

NLP Techniques for Water Quality Analysis in Social Media Content

Muhammad Asif Ayub¹, Khubaib Ahmad¹, Kashif Ahmad²,
Nasir Ahmad¹, Ala Al-Fuqaha²

¹ Department of Computer Systems Engineering, University of Engineering and Technology, Peshawar, Pakistan.

² Division of Information and Computing Technology, College of Science and Engineering, Hamad Bin Khalifa University, Qatar Foundation, Doha, Qatar.

{khubaibtakkar,asifayub836}@gmail.com,{kahmad,aalfuqaha}@hbku.edu.qa,n.ahmad@uetpeshawar.edu.pk

ABSTRACT

This paper presents our contributions to the MediaEval 2021 task namely "WaterMM: Water Quality in Social Multimedia". The task aims at analyzing social media posts relevant to water quality with particular focus on the aspects like watercolor, smell, taste, and related illnesses. To this aim, a multimodal dataset containing both textual and visual information along with meta-data is provided. Considering the quality and quantity of available content, we mainly focus on textual information by employing three different models individually and jointly in a late-fusion manner. These models include (i) Bidirectional Encoder Representations from Transformers (BERT), (ii) Robustly Optimized BERT Pre-training Approach (XLM-RoBERTa), and a (iii) custom Long short-term memory (LSTM) model obtaining an overall F1-score of 0.794, 0.717, 0.663 on the official test set, respectively. In the fusion scheme, all the models are treated equally and no significant improvement is observed in the performance over the best performing individual model.

1 INTRODUCTION

In recent years, social media has emerged as a valuable tool and platform to discuss and convey concerns over different challenges and daily life issues [1]. The literature covers a diversified list of societal, environmental, and technological topics, such as racism and hate speech [6], public health [7], natural disasters and rehabilitation [8], and technological conspiracies [4], discussed in social media outlets. More recently, there have been debates in social networks on environmental issues especially the quality of air and drinking water in different parts of the world. The discussions generally revolve around the topics like strange color, smell, bad taste, and diseases caused by polluted water. This information could help in several ways. For instance, it can serve as valuable feedback for public authorities on the water distribution network. However, extracting information from such informal sources is very challenging. It is possible that social media posts containing water-quality-related keywords do not represent discussions on polluted water. In this regard, Machine Learning (ML) and Natural Language Processing (NLP) techniques could be employed to automatically analyze and filter out irrelevant posts. In order to explore the potential of ML and NLP techniques in this challenging problem, a task namely "WaterMM: Water Quality in Social Multimedia" has been introduced in the benchmark MediaEval 2021 competition [2].

This paper provides a detailed description of the methods proposed by team CSE-Innoverts for the water quality analysis represented in the MediaEval task. The dataset provided for the task covers multi-modal information including textual, visual, and meta-data. However, images are available for very few posts. Moreover, the majority of the available images are not relevant. Thus, we mainly focus on textual information by proposing four different solutions as detailed in Section 2.

2 PROPOSED APPROACHES

In total, we submitted 4 different runs by employing three different Neural Networks (NNs) architectures, namely BERT [3], XLM-RoBERTa [5], and LSTM, individually and jointly in a late fusion scheme. Run 1 is based on the late fusion where we jointly employed the models by aggregating the classification scores obtained with the individual models. Figure 1 provides the block diagram of the proposed methodology for Run 1. Run 2, Run 3, and Run 4 are based on the individual models namely BERT, XLM-RoBERTa, and LSTM, respectively. The details of the individual model based solutions are provided below.

- **BERT-based Solution (Run 2):** In this proposed solution, we rely on a pre-trained BERT model, which is fine-tuned on the data development set provided by the task organizers. Before proceeding with fine-tuning the model, necessary pre-processing is performed, using *Tensorflow* libraries, to bring the data in the required form to be used for training the model. Since it is a binary classification task, we used *Binary Cross entropy* loss function with *Adaptive Moments (Adam)* optimizer.
- **XLM-RoBERTa-based Solution (Run 3):** In this approach, we rely on the multilingual pre-trained XLM-RoBERTa model that is fine-tuned on the development set. As a first step, the input text is tokenized in the pre-processing phase. A pre-trained model is then fine-tuned on the pre-processed data using Adam optimizer with a binary cross-entropy loss function.
- **LSTM-based Solution (Run 4):** In this approach, we rely on a custom LSTM model. The model is composed of three layers including an input, LSTM, and output layer. We used this model as a baseline for our experiments. However, the model obtained encouraging results on the development and was thus utilized in the fusion scheme.

We also cleaned the data before feeding into the models by removing URLs, account handles, emojis, and unnecessary punctuation.




Figure 1: Block diagram of the proposed methodology.

Moreover, in all the proposed solutions, we used an up-sampling technique to balance the dataset.

3 RESULTS AND ANALYSIS

3.1 Evaluation Metric

For the evaluation of the proposed methods, we used four different metrics, namely (i) accuracy, (ii) micro precision, (iii) micro recall, and (iv) micro F1-score. Precision, recall, and f1-scores are the official metrics while accuracy has been used as an additional metric for the evaluation of the methods on the development set.

3.2 Experimental Results on the Development Set

Table 1 provides the experimental results of our proposed solutions on the development set. To this aim, a separate validation set composed of 1,810 samples is used. Run 1 represents our fusion-based solutions while Run 2, Run 3, and Run 4 represent our solutions based on the individual models namely BERT, RoBERTa, and LSTM, respectively. On the development set, overall better results are obtained with the BERT-based solution obtaining an overall F1-score and accuracy of 0.950 and 0.929, respectively. The least performance in terms of F1-score and accuracy are observed for RoBERTa.

Table 1: Evaluation of our proposed solutions on the development set in terms of precision, recall, f1-score, and accuracy.

Runs	Precision	Recall	F1-Score	Accuracy
Run 1	0.950	0.925	0.938	0.914
Run 2	0.949	0.950	0.950	0.929
Run 3	0.862	0.900	0.881	0.836
Run 4	0.885	0.947	0.915	0.885

3.3 Experimental Results on the Test Set

Table 2 provides the official results on the test set in terms of precision, recall, and f1-score. Overall better results are obtained for BERT among the individual model-based solutions while the least scores are observed for the LSTM based solution. However, interestingly, no significant improvement in the performance for the fusion-based solution over the best-performing individual models-based solution has been observed. One of the possible reasons could

be the low-performing models as all the models are treated equally by simply aggregating the obtained posterior probabilities. This limitation could be addressed by using merit-based fusion where weights are assigned to the contributing models based on the performance of the model.

Table 2: Evaluation of our proposed solutions on the test set in terms of micro precision, recall, and f1-score.

Runs	Precision	Recall	F1-Score
Run 1	0.732	0.866	0.794
Run 2	0.732	0.866	0.794
Run 3	0.606	0.877	0.717
Run 4	0.565	0.801	0.663

4 CONCLUSIONS AND FUTURE WORK

The quantity and quality of the images associated with the social media posts were not good enough to contribute to the task. Thus, we focused on the textual information only by employing several NNs based solutions. In total, four different solutions including a fusion and three individual models based solutions. In the current implementation, we used a simple fusion mechanism by simply aggregating the posterior probabilities obtained with each individual model.

In the future, we aim to employ more sophisticated fusion schemes by assigning merit-based weights to the contributing models. We also aim to make use of the additional information available in the form of metadata in our future fusion-based solutions.

REFERENCES

- [1] Kashif Ahmad, Konstantin Pogorelov, Michael Riegler, Nicola Conci, and Pal Halvorsen. 2019. Social media and satellites: Disaster event detection, linking and summarization. *MULTIMEDIA TOOLS AND APPLICATIONS* 78, 3 (2019), 2837–2875.
- [2] Stelios Andreadis, Ilias Gialampoukidis, Aristeidis Bozas, Anastasia Moutzidou, Roberto Fiorin, Francesca Lombardo, Anastasios Karakostas, Daniele Norbiato, Stefanos Vrochidis, Michele Ferri, and Ioannis Kompatsiaris. 2021. WaterMM: Water Quality in Social Multi-media Task at MediaEval 2021. In *Proceedings of the MediaEval 2021 Workshop, Online*.
- [3] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [4] Abdullah Hamid, Nasrullah Shiekh, Naina Said, Kashif Ahmad, Asma Gul, Laiq Hassan, and Ala Al-Fuqaha. 2020. Fake news detection in social media using graph neural networks and NLP Techniques: A COVID-19 use-case. *arXiv preprint arXiv:2012.07517* (2020).
- [5] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* (2019).
- [6] Ariadna Matamoros-Fernández and Johan Farkas. 2021. Racism, Hate Speech, and Social Media: A Systematic Review and Critique. *Television & New Media* 22, 2 (2021), 205–224.
- [7] Salman Bin Naeem, Rubina Bhatti, and Aqsa Khan. 2021. An exploration of how fake news is taking over social media and putting public health at risk. *Health Information & Libraries Journal* 38, 2 (2021), 143–149.

[8] Naina Said, Kashif Ahmad, Michael Riegler, Konstantin Pogorelov, Laiq Hassan, Nasir Ahmad, and Nicola Conci. 2019. Natural disasters

detection in social media and satellite imagery: a survey. *Multimedia Tools and Applications* 78, 22 (2019), 31267–31302.