

# Deep Models for Visual Sentiment Analysis of Disaster-related Multimedia Content

Khubaib Ahmad<sup>1</sup>, Muhammad Asif Ayub<sup>1</sup>, Kashif Ahmad<sup>2</sup>,  
Ala Al-Fuqaha<sup>2</sup>, Nasir Ahmad<sup>1</sup>

<sup>1</sup> Department of Computer Systems Engineering, University of Engineering and Technology, Peshawar, Pakistan.

<sup>2</sup> 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 a solutions for the MediaEval 2021 task namely "Visual Sentiment Analysis: A Natural Disaster Use-case". The task aims to extract and classify sentiments perceived by viewers and the emotional message conveyed by natural disaster-related images shared on social media. The task is composed of three sub-tasks including, one single label multi-class image classification subtask, and, two multi-label multi-class image classification subtasks. Both the multi-label classification tasks cover different sets of labels. In our proposed solutions, we mainly rely on two different state-of-the-art models namely, Inception-v3 and VggNet-19, pre-trained on ImageNet. Both the pre-trained models are fine-tuned for each of the three subtasks using different strategies. Overall encouraging results are obtained on all of the three subtasks. On the single-label classification subtask (i.e. subtask 1), we obtained the weighted average F1-scores of 0.540 and 0.526 for the Inception-v3 and VggNet-19 based solutions, respectively. On the multi-label classification tasks i.e., subtask 2 and subtask 3, the weighted F1-scores of our Inception-v3 based solutions are 0.572 and 0.516, respectively. Similarly, the weighted F1-scores of our VggNet-19 based solution on the subtask 2 and subtask 3 are 0.584 and 0.495, respectively.

## 1 INTRODUCTION

Over the last few years, natural disasters analysis in social media outlets has been one of the active areas of research. During this time several interesting solutions exploring different aspects of natural disasters have been proposed [9]. Some key aspects of natural disasters explored in the literature include disaster detection [8], disaster news dissemination [1], and disasters severity and damage assessment [2, 7]. Some efforts on the sentiment analysis of natural disaster-related social media posts have also been reported. However, most of the efforts made in this regard are based on textual information [3]. More recently, Hassan et al. [5] introduced the concept of visual sentiment analysis of natural disaster-related images by proposing a deep sentiment analyzer. However, the topic is very challenging and there are several aspects of visual sentiment analysis of natural disaster-related visual content that yet need to be explored. As part of their efforts to further explore the topic, the authors proposed a task namely "Visual Sentiment Analysis: A Natural Disaster Use-case Task at MediaEval 2021" [6].

This paper provides the details of the solutions proposed by team CSE-Innoverts for the visual sentiment analysis task. The task is composed of three sub-tasks including a (i) single-label multi-class classification task with three labels, a (ii) multi-label multi-class classification task with seven labels, and a (iii) multi-label multi-class classification task with 11 labels. In the first subtask, the participants need to classify an image into *Negative*, *Positive*, and *Neutral* sentiments. In the second subtask, the proposed solution aims to differentiate among *Joy*, *sadness*, *fear*, *disgust*, *anger*, *surprise*, and *neutral*. The final subtask is composed of 11 labels including *anger*, *anxiety*, *craving*, *empathetic pain*, *fear*, *horror*, *joy*, *relief*, *sadness*, and *surprise*.

## 2 PROPOSED APPROACH

### 2.1 Methodology for Single-label Classification task (subtask 1)

For the first task, we mainly rely on two different Convolutional Neural Networks (CNNs) architectures namely Inception V-3 and VggNet based on their proven performances in similar tasks [10]. Since the available dataset is not large enough to train the models from the scratch, we fine-tuned the existing models pre-trained on ImageNet dataset [4]. In the literature, generally, the models pre-trained on ImageNet and Places dataset [11] are fine-tuned for image classification tasks. However, our choice for the current implementation is based on the better performance of the models pre-trained on the ImageNet dataset in similar tasks [10]. In this work, the models are fine-tuned for 50 epochs using Adam optimizer with a learning rate of 0.0001.

It is important to mention that the provided dataset is imbalanced with a large number of *negative* samples while fewer samples are available in the *neutral* class. Before fine-tuning the models, we applied an up-sampling technique to balance the dataset. Moreover, some data augmentation techniques are also employed to further increase the training samples by cropping, rotating, and flipping the image patches.

### 2.2 Methodology for Multi-label Classification tasks (subtask 2 and subtask 3)

We used the same strategy of fine-tuning the existing pre-trained state-of-the-art models for the subtask 2 and subtask 3. However, to deal with the multi-label classification, several changes are made. For instance, the top layers of the models are extended to support

**Table 1: Evaluation of our proposed solutions on the development set in terms of F1-score.**

Runs	Binary Accuracy			Weighted F1-score		
	Subtask 1	Subtask 2	Subtask 3	Subtask 1	Subtask 2	Subtask 3
Run 1	0.722	0.664	0.675	0.714	0.588	0.479
Run 2	0.750	0.710	0.628	0.666	0.535	0.479

the multi-label classification tasks. Moreover, the sigmoid Cross-Entropy loss function is used to deal with every CNN output vector component independently.

Similar to subtask 1, the distribution of the samples in the sentiment categories covered in subtask 2 and subtask 3 is not balanced. To this aim, the same strategy of up-sampling the minority classes is used to balance the dataset. Moreover, the data augmentation techniques are also employed in these subtasks.

### 3 RESULTS AND ANALYSIS

#### 3.1 Evaluation Metric

We used two different metrics for the evaluation of the proposed solutions. On the test set, the evaluations are carried out in terms of weighted F1-score, which is the official evaluation metric of the task. On the development set, we used binary accuracy as an additional metric along with the weighted F1-score. For computing the scores in the multi-label classification task, we used the default threshold (i.e., 0.5).

#### 3.2 Experimental Results on the development set

Table 1 provides the experimental results of our proposed solutions on the development set in terms of F1-score. It is important to note that our validation set in these experiments is composed of 487 samples. As can be seen, overall better results are obtained on the single-label classification subtask 1, which is composed of three classes only. As we go deeper in the sentiment categories/classes hierarchy the performance of the algorithms decreases as the inter-class variation decreases.

As far as the performance of the models is concerned, Inception-v3 has significant improvements over VggNet-19 on subtask 1 and subtask 2 while comparable results are obtained on subtask 3.

#### 3.3 Experimental Results on the test set

Table 2 presents the official results of our proposed solutions on the test set. Surprisingly, overall better results are obtained on a multi-label classification task subtask 2 for both the models. On the other hand, similar to the development set, the least performance is observed for both models on subtask 3. As far as the performance of the models is concerned, Inception-v3 based solution outperformed the VggNet-19 based solution on subtask 1 and subtask 3 while comparable results are obtained on subtask 2.

### 4 CONCLUSIONS AND FUTURE WORK

The challenge is composed of three tasks including a single-label and two multi-label image classification tasks with different sets

**Table 2: Evaluation of our proposed solutions on the test set in terms of weighted F1-score.**

Runs	Subtask 1	Subtask 2	Subtask 3
Run 1	0.540	0.572	0.516
Run 2	0.526	0.584	0.495
Highest Score	0.771	0.627	0.583

of labels. The first task aims to cover the conventional three categories/labels generally used to represent sentiments. The other two tasks aim to cover sets of labels more specific to natural disasters. These three sets of labels allow to explore different aspects of the domain, and the task’s complexity increases by going deeper in the sentiments hierarchy. For all the tasks, we rely on two state-of-the-art deep architectures namely Inception-v3 and VggNet-19. To this aim, the models pre-trained on the ImageNet dataset are fine-tuned on the development dataset.

In the current implementations, we rely on object-level information only by employing the models pre-trained on ImageNet dataset. We believe, scene-level features could also contribute to the task. In the future, we aim to jointly utilize both object and scene-level information for better performance on all the tasks. Moreover, we aim to employ merit-based fusion schemes by considering the contribution of the individual models to the tasks.

### REFERENCES

- [1] Kashif Ahmad, Michael Riegler, Konstantin Pogorelov, Nicola Conci, Pål Halvorsen, and Francesco De Natale. 2017. Jord: a system for collecting information and monitoring natural disasters by linking social media with satellite imagery. In *Proceedings of the 15th International Workshop on Content-Based Multimedia Indexing*. 1–6.
- [2] Firoj Alam, Ferda Ofli, Muhammad Imran, Tanvirul Alam, and Umair Qazi. 2020. Deep Learning Benchmarks and Datasets for Social Media Image Classification for Disaster Response. In *2020 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. IEEE, 151–158.
- [3] Ghazaleh Beigi, Xia Hu, Ross Maciejewski, and Huan Liu. 2016. An overview of sentiment analysis in social media and its applications in disaster relief. *Sentiment analysis and ontology engineering* (2016), 313–340.
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 248–255.
- [5] Syed Zohaib Hassan, Kashif Ahmad, Ala Al-Fuqaha, and Nicola Conci. 2019. Sentiment analysis from images of natural disasters. In *International Conference on Image Analysis and Processing*. Springer, 104–113.
- [6] Syed Zohaib Hassan, Kashif Ahmad, Michael Riegler, Steven Hicks, Nicola Conci, Pal Halvorsen, and Ala Al-Fuqaha. 2021. Visual Sentiment Analysis: A Natural Disaster Use-case Task at MediaEval 2021. In *Proceedings of the MediaEval 2021 Workshop Online*.
- [7] Nayomi Kankanamge, Tan Yigitcanlar, Ashantha Goonetilleke, and Md Kamruzzaman. 2020. Determining disaster severity through social media analysis: Testing the methodology with South East Queensland Flood tweets. *International journal of disaster risk reduction* 42 (2020), 101360.
- [8] Naina Said, Kashif Ahmad, Nicola Conci, and Ala Al-Fuqaha. 2021. Active learning for event detection in support of disaster analysis applications. *Signal, Image and Video Processing* (2021), 1–8.

- [9] 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.
- [10] Naina Said, Konstantin Pogorelov, Kashif Ahmad, Michael Riegler, Nasir Ahmad, Olga Ostroukhova, Pål Halvorsen, and Nicola Conci. 2018. Deep Learning Approaches for Flood Classification and Flood Aftermath Detection.. In *MediaEval*.
- [11] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. 2017. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence* 40, 6 (2017), 1452–1464.