

The Multimedia Satellite Task at MediaEval 2017

Emergency Response for Flooding Events

Benjamin Bischke^{1, 2}, Patrick Helber^{1, 2}, Christian Schulze¹,

Venkat Srinivasan³, Andreas Dengel^{1, 2}, Damian Borth¹

¹ German Research Center for Artificial Intelligence (DFKI), Kaiserslautern, 67663, Germany

² Technical University of Kaiserslautern, Germany

³ Department of Computer Science, Virginia Tech, VA 24061, USA

benjamin.bischke@dfki.de,patrick.helber@dfki.de,christian.schulze@dfki.de

svenkat@vt.edu,andreas.dengel@dfki.de,damian.borth@dfki.de

ABSTRACT

This paper provides a description of the MediaEval 2017 Multimedia Satellite Task. The primary goal of the task is to extract and fuse content of events which are present in Satellite Imagery and Social Media. Establishing a link from Satellite Imagery to Social Multimedia can yield to a comprehensive event representation which is vital for numerous applications. Focusing on natural disaster events in this year, the main objective of the task is to leverage the combined event representation with the context of emergency response and environmental monitoring. In particular, our task focuses this year on flooding events and consists of two subtasks. The first *Disaster Image Retrieval from Social Media* subtask requires participants to retrieve images from Social Media which show a direct evidence of the flooding event. The second task *Flood Detection in Satellite Images* aims to extract regions in satellite images which are affected by a flooding event. Extracted content from both tasks can be fused by means of the geographic information. The task seeks to go beyond state-of-the-art flooding map generation towards recent approaches in Deep-Learning while augmenting the satellite information at the same time with rich social multimedia.

1 INTRODUCTION

Recent advances in earth observation are opening up a new exciting area for exploration of satellite image data. Programs like ESA Copernicus, NASA Landsat, and private companies like Planet Labs or Digital Globe provide access to such imagery, for the first time. Large-scale datasets such as the EuroSAT-Dataset [4] or the ImageCLEFremote-Dataset [2] have emerged from these programs and encourage research in this direction to extract meaningful insights from this new data source. A proper analysis of these satellite images has potential to change how agriculture, urbanization and environmental monitoring will be done in the future. Hand in hand with this development, the Multimedia Satellite Task at MediaEval 2017 addresses natural disaster and environmental monitoring, allowing to raise situational awareness for such events. According to the United Nations Office for the Coordination of Humanitarian Affairs¹, flooding events represent currently the most often observed natural disaster type on our planet. Given this significant

¹<http://reliefweb.int/disasters>

importance for this natural disaster type, our Multimedia Satellite Tasks specifically focuses on flooding events in this year.

One challenge when solely relying on remote sensing is the sparsity problem of satellite data over time. Due to the delayed receiving time of satellite imagery and low temporal revisit time of a particular location by satellites, locations are often sparsely sensed with missing information. In the context of natural disaster monitoring, where effects are often present at multiple locations at the same time, missing information represents a crucial problem since humanitarian organizations and rescuer efforts need to rely on up-to-date disaster maps.

In order to overcome this problem and provide an accurate and comprehensive view of the event, the objective of this task is to fuse satellite imagery with real-time multimedia content from Social Media. Our approach is motivated by previous work in [1, 3, 7] which demonstrated the contextual enrichment of remote-sensed events in satellite imagery by leveraging contemporary content from Social Media. Our multimedia satellite task constitutes a combination of satellite image processing and social media retrieval, where the particular challenges are addressed in two separate subtasks. Task participants are required to retrieve images which provide direct evidence of flooding event from a given set of Flickr images. Beyond that, participants quantify the geospatial impact of the flooding events in the corresponding satellite images in form of segmentation masks.

2 TASK DETAILS

In the following, we define two tasks for our challenge.

Disaster Image Retrieval from Social Media.

The goal of the first subtask is to retrieve all images which show direct evidence of a flooding event from social media streams, independently of a particular event. The objective is to design an algorithm that given any collection of multimedia images and their metadata (e.g., YFCC100M, Twitter, Wikipedia, news articles) is able to identify those images that are related to a flooding event. Please note, that only those images which convey a visual evidence of a flooding event will be considered as True Positives. Specifically, we define images showing “unexpected high water levels in industrial, residential, commercial and agricultural area” as images providing evidence of a flooding event.

The main challenges of this task lie in the proper discrimination of the water levels in different areas (e.g., images showing a lake vs. showing a flooded street) as well as the consideration of different

Metadata	image_id, image_url, date_taken, date_uploaded, user_nsid, user_nickname, title, description, user_tags, capture_device, latitude, longitude, license_url, license_name
Visual Features	AutoColorCorrelogram, EdgeHistogram, Color and Edge Directivity Descriptor (CEDD), Color-Layout, Fuzzy Color and Texture Histogram (FCTH), Joint Composite Descriptor (JCD), Gabor, ScalableColor, Tamura

Table 1: Details of provided metadata information and visual features for the DIRSM-Dataset

types of flooding events (e.g., coastal flooding, river flooding, pluvial flooding). Participants are allowed to submit 5 runs:

- Required run 1: using visual data only
- Required run 2: using metadata only
- Required run 3: using metadata-visual data only fused without resources other than those provided by the organizers
- General run 4, 5: everything automated allowed, including using data from external sources (e.g. Twitter, Flickr)

Flood-Detection in Satellite Images.

The aim of the second subtask is to develop a method that is able to identify regions in satellite imagery which are affected by a flooding. Participants are given a set of satellite image patches for multiple instances of flooding events along with corresponding segmentation masks for the flooding to train their models. Participants report for the unseen image patches a segmentation masks of the flooded area. Participants are allowed to submit 5 runs:

- Required run 1, 2, 3: using satellite data only
- General run 4, 5: everything automated allowed, including using data from external sources

3 DATA

Disaster Image Retrieval from Social Media Dataset.

The dataset for the first subtask consists of 6,600 Flickr images. All images were extracted from the YFCC100M-Dataset [6] which are shared under Creative Commons licenses. The dataset contains one image per user to avoid a bias towards content from same locations and the actively content-sharing users.

Images with the tags of *flooding*, *flood* and *floods* were selected and additionally refined by human annotators according to the strength of the evidence of flooding that they depict: very strong non-evidence of a flooding (0), non-evidence of a flooding (1), direct evidence of a flooding (4), very strong direct evidence of a flooding (5), or with “don’t know” answer (3). The definition of relevance was available to the annotators in the interface during the entire process. The annotation process was not time restricted. The scores were collected from two annotators and the final ground truth label was determined as flooding if both annotators rated the image with 4 or 5 and as non flooding for scores of 0 or 1. To cover a broader diversity of images, we injected additional distractor images in the dataset.

For each image, image metadata from YFCC100M and visual feature descriptors are provided to participants. Visual features were extracted with the open-source LIRE library² using default parameter settings. A overview of the provided features is given

in Table 1. The dataset is separated with a ratio of 80/20 into the following two sets:

- **Development-Set** contains 5,280 images, along with features and class labels (1=evidence of a flooding event and 0=no evidence)
- **Test-Set** contains 1,320 images and features

Flood-Detection in Satellite Images Dataset.

The dataset for the second subtask consists satellite image patches which have been derived from Planet’s 4-band satellites [5]. The imagery has a ground-sample distance (GSD) of 3.7 meters and an orthorectified pixel size of 3 meters. The data was collected from eight different flooding events between 01.06.2016 and 01.05.2017. The image patches have the shape of 320 x 320 x 4 pixels and are provided in the GeoTiff format. All image scenes have been projected in the UTM projection using the WGS84 datum (EPSG:3857). Each image patch contains four channels with *Red*, *Green*, *Blue*, and *Near Infrared* band information. Pixel values are represented in a 16 bit digital number format. The dataset is separated as follows:

- **Development-Set** contains 462 image patches from six locations. For each image patch we provide a segmentation mask of the flooded area, extracted by human annotators (0=background, 1= flooded area).
- **Test-Set-1** contains unseen patches extracted from the same region which are present in the development set.
- **Test-Set-2** contains unseen patches extracted from a different region which are not present in the dev-set.

4 EVALUATION

Disaster Image Retrieval from Social Media.

The official metric for evaluating the correctness of retrieved images from Social Media is Average Precision at k (AP@k) at various cutoffs, k=50,100, 200, 300, 400, 500. The metric measures the number of relevant images among the top k retrieved results and takes the rank into consideration.

Flood-Detection in Satellite Images.

In order to assess performance of generated segmentation masks for flooded areas in the satellite image patches, the intersection-over-union metric (Jaccard Index), is used for the official evaluation: $IoU = TP / (TP + FP + FN)$, where TP, FP, and FN are the numbers of true positive, false positive, and false negative pixels, respectively, determined over the whole test set. The metric measures the accuracy for the pixel-wise classification.

ACKNOWLEDGMENTS

We would like to thank Planet for providing us with high resolution satellite images for this task. Additionally, this work was supported BMBF project MOM (Grant 01IW15002).

²LIRE, <http://www.lire-project.net/>

REFERENCES

- [1] Kashif Ahmad, Michael Riegler, Ans Riaz, Nicola Conci, Duc-Tien Dang-Nguyen, and Pål Halvorsen. 2017. The JORD System: Linking Sky and Social Multimedia Data to Natural Disasters. In *Proceedings of the 2017 ACM on International Conference on Multimedia Retrieval*. ACM, 461–465.
- [2] Helbert Arenas, Md Bayzidul Islam, and Josiane Mothe. 2017. Overview of the ImageCLEF 2017 Population Estimation (Remote) Task. (2017).
- [3] Benjamin Bischke, Damian Borth, Christian Schulze, and Andreas Dengel. 2016. Contextual enrichment of remote-sensed events with social media streams. In *Proceedings of the 2016 ACM on Multimedia Conference*. ACM, 1077–1081.
- [4] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. 2017. EuroSAT: A Novel Dataset and Deep Learning Benchmark for Land Use and Land Cover Classification. *arXiv preprint arXiv:1709.00029* (2017).
- [5] Planet Team. 2017. Planet Application Program Interface: In Space for Life on Earth, San Francisco, CA. (2017).
- [6] Bart Thomee, David A. Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. 2016. YFCC100M: The New Data in Multimedia Research. *Commun. ACM* 59, 2 (Jan. 2016), 64–73. <https://doi.org/10.1145/2812802>
- [7] Alan Woodley, Shlomo Geva, Richi Nayak, and Timothy Campbell. 2016. Introducing the Sky and the Social Eye. In *Working Notes Proceedings of the MediaEval 2016 Workshop*, Vol. 1739. CEUR Workshop Proceedings.