People and vehicles in danger - A fire and flood detection system in social media

Panagiotis Giannakeris, CERTH-ITI, Thessaloniki, Greece, giannakeris@iti.gr
Konstantinos Avgerinakis, CERTH-ITI, Thessaloniki, Greece, koafgeri@iti.gr
Anastasios Karakostas, CERTH-ITI, Thessaloniki, Greece, akarakos@iti.gr
Stefanos Vrochidis, CERTH-ITI, Thessaloniki, Greece, stefanos@iti.gr
Ioannis Kompatsiaris, CERTH-ITI, Thessaloniki, Greece, ikom@iti.gr

***Abstract*—This paper presents a novel warning system framework for detecting people and vehicles in danger. The system was tested in several images compiled from Flickr and other social media sources and is highly suggested to get integrated in future warning surveillance and safety systems for preventing or solving crisis events. The proposed framework recruits State-of-the-Art deep learning technologies so as to solve a series of image processing and machine learning challenges and provides a near real-time localization solution for detecting and scoring severity safety levels of people and vehicles in flood and fire images.**

***Index Terms*—computer vision, deep learning, image segmentation, object detection**

I. Introduction

Nowadays, it is well known that social media analytics play a great role on several aspects of daily life. This technology proved that the analysis of diverse information such as group and individual behaviour in public and private places, market upside-downs and customer preferences, emergency or crisis detection in urban and rural environments, can provide meaningful metadata so as to solve or optimize the solution of current societal problems.

More specifically, image and video capturing by the use of mobile and wearable cameras in conjunction with the sharing culture that the internet brought in human societies has inspired the deployment of several computer vision and image understanding technologies which endeavour to solve several societal problems with security and safety domains to be one of the most serious amongst others. Millions of images are daily uploaded on social media, while a great deal of them might include the existence of a crisis or emergency event. Taken this into account and inspired from the recent advance in image understanding, we suggest a novel framework that combines several technologies so as to detect and score the danger that people and vehicles might be in fire and flood scenarios. As far as we are concerned, while there might be a lot of works that have focused either holistically (i.e. image classification) or locally (i.e. image segmentation) on the existence of fire and flood detection and separate with people detection, there is none that have proposed a combined severity level estimation scoring system as ours.

Our framework has been evaluated in benchmark fire and flood classification and localization datasets, while we have also crawled a great deal of emergent images from Flickr so as to test our algorithm in a real case environment. The dataset is a collection of Fire-Flood Flickr (3F) annotated images. The proposed crisis management solution is very important for the authorities since it enhances the crisis situational awareness by using the automatic analysis of multiple files especially during the crisis.

II. Related work

Deep Convolutional Neural Networks [1] constitute the State-of-the-Art in **image classification**. After the initial interest, investigation on deeper models began to take place with [2], where smaller stacked kernels started to be investigated and advanced with Inception [3] and ResNet [4] architectures. As far as security and safety domains are concerned we can encounter flood classification in [5], [6] and fire classification in [7], [8].

Semantic image segmentation SoA has also tend to use deep CNNs as well [9], [10] by simply changing the objective of the classifier and label each pixel in the image individually, leading to a classification mask for the whole image instead of a recognition class. As far as security and safety domains are concerned, we scarcely find a technique that uses a deep CNN, as there are no groundtruth available masks and the training of these models is infeasible. A worth-to-note technique which performs fire detection in social images with the use of color and texture attributes was presented in [11].

Object detection has numerous applications: autonomous vehicles, smart video surveillance, facial detection, ambient assisted living, etc. Naturally, deep CNN architectures were thoroughly examined for this. Early works such as [12] include multi scale bounding box proposal generation techniques like Selective Search [13], as a feeding mechanism of candidate boxes to deep classifiers. The trend later became to incorporate this function into single shot object detectors, using end-to-end deep architectures [14], [15], [16]. Those models achieved a better trade-off between accuracy and speed. In a previous work of ours [17] we have proposed a novel scheme to detect vehicles and pedestrians from traffic surveillance cameras. The same framework has also been deployed in UA-DETRAC vehicle detection dataset [18], achieving a really high detection rate.

978-1-5386-0951-4/18/$31.00 ©2018 IEEE

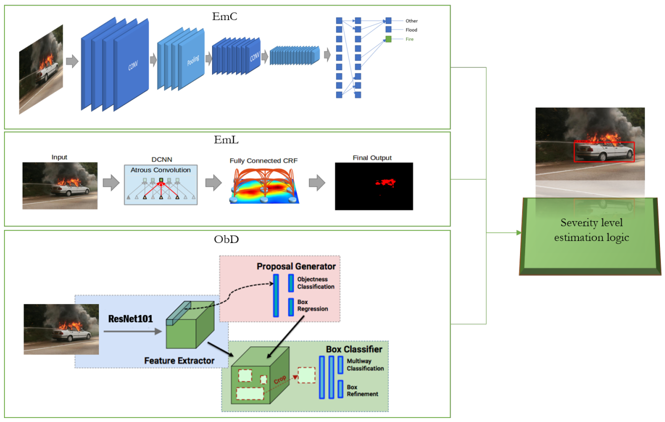

Fig. 1: Overall diagram of our unified framework

III. Methodology

To begin with the analysis of our unified framework, we need to clarify that our work mainly focuses on the detection of people and vehicles that undergo danger in fire and flood scenarios in social media images and other visual sources. For that purposes, we gathered a great deal of visual content and deployed the following set of computer vision techniques: a) **Image classification** so as to determine whether an image contains an emergent event or not; b) **Emergency localization** in order to detect the regions where fire and flood pixels exist in flood and fire pictures; c) **Object detection** so as to find people and vehicles that exist in the image and d) **Severity level estimation** to define the danger that the people and vehicles undergo based on their proximity to the emergent event. The overall diagram is depicted in Figure 1.

A. Emergency classification

The Emergency Classification **(EmC)** component is based on State-of-the-Art image classification techniques and is used so as to determine which images contain an emergency event. Inspired from the recent success that deep learning showed in image understanding [2] and scene recognition [19], we chose to fine-tune the pre-trained parameters of the $VGG-16$ on $Places365$ dataset so as to leverage useful distinctions between various visual clues that relate to generic scenery images. A set of amendments were performed on its architecture so as to fit it to our purposes: Initially we removed the final Fully Connected (FC) layer and replaced it with a new FC layer with a width of 3 freezing the weights up to the previous layer and finally we deployed a softmax classifier so as to enable multi-class recognition. More specifically the EmC results into three-class image recognition: ”Fire”, ”Flood” and ”Other”, where ”Other” may represent any theme except for fire and flood events, e.g. scenes of interior, forests, snowy mountains, crowded streets, urban life etc.

The EmC results are integrated in the framework to indicate the existence of fire and flood events in a holistic manner and the component's purpose is to give an early indication and a first segment of solid information about the existence of an emergency in the image. This information, taken from an initial observation of the whole image, is rather useful to be integrated into the severity level analysis as shown later.

B. Emergency localization

Simultaneously with EmC, we deploy an Emergency Localization **(EmL)** component, which is responsible to semantically segment the regions where fire and flood pixels exist in case EmC's result indicates an emergency situation. Inspired from the recent success that semantic image segmentation achieved in [20], we chose to adopt DeepLab and an architecture of ”atrous convolution”, which uses convolution with up-sampled filters, so as to solve the emergency localization problem in images. We chose DeepLab because ”atrous convolution” allows a wider reception field of the convolution filters, leading to richer context representations, while it also combines the result feature vectors of the final convolutional layer with a fully connected Conditional Random Field (CRF) which provides refined segmentation masks as it includes neighboring context on its calculations.

C. Object detection

The Object Detector **(ObD)** component is responsible to provide a set of bounding boxes of the persons and vehicles in social media images, as well as their immediate surroundings. Groups of people or individuals are detected as persons, while vehicles may contain one of the following categories: cars, trucks, buses, bicycles and motorcycles. The basis of our object detection component is inspired from Faster R-CNN [14], pretrained on COCO dataset [21], with some alterations so as to make it fit to our emergency event purposes. More specifically, based on [22], we chose to deploy $ResNet101$ feature extractor so as to extract deep features and then use an Region of Interest (RoI) pooling scheme to classify candidate boxes. Unfortunately, as far as we are concerned, there is no fire/flood emergency image dataset to depict people or vehicles in immediate danger to train our detector, so we opt to train our model only in COCO and keep only the object classes that we are interested in.

D. Severity level estimation

Our unified framework is completed with the severity level estimation component which combines EmC, EmL and ObD results so as to define a severity level label for each in-danger bounding box of the gathered social images: (a) 'Safe target', (b) 'Target possibly in danger', (c) 'Target in danger', which can also be interpreted as a qualitative risk assessment scale of three levels: 'Low', 'Medium', or 'High' respectively.

The possible outcomes of the system logic are described here:

- Low risk for an emergency event we have when the EmC classifies the candidate image as 'other'. All the detected bounding boxes from the ObjD are declared 'Safe targets'.
- Medium risk we have when an image is classified as emergency from the EmC component (i.e. fire/flood).

All targets that are detected from ObD are automatically characterized as 'Possibly in danger'.
- Elevation to high risk we have when a bounding box detected by ObD coincides with EmL emergency masks (i.e. fire/flood).

IV. Experimental work

To evaluate each component of our framework, we made several experiments on benchmark fire and flood datasets. For Emergency Classification (EmC) we use the MediaEval's Disaster Image Retrieval from Social Media (DIRSM) dataset [6], while for Emergency Localization (EmL) we use the BowFire dataset [11]. As far as severity level estimation is concerned, we didn't find an available dataset for emergency situations on the literature, thus we decided to create our own one with images from Flickr and other social media sources, annotate it and named it as Fire-Flood Flickr (3F) emergency dataset.

A. 3F-emergency dataset

Due to the lack of emergency dataset for severity level estimation, we contribute to the literature the 3F-emergency dataset, which is a collection of social media images for fire and flood emergency scenarios. For compilation purposes, we also used annotated images of flood [6] and fire [11] events, in addition to our retrieved and annotated fire and flood images from Flickr.

More specifically, the expanded 3F-emergency dataset is enriched with DIRSM [6] flood images, BowFire [11] and UIA-CAIR [23] fire images. Our initial 3F-emergency dataset is composed of 6K images from Flickr and in the expanded version another 6K images are added. The goal of this compilation is to create a large enough dataset so as to train the CNN models as more accurate as it can be. For collecting visual content from Flickr, we used Flickr API [24] which lets users search the database with text or tag queries and returns image IDs. We retrieved only images accompanied with a Creative Commons compatible license, while the search queries we used were: 'flood', 'flooded street', 'fire', 'burning vehicle', 'fire accident', 'flood emergency', 'fire explosion', 'wildfire', etc. We also used queries such as 'busy street', 'forests', 'mountains', 'sunrise', 'beach', etc to get images for the 'Other' class.

While all the above pictures accompanied with a class label, only BowFire are complemented with groundtruth masks that can be used for Emergency Localization purposes. Unfortunately, BowFire dataset contain pictures only for fire events, so we led to the temporary solution to use VideoWaterDB [25] which contain different video shots from water scenarios, so as to simulate a flood scenario event. Another unfortunate fact in our localization experiments is that that the groundtruth masks that BowFire provides contains only fire pixels, no matter if flood or large water bodies exists as well, which usually happen in real case scenarios (i.e. a fire-fighter in a flooded city extinguishing a fire that burns a building). That led us to the unfortunate position to remove these images from the emergency localization training, as they would erroneously reduce the accuracy rates when pixels were determined as background and not as flood. Additionally, we added some of the 'Other' type images from the 3F-Dataset to use as background instances. Their ground truth mask was easily created by tagging all the pixels as background. Free-copyright results on 3F-Dataset are provided in [26].

B. System evaluation

For system evaluation purposes, we provide quantitative results for image classification and image segmentation and qualitative results for severity level estimation. Our models were trained in generic datasets and fine-tuned in emergency related ones so as to have a better representation and generalization system.

Image classification evaluation took place in MediaEval's Disaster image retrieval from social media (DIRSM) dataset, where flood and other type of images were provided. A 10 fold cross validation was followed to evaluate Emergency Classification (EmC) component. Recognition accuracy results and comparison with State-of-the-Art are provided in Table I, where we can see that EmC outperforms all image classification methods that were presented in MediaEval's Multimedia Satelite Task 2017, scoring 1.77% higher from the second rival.

A separate classification evaluation was also took place on our 3F-emergency dataset, where from a total of 12423 images, 2485 were selected randomly (20%) as a test set and the remaining were used to train EmC model. Here, the classes were 3: 'fire', 'flood' and 'other', while there were some examples where some of the two coincide and the discrimination was quite difficult to tell. Nevertheless, our framework achieved a mean accuracy recognition rate that reached the 87.32%, with 83.7% 'other' class achieving the lowest score, fire the highest with 93.3% and 88.96% for flood.

Image localization evaluation took place on BowFire [11] and VideoWaterDB [25] datasets for fire and flood segmentation respectively. Comparisons regarding the fire segmentation results took place on BowFire by computing recall and precision metrics and are depicted in Table II. As far as recall is concerned, we can observe that our results are really close to [35] and tied for second place with [36] outperforming the rest, meaning that we found a great deal of pixels that were groundtruthed as fire. On the other hand, as far as precision is concerned, we didn't achieve as well as we expected, as there were a great deal of background pixels that misclassified as fire, leading to lower precision rates than other SoA techniques. These false alarms however, can be alleviated in our warning framework, as the use of EmC component can eliminate a great deal of images that do not contain a threat, which can eventually increase the precision rate on the final severity level estimation.

Image localization evaluation has also been performed on the whole collected dataset (i.e. flood, fire, other) and precision, recall and IoU results are aggregated in Table III. It is noteworthy that we achieved very good precision and recall scores for 'flood' and 'other' classes as they are much more

TABLE I: Image classification results on DIRSM Dataset and comparison with SoA

Authors	Ours	[27]	[6]	[28]	[29]	[30]	[31]	[5]
Accuracy	97.50%	87.88%	92.27%	95.11%	70.16%	87.87%	95.73%	95.71%

TABLE II: Fire localization results on BowFire Dataset and comparison with SoA

Authors	Precision	Recall	F1-score
Celik [32]	52%	68%	53%
Rossi [33]	$< 40\%$	$20\% - 30\%$	$< 30\%$
Rudz [34]	63%	$< 50\%$	$50\% - 60\%$
Chino [11]	50%	$60\% - 70\%$	$50\% - 60\%$
Chen [35]	37%	84%	45%
Avalhais [36]	62%	77%	63%
Zhang [37]	50%	31%	29%
Ours	39%	77%	52%

TABLE III: Localization results on our custom Dataset

Metric	Backgr	Flood	Fire
Precision	97%	92%	39%
Recall	94%	92%	77%
IoU	92%	85%	35%
MeanIoU	71%		

rigid than 'fire' and much harder to be accurately found. IoU score for each class is calculated dividing TP instances by the sum of TP, FP, and FN instances, leading eventually into a decent meanIoU 71.0% segmentation accuracy.

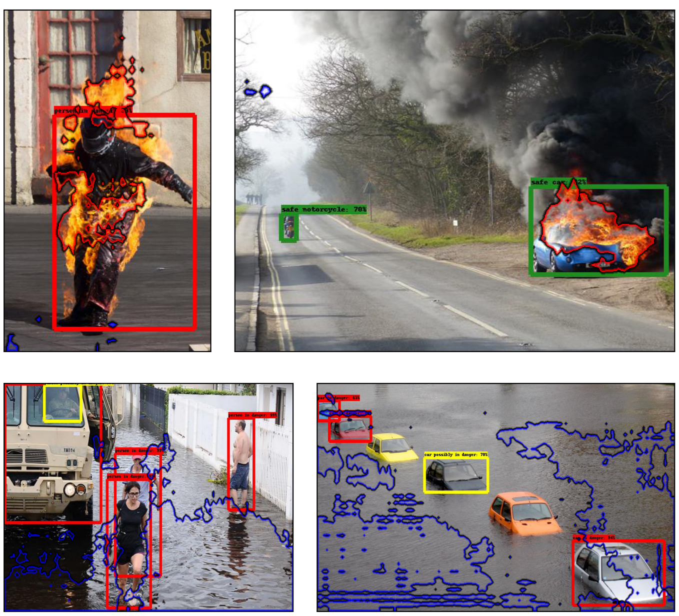

Fig. 2: Qualitative results of people and vehicles in danger

Severity level estimation in this work is evaluated with qualitative results and a set of successful and failure images are provide in Figure 2. We ask the reader to follow our github repository [26] to download extra outcomes of our framework.

We visualize the severity level of danger in the resulting candidate target bounding boxes from the ObD component by using a three color palette to draw them: (a) Green for 'Safe' targets, (b) yellow for 'Possibly in danger' targets, and (c) red for targets classified as being 'in danger'. A second color pallet is used to draw the embossed regions that result from EmL component and are colored as red for fire and blue for flood regions.

Analysing the qualitative results in Figure 2, we can see that our framework can work very well in very demanding situations such as the top-left picture, where a person fired up can be easily isolated from the background environment which is quite irrelevant with the emergent event. A successfully captured flood event is depicted in the bottom left picture, where we can see the people who are in the water obtain an 'in danger' label contrary to the one who is in the car and far from the flood and is labelled as 'possibly in danger'.

Analysing now the failure cases, we can see that in some cases we might have a good EmL mask, but fail to recognise the picture as emergent using EmC, giving an erroneous 'safe' label, like the car which is on fire on the top-right picture. However, examining a plethora of testing samples it can be confirmed that this failure rarely appears ([26]). Other, more frequent cases of failure are showcased in the bottom right picture, where a series of flooded cars is depicted. As we can see there might be some cases where the ObD may not find all the targets or the EmL mask is not so well formed, leading to missing or erroneous labels. This was very usual in flood scenarios, where the water covers a great deal of the object or the object occludes the water, leading to bad bounding boxes and segmentation masks, contrary to fire events where the fire usually occludes the target and not vice versa.

Overall, on the most of the test samples that were examined, rarely a target in danger did not get at least a 'Medium' level tag. The most cases of inconsistency and confusion happened frequently between 'Medium' and 'High' level tags, because EmL did not work so well in flood cases, which is mainly attributed to the lack of groundtruth masks to train the model. Flood detection worked very well in the provided dataset but it undoubtedly needs much more data to generalize it in more emergent situations.

V. Conclusion

We conclude that our work could contribute greatly to crisis management procedures during and after such events and can be integrated in crisis management and decision support systems. Furthermore, we hope that with this work we will manage to arouse interest towards the creation of larger scale datasets that will be used to analyse emergency situations.

Acknowledgment

This work was supported by beAWARE [38] and EOPEN [39] project partially funded by the European Commission under grant agreement No 700475 and 776019.

REFERENCES

[1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.

[2] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *CoRR*, vol. abs/1409.1556, 2014. [Online]. Available: http://arxiv.org/abs/1409.1556

[3] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, A. Rabinovich *et al.*, "Going deeper with convolutions." Cvpr, 2015.

[4] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.

[5] B. Bischke, P. Bhardwaj, A. Gautam, P. Helber, D. Borth, and A. Dengel, "Detection of flooding events in social multimedia and satellite imagery using deep neural networks," in *Working Notes Proc. MediaEval Workshop*, 2017, p. 2.

[6] K. Avgerinakis, A. Moumtzidou, S. Andreadis, E. Michail, I. Gialampoukidis, S. Vrochidis, and I. Kompatsiaris, "Visual and textual analysis of social media and satellite images for flood detection@ multimedia satellite task mediaeval 2017," in *Multimedia satellite task MediaEval 2017*, 2017.

[7] K. Muhammad, J. Ahmad, and S. W. Baik, "Early fire detection using convolutional neural networks during surveillance for effective disaster management," *Neurocomputing*, 2017.

[8] J. Sharma, O.-C. Granmo, M. Goodwin, and J. T. Fidje, "Deep convolutional neural networks for fire detection in images," in *International Conference on Engineering Applications of Neural Networks*. Springer, 2017, pp. 183–193.

[9] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 3431–3440.

[10] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.

[11] D. Y. Chino, L. P. Avalhais, J. F. Rodrigues, and A. J. Traina, "Bowfire: detection of fire in still images by integrating pixel color and texture analysis," in *Graphics, Patterns and Images (SIBGRAPI), 2015 28th SIBGRAPI Conference on*. IEEE, 2015, pp. 95–102.

[12] R. Girshick, "Fast r-cnn," *arXiv preprint arXiv:1504.08083*, 2015.

[13] J. R. Uijlings, K. E. Van De Sande, T. Gevers, and A. W. Smeulders, "Selective search for object recognition," *International journal of computer vision*, vol. 104, no. 2, pp. 154–171, 2013.

[14] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in *Advances in neural information processing systems*, 2015, pp. 91–99.

[15] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 779–788.

[16] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," in *European conference on computer vision*. Springer, 2016, pp. 21–37.

[17] K. Avgerinakis, P. Giannakeris, A. Briassouli, A. Karakostas, S. Vrochidis, and I. Kompatsiaris, "Intelligent traffic city management from surveillance systems (certh-iti)," in *IEEE Smart World 2017, NVIDIA AI city challenge*, 2017.

[18] L. Wen, D. Du, Z. Cai, Z. Lei, M. Chang, H. Qi, J. Lim, M. Yang, and S. Lyu, "DETRAC: A new benchmark and protocol for multi-object tracking," *CoRR*, vol. abs/1511.04136, 2015. [Online]. Available: http://arxiv.org/abs/1511.04136

[19] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva, "Learning deep features for scene recognition using places database," in *Advances in neural information processing systems*, 2014, pp. 487–495.

[20] L. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs," *CoRR*, vol. abs/1606.00915, 2016. [Online]. Available: http://arxiv.org/abs/1606.00915

[21] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *European conference on computer vision*. Springer, 2014, pp. 740–755.

[22] J. Huang, V. Rathod, C. Sun, M. Zhu, A. Korattikara, A. Fathi, I. Fischer, Z. Wojna, Y. Song, S. Guadarrama *et al.*, "Speed/accuracy trade-offs for modern convolutional object detectors," in *IEEE CVPR*, 2017.

[23] UIA-CAIR, "Fire-Detection-Image-Dataset," https://github.com/cair/Fire-Detection-Image-Dataset.

[24] Flickr, "Flickr API," https://www.flickr.com/services/api/.

[25] P. Mettes, R. T. Tan, and R. C. Veltkamp, "Water detection through spatio-temporal invariant descriptors," *Computer Vision and Image Understanding*, vol. 154, pp. 182–191, 2017.

[26] "PVD Framework github page," https://github.com/ppgiannak/pvd-framework.

[27] 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, and R. d. S. Torres, "Data-driven flood detection using neural networks."

[28] K. Ahmad, P. Konstantin, M. Riegler, N. Conci, and P. Holversen, "Cnn and gan based satellite and social media data fusion for disaster detection," in *Working Notes Proc. MediaEval Workshop*, 2017, p. 2.

[29] L. Lopez-Fuentes, J. van de Weijer, M. Bolaños, and H. Skinnemoen, "Multi-modal deep learning approach for flood detection."

[30] M. S. Dao, Q. N. M. Pham, D. Nguyen, and D. Tien, "A domain-based late-fusion for disaster image retrieval from social media," 2017.

[31] S. Ahmad, K. Ahmad, N. Ahmad, and N. Conci, "Convolutional neural networks for disaster images retrieval."

[32] T. Celik and H. Demirel, "Fire detection in video sequences using a generic color model," *Fire Safety Journal*, vol. 44, no. 2, pp. 147–158, 2009.

[33] L. Rossi, M. Akhloufi, and Y. Tison, "On the use of stereovision to develop a novel instrumentation system to extract geometric fire fronts characteristics," *Fire Safety Journal*, vol. 46, no. 1-2, pp. 9–20, 2011.

[34] S. Rudz, K. Chetehouna, A. Hafiane, H. Laurent, and O. Séro-Guillaume, "Investigation of a novel image segmentation method dedicated to forest fire applications," *Measurement Science and Technology*, vol. 24, no. 7, p. 075403, 2013.

[35] C.-S. Chen, C.-W. Yeh, and P.-Y. Yin, "A novel fourier descriptor based image alignment algorithm for automatic optical inspection," *Journal of Visual Communication and Image Representation*, vol. 20, no. 3, pp. 178–189, 2009.

[36] L. P. Avalhais, J. Rodrigues, and A. J. Traina, "Fire detection on unconstrained videos using color-aware spatial modeling and motion flow," in *Tools with Artificial Intelligence (ICTAI), 2016 IEEE 28th International Conference on*. IEEE, 2016, pp. 913–920.

[37] J. Zhang, X. Wang, and M. Lv, "Flame image segmentation algorithm based on background subtraction," in *PIAGENG 2013: Image Processing and Photonics for Agricultural Engineering*, vol. 8761. International Society for Optics and Photonics, 2013, p. 876112.

[38] "beAWARE Project," http://beaware-project.eu.

[39] "EOPEN Project," http://eopen-project.eu/.