learning and involves assigning categories to documents which can be documents, web pages etc. In general it involves the following steps: a) document collection, b) document preprocessing, which converts the original text data in a data-mining-ready structure, c) Text representation [23], which models documents and transforms them into numeric vectors. The most commonly model is the BoW model which can use different term weighting schemas such as the Boolean, the Term Frequency (TF), and the Term Frequency Inversed Document Frequency (TF-IDF). A more recent model version is word2vec [11], which comprises two-layer neural networks trained to reconstruct linguistic contexts of words and produce eventually word embeddings, and d) Serving the feature vector as input to a classifier (i.e. SVM, Naïve Bayes or Random Forests (RFs) [3]) which is tuned in order to achieve maximum performance.

Recently, text classification techniques that consider the characteristic features of short texts appearing in many areas such as Instant Messages and Twitter were developed. These texts are usually noisier, less topic-focused, shorter, and they contain many non-standard terms. Methods handling this type of texts include semantic analysis, semi-supervised short text classification, ensemble models for short text [18] and feature selection [2]. Moreover, some techniques target specifically Twitter due to its extensive use, e.g. [16] that considers the emoticons, the URLs existing in the text, the number of retweets and other features. However, given that these methods do not focus on the content are usually used for classifying text in more generic classes such as news, events and opinions.

3 PROPOSED FRAMEWORK

The proposed architecture describes a social media monitoring pipeline that collects effectively and in real-time posts from social media and specifically Twitter and classifies them as relevant or irrelevant to a natural disaster event. The classification considers both textual information and visual information (if available). The complete flow of the social media monitoring architecture is demonstrated in 1. The proposed framework involves the classification of each tweet in order to determine its relevancy to a specified natural disaster, i.e. flood. Two modalities are considered during classification; textual and visual. Visual classification is language-independent, since only visual characteristics are taken into account and in this manner it can be applied to any image retrieved regardless of the language of the tweet. Thus, if an image is uploaded along with the tweet, its visual features are extracted and it is then fed to a classifier (Section 3.1). In case the crawled tweet does not include an image or the result of the visual classifier is negative, the actual text is used to estimate the relevancy by using a text classifier (Section 3.2). In the sequel, we present an overview of the visual and textual classification modules.

3.1 Social media image classification framework

In the employed framework, we train a 22-layer GoogLeNet network [19] on 5055 ImageNet concepts [13], which are a subset of the 12,988 ImageNet concepts. Then, this network is applied on

Figure 1: Social media monitoring module architecture.

the TRECVID SIN 2013¹ development dataset and the output of the last pooling layer with dimension 1024 is used for global image representation. We use the annotated dataset for training and validating an SVM classifier. The SVM classifier can be tuned by setting different t and g values in order to achieve maximum performance. The parameter t in SVM classifiers defines the kernel type, while g stands for the gamma in the kernel function.

3.2 Social media text classification framework

Image classification is supported by text classification to assess the relevance of a social media post. In the employed framework, we evaluated several methods belonging to the traditional text classification, as well as the Jaccard similarity method. Jaccard similarity method [7] is a measure that concludes whether a document belongs to a specific class by calculating its similarity with the instances that are known to belong to the class. The maximum of the similarity calculated between the query document and the set of documents belonging to the class of interest is compared to a threshold value that is defined empirically in order to decide whether the query document belongs or not to the specific class. Similar measures are the Manhattan distance, the cosine similarity, and the Euclidean distance.

Thus, for the traditional text classification we approach each of the aforementioned steps as follows:

- (1) Collection of short text messages from Twitter Streaming API².
- (2) Pre-processing of the collected text by either applying DBpedia Spotlight that automatically annotates it with respective DBpedia resources [4] which can also be used as high-level textual features or simple removing punctuation, non-characters, stop words and word stemming.
- (3) Text representation using Term Frequency (TF), TFIDF and word2vec. Various experiments were realized for different feature length and n-gram values for the BoW methods, and different corpus and vector dimensions for the word2vec method.
- (4) Serving of feature vector as input to a classifier (i.e. SVM, Naïve Bayes or Random Forests (RFs)) which is tuned in order to achieve maximum performance.

¹<http://www-nlpir.nist.gov/projects/tv2013/tv2013.html>

²<https://developer.twitter.com>

Table 1: Statistics of MediaEval 2017 dataset

	Annotation for concept 'flood'		Sum
	True	False	
Train set	1280	2240	3520
Validation set	640	1120	1760
Total Records	1920	3360	5280

Regarding the Jaccard similarity approach, after collecting and preprocessing the data in the same manner as previously, we calculate the Jaccard similarity coefficient between the new text description and each positively annotated text description, using $J(W_q, W_{t_n}) = \frac{|W_q \cap W_{t_n}|}{|W_q \cup W_{t_n}|}$, where W_q stands for the set of terms of the new text description, and W_{t_n} for the set of terms of the n text description of the positively annotated dataset tests. Then, the maximum value of the Jaccard similarity coefficients is compared to a threshold defined empirically in order to determine whether the new text description will be considered as positive or not.

In the following, we examine the performance of the proposed framework and we tune the involved parameters.

4 EVALUATION

This section describes the datasets where the text and image classification modules are evaluated, and the experiments conducted. The proposed framework is evaluated for the flood event, but it can be extended to any other events.

4.1 Dataset Description

The dataset that was used for developing and evaluating both the visual and textual classification modules, is the MediaEval 2017 dataset³. It was provided within the context of the Disaster Image Retrieval from Social Media (DIRSM) subtask whose goal was to identify all images that show direct evidence of a flooding event. Along with the dataset a set of visual descriptors were also precomputed and provided to the contestants which were evaluated during the building of the visual classifier. Table 1 contains the statistics of the MediaEval2017 dataset.

4.2 Experiments

The evaluation of the visual classification part is done by using precision, recall, and F-score. These metrics are calculated in every run in order to decide the best performing classification method.

4.2.1 Social media image classification. In order to find the best performing feature and classifier for identifying images that contain evidence of flood, several features are compared and the parameters of SVM classifiers are tuned in order to maximise their performance. In detail, the features provided by the Multimedia-Satellite challenge were tested (i.e. acc, fcth, jcd, cedd, eh, sc, cl, and tamura) and the DCNN-based features that were produced from our framework. SVM classifiers were trained for each feature for different parameters and results showed that the proposed DCNN feature outperformed almost all of them (see Table 2). The SVM parameters

³<https://multimediaeval.github.io/2017-Multimedia-Satellite-Task/>

that were tuned were t that defines the type of kernel function and can be set to linear, polynomial, radial and sigmoid and g that is the gamma parameter in the kernel function. Consequently, the best results were obtained for the DCNN-based features for $t = 1$ (polynomial function) and $g = 0.5$ or $g = 0.03125$. Figure 3 depicts the top 18 results returned by the classifier.

Table 2: Evaluation of visual features and classifiers.

Descriptors	SVM parameters		Precision	Recall	F-score
	t	g			
acc	1	0,03125	0,5359	0,1516	0,2363
acc	1	0,5	0,4830	0,1328	0,2083
cedd	1	0,03125	0,6085	0,5391	0,5717
cedd	1	0,5	0,5925	0,3953	0,4742
cl	1	0,03125	0,6115	0,3641	0,4564
cl	1	0,5	0,5957	0,3016	0,4004
eh	1	0,03125	0,6682	0,4688	0,5510
eh	1	0,5	0,6605	0,4469	0,5331
fcth	1	0,03125	0,5956	0,4625	0,5207
fcth	1	0,5	0,5000	0,2578	0,3402
jcd	1	0,03125	0,6388	0,5250	0,5763
jcd	1	0,5	0,6025	0,3719	0,4599
sc	1	0,03125	1,0000	0,0016	0,0031
sc	1	0,5	0,2500	0,0016	0,0031
tamura	1	0,03125	0,5246	0,0500	0,0913
tamura	1	0,5	0,3913	0,0141	0,0271
dcnn-based	1	0,03125	0,8195	0,8016	0,8104
dcnn-based	1	0,5	0,8195	0,8016	0,8104

4.2.2 Social media text classification. In the sequel, the traditional text classification methods are evaluated. In all cases, three classifiers, namely SVM, Naïve Bayes, and Random Forests, are tested for a set of parameters. Specifically, for the case of SVM the penalty parameter and the kernel type is tested, for the Naïve Bayes the additive smoothing parameter is adjusted and finally for the RFs the parameters that are tested are the number of trees in the forest and the number of features used for best split. For the remaining parameters, default values are used. Moreover, regarding the methods using TF and TF-IDF representation, different n -gram values and min_df values are considered during text vectorisation. The min_df value affects the size of the feature length since it ignores terms that have a document frequency strictly lower than the given threshold when building the vocabulary. Table 3 contains the best results of the TF and the TF-IDF representation methods for each different classifier for different text inputs.

Regarding the word2vec methodology several runs were performed for a set of different parameters including the *size* parameter that defines the dimension of the feature vector, the *window* parameter that defines the maximum distance between the current and predicted word within a sentence and the *sg* parameter that defines the training algorithm that can either skip-gram or CBOW. Specifically, the values that are tested for the aforementioned parameters are the following: *size* = {100, 200, 300, 400, 500}, *window* = {2,3} and *sg* = {0,1}.

Table 4 contains the best results of the word2vec method. The sizes of the corpora used are 6,600 records for the mediaEvalFloods_corpus, and around 830,000 records for the twitterFloods_corpus which was produced by crawling tweets that include the keyword “flooding”. After a careful observation of the table, we conclude that the larger corpus achieves better performance. The best runs from all tables along with information concerning the text representation parameters and the classifier parameters can be found in Table reftab:bestparams. From the table we can deduce that the best performing method according to F-score is the TF method, while the second best performing method is the word2vec which has slightly lower F-score but significantly better precision.

Finally we evaluated for text classification uses the Jaccard Similarity Coefficient. Figure 2 depicts the F-score values for different text input and different values of the e threshold parameter. After a careful observation of the figure, we conclude that for the Jaccard similarity with DBpedia concepts slightly improves the classification performance. Moreover, it is evident that the method has good results for very low values of the e parameter, i.e. around 0.1 and drops significantly after that value. The main disadvantage of the Jaccard method is that it is rather slow compared to the other methods during the evaluation of new incoming text, given that this text must be compared against all positively annotated texts in order to determine its relevancy with them. However, since the text classification method is part of the social media monitoring pipeline that is triggered very regularly (usually around 1 second), it is not considered an optimal solution.

4.2.3 Social media visual and textual classification. In this section we present the evaluation of the total framework, which is a two-stage approach that considers both visual features and textual features. Two simple ways of combining the two modalities are examined that arise from the boolean operations AND and OR. The first is based on the boolean operation OR and is a sequential mode where, the result of the visual classifier determines initially the relevancy and in case it is negative the result of the textual classifier is examined as well in order to determine the relevancy of tweet. The aim of using a sequential mode is to avoid as much as possible running unnecessary processes (i.e. two classifiers) and thus consuming more time. The second is based on the boolean operation AND and evaluates simultaneously both modes which should be both positive in order for the tweet to be considered as relevant.

Table 6 contains the evaluation results of both combined methods and the single modalities as well in order to allow fast comparison. We can deduce that the combined mode based on the OR boolean operation performs better given that the F-score is higher compared to the single modes and the simultaneous combined mode. After a careful observation of the table, it is evident that there is a significant decrease in the precision of the framework which is balanced by a significant increase in the recall. On the other hand the simultaneous combined mode has slightly better F-score than the single modes but has a much higher precision and lower recall which is expected given that the two modes are combined in an AND boolean way. However, a disaster monitoring systems would favor higher recall given that it is more important identifying more

Table 3: Evaluation of TF and TF-IDF representation method.

n-gram	Text-input	Classifier	TF			TF-IDF		
			Precision	Recall	Fscore	Precision	Recall	Fscore
1	Without stop words	SVM	0,72655	0,85938	0,78740	0,71120	0,82344	0,76322
		Naive Bayes	0,78522	0,71406	0,74795	0,76311	0,65938	0,70746
		RandomForest	0,74098	0,89844	0,81241	0,73359	0,89063	0,80452
	Without stop words & with stemming	SVM	0,73986	0,48438	0,58546	0,69727	0,43906	0,53883
		Naive Bayes	0,78776	0,60313	0,68319	0,77186	0,56563	0,65284
		RandomForest	0,76087	0,76563	0,76324	0,76973	0,74688	0,75813
	DBPedia concepts	SVM	0,66185	0,35781	0,46450	0,66462	0,33750	0,44767
		Naive Bayes	0,72624	0,50156	0,59335	0,69752	0,48281	0,57064
		RandomForest	0,66983	0,66250	0,66614	0,69581	0,59688	0,64256
	Without stop words	SVM	0,72267	0,84688	0,77986	0,70572	0,80938	0,75400
		Naive Bayes	0,77634	0,74844	0,76213	0,76056	0,67500	0,71523
		RandomForest	0,74804	0,89531	0,81508	0,74443	0,88750	0,80969
2	Without stop words& with stemming	SVM	0,74485	0,45156	0,56226	0,69825	0,43750	0,53794
		Naive Bayes	0,78044	0,61094	0,68536	0,77419	0,56250	0,65158
		RandomForest	0,76837	0,75156	0,75987	0,76056	0,75938	0,75997
	DBPedia concepts	SVM	0,66382	0,36406	0,47023	0,66875	0,33438	0,44583
		Naive Bayes	0,72955	0,50156	0,59444	0,70000	0,48125	0,57037
		RandomForest	0,67742	0,65625	0,66667	0,66831	0,63594	0,65172

Table 4: Evaluation of word2vec representation method.

Text input	Corpus	Vector dimension	Words windows	Training algorithm	Precision	Recall	Fscore
without stop words	mediaEvalFloods_corpus	100	3	1	0,75835	0,74531	0,75177
text with stop words removed	twitterFloods_corpus	200	3	0	0,79341	0,82813	0,81040
without stop words & with stemming	mediaEvalFloods_corpus	100	3	1	0,76167	0,71406	0,73710
without stop words & with stemming	twitterFloods_corpus	200	3	0	0,77647	0,82500	0,80000
DBPedia concepts	mediaEvalFloods_corpus	100	2	1	0,75455	0,77813	0,76615
DBPedia concepts	twitterFloods_corpus	[100 - 500]	[2,3]	[0,1]	0,86667	0,02031	0,03969

Table 5: Best parameters from TF, TF-IDF, word2vec text classification methods.

Method	Text Input	Parameters	Classifiers & parameters	Precision	Recall	Fscore
TF	without stop words	n-gram = 2, min_df = 0,003, features length = 1068	Random Forest Features num: auto Number of trees: 200	0,74804	0,89531	0,81508
TF-IDF	without stop words	n-gram = 2, min_df = 0,003, features length = 1068	Random Forest Features num: auto Number of trees: 500	0,74443	0,88750	0,80969
word2vec	without stop words	corpus = twitterFloods_corpus, vector dimension = 200, words window = 3, training algorithm = 0	SVM Penalty parameter: 5.0 Kernel type: rbf	0,79341	0,82813	0,81040

cases that require attention than ignoring such messages in favor of a higher precision.

5 CONCLUSIONS

This work presents an original framework that assesses the relevance of a social media post to a target event, such as floods. Our framework contributes to the crisis management procedures before,

Table 6: Evaluation results of single modalities and combined modalities.

Method	Precision	Recall	Fscore
Single mode using text data (word2vec)	0,79341	0,82813	0,81040
Single mode using visual data	0,81949	0,80156	0,81043
Sequential combined mode (OR)	0.73659	0.94375	0.82740
Simultaneous combined mode (AND)	0.91245	0.73281	0.81282

Figure 2: Evaluation of Jaccard Similarity method. **Figure 3: Visualisation of DCNN-based visual classifier validation results.**

during and after the crisis and can be integrated in crisis management and decision support systems. Experiments on the MediaEval benchmark dataset and crawled posts from Twitter involving visual and textual information have shown that the best performing method in the case of image classification is the use of DCNN-based features together with an SVM classifier, while in the case of text classification stand out the TF-IDF method for text representation and Random Forests as classifier. As a future work, we plan to fuse visual and textual features, as well as combine traditional text classification techniques with short text classification techniques that additional social media metadata. Finally, we plan to evaluate the framework on other events including fires and heatwave as well as in order languages.

ACKNOWLEDGMENTS

This work was partially supported by the European Commission under contracts H2020-700475 beAWARE and H2020-776019 EOPEN.

REFERENCES

- [1] Flavia Sofia Acerbo and Claudio Rossi. 2017. Filtering Informative Tweets During Emergencies: A Machine Learning Approach. In *Proceedings of the First CoNEXT Workshop on ICT Tools for Emergency Networks and Disaster Relief (I-TENDER '17)*. ACM, New York, NY, USA, 1–6. <https://doi.org/10.1145/3152896.3152897>
- [2] Charu C Aggarwal and ChengXiang Zhai. 2012. *Mining text data*. Springer Science & Business Media.
- [3] Leo Breiman. 2001. Random forests. *Machine learning* 45, 1 (2001), 5–32.
- [4] Joachim Daiber, Max Jakob, Chris Hokamp, and Pablo N Mendes. 2013. Improving efficiency and accuracy in multilingual entity extraction. In *Proceedings of the 9th International Conference on Semantic Systems*. ACM, 121–124.
- [5] Muhammad Imran, Carlos Castillo, Fernando Diaz, and Sarah Vieweg. 2015. Processing social media messages in mass emergency: A survey. *ACM Computing Surveys (CSUR)* 47, 4 (2015), 67.
- [6] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. 2014. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the 22nd ACM international conference on Multimedia*. ACM, 675–678.
- [7] Vijaymeena M K and Kavitha K. 2016. A Survey on Similarity Measures in Text Mining. 3 (03 2016), 19–28.
- [8] Foteini Markatopoulou, Vasileios Mezaris, and Ioannis Patras. 2015. Cascade of classifiers based on binary, non-binary and deep convolutional network descriptors for video concept detection. In *Image Processing (ICIP), 2015 IEEE International Conference on*. IEEE, 1786–1790.
- [9] Raina M Merchant, Stacy Elmer, and Nicole Lurie. 2011. Integrating social media into emergency-preparedness efforts. *New England Journal of Medicine* 365, 4 (2011), 289–291.
- [10] Stuart E Middleton, Lee Middleton, and Stefano Modafferi. 2014. Real-time crisis mapping of natural disasters using social media. *IEEE Intelligent Systems* 29, 2 (2014), 9–17.
- [11] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*. 3111–3119.
- [12] Anastasia Mourtzidou, Symeon Papadopoulos, Stefanos Vrochidis, Ioannis Kompatsiaris, Konstantinos Kourtidis, George Hloupis, Ilias Stavrakas, Konstantina Papachristopoulou, and Christodoulos Keratidis. 2016. Towards Air Quality Estimation Using Collected Multimodal Environmental Data. In *International Workshop on the Internet for Financial Collective Awareness and Intelligence*. Springer, 147–156.
- [13] Nikiforos Pittaras, Foteini Markatopoulou, Vasileios Mezaris, and Ioannis Patras. 2017. Comparison of fine-tuning and extension strategies for deep convolutional neural networks. In *International Conference on Multimedia Modeling*. Springer, 102–114.
- [14] MV Sangameswar, M Nagabhushana Rao, and S Satyanarayana. 2017. An algorithm for identification of natural disaster affected area. *Journal of Big Data* 4, 1 (2017), 39.
- [15] Hitoshi Sato, Kunihiro Takeda, Kazuhiro Matsumoto, Hirokazu Anai, and Yuzuru Yamakage. 2016. Efforts for Disaster Prevention/Mitigation to Protect Society from Major Natural Disasters. *FUJITSU Sci. Tech. J* 52, 1 (2016), 107–113.
- [16] P Selvaperumal and A Suruliandi. 2014. A short message classification algorithm for tweet classification. In *Recent Trends in Information Technology (ICRTIT), 2014 International Conference on*. IEEE, 1–3.
- [17] Karen Simonyan and Andrew Zisserman. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. *CoRR* abs/1409.1556 (2014). [arXiv:1409.1556](http://arxiv.org/abs/1409.1556) <http://arxiv.org/abs/1409.1556>
- [18] Ge Song, Yunming Ye, Xiaolin Du, Xiaohui Huang, and Shifu Bie. 2014. Short Text Classification: A Survey. *Journal of Multimedia* 9, 5 (2014).
- [19] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1–9.

- [20] Nataliya Tkachenko, Stephen Jarvis, and Rob Procter. 2017. Predicting floods with Flickr tags. *PloS one* 12, 2 (2017), e0172870.
- [21] Brandon Truong, Cornelia Caragea, Anna Squicciarini, and Andrea H Tapia. 2014. Identifying valuable information from twitter during natural disasters. *Proceedings of the Association for Information Science and Technology* 51, 1 (2014), 1–4.
- [22] Si Si Mar Win and Than Nwe Aung. 2017. Target oriented tweets monitoring system during natural disasters. In *Computer and Information Science (ICIS), 2017 IEEE/ACIS 16th International Conference on*. IEEE, 143–148.
- [23] Jun Yan. 2009. *Text Representation*. Springer US, Boston, MA, 3069–3072. https://doi.org/10.1007/978-0-387-39940-9_420
- [24] Jie Yin, Andrew Lampert, Mark Cameron, Bella Robinson, and Robert Power. 2012. Using social media to enhance emergency situation awareness. *IEEE Intelligent Systems* 27, 6 (2012), 52–59.