



Automatic detection of passable roads after floods in remote sensed and social media data

Kashif Ahmad^{a,*}, Konstantin Pogorelov^b, Michael Riegler^c, Olga Ostroukhova^d, Pål Halvorsen^c, Nicola Conci^e, Rozenn Dahyot^a

^a Adapt Centre, School of Computer Science and Statistics, Trinity College Dublin, Ireland

^b Simula Research Laboratory, Oslo, Norway

^c Simula Metropolitan Center for Digital Engineering and University of Oslo, Norway

^d Research Institute of Multiprocessor Computation Systems n.a. A.V. Kalyaev, Russia

^e University of Trento, Italy

ARTICLE INFO

Keywords:

Flood detection
Convolutional neural networks
Natural disasters
Social media
Satellite imagery
Multimedia indexing and retrieval

ABSTRACT

This paper addresses the problem of floods classification and floods aftermath detection based on both social media and satellite imagery. Automatic detection of disasters such as floods is still a very challenging task. The focus lies on identifying passable routes or roads during floods. Two novel solutions are presented, which were developed for two corresponding tasks at the MediaEval 2018 benchmarking challenge. The tasks are (i) identification of images providing evidence for road passability and (ii) differentiation and detection of passable and non-passable roads in images from two complementary sources of information. For the first challenge, we mainly rely on object and scene-level features extracted through multiple deep models pre-trained on the ImageNet and Places datasets. The object and scene-level features are then combined using early, late and double fusion techniques. To identify whether or not it is possible for a vehicle to pass a road in satellite images, we rely on Convolutional Neural Networks and a transfer learning-based classification approach. The evaluation of the proposed methods is carried out on the large-scale datasets provided for the benchmark competition. The results demonstrate significant improvement in the performance over the recent state-of-art approaches.

1. Introduction

Natural disasters, such as floods, earthquakes and storms, may cause significant damage to both human life and infrastructures. In such adverse events, an instant access to relevant information may certainly help in the rescue operations, which will ultimately help in mitigating the damage [1,2]. Having an idea of the damage inflicted due to a disaster, government and non-government organizations could allocate their resources to the affected areas accordingly. In such situations, especially in flood events, information about the passability of the roads, i.e., it is possible to travel through the affected areas, is a crucial element for emergency response and deployment of the resources for rescue operations.

In this respect, social media emerged as an important source of information and has been proved very effective in emergency situations, where news agencies could not provide information at all or in time [3, 4]. In [5], several situations have been reported, where news agencies with conventional sources of information failed to report in a timely fashion, simply due to the insufficient number of reporters covering the world. On the other hand, due to wide geographical coverage and high

spatial and multi-spectral resolutions, satellite imagery has been widely used in the past for the analysis of natural disasters and their impact on the environment [6,7].

The joint use of social media and remotely sensed information, which are expected to complement each others, has been already investigated in the literature for the analysis of natural disasters and their potential impact on the environment [1,8–10]. While the satellite imagery provide a birds' eye view of disaster covering a large geographic area, social media provides more specific and detailed information as shown in Fig. 1.

To encourage research in this area, as well as to compare the performance of different software solutions, a task at the MediaEval benchmarking initiative has been proposed for two consecutive years i.e., MediaEval 2017 [12] and MediaEval 2018 [11], with particular attention to the problem of flood detection. In MediaEval 2017, the challenge aimed at flood detection in images from social media and satellite. In Mediaeval 2018, the challenge targeted analysis of social media and satellite imagery for detection of passable roads in flood-affected regions. The 2018 challenge is composed of two parts:

* Corresponding author.

E-mail address: ahmadka@tcd.ie (K. Ahmad).



Fig. 1. A joint view of floods from social media and satellite imagery. Satellite imagery provide a bird's eye view covering a large area while social media provides more specific and detailed information [11].

1. **FCSM: Flood Classification for Social Multimedia.** This task is further divided in two sub-tasks aiming to predict:

- (a) whether there is evidence of a flood in a given social media image or not and,
- (b) if evidence of flood exists in the image, whether it is possible to pass through the flooded road (passability).

2. **FDSI: Flood Detection in Satellite Imagery.** This task aims to analyse the roads from satellite images, and predict whether or not it is possible for a vehicle to pass a road.

This paper addresses in particular the MediaEval 2018 challenge. For FCSM 1, we rely on an ensemble of deep models using three different fusion techniques. For FDSI 2, we rely on a Convolutional Neural Network (CNN) architecture and a transfer learning-based classification approach to identify passable roads in satellite imagery. The main contributions of the work can be summarized as:

- (i) We analyse flood-related images from social media and satellite to identify passable roads. In particular, we aim to analyse (a) whether the images provide evidence for road passability and (b) to differentiate between images showing passable vs. non passable roads.
- (ii) On the FCSM task, being an important ingredient of several multimedia analysis applications frameworks, we analyse the performances of four models from three well-known architectures, pre-trained on object and places datasets, both individually and in different combinations, using three different fusion techniques: early, late and double fusion. We believe such rigorous analysis will set a benchmark for the future research on the topic.
- (iii) For the analysis of satellite imagery in the FDSI task, we propose a CNN- and a transfer learning-based classification approach to identify passable roads in satellite imagery.
- (iv) We have performed experiments on a challenging benchmark dataset provided for the competition, and show that better scores are achieved compared to the recent literature.

The rest of the paper is organized as follows. Related work is discussed in Section 2, followed by a detailed description of the MediaEval-2018 challenge providing an overview of the other proposed solutions and the dataset used for the evaluation in Section 3. A detailed description of our proposed solutions is provided in Section 4. We show experimentally how our approach performs (Section 5), and we conclude and provide potential future research directions in Section 6.

2. Related work

The literature on natural disasters detection and analysis can be roughly divided into two parts, namely (i) disaster detection in social

media, and (ii) disaster detection in satellite imagery. In the next subsections, we provide a detailed review of the relevant literature in both domains.

2.1. Disasters detection in social media images

In recent years, several applications have emerged to make use of data posted on social media platforms in combination to other media streams available (e.g. Google Street view imagery, OpenStreetMap map) allowing 3D reconstruction of cities [13], inferring pedestrian trajectories [14] and visualizing sentiments in 3D [15].

Geo-located and time-stamped data available in form of text and visual content on social media have also been widely utilized for disaster events analysis to gather useful information to be used in rescue and rehabilitation [1]. To this aim, most of the approaches rely on two types of complementary information including visual contents and the additional information associated with images in the form of metadata, such as user tags, geo-location and temporal information. For instance, in [16], users' tags and other useful information from metadata are jointly utilized with visual features in an early fusion scheme. In [8], visual features extracted through deep models, pre-trained on ImageNet [17], are complemented by textual information, such as users tags, geo-location and temporal information along with textual description. Both textual and visual features are evaluated individually and jointly by concatenating the relevant feature vectors.

Existing pre-trained models are also used in [18], where five different models from four state-of-the-art deep architectures, namely AlexNet [19], GoogleNet [20], VggNet [21] and ResNet [22], pre-trained on the large-scale ImageNet and Places datasets [23], were used. The basic insight of the paper was to combine object and scene-level features for the flood classification task. Individual Support Vector Machines (SVMs) are then trained on the features extracted through each model, followed by a fusion phase where three different late fusion techniques are used to combine the scores obtained through the individual classifiers along with a Random Forest classifier trained on textual features. Object and scene-level features are also used in [24–26], for the classification of flooded and non-flooded images in social media.

Tkachenko et al. [27] rely on hand-crafted visual features, such as the colour and edge directivity descriptor (CEDD) [28], colour layout [29], and Gabor wavelets [30]. A more sophisticated solution has been proposed for textual information (i.e., description, title and users' tags) relying on word embeddings trained on the entire YFCC100m dataset [31]. Each textual feature is extracted separately, and then concatenated to form a single feature vector. Moreover, to translate users' tags into English, a machine translation technique has been

employed. In [32], handcrafted visual features are concatenated into a single feature vector followed by dimensionality reduction and classification. Term Frequency Inverse Document Frequency (TFIDF) [33] are measured for users' tags to represent textual features. In [34], handcrafted visual features along with textual information are used for the classification of flood-related images.

An active learning framework intending to collect, filter and analyse social media contents for natural disasters has been proposed in [35]. For data collection, a publicly available system, namely AIDR [36], has been used to crawl social media platforms, followed by a crowd-sourcing activity for data annotation. A pre-trained model [21] is then fine-tuned on the annotated images for classification purposes.

2.2. Disaster events detection in satellite imagery

Being one of the most valuable sources of information for disaster analysis [1,37,38], a growing portion of research also aims at the detection and classification of natural disaster events in satellite imagery. Liu et al. [39] proposed a deep architecture along with a wavelet transformation-based pre-processing scheme for the identification of areas affected by disasters in satellite imagery. Amit et al. [40] propose a CNN-based deep architecture composed of five weighted layers for landslides and flood detection in satellite imagery.

Benjamin et al. [8] approach the flood detection in satellite imagery as an image segmentation task where a CNN-based framework with three different training strategies has been adopted. In order to remove a location bias due to local changes in images due to lighting conditions and other atmospheric distortions, the individual components of the provided satellite imagery, i.e., RGB and IR, are normalized before training the model. In [10], authors exploit the diversity of different CNNs, which are mainly based on dilated convolution [41] and deconvolution [42], in a fusion framework. Initially, binary maps obtained with the individual models are concatenated, which are then used to train SVMs for analysing which individual model and in what cases is better. Finally, SVMs are trained on features/maps obtained with the combination of the best models to predict the final binary maps of the test images.

Ahmad et al. [18] tackle flood detection in satellite imagery as a generative problem where a framework based on Adversarial Generative Networks (GANs) has been proposed. The framework mainly relies on a GANs architecture, namely V-GAN [43], originally developed for the retinal vessel segmentation. In order to adopt the architecture for the flood detection task, the top layer of the generative network is extended with a threshold mechanism to generate a binary segmentation mask of the flooded regions in satellite imagery. In another work from the same authors [1], the input layer is modified to support 4 channel input images (i.e., RGB and IR) and several experiments are conducted to evaluate the performance of RGB and IR components, individually and jointly. In [27], different indices, namely Land Water Index (LWI), Normalized Difference Vegetation Index (NDVI) and Normalized Difference Water Index (NDWI) are selected from the spectral images. Subsequently, two different strategies based on supervised classification and unsupervised clustering techniques are then adopted for the identification of flooded regions in satellite imagery. On the other hand, Avgrinak et al. [25] rely on Mahalanobis distance [44] and some morphological operations for the task.

3. MediaEval 2018 challenge on multimedia satellite task: Emergency response for flooding events

3.1. Overview of the presented/submitted works

The problem posted in the MediaEval 2018 challenge is different from the existing state-of-the-art, and focuses on a different perspective of detecting road passability in flood-affected areas from social media and satellite imagery.

Similar to the challenge posted in MediaEval 2017 [12], in MediaEval 2018, two tasks have been included in the challenge. The majority of the solutions proposed for the challenge rely on deep architectures. Fen et al. [45] use an ensemble of deep models pre-trained on Imagenet dataset along with textual features extracted through *fasttext* [46] for the FCSM task. In [47], the performance of several deep models is evaluated in a framework with double-ended classifier and compact loss function treating the two sub tasks as one class classification problem by individually training the models on the provided dataset. Moreover, data augmentation techniques are used to increase the number of training samples. On the other hand, textual features are represented through embedding, initialized with Glove [48]. Another deep architecture method is presented in [49], where an existing deep model pre-trained on ImageNet is fine-tuned on the provided dataset. The Bag of visual Words (BoW) model over the textual information is also used for the FCSM task. For the FDSI task, image patches of size 50×50 pixels are extracted around each of the two given points. Visual features are then extracted as RGB histograms with 16 bins per channel followed by training and SVMs based classification of the test patches.

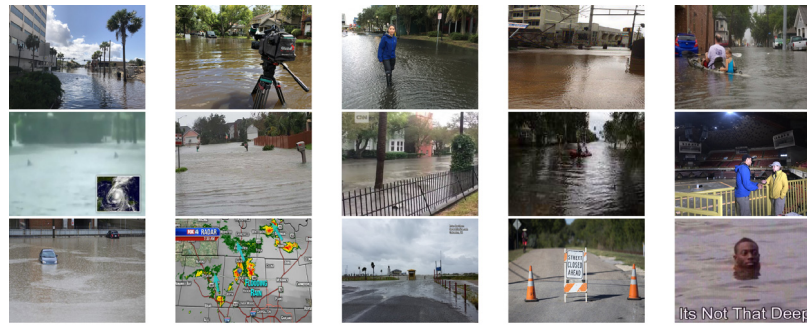
In [50], a deep architecture is proposed in the FCSM task, containing two separate deep architectures (VggNet) for the evidence and passability sub-tasks. Each image is processed by both networks, aiming to predict evidence for the road passability and differentiating in passable and non-passable images. Moreover, stacked auto-encoders are used for the early fusion of textual and visual features. For the FDSI task, a ResNet model pre-trained on ImageNet is fine-tuned on the satellite imagery provided for the challenge. The authors in [51], utilize scene-level features extracted through a deep model pre-trained on places dataset for the evidence of road passability. For the second sub-task, to differentiate in passable and non-passable roads, both object and scene-level features are extracted through deep models. Moreover, the presence of a boat is also used as an indication of the non-passable roads. Hanif et al. [52] adopted an ensemble framework to jointly utilize local features and global visual features for the FCSM task. Global features are extracted through several features descriptors, while a CNN-based local feature descriptor is used for the extraction of local features. On the other hand, textual information is represented through frequency-inverse document frequency (TF-IDF). An ensemble framework is then used to jointly utilize these features for the FCSM task.

3.2. Dataset


The dataset for evaluation is provided in the MediaEval 2018 benchmark competition on Multimedia and Satellite task. In the challenge, two different collections containing images and associated meta-data from social media and satellite imagery have been provided for the FCSM and FDSI tasks, respectively.

For the FCSM task, participants were provided with a set of tools to download the tweets along with the associated images from three hurricanes, namely hurricane Harvey, hurricane Maria, and hurricane Irma, hitting Southeast Texas, Puerto Rico, and Florida, respectively, in 2017. The dataset is collected during the events from Twitter using keywords, such as flooding and floods. The dataset also contains additional information, such as rainfall and climate predictions. The development and test sets are provided, separately. The development set is composed of 7,387 tweets and associated images, while the test set contains a total of 3683 tweets/images. In Table 1, we provide the number of negative and positive samples available in the dataset for each sub-task of the FCSM task. The tweets and associated images are annotated through human annotators via crowdsourcing, conducted on Figure Eight.¹ The ground truth is provided in two separate files, one for each of the sub-tasks. Moreover, the participants are also provided with a set of handcrafted visual feature.

¹ <https://www.figure-eight.com>.



(a) Sample images from FCSM dataset.



(b) Sample images from FDSI dataset depict passable (top line) and non-passable (bottom line) roads.

Fig. 2. Sample images from the datasets used for validation.

Table 1

Details of the dataset in terms of the number of positive and negative samples for both sub-tasks.

Sub-task	No. of positive samples	No. of negative samples
Evidence	2193	3877
Passability	976	1217

For FDSI task, satellite image patches of flooded areas from the three disaster events have been provided. The FDSI dataset covers image patches from DigitalGlobe taken by the satellite World-View3 (0.3m resolution). The dataset is provided in two files, i.e., the one containing the cropped satellite images of flooded areas, and the ground-truth labels for road passability given two points of the road network. The dataset is provided in two sets (development and test sets) containing 1,438 and 225 image patches, respectively. Fig. 2 reports some sample images from both datasets.

4. Our proposed solutions for flood classification and detection

In this section, we present our proposed solutions for the tasks. First, we describe the methodology proposed for the FCSM task, and then we provide the details of the methodology chosen for the FDSI task.

4.1. Methodology for FCSM task

In the first task, the objective is the analysis of images from social media providing direct evidence for passability of roads using conventional means of transportation. The task can be divided into two sub-tasks, namely (i) identification of images providing an evidence for road passability and (ii) differentiating passable and non-passable images among the ones providing an evidence for the road passability identified in the first step. Both steps are carried out sequentially. In Fig. 3, we provide the flowchart of the first challenge showing the two sub-tasks of the FCSM task.

Fig. 4 shows the block diagram of the methodology we adopted for the FCSM task. The proposed method is mainly composed of three phases, namely, (i) feature extraction, (ii) classification, and (iii) fusion.

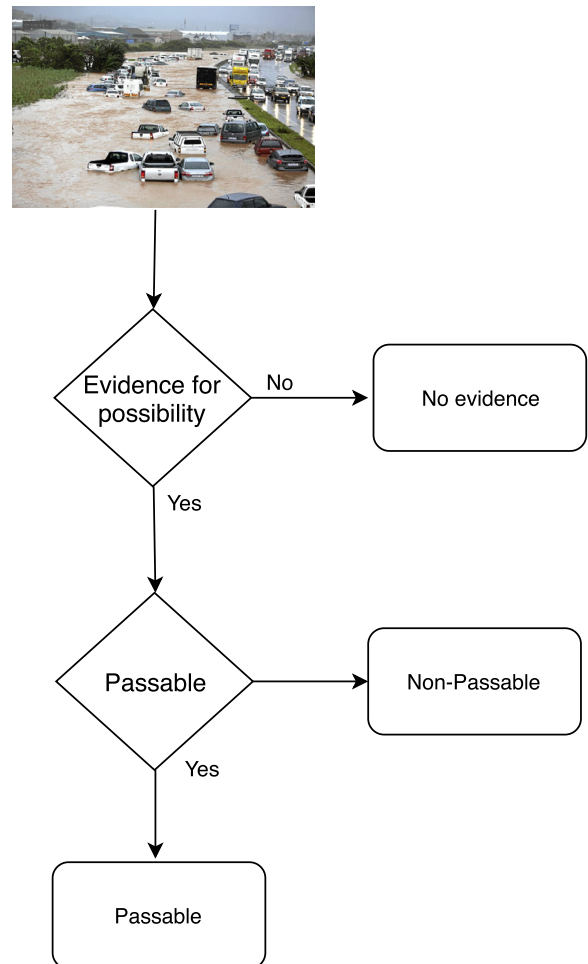


Fig. 3. Flow chart of the first task showing the two sub-tasks of the FCSM task.

For feature extraction, we rely on four different models, pre-trained on ImageNet and Places datasets, from three state-of-the-art deep architectures, namely AlexNet [19], VggNet [21] and ResNet [22]. For classification we adopt a SVM. The basic motivation for the feature extraction through these deep models pre-trained on object and places datasets comes from our previous experience [26,53], where object and scene-level features showed better performance when jointly utilized. In the final phase, we use three different fusion schemes to combine the capabilities of the four models in the FCSM task. In the next subsections, we provide a detailed description of each component.

4.1.1. Feature extraction and classification

In this phase, we rely on four different deep models for feature extractions. Two of the models (ResNet and VggNet) are pre-trained on ImageNet [17], and the other two are pre-trained on the Places dataset [23].

Features are extracted from the last fully connected layer of each model (i.e., Fc-7 for AlexNet and VggNet, Fc-1000 for ResNet). For VggNet, we use the architecture with 19 layers, and for ResNet the configuration with 50 layers. We use the models as feature descriptors without any re-training and fine-tuning. For feature extraction with all models, we used the Caffe toolbox.² After feature extraction, SVMs are trained on features extracted through each individual model. For classification, we use the default parameters using the *Fit-Multiclass* model from the MathWorks³ toolbox.

4.1.2. Fusion

To jointly exploit the capabilities of the individual models in the FCSM task, we rely on three different fusion techniques: early, late, and double fusion. In the early fusion, we concatenate the features extracted through the different models. For the late fusion, in the current implementation, we simply average the results obtained by the individual models. For the third fusion technique, we combine the results obtained from the first two techniques in an additional late fusion step by averaging their scores.

4.2. Methodology for FDSI task

For the FDSI sub-task, we have also opted for a CNN and a transfer learning-based classification approach, previously validated in a different application domain in our work [54]. In fact, we initially tried to apply a GAN-based approach introduced in our previous works for the flood detection satellite imagery [1] and medical imagery [55,56]. We conducted an exhaustive set of experiments, but we unfortunately could not achieve sufficiently accurate roads passability detection performance. The reason for that is the limited size of the dataset (only 1,437 samples were provided in the development set). This, in combination with the large variety of landscapes, road types, types of obstacles and weather conditions, etc., prevents the GAN-based approach from adequate training and finding key visual features required to reliably distinguish between flooded and non-flooded roads.

Fig. 5 provides the block diagram of the methodology proposed for the FDSI task. This approach is based on the Inception v3 architecture [20] pre-trained on the ImageNet dataset [17] and the retraining method described in [57].

For this work, we froze all the basic convolutional layers of the network and only retrained the top fully connected layer with softmax activation after random initialization of its weights. The new fully connected layer was retrained using the RMSprop [58] optimizer, which allows an adaptive learning rate during the training process.

As the input for the CNN model, we used the image patches, extracted from the full images using the provided coordinates of the target road end points of the road-segment. Visual inspection of the

Table 2

Evaluation of the individual models in terms of accuracy on the validation set.

Models	Performance (accuracy in %)	
	Evidence sub-task	Passability sub-task
AlexNet (Places)	86.19	71.00
VggNet (Places)	86.79	71.85
VggNet (ImageNet)	86.79	69.85
ResNet (ImageNet)	85.05	69.28

generated roads patches from the training dataset, showed relatively good coverage for the road-related areas and enough coverage of the neighbourhood areas, and give enough visual information for the following CNN-based analysis and classification.

Moreover, in order to increase the number of training samples, we also performed various augmentation operations on the images. Specifically, we performed horizontal and vertical flipping, and change of brightness in the interval of $\pm 40\%$.

After the model has been retrained, we used it as a multi-class classifier that provides the probability value for each of two classes: passable and non-passable. The final passability detection is done via the selection of the class with higher probability. In case of equal class probabilities, we mark the road patch as non-passable.

5. Evaluation and results

5.1. Experimental results of the FCSM task

For the evaluation of the FCSM task, we used two different metrics, namely (i) accuracy and (ii) mean F1 score, which is the official metrics of the competition. The accuracy is used to evaluate the performance of the sub-tasks independently while the official metric of the challenge (i.e., mean F1 score given by Eq. (1)) is used for comparison purposes. The basic motivation for using the two different metrics is to evaluate the performance on the sub-tasks both independently as well as when the evaluation of the second task is conditioned on the first sub-task of the evidence challenge.

$$\text{Mean_F1} = (\text{F1}(\text{evidence AND passable}) + \text{F1}(\text{evidence AND non_passable}))/2 \quad (1)$$

For the FDSI task, we rely on the official metric F1 Score, computing the harmonic mean between precision and recall for the non-passable roads.

The rest of the section provides a detailed description of the conducted experiments, the results achieved, and the comparisons against the state-of-the-art on the FCSM task.

5.1.1. Evaluation of the individual models

Table 2 provides the experimental results of our first experiment where we evaluate the performance of the individual deep models on both sub-tasks, namely (i) identification of images providing evidence for road passability, and (ii) differentiation between images showing passable vs. non passable roads, in terms of accuracy per class on both sub-tasks. The evaluation is carried out on the development set allocating 60% and 40% images for training and testing, respectively. Overall, better results are obtained on the first sub-task. However, the performances in differentiating between passable and non-passable roads are generally lower, demonstrating a significantly higher complexity of the task. Although there is no significant difference in the performance of the models in identifying the images providing an evidence of road passability, slightly better results are obtained using VggNet pre-trained on the Places dataset. Moreover, in differentiating between passable and non-passable roads, the models pre-trained on the Places dataset outperform the ones pre-trained on the object dataset, showing the importance of scene-level features in this particular task.

² <http://caffe.berkeleyvision.org/>.

³ <https://www.mathworks.com>.

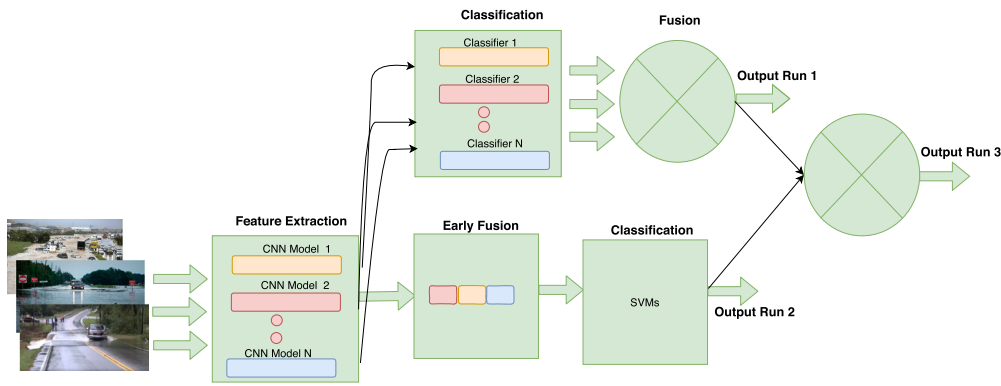


Fig. 4. Block diagram of the proposed methodology of the three different techniques used for the FCSM task. The input images are representing road segments, which are then processed using three pipelines, namely (i) late fusion (upper block), (ii) early fusion (the lower block) and (iii) double fusion coming both blocks at the final block.

Table 3
Evaluation of the individual models in terms of per class accuracy on the validation set.

Models	Evidence sub-task (accuracy in %)		Passability sub-task (accuracy in %)	
	Evidence class	No-evidence class	Passable class	Non-passable class
AlexNet (Places)	78.11	90.80	57.33	81.25
VggNet (Places)	80.00	90.71	57.66	82.50
VggNet (ImageNet)	79.88	90.77	55.66	80.50
ResNet (ImageNet)	77.41	89.49	52.00	82.25

Table 4
Evaluation results of the fusion experiment in terms of accuracy on the validation set.

Methods	Performance (accuracy in %)	
	Evidence sub-task	Passability sub-task
Early fusion	88.81	77.00
Late fusion	90.36	76.00
Double fusion	96.43	80.71

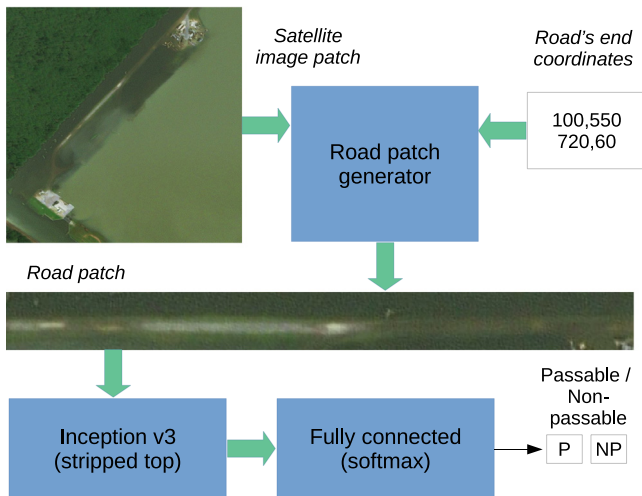


Fig. 5. Block diagram of the proposed satellite image processing methodology used for the FDSI task.

In order to analyse the performance of the models on individual classes of the dataset, we also provide the results of the models in terms of per-class accuracy in Table 3. In both sub-tasks, the accuracy on the negative samples (i.e., images with no evidence and non-passable images) is generally high. Significant variations in the performances of the individual model can be observed on positive samples (i.e., images providing evidence for passability and images of passable roads). The variations in the performances of the individual models provide basis for our second experiment, where we use three different fusion techniques to combine the capabilities of the individual models for the potential improvement in the performances.

5.1.2. Results of the fusion experiments

Table 4 provides the experimental results of our fusion techniques on both sub-tasks in terms of accuracy on the validation set. As can be seen, fusion contributes to significantly improving the performances of the individual models. The double fusion method, which combines the

results of both early and late fusion, secures a significant improvement of 6.03% and 3.71% over the best single fusion methods, on evidence and passability sub-tasks, respectively.

5.1.3. Comparison against state-of-the-art on the FCSM tasks

Table 5 provides the comparisons of our proposed methods against the existing state-of-the-art. It is important to mention that in the benchmark competition teams could submit up to five runs: visual information only for Run 1, textual information only in Run 2, combination of textual and visual features in Run 3, and two Runs (4 and 5) with no restriction. In our method, we rely on visual information only. Our first Run is based on late fusion, while in the fourth Run, we concatenate the features extracted through the individual models (early fusion). Our final Run is based on the double fusion. The conducted experiments show that the proposed solution outperforms the available state-of-the-art in all four configurations. Our best run with double fusion exhibits a significant improvement over most of the methods, and achieves comparable results with method proposed by Anastasia et al. [50], which demonstrated the significance of combining multiple deep models for the classification purposes.

5.2. Experimental results of the FDSI task

For the experimental setup of the FDSI task, we decided to perform only two mandatory runs, which rely on the task-provided training data only. Due to a limited amount of training samples available, training the deep architecture from the scratch is not possible. Thus, we decided to perform two types of training for our transfer-learning detection approach.

First, we implemented a pipeline for classification that differs from common procedures. This process was involving all the training samples



Fig. 6. Sample road patch images depicting the classification results of the validation dataset.

Table 5

Comparisons against other methods from the benchmarking competition on the FCSM dataset. In the competition up-to 5 runs were allowed.

Methods	Performance (F1 measure)				
	Visual	Textual	Multi-modal	Run 4 (visual)	Run 5 (visual)
Feng et al. [45]	64.35	32.81	59.49	52.16	51.59
Armin et al. [49]	20.00	24.00	–	17.00	35.00
Anastasia et al. [50]	66.65	30.17	66.43	55.12	54.48
Zhao et al. [51]	63.88	12.86	–	63.13	63.89
Hanif et al. [52]	45.04	31.15	45.56	–	–
Dias et al. [59]	62.82	60.92 (visual)	62.93 (visual)	62.91	64.84
Our method	63.58 (late fusion)	–	–	60.59 (early fusion)	65.03 (double fusion)

into the training process as both training and validation sets. Usually, for classification tasks, this would result into over-fitting of the model and inability to correctly classify the test samples. However, for this specific classification task, we used the following steps to prevent possible over-fitting. First, we limit the number of training epochs to 100 (selected by the experimental evaluation). Next, the used training data augmentation scheme allow us to generate the significant amount of additional training data. Lastly, the provided training set contains the flooding samples with a high variety of road shapes and conditions. As the result, the last layers of the model were correctly retrained without over-fitting. Fig. 6 depicts some samples of the classification output for the validation dataset in terms of true positive, true negative, false positive, and false negative.

The official F1-Score metric (see Table 6) on the non-passable road class for the first “All-train” Run is 62.30%. To verify our idea of the feasibility of using all the training data for both training and validation, we also performed a normal network training with a random 50/50 development/validation data split. This second *Half-trained* Run resulted in a F1-Score of 61.02%, which is slightly lower comparing to the *All-trained* Run. This is confirming the validity of our idea of using the complete training dataset together with heavy data augmentation to improve road patches classification performance.

We also provide comparisons of our method for FDSI task against the state-of-the-art in terms of F1 score in Table 7. Our method outperforms the state-of-the-art in both Run 1 and Run 2 with an improvement of 5.30% and 4.02%, respectively.

Table 6
Evaluation of our proposed approach for the FDSI task in terms of F1 scores.

Run	Method	F1 score
1	All-train	62.30%
2	Half-train	61.02%

Table 7
Comparisons against other state-of-the-art methods from the benchmarking competition on FDSI dataset. In the competition, up-to 5 runs were allowed.

Methods	Performance (F1 measure)				
	Run 1	Run 2	Run 3	Run 4	Run 5
Armin et al. [49]	57.00	32.00	39.00	56.00	57.00
Anastasia et al. [50]	56.45	–	–	–	–
Our method	62.30	61.02	–	–	–

6. Conclusions and future work

In this paper, we addressed a challenging problem of detecting the passability of roads in social media and satellite imagery. In the social media image analysis, we mainly relied on deep features extracted through different pre-trained deep models individually and jointly, through fusion. We observed that better results can be achieved for the models pre-trained on places dataset compared to the ones pre-trained on objects dataset, showing the importance of the scene-level information in the task. We also observed that the object-level information well complements the scene-level features when jointly utilized. Among the fusion methods, double fusion combines the capabilities of both early and late fusions and ultimately leads to better results. Considering the improvement achieved thanks to the fusion, we believe that, in the future, some optimization methods can be adopted to assign more accurate weights to the individual models.

In the satellite sub-task, we found that a simple image segmentation approach is of no help, and we implemented a task-oriented CNN and transfer learning-based approach. This approach was able to classify image patches with roads and achieved an F1-Score of 62.30% for the non-passable road class.

As a future research direction we expect that a combination of CNN- and GAN-based approaches pre-trained on the existing annotated road network and flooded areas datasets could provide more accurate results. In fact, although our experiments with GANs were not satisfactory, we believe that this issue could be overcome by the availability of a richer dataset. Apart from the improved segmentation, we plan to implement a resolution independent version of our satellite imagery analysis algorithm suitable for utilizing the open-sourced real-time satellite data sources, and to further improve the flooding detection performance by adding freely available Digital Elevation Model (DTM) as an additional layer of input information. DTM combined with the whole road network definition using predefined maps, like, for example, Open Street Map will allow us to generate a live road network status view and to predict the possible flooding expansion to the non-flooded roads.

Acknowledgements

This research is partly supported by the ADAPT Centre for Digital Content Technology, which is funded under the Science Foundation Ireland Research Centres Programme, Ireland (Grant 13/RC/2106) and is co-funded under the European Regional Development Fund.

Conflict of interest

The authors declare that there is no conflict of interest.

References

- [1] K. Ahmad, K. Pogorelov, M. Riegler, N. Conci, P. Halvorsen, Social media and satellites, *Multimedia Tools Appl.* (2018) 1–39.
- [2] K. Ahmad, M. Riegler, K. Pogorelov, N. Conci, P. Halvorsen, F. De Natale, 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, ACM, 2017*, p. 12.
- [3] T. Brouwer, D. Eilander, A. Van Loenen, M.J. Booij, K.M. Wijnberg, J.S. Verkade, J. Wagemaker, Probabilistic flood extent estimates from social media flood observations., *Nat. Hazards Earth Syst. Sci.* 17 (5) (2017).
- [4] J. Yin, A. Lampert, M. Cameron, B. Robinson, R. Power, Using social media to enhance emergency situation awareness, *IEEE Intell. Syst.* 27 (6) (2012) 52–59.
- [5] B. Stelter, N. Cohen, Citizen journalists provided glimpses of Mumbai attacks, *N.Y. Times* 30 (2008).
- [6] J. Kansas, J. Vargas, H.G. Skatter, B. Balicki, K. McCullum, Using landsat imagery to backcast fire and post-fire residuals in the boreal shield of saskatchewan: implications for woodland caribou management, *Int. J. Wildland Fire* 25 (5) (2016) 597–607.
- [7] R. Lagerstrom, Y. Arzhaeva, P. Szul, O. Obst, R. Power, B. Robinson, T. Bednarz, Image classification to support emergency situation awareness, *Front. Robot. AI* 3 (2016) 54.
- [8] B. Bischke, P. Bhardwaj, A. Gautam, P. Helber, D. Borth, A. Dengel, Detection of flooding events in social multimedia and satellite imagery using deep neural networks, in: *Proceedings of the Working Notes Proceeding MediaEval Workshop, Dublin, Ireland, 2017*, pp. 13–15.
- [9] M. Jing, B. Scotney, S. Coleman, et al., Flood event image recognition via social media image and text analysis, in: *Signals and Systems Conference, ISSC, 2016*, pp. 4–9.
- [10] K. Nogueira, S.G. Fadel, Í.C. Dourado, R.d.O. Werneck, J.A. Muñoz, O.A. Penatti, R.T. Calumby, L.T. Li, J.A. dos Santos, R.d.S. Torres, Exploiting convnet diversity for flooding identification, *IEEE Geosci. Remote Sens. Lett.* 15 (9) (2018) 1446–1450.
- [11] B. Bischke, P. Helber, Z. Zhao, J. de Bruijn, D. Borth, The multimedia satellite task at mediaeval 2018: Emergency response for flooding events, in: *Proceedings of the MediaEval 2018 Workshop, Sophia-Antipolis, France, Oct. 29–31, 2018*.
- [12] B. Bischke, P. Helber, C. Schulze, S. Venkat, A. Dengel, D. Borth, The multimedia satellite task at mediaeval 2017: Emergency response for flooding events, in: *Proceedings of the MediaEval 2017 Workshop, Dublin, Ireland, Sept. 13–15, 2017*.
- [13] N. Snavey, S.M. Seitz, R. Szeliski, Modeling the world from internet photo collections, *Int. J. Comput. Vis.* (80) (2008) 189–210, <http://dx.doi.org/10.1007/s11263-007-0107-3>.
- [14] A. Bulbul, R. Dahyot, Populating virtual cities using social media, *Comput. Animat. Virtual Worlds* 28 (2017) 1–10, <http://dx.doi.org/10.1002/cav.1742>.
- [15] A. Bulbul, R. Dahyot, Social media based 3d visual popularity, *Comput. Graph.* 63 (2017) 28–36, <http://dx.doi.org/10.1016/j.cag.2017.01.005>.
- [16] Y. Yang, H.-Y. Ha, F. Fleites, S.-C. Chen, S. Luis, Hierarchical disaster image classification for situation report enhancement, in: *Information Reuse and Integration, IRI, 2011 IEEE International Conference on, IEEE, 2011*, pp. 181–186.
- [17] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei -Fei, Imagenet: A large-scale hierarchical image database, in: *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on, IEEE, 2009*, pp. 248–255.
- [18] K. Ahmad, K. Pogorelov, M. Riegler, N. Conci, H. Pal, Cnn and gan based satellite and social media data fusion for disaster detection, in: *Proceedings of the MediaEval 2017 Workshop, Dublin, Ireland, 2017*.
- [19] A. Krizhevsky, I. Sutskever, G.E. Hinton, Imagenet classification with deep convolutional neural networks, in: *Advances in Neural Information Processing Systems, 2012*, pp. 1097–1105.
- [20] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, Z. Wojna, Rethinking the inception architecture for computer vision, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016*, pp. 2818–2826.
- [21] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition, 2014, arXiv preprint [arXiv:1409.1556](https://arxiv.org/abs/1409.1556).
- [22] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016*, pp. 770–778.
- [23] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, A. Oliva, Learning deep features for scene recognition using places database, in: *Advances in Neural Information Processing Systems, 2014*, pp. 487–495.
- [24] S. Ahmad, K. Ahmad, N. Ahmad, N. Conci, Convolutional neural networks for disaster images retrieval, in: *Proceedings of the MediaEval 2017 Workshop, Sept. 13–15, 2017, Dublin, Ireland, 2017*.
- [25] K. Avgerinakis, A. Mouttazidou, S. Andreadis, E. Michail, I. Gialampoukidis, S. Vrochidis, I. Kompatsiaris, Visual and textual analysis of social media and satellite images for flood detection@ multimedia satellite task MediaEval 2017, in: *Proceedings of the Working Notes Proceeding MediaEval Workshop, Dublin, Ireland, 2017*, pp. 13–15.
- [26] K. Ahmad, A. Sohail, N. Conci, F. De Natale, A comparative study of global and deep features for the analysis of user-generated natural disaster related images, in: *2018 IEEE 13th Image, Video, and Multidimensional Signal Processing Workshop, IVMSP, IEEE, 2018*, pp. 1–5.

- [27] N. Tkachenko, A. Zubiaga, R. Procter, Wisc at mediaeval 2017: Multimedia satellite task, in: Working Notes Proceedings MediaEval Workshop, 2017, p. 2.
- [28] S.A. Chatzichristofis, Y.S. Boutalis, CEDD: color and edge directivity descriptor: a compact descriptor for image indexing and retrieval, in: International Conference on Computer Vision Systems, Springer, 2008, pp. 312–322.
- [29] E. Kasutani, A. Yamada, The MPEG-7 color layout descriptor: a compact image feature description for high-speed image/video segment retrieval, in: Image Processing, 2001. Proceedings. 2001 International Conference on, vol. 1, IEEE, 2001, pp. 674–677.
- [30] Y. Bai, L. Guo, L. Jin, Q. Huang, A novel feature extraction method using pyramid histogram of orientation gradients for smile recognition, in: Image Processing, ICIP, 2009 16th IEEE International Conference on, IEEE, 2009, pp. 3305–3308.
- [31] B. Thomee, D.A. Shamma, G. Friedland, B. Elizalde, K. Ni, D. Poland, D. Borth, L.J. Li, YFCC100M: the new data in multimedia research, *Commun. ACM* 59 (2) (2016) 64–73.
- [32] M. Hanif, M.A. Tahir, M. Khan, M. Rafi, Flood detection using social media data and spectral regression based kernel discriminant analysis, in: Proceedings of the MediaEval 2017 Workshop, Sept. 13–15, 2017, Dublin, Ireland, 2017.
- [33] G. Salton, C. Buckley, Term-weighting approaches in automatic text retrieval, *Inf. Process. Manage.* 24 (5) (1988) 513–523.
- [34] D. Minh-Son, P. Quang-Nhat-Minh, D.N. Duc-Tien, A domain-based late-fusion for disaster image retrieval from social media, in: Proceedings of the MediaEval Workshop. Dublin Google Scholar, 2017.
- [35] F. Alam, F. Ofli, M. Imran, Processing social media images by combining human and machine computing during crises, *Int. J. Hum. Comput. Interact.* 34 (4) (2018) 311–327.
- [36] M. Imran, C. Castillo, J. Lucas, P. Meier, S. Vieweg, AIDR: Artificial intelligence for disaster response, in: Proceedings of the 23rd International Conference on World Wide Web, ACM, 2014, pp. 159–162.
- [37] K. Ahmad, M. Riegler, A. Riaz, N. Conci, D.T. Dang-Nguyen, P. Halvorsen, 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, 2017, pp. 461–465.
- [38] V. Klemas, Remote sensing of floods and flood-prone areas: an overview, *J. Coast. Res.* 31 (4) (2014) 1005–1013.
- [39] Y. Liu, L. Wu, Geological disaster recognition on optical remote sensing images using deep learning, *Procedia Comput. Sci.* 91 (2016) 566–575.
- [40] S.N.K.B. Amit, S. Shiraishi, T. Inoshita, Y. Aoki, Analysis of satellite images for disaster detection, in: Geoscience and Remote Sensing Symposium, IGARSS, 2016 IEEE International, IEEE, 2016, pp. 5189–5192.
- [41] F. Yu, V. Koltun, Multi-scale context aggregation by dilated convolutions, 2015, arXiv preprint arXiv:1511.07122.
- [42] V. Badrinarayanan, A. Kendall, R. Cipolla, Segnet: A deep convolutional encoder-decoder architecture for image segmentation, *IEEE Trans. Pattern Anal. Mach. Intell.* (12) (2017) 2481–2495.
- [43] J. Son, S.J. Park, K.-H. Jung, Retinal vessel segmentation in fundoscopic images with generative adversarial networks, 2017, arXiv preprint arXiv:1706.09318.
- [44] R. De Maesschalck, D. Jouan-Rimbaud, D.L. Massart, The mahalanobis distance, *Chemometr. Intell. Lab. Syst.* 50 (1) (2000) 1–18.
- [45] Y. Feng, S. Shebotnov, C. Brenner, M. Sester, Ensembled convolutional neural network models for retrieving flood relevant tweets, in: Proceedings of the MediaEval 2018 Workshop, Sophia-Antipolis, France, Oct. 29–31, 2018.
- [46] P. Bojanowski, E. Grave, A. Joulin, T. Mikolov, Enriching word vectors with subword information, *Trans. Assoc. Comput. Linguist.* 5 (2017) 135–146.
- [47] L. Lopez-Fuentes, A. Farasin, H. Skinnemoen, P. Garza, Deep Learning models for passability detection in flooded roads, Sophia-Antipolis, France, Oct. 29–31, 2018.
- [48] J. Pennington, R. Socher, C. Manning, Glove: Global vectors for word representation, in: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP, 2014, pp. 1532–1543.
- [49] A. Kirchkopf, D. Slijepcevic, M. Zeppelzauer, M. Seidl, Detection of road passability from social media and satellite images, in: Proceedings of the MediaEval 2018 Workshop, Sophia-Antipolis, France, Oct. 29–31, 2018.
- [50] A. Moumtzidou, P. Giannakeris, S. Andreadis, A. Mavropoulos, G. Meditskos, I. Gialampoukidis, K. Avgerinakis, I. Kompatsiaris, A multimodal approach in estimating road passability through a flooded area using social media and satellite images, in: Proceedings of the MediaEval 2018 Workshop, Sophia-Antipolis, France, Oct. 29–31, 2018.
- [51] Z. Zhengyu, M. Larson, N. Oostdijk, Exploiting local semantic concepts for flooding-related social image classification, in: Proceedings of the MediaEval 2018 Workshop, Sophia-Antipolis, France, Oct. 29–31, 2018.
- [52] M. Hanif, M. Tahir, M. Rafi, Detection of passable roads using ensemble of global and local features, in: Proceedings of the MediaEval 2018 Workshop, Sophia-Antipolis, France, Oct. 29–31, 2018.
- [53] K. Ahmad, M.L. Mekhalfi, N. Conci, F. Melgani, F.D. Natale, Ensemble of deep models for event recognition, *ACM Trans. Multimedia Comput. Commun. Appl. OMM* 14 (2) (2018) 51.
- [54] K. Pogorelov, S.L. Eskeland, T. de Lange, C. Griwodz, K.R. Randel, H.K. Stensland, D.T. Dang-Nguyen, C. Spampinato, D. Johansen, M. Riegler, et al., A holistic multimedia system for gastrointestinal tract disease detection, in: Proceedings of the 8th ACM on Multimedia Systems Conference, ACM, 2017, pp. 112–123.
- [55] K. Pogorelov, O. Ostroukhova, M. Jeppsson, H. Espeland, C. Griwodz, T. de Lange, D. Johansen, M. Riegler, P. Halvorsen, Deep learning and hand-crafted feature based approaches for polyp detection in medical videos, in: 2018 IEEE 31st International Symposium on Computer-Based Medical Systems, CBMS, IEEE, 2018, pp. 381–386.
- [56] K. Pogorelov, O. Ostroukhova, A. Petlund, P. Halvorsen, T. de Lange, H.N. Espeland, T. Kupka, C. Griwodz, M. Riegler, Deep learning and handcrafted feature based approaches for automatic detection of angiectasia, in: Biomedical & Health Informatics, BHI, 2018 IEEE EMBS International Conference on, IEEE, 2018, pp. 365–368.
- [57] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, T. Darrell, DeCAF: A deep convolutional activation feature for generic visual recognition, in: Proceedings of the Internal Conference on Machine Learning, ICML, vol. 32, 2014, pp. 647–655.
- [58] T. Tieleman, G. Hinton, Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude, COURSE: Neural Netw. Mach. Learn. 4 (2) (2012).
- [59] D. Dias, U. Dias, Flood detection from social multimedia and satellite images using ensemble and transfer learning with CNN architectures, in: Proceedings of the MediaEval 2018 Workshop, Sophia-Antipolis, France, Oct. 29–31, 2018.