

Visual Sentiment Analysis: A Natural Disaster Use-case Task at MediaEval 2021

Syed Zohaib Hassan¹, Kashif Ahmad², Michael A. Riegler¹,
Steven Hicks¹, Nicola Conci³, Pål Halvorsen¹, Ala Al-Fuqaha²

¹SimulaMet, Norway

²Information and Computing Technology (ICT) Division, College of Science and Engineering, Hamad Bin Khalifa University, Doha 34110, Qatar

³University of Trento, Italy

syed@simula.no, kahmad@hbku.edu.qa, michael@simula.no, steven@simula.no,
nicola.conci@unitn.it, paalh@simula.no, aalfuqaha@hbku.edu.qa

ABSTRACT

The *Visual Sentiment Analysis* task is being offered for the first time at MediaEval. The main purpose of the task is to predict the emotional response to images of natural disasters shared on social media. Disaster-related images are generally complex and often evoke an emotional response, making them an ideal use case of visual sentiment analysis. We believe being able to perform meaningful analysis of natural disaster-related data could be of great societal importance, and a joint effort in this regard can open several interesting directions for future research. The task is composed of three sub-tasks, each aiming to explore a different aspect of the challenge. In this paper, we provide a detailed overview of the task, the general motivation of the task, and an overview of the dataset and the metrics to be used for the evaluation of the proposed solutions.

1 INTRODUCTION

An enormous amount of multimedia content is generated and shared over the internet daily, especially after the introduction of several social media platforms, such as Twitter, Instagram, and Facebook. Among other types of content, visual content is used extensively to deliver a specific message to the viewer, implying the common phrase “a picture is worth a thousand words” [7]. Various emotions can be expressed in visual content, where extracting and analyzing them could be of great value in different application domains, such as education, entertainment, advertisement, and journalism. However, it is not clear which part of an image evoke certain emotions and, more importantly, how the underlying sentiments can be derived from a scene by an automatic algorithm and how the sentiments can be expressed. This opens an interesting line of research to interpret emotions and sentiments perceived by users viewing visual content.

Disaster-related images can be challenging to interpret, especially when analyzing them using current visual analysis techniques. These images generally contain important details in both the background and foreground, making no single part the main point of interest [6]. Developing a visual sentimental analysis pipeline for such a use-case is difficult but could greatly serve our society. Example use cases include having news agencies fact check and provide

their audience with relevant information in case of adverse events. It can also be used to assess the damage caused by a disaster and use the relevant imagery to create awareness and raise funds for humanitarian activities. We note that the resultant solutions are not intended to manipulate public emotions rather are concerned with estimating disaster severity and supporting disaster relief. For instance, humanitarian organizations can benefit from these solutions to reach a wider audience by communicating visual content that best demonstrates the evidence of a certain event [5].

In order to facilitate future work in the domain, a large-scale dataset is collected, annotated, and made publicly available. For the annotation of the dataset, a crowd-sourcing activity with a large number of participants has been conducted.

2 RELATED WORK

The literature reports several interesting works aiming to minimize and mitigate the loss of lives and infrastructure damage. Different types of solutions rely on ground sensors and satellite imagery for effective disaster handling [1]. With the emergence of social media, more human-centric approaches have arisen for effective response management to natural disasters. There are also some efforts on sentiment analysis of disaster-related text. For instance, IN-SPIRE [4] and PUISE [8] are two frameworks dealing with the sentimental analysis of text data shared on social media. There has been a great deal of work done on textual sentimental analysis of disaster-related textual content. However, the concept of extracting sentiments from the visual content is quite new. There are also some efforts on visual sentiment analysis of natural disaster-related multimedia content. For instance, in [5, 9] a deep visual sentiment analysis is introduced along with a benchmark dataset visual sentiment analysis of natural disaster analysis. We believe the task will help in developing a task force to further explore this interesting aspect of natural disaster-related visual content.

3 DATASET DETAILS

As a first step, in creating this benchmark dataset [5], we crawled images related to natural disasters from the social media platforms, such as Twitter, Google API, and Flickr. During this process, special consideration was put on the fact that the crawled images meet the licensing policy of free usage and sharing. A keyword search was the basic technique used to crawl these images, for example, searching by type of disaster, such as *floods*, *wildfires*, *earthquake*,

landslides, and *storms*, or some specific recent natural disasters that happened in different parts of the world, such as *floods in Pakistan* etc.

After data collection, the next step was data annotation, which was crowd-sourced using Microworkers. The challenging part for crowd-sourcing was designing the task and selecting proper keywords to represent different categories of the annotated data. The three keywords namely *Positive*, *Negative*, and *Neutral* are the most widely used sentiment labels in the literature. However considering the applications and the complexity of sentimental analysis of disaster-related images, it is quite complicated to interpret and represent the sentiment and emotion in these three categories. We rather need a more specific and larger set of sentiment categories/labels to interpret the emotions associated with natural disasters. There is an interesting recent work in psychology on emotions and sentiments representation, where 27 different emotion categories are reported. Based on this study, we set up four different types of annotations, which are expected to better represent the sentiments and emotions associated with natural disasters by covering different aspects of natural disasters. These sets of labels are described below.

- (1) Positive, negative, and neutral.
- (2) Relax/calm, normal and simulated/excited.
- (3) Joy, sadness, fear, disgust, anger, surprise, and neutral.
- (4) Anger, anxiety, craving, empathetic pain, fear, horror, joy, relief, sadness, and surprise.

As a first step towards visual sentiment analysis of natural disaster multimedia content, the scope of the task is limited to images from different types of natural disasters. Overall, 4,003 images were selected for the crowd-sourcing study. To ensure quality and consistency, each image was annotated by five different persons. The study was disseminated through multiple channels. A total of 10,010 responses were received during the study, and 2,338 participants participated from 98 different countries.

4 TASK DESCRIPTION

It is the first time we are hosting this task in MediaEval. The task is closely related to previous tasks namely "The 2017-2019 Multimedia Satellite Task: Emergency Response for Flooding Events" [3] and "The 2020 Flood-related Multimedia Task" [2]. However, the goals, challenges, images, and the types of natural disasters covered in the dataset are different from the previous tasks.

Images from the dataset provided for the task are filtered to check whether there is a majority consensus among the annotations provided by the crowd-sourcing participants. The resulted dataset consists of 3,631 images. The development set is composed of 2,432 images, and the test set contains 1,199 images. The task is divided into three subtasks, which are described below.

4.1 Subtask 1

This is a multi-class single label classification task, where the images are arranged in three different classes, namely *positive*, *negative*, and *neutral*. There is a strong imbalance towards the negative class, given the nature of the topic.

4.2 Subtask 2

This is a multi-class multi-label image classification task, where the participants are provided with multi-labeled images. The multi-label classification strategy, which assigns multiple labels to an image, better suits our visual sentiment classification problem and is intended to show the correlation of different sentiments. In this task, seven classes, namely *joy*, *sadness*, *fear*, *disgust*, *anger*, *surprise*, and *neutral*, are covered.

4.3 Subtask 3

This task is also multi-class multi-label, however, a wider range of sentiment classes are covered. Going deeper in the sentiment hierarchy, the complexity of the task increases. The sentiment categories covered in this task include *anger*, *anxiety*, *craving*, *empathetic pain*, *fear*, *horror*, *joy*, *relief*, *sadness*, and *surprise*.

5 EVALUATION

The proposed solutions will be evaluated using a weighted F_1 score. It is a more preferable metric when dealing with imbalanced datasets as it is more sensitive to data distribution. It also provides a good balance between precision and recall as it is a harmonic mean of these two metrics. The weighted F_1 score can be calculated using the following four equations.

$$Precision_i = \frac{\text{True Positives}_i}{(\text{True Positives}_i + \text{False Positives}_i)} \quad (1)$$

$$Recall_i = \frac{\text{True Positives}_i}{(\text{True Positives}_i + \text{False Negatives}_i)} \quad (2)$$

$$F_1 \text{ Score}_i = \frac{2 * Precision_i * Recall_i}{(Precision_i + Recall_i)} \quad (3)$$

$$Weighted F_1 = \frac{\sum_{i=1} (F_1 \text{ Score}_i * c_i)}{\sum_{i=1} (c_i)} \quad (4)$$

Equation 1, 2, and 3 shows the calculation of precision, recall and F_1 score for each class, respectively. Equation 4 shows the calculation weighted F_1 score where c_i is the number of instances of a particular class.

6 DISCUSSION AND OUTLOOK

Though the literature reports some initial efforts on the topic, several important aspects need to be addressed yet. We believe that extracting and representing sentiments from disaster-related visual content will benefit several stakeholders, such as news agencies, Government and non-government, and other humanitarian organizations. From a research point of view, visual sentiment analysis of natural disaster-related visual content is not limited to object recognition. Instead, it requires a more complex identification analysis and establishing a connection among salient objects, scenes, expressions, and color schemes. We hope the task will result in a joint effort towards this important topic, and it also will open new research directions.

REFERENCES

- [1] Ahsan Adeel, Mandar Gogate, Saadullah Farooq, Cosimo Ieracitano, Kia Dashtipour, Hadi Larjani, and Amir Hussain. 2018. A New Efficient Alert Model for Disaster Management. *Geological Disaster Monitoring Based on Sensor Networks* (2018), 57–66.

- [2] Stelios Andreadis, Ilias Gialampoukidis, Anastasios Karakostas, Stefanos Vrochidis, Ioannis Kompatsiaris, Roberto Fiorin, Daniele Norbiato, and Michele Ferri. 2020. The flood-related multimedia task at mediaeval 2020. In *Proceedings of the MediaEval 2020 Workshop, Online*. 14–15.
- [3] Benjamin Bischke, Patrick Helber, Christian Schulze, Venkat Srinivasan, Andreas Dengel, and Damian Borth. 2017. The Multimedia Satellite Task at MediaEval 2017.. In *MediaEval*.
- [4] Michelle Gregory, Deborah Payne, Dave McColgin, Nick Cramer, and V DouglasLove. 2007. Visual Analysis of Weblog Content. *INTERNATIONAL CONFERENCE ON WEB AND SOCIAL MEDIA* (2007).
- [5] Syed Zohaib Hassan, Kashif Ahmad, Steven Hicks, Pål Halvorsen, Ala Al-Fuqaha, Nicola Conci, and Michael Riegler. 2020. Visual sentiment analysis from disaster images in social media. *arXiv preprint arXiv:2009.03051* (2020).
- [6] Imran Khan, Kashif Ahmad, Namra Gul, Talhat Khan, Nasir Ahmad, and Ala Al-Fuqaha. 2021. Explainable Event Recognition. *arXiv preprint arXiv:2110.00755* (2021).
- [7] 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.
- [8] Marc A Smith and Andrew T Fiore. 2001. Visualization components for persistent conversations. *Proceedings of the SIGCHI conference on Human factors in computing systems* (2001), 136–143.
- [9] Nicola Conci Syed Zohaib, Kashif Ahmad and Ala Al-Fuqaha. 2019. Sentiment Analysis from Images of Natural Disasters. *International Conference on Image Analysis and Processing* (2019).