Deep Learning for Social Media Analysis in Crises Situations

Mehdi Ben Lazreg Morten Goodwin Ole-Christoffer Granmo

University of Agder, Grimstad, Norway

Abstract

Social media has become an important open communication medium during crises. This has motivated much work on social media data analysis for crises situations using machine learning techniques but has mostly been carried out by traditional techniques. Those methods have shown mixed results and are criticised for being unable to generalize beyond the scope of the designed study. Since every crisis is special, such retrospect models have little value. In contrast, deep learning shows very promising results by learning in noisy environments such as image classification and game playing. It has, therefore great potential to play a significant role in the future social media analysis in noisy crises situations. This position paper proposes an approach to improve the social media analysis in crises situations to achieve better understanding and decision support during a crisis. In this approach, we aim to use Deep Learning to extract features and patterns related to the text and concepts available in crisis related social media posts and use them to provide an overview of the crisis.

Keywords: deep learning, social media, crises situations

1 Introduction

A vast variety of natural and human-caused crises occur around the world. The diversity and immediacy of these crises cause severe challenges not only for the people affected and responders, but also to the research community. Some of the unsolved research challenges include:

- How can machine learning be used to detect a crisis as soon as it occurs from external sources including social media?
- How can machine learning use crisis related social media data to acquire information about a crisis' status and progression?
- How can artificial intelligence support the responders in making the correct decisions during a crisis?

Social media plays a pivotal role in most crises today, from getting life signs from people affected to communicating with responders. However, most research using social media in crises situations are one-off solutions with a specialised technique or addressed area [1]. A one-off solution means that the research focuses on finding a technique that yields the best result in a specific crisis. The technique fails to generalize beyond the study. This finding has several implications due to the distinct nature of crisis. Firstly, crises are diverse, and range from natural, technological, financial to political crises, and even previously unforeseen types of crises. Secondly, crises evolve through time. Different aspects of a crisis change as time passes, which means what is learnt earlier on may not be applicable later in the crisis. Finally, crises are unpredictable in nature as unexpected event may occur.

Deep Learning (DL) has the potential to improve social media analysis in crises situations because of its ability to learn patterns from unlabelled data [17]. This property has enabled DL to produce breakthroughs in the domain of image, text and speech recognition. Moreover, DL has the ability to generalize learnt patterns beyond data similar to the training data, which can be advantageous while dealing with social media analysis in crises situations. Despite the breakthroughs brought by DL, improvements are still to be made the further optimise it and improve its performance [14]. This paper proposes to explore the uninvestigated area of how the emerging advantages of DL can be expanded upon to address the pertinent challenges of evolving crisis analytics for social media.

This paper is organised as follows. In section 2 we discuss the use of social media in crises situations with a special focus on the use of machine learning. Section 3 continues with Deep Learning, and section 4 proposes an approach for applying Deep Learning to for social media analysis. Finally, section 5 concludes.

2 Machine Learning in Social Media Analysis in Crises Situations

Social media has become an open crises communication medium. As an example, during the tsunami in the Philippines in 2012, 558126 tweets were produced in 8 different languages in the course of the seven first hours following the crises [2]. Similarly, 20000 tweets/day were registered midst the 2012 Sandy storm in New York, and 5000 tweets/second were reported during the 2011 earthquake in Virginia (US) [2].

There is no doubt that valuable, high throughput data is produced on social media only seconds after a crisis occurs. However, processing and inferring valuable knowledge from such data are difficult for several reasons. The messages are typically brief, informal, and heterogeneous (mix of languages, acronyms, and misspellings) with varying quality, and it is often required to know the context of the message to understand its meaning. Moreover, posts on other mundane events are also part of the data, which introduces additional noise for training. To address the challenges of detecting and classifying a crisis in heterogeneous data, supervised and unsupervised machine learning techniques were used.

2.1 Supervised learning

In order to classify a social media message as part of one particular crisis event, several features related to the message need to be used, including the nature of the message (factual, emotional or subjective), the information provided, the information source, credibility, time and location. Note that some of the features can be automatically extracted, but others need human labelling.

From the message examples, the supervised learning algorithm learns a predictive function (representing the relation between the features and particular crises) so that it can classify any new unknown message as part of one of the categories of crises. Several approaches have been applied for this: Naïve Bayes and Support Vector Machine (SVM) [5][6], Random Forests [7], and Logistic Regression [8]. Further, to mitigate the complexity of social media messages, some research focuses on only analysing tweets with certain tags [4]. As example, the tag “breaking news” and “news” can be used to identify breaking news tweets [9]. In the same way, the occurrence of the words “earthquake” and “landslide” were used as features in a SVM classifier to classify earthquakes [8] and landslides [10] respectively.

However useful in reducing the complexity of the data, this approach neglects the potential valuable information contained in the text such as the status of the crisis, potential victims, needed resources, and so on. In a supervised approach, (human made) labels are necessary for training the classifiers, but they might be highly difficult to obtain, especially in case of multi-language messages or context knowledge [4]. Furthermore, such labels are not always reliable and may not be available at the time of the crisis. Moreover, reusing a classifier trained on data from previous disasters may not perform well in practise and intuitively returns a loss of accuracy even if the crises have a lot in common. To solve this issue research has been carried out to use unsupervised learning techniques.

2.2 Unsupervised learning

Unsupervised methods are used to identify patterns in unlabelled data. They are most useful when the information seekers do not know specifically what information to look for in the data – which is the case in many crises situations. An example is grouping tweets into stories (clusters of tweets) after a keyword filter [11]. This method reduces the number of social media messages to be handled by humans since it groups equivalent messages together. Another application using unsupervised learning identifies events related to public and safety with a spatio-temporal clustering approach

[12]. In addition to strictly clustering elements into groups, soft clusters have been used to allow items to simultaneously belong to several clusters with variant degrees. In this methods, the tweets similarity is based on words they contain and the length of the tweets [13]. The approach was applied on data from the Indonesia earthquake (2009) and it detected aspects related to the crisis (relief, deaths, missing persons, and so on).

3 Deep Learning

Imitating the efficiency of the human brain has been a huge challenge for the artificial intelligence field. The emergence of DL has fuelled a paradigm shift and made it a more achievable goal. DL is a machine learning technique with its roots in Neural Networks that allows learnt models composed of multiple processing layers so that the knowledge state has several layers of abstraction. DL is particularly valuable because it is shown to find complex structures in large data using an algorithm to update its internal representation of each layer in a way that other state-of-the-art algorithms are not [14]. Conventional supervised machine learning techniques require careful engineering to transform raw data into suitable features for classification of inputs, whereas, DL fed with raw data, discovers the representation and features needed for detection and classification. DL techniques with back propagation and deep convolutional nets have brought breakthroughs in image processing while recurrent nets have brought amazing advancements in sequential data analysis such as text and speech recognition [15]. DL is used to analyse X-ray images to detect potential diseases [16], and to recognise handwriting [17]. DL is applied to online tasks perhaps most notably, it was the first AI machine to beat a human expert in the game of Go – by most AI scholars considered one of the most complex games for artificial intelligence [18]. Further, DL is successfully applied to text mining to organise text documents in databases by topic [19], to analyse costumers review on a given product and deduce what they think about it [20]. Moreover, it has been explored in chemical text mining to recognise drugs and chemical compounds [21], and sentiment analysis [22]. However, DL has to a very little degree been explored for crises management [4].

Despite the breakthroughs brought by DL, using DL for unsupervised learning has not been much explored and was for a long time overshadowed by the success of supervised learning [14]. Unsupervised learning is important to explore since, in that paradigm, the AI machines discover structures by observing the data without being told what each feature in the data represents.

4 Social Media Analysis

This paper propose an approach to improve the social media analysis in crises situations to achieve better understanding and decision support during a crisis. The approach is summarised in Figure 1 and consists of moving from low to high level of abstraction. We plan to proceed from: Using DL to transform non-standard words into their canonical form (1). Then, understand the semantic of the text (2). Finally, provide an overview of a crisis development (3).

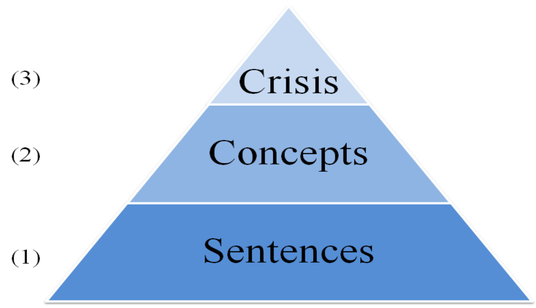

Figure 1: Overview of the proposed approach

4.1 Analysis on words and sentences

The machine needs to recognise words in the sentence to understand it. In social media text this task is challenging since the messages typically contain misspellings, abbreviations, deletions, and phonetic spellings. Traditional supervised learning approaches are extremely dependent on the correctness of the training set. To learn diverse functionality, training sets that represent each category of the data are required. The number of ways a word can be misspelled is huge, which means that traditional approaches fall short [23]. Unsupervised approaches can find similarities between spelling variations of words using clusters containing the correct word as well as its different misspellings. Nevertheless, unsupervised learning also falls short because they depend on rigid metrics for similarity that influence the clusters [23].

To address this challenge, we plan to use DL to deduct high-level abstractions (e.g. meaning) from lower levels (e.g. letters or sub-string). A new configuration at low abstraction level may then lead to a better representation of words and similarities between different spellings. The assumption is that even with limited training sets, a new example could be meaningfully represented using the low-level abstractions. We will look at ways to modify the basic DL algorithm so it can detect features that relate non-standard versions of words or phrases, which will be used to deduct the correct text. For this first phase, we will train and test the developed algorithm on collected texts from crisis events on Twitter. The text contains ungrammatical sentences, non-standard and misspelled words. The algorithm will be evaluated based on correctness metrics and error rates of non-standard words that the algorithm fails to recognise. Due to DL ability to be applied in noisy environments, we predict that this approach will outperform state-of-the-art in the field of text mining.

4.2 Identify and understand concepts

After word and sentence analysis is carried out, the proposed approach moves to a higher level of abstraction: Classifying concepts from multiple social media messages. Explicitly, this means identifying what a writer tries to express (e.g. informing about a situation, crying out for help, express explicit needs, and so on) and understanding concepts from the message. The state-of-the-art in this area mostly centres on supervised learning techniques by training the algorithm on a set of text on each topic to learn a predictive function, which in turn is used to classify a new topic into a previously learnt topic [24]. A limitation of this approach is the scope of predefined topics: If a text about an unforeseen topic is presented to the algorithm, such as a new crisis, it will wrongly classify it as one of the existing topics. A challenge is that crises are diverse, and the number of topics discussed in social media during a single crisis is big, dynamic, and changing. The complexity of social media data makes getting human labelled data for each topic very expensive and time-consuming. Adversely, unsupervised techniques try to look for co-occurrences of terms in the text as a metric of similarity [26] and inferring the word distribution in the set of word the text contains and using their frequencies for document clustering [27].

Using DL to understand a sentence or a document is an ongoing research topic in which progress is still to be made. DL has been used to predict the next word on a sequence of semantically related words [14], and this ability suggests that DL have learnt a semantic representation of the words. DL also has some success in predicting the next character in a sequence of characters which is used to generate text, and in machine translation [14]. To address this challenge we aim to use DL on the normalised text to automatically learn distinct features for each discussed concept. Further, we intend to find a strategy to decompose a document into concepts, segmenting the text into semantically meaningful atomic units. By identifying the underlying concepts of a document or a sentence, a deeper understanding is established. This lies the foundation for building crisis understanding. We will investigate ways to improve unsupervised DL for concept discovery. A semi-supervised method can be used to improve the performance of the data representation. When the model is actually able to represent the unlabelled data, labels can be added to transform the problem from an unsupervised to a supervised learning problem. For this model, we will gather data from social media platforms on crises events, including Twitter, transform the noisy data into a workable form using the approach described in Section 4.1, and use this data to train and test our model. We will empirically verify the result of each test, which will help us understand how the model performed. We will base our evaluation on established correctness metrics and error rates of topic or concepts that the algorithm fails to recognise.

4.3 Crisis understanding

Even though the crisis data is valuable with high throughput, it is small compared to the 7Gb/min of data produced by Twitter alone [3]. Hence, this topic integrates two areas: DL and big data (large and complex datasets that cause traditional data processing application to be inadequate). The aim of this phase is to detect crisis patterns in social media text that can be used to retrieve crisis related messages from big social media data in a way that it gives an overview of its status.

The state-of-the-art in the use of machine learning on social media in a crisis (see Section 2) are one-off solutions with specialised techniques or addressed areas, including the labels needed for training the classifiers which are not always reliable or available. Also, reusing a classifier trained on data from previous disasters may not perform well in practise and intuitively returns a loss of accuracy even if the crises have a lot in common.

One of the key features of DL is the analyses of a big amount of unsupervised data [25], which makes it valuable for big data analytics with unlabelled and uncharacterized data. DL can be used to address the important problems including extracting patterns from massive data, and information retrieval. It can provide a generic solution that infers similarity and dissimilarity patterns between different crises. Nonetheless, DL algorithms can become computationally expensive when dealing with high dimensional data due to its deep layered hierarchy and number of parameters to learn. The computing expensiveness becomes more of a problem in social media where the data is streaming rapidly and changing fast. Methods for incremental learning have been developed to deal with this challenge that includes the use of DL [28]. To address this challenge question, we will use and expand incremental DL to infer information about rare events (crises data) in a mix of a massive and diverse data which, in this stage, include the concepts learnt previously and metadata (including time, location and writer of a tweet). The result will be presented in a spatiotemporal overview of the crisis. A spatiotemporal overview presents the statue of crisis in a location at different points in time, an approach which to very little degree has been investigated [4]. The overview will be used as a decision support system to help take the most appropriate actions to resolve the crisis. We will present the model with a set of crisis practitioners that will test it and provide inputs (in the form of a survey) on how helpful this model would be in the case of crisis. Correctness metrics and error rates of the algorithm will also inform of its abilities.

5 Conclusion

This position paper presents an overview of machine learning techniques used for social media analysis in crises situations today. The current approaches, based on traditional machine learning techniques, are heavily criticised for being one-off studies which cannot be generalized. Since every crisis is special, such retrospect models have little value.

Deep Learning (DL) has the potential to mitigate this problem since it has been shown to very good at generalizing. The paper presents a possible approach for applying DL to crises analysis. The model starts with normalizing social media data. This includes mapping noisy words to the original word. Further, the normalized text can be used to deduce the concepts and the topic of a cluster of texts. Finally, the texts related to a crisis situation are retrieved and a spatiotemporal representation of the crises of the crises is produced based on those texts. In this way, DL can be used to offer a decision support system to crises responders to help them better understand a crisis situation and produce more efficient decisions than traditional machine learning techniques.

References

[1] Maresh-Fuehrer, M.M. and R. Smith, Social media mapping innovations for crisis prevention, response, and evaluation. Computers in Human Behavior, 2016. 54: p. 620-629.

[2] Zielinski, A., et al. Social media text mining and network analysis for decision support in natural crisis management. in ISCRAM 2013 Conference Proceedings - 10th International Conference on Information Systems for Crisis Response and Management. 2013.

[3] Valkanas, G., et al. Mining Twitter Data with Resource Constraints. in Web Intelligence (WI) and Intelligent Agent Technologies (IAT), 2014 IEEE/WIC/ACM International Joint Conferences on. 2014.

[4] Imran, M., et al., Processing Social Media Messages in Mass Emergency: A Survey. CoRR, 2014. abs/1407.7071.

[5] Yin, J., et al., Using social media to enhance

emergency situation awareness. IEEE Intelligent Systems, 2012. 27(6): p. 52-59.

[6] Power, R., et al., Emergency situation awareness: Twitter case studies, in Information Systems for Crisis Response and Management in Mediterranean Countries. 2014, Springer. p. 218-231.

[7] Imran, M., et al. Aidr: Artificial intelligence for disaster response. in Proceedings of the companion publication of the 23rd international conference on World wide web companion. 2014. International World Wide Web Conferences Steering Committee.

[8] Ashktorab, Z., et al., Tweedr: Mining twitter to inform disaster response. Proc. of ISCRAM, 2014.

[9] Phuvipadawat, S. and T. Murata. Breaking news detection and tracking in Twitter in Web Intelligence and Intelligent Agent Technology (WI-IAT). IEEE/WIC/ACM International Conference on. 2010.

[10] Musaev, A., D. Wang, and C. Pu. LITMUS: Landslide detection by integrating multiple sources. in 11th International Conference Information Systems for Crisis Response and Management (ISCRAM). 2014.

[11] Rogstadius, J., et al., CrisisTracker: Crowdsourced social media curation for disaster awareness. IBM Journal of Research and Development, 2013. 57(5): p. 4: 1-4: 13.

[12] Berlingerio, M., et al. SaferCity: a system for detecting and analyzing incidents from social media. in Data Mining Workshops (ICDMW), 2013.

[13] Kireyev, K., L. Palen, and K. Anderson. Applications of topics models to analysis of disaster-related twitter data. in NIPS Workshop on Applications for Topic Models: Text and Beyond. 2009. Canada: Whistler.

[14] LeCun, Y., Y. Bengio, and G. Hinton, Deep learning. Nature, 2015. 521(7553): p. 436-444.

[15] Deng, L., et al. Recent advances in deep learning for speech research at Microsoft. in Acoustics, Speech and Signal Processing (ICASSP), 2013.

[16] Anavi, Y., et al. A comparative study for chest radiograph image retrieval using binary texture and deep learning classification. 37th Annual International Conference of the IEEE in Engineering in Medicine and Biology Society (EMBC), 2015.

[17] Arel, I., D.C. Rose, and T.P. Karnowski, Deep Machine Learning - A New Frontier in Artificial Intelligence Research. IEEE Computational Intelligence Magazine, 2010. 5(4): p. 13-18.

[18] Silver, D., et al., Mastering the game of Go with deep neural networks and tree search. Nature, 2016. 529(7587): p. 484-489.

[19] Kalaimani, E., E. Kirubakaran, and P. Anbalagan, Classification and topic discovery of web documents using semantic deep learner based on semantic smoothing model. International Journal of Applied Engineering Research, 2015. 10(2): p. 2575-2598.

[20] Lu, Y., et al., CHEMDNER system with mixed conditional random fields and multi-scale word clustering. Journal of Cheminformatics, 2015. 7.

[21] Wang, Y., et al., Word vector modeling for sentiment analysis of product reviews, in Communications in Computer and Information Science. 2014. p. 168-180.

[22] Hailong, Z., G. Wenyan, and J. Bo. Machine learning and lexicon based methods for sentiment classification: A survey. in Proceedings - 11th Web Information System and Application Conference, WISA. 2014.

[23] Acharyya, S., et al., Language independent unsupervised learning of short message service dialect. International Journal on Document Analysis and Recognition (IJDAR), 2009. 12(3): p. 175-184.

[24] Aggarwal, C.C. and H. Wang, Text Mining in Social Networks, in Social Network Data Analytics, C.C. Aggarwal, Editor. 2011, Springer US: Boston, MA. p. 353-378.

[25] Najafabadi, M.M., et al., Deep learning applications and challenges in big data analytics. Journal of Big Data, 2015. 2(1): p. 1-21.

[26] Dumais, S.T., Latent semantic analysis. Annual review of information science and technology, 2004. 38(1): p. 188-230.

[27] Blei, D.M., A.Y. Ng, and M.I. Jordan, Latent dirichlet allocation. the Journal of machine Learning research, 2003. 3: p. 993-1022.

[28] Zhou, G., K. Sohn, and H. Lee, Online incremental feature learning with denoising autoencoders. in International Conference on Artificial Intelligence and Statistics, 2012