

A MULTIMODAL SUPERVISED MACHINE LEARNING APPROACH FOR SATELLITE-BASED WILDFIRE IDENTIFICATION IN EUROPE

Angelica Urbanelli^{1*}, Luca Barco^{1*}, Edoardo Arnaudo^{1,2}, Claudio Rossi¹

1. LINKS Foundation, *AI, Data & Space (ADS)*, Torino (TO), Italy

2. Politecnico di Torino, *Dipartimento di Automatica e Informatica (DAUIN)*, Torino (TO), Italy

ABSTRACT

The increasing frequency of catastrophic natural events, such as wildfires, calls for the development of rapid and automated wildfire detection systems. In this paper, we propose a wildfire identification solution to improve the accuracy of automated satellite-based hotspot detection systems by leveraging multiple information sources. We cross-reference the thermal anomalies detected by the Moderate-resolution Imaging Spectroradiometer (MODIS) and the Visible Infrared Imaging Radiometer Suite (VIIRS) hotspot services with the European Forest Fire Information System (EFFIS) database to construct a large-scale hotspot dataset for wildfire-related studies in Europe. Then, we propose a novel multimodal supervised machine learning approach to disambiguate hotspot detections, distinguishing between wildfires and other events. Our methodology includes the use of multimodal data sources, such as the ERSI annual Land Use Land Cover (LULC) and the Copernicus Sentinel-3 data. Experimental results demonstrate the effectiveness of our approach in the task of wildfire identification.

Index Terms— machine learning, computer vision, earth observation, hotspot disambiguation.

1. INTRODUCTION

The world is experiencing a surge in catastrophic natural events, including wildfires, floods, and storms, the frequency and magnitude of which are exacerbated by climate change effects. The rising global average temperature is expected to increase the risk of fires across a wide range of latitudes, leading to more frequent and impacting events and consequently to the release of a significant amount of smoke, carbon dioxide and heat, which further drives up climate change [1]. Automated wildfire detection systems are crucial to identify fires in their early stage, allowing a prompt response and a consequent reduction of the impacts on human ecosystems, especially in remote areas that are sparsely populated and urbanized.

In this context, the Moderate-resolution Imaging Spectroradiometer (MODIS), and the Visible Infrared Imaging Radiometer Suite (VIIRS) provide operational services aimed at detecting thermal anomalies. On top of these systems, various solutions have been developed to convert raw measurements into *hotspot detections*, i.e., thermal anomalies that are likely to be caused by fires. Existing *hotspot detections* services implement algorithms that can adapt to different instruments and resolutions, filtering out as many *false positive* as possible. However, the accuracy of such services is not ideal for integrating their outputs in operational chains aimed at the automatic detection of wildfires.

In this work, we propose to tackle this problem through an *hotspot disambiguation* task, introducing a novel supervised machine learning approach leveraging multiple data sources to increase the accuracy of wildfire detection from satellite-based thermal anomalies in Europe. Moreover, we openly release both the code of the proposed approach as well as the large-scale dataset we use for the performance evaluation ¹.

This paper is organized as follows. In Section 2, we present a review of the related works in the field of automatic fire detection using satellite observations, while in Section 3 we detail the employed data sources and how we construct our large-scale hotspot dataset. In Section 4, we describe the proposed methodology, together with the machine learning models used, while in Section 5 we present our experimental settings and results. Lastly, in Section 6, we draw our conclusions and discuss potential future works.

2. RELATED WORKS

In the context of fire detection, significant efforts have been carried out to create scalable approaches able to extract thermal anomalies from raw satellite data. These works, carried out by *Giglio et al.* can adapt to different satellites, resolutions, sensors, and data sources, e.g., MODIS and VIIRS [2, 3], providing collections of hotspots that can be used for downstream tasks, such as alerting or mapping services. Since there is no manual or automatic validation provided

¹This work was carried out in the context of the H2020 projects: SAFERS (GA n.869353) and OVERWATCH (GA n.101082320)

* Equal contribution.

¹Code and dataset are available at the following URL: <https://github.com/links-ads/hotspot-disambiguation>

with the outputs of these algorithms, some studies compare the output of MODIS and VIIRS hotspot detection algorithms using various data sources, namely Landsat feeds [4], Geostationary Earth Observation satellites with a lower revisit time [5], the ESA Climate Change Initiative Land Cover and Global Annual Burned Area Maps [6], and the Copernicus Global Land Service Land Cover [7]. While a manual validation approach is not viable nor feasible without relying on data from other satellites or alternative algorithms, the automatic validation of currently detected hotspots for wildfires identification remains an open challenge.

In this work, we leverage on the European Forest Fire Information System (EFFIS), which is the most important European database of burned areas, to extract the MODIS and VIIRS hotspots that correspond to real wildfires and use such ground truth to train different supervised machine learning models using multimodal data, namely the ESRI annual Land Use Land Cover (LULC), Sentinel-3 Sea and Land Surface Temperature Radiometer (SLSTR) and Ocean and Land Colour Instrument (OLCI). We compare the performance of both classical machine learning models and more recent deep learning architectures, obtaining substantial improvements in wildfire classification accuracy.

3. DATASET

3.1. Data sources

We build our wildfire hotspot dataset by harmonizing several data sources, specifically: MODIS, VIIRS, EFFIS, Sentinel-3 SLSTR and OLCI, LULC.

MODIS is an instrument developed by NASA, collecting data across 36 spectral bands with a spatial resolution varying within the range of 250 to 1000 meters. It is able to identify areas that are hotter than their surroundings and flag them as active fires. In this study, we exploit the MCD14ML data collection, which has a spatial resolution of 1km. Similarly, VIIRS is a NASA instrument designed to be complementary to MODIS. It uses similar fire detection algorithms but with a higher spatial resolution, specifically VNP14 and VNP14IMG with 750 m/pixel and 350 m/pixel, respectively.

EFFIS instead is the reference European service providing critical information about wildfire risks and impacts, allowing for more effective prevention, mitigation, and response efforts. Since 2015, it is part of the Copernicus Emergency Management System². The EFFIS Burnt Area product contains a list of manually validated wildfires events detected from the daily processing of MODIS and Sentinel-2 satellite imagery, at 250m and 20m spatial resolution, respectively. Wildfires registered in this database must have a minimum area of 30 hectares, thus excluding smaller fires. In addition, the database also provides a rough estimate of the start and end time for each event, but they may not correspond to the

actual timeline, especially considering its end. For this reason, we estimate the extinction date of the fire considering the first date on which there are less than 2 hotspots in that area.

Sentinel-3 is a satellite mission that provides global observations of Earth’s oceans, land surfaces, and atmosphere to support environmental and climate monitoring. The satellite was launched on 16/02/2016 and hosts four instruments, including the SLSTR and the OLCI. The SLSTR instrument has 9 spectral bands and 2 additional active fire bands useful for fire monitoring. It has a spatial resolution of 1 km, with a revisit time of less than half a day. The OLCI instrument has 21 spectral bands that measure biological information, including the land cover. It has a spatial resolution of 300m with a revisit time of less than 2 days with 2 satellites. Data are gathered with a spatial resolution of 300m. A bicubic resampling algorithm is used for upsampling the 1 km resolution SLSTR bands.

Lastly, LULC is a collection of annual maps of land cover and land use provided by ESRI, and derived from Sentinel-2 imagery at 10m resolution by exploiting deep learning models. These maps comprise 9 land cover classes, namely: *water, trees, flooded vegetation, crops, built area, bare ground, snow/ice, clouds and rangeland* [8]. For this analysis, the map of 2018 has been used with a spatial resolution of 300m to be compliant with Sentinel-3 data.

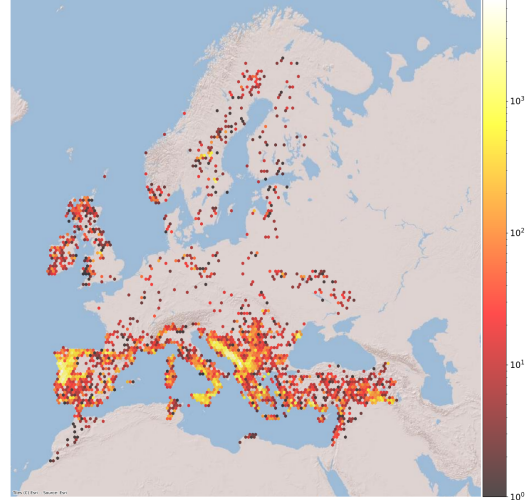


Fig. 1. Distribution of hotspots intersecting with EFFIS burned areas. The concentration of points is visually depicted through a logarithmic color scale.

3.2. Cross-Referencing

We construct our hotspot dataset by cross-referencing MODIS and VIIRS thermal anomalies extracted from the algorithm by Giglio *et al.* on the Mediterranean region with the EFFIS wildfire database, obtaining a collection of more than 9.5M hotspots from January 2012 to August 2022. Specifically, we

²www.copernicus.eu

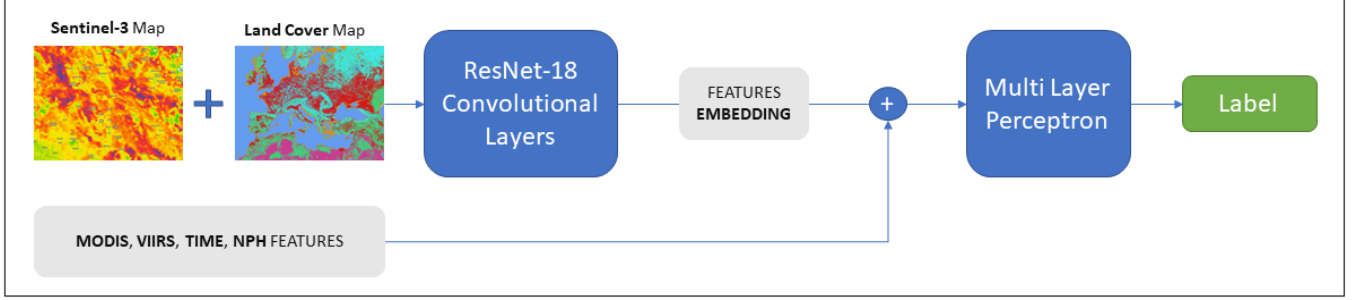


Fig. 2. Architecture schema for the end-to-end solution (RN-E2E).

consider a hotspot to be positive if it co-occurs in both space and time with the corresponding burned areas, indicating a strong association between the detected thermal anomalies and the documented fire events in the EFFIS database. Figure 1 shows the distribution of positive hotspots, with areas of pronounced concentration, whereas others display moderate dispersion and regions devoid of fire activity are also observed.

3.3. Data Preprocessing and Integration

To cope with the limited resources and computational capacity, we train our models on a subset of 500k data points, starting the selection from February 2016 due to Sentinel-3 data availability. These points are selected performing a filter on the timestamp, then a stratified sampling approach is employed to maintain the original dataset’s geographical distribution and label proportions. Given the imbalanced proportion of positive and negative instances, amounting to 3% and 97% respectively, we undersample the negative class targeting a less skewed distribution, namely 10% of positive examples with 90% of negative examples, while trying to ensure the geographical distribution balancing. In addition to the other data sources (i.e., LULC, S3) we compute the Number of Previous Hotspots (NPH) occurring in the same place 12, 24 and 36 hours before, as well as time-related features that provide insights into the weekly periodicity of the event (*Time* features). Among the available MODIS and VIIRS features, we select the Fire Radiative Power (FRP) and the following bands: t_{21} , t_{31} , t_{m13} , t_{m15} , t_{i4} , t_{i5} . We keep all the bands from Sentinel-3 and Land Cover.

4. METHODOLOGY

Baselines. First, we conduct our experiments on classical machine learning algorithms, namely Logistic Regression (LR), Multi Layer Perceptron (MLP), and XGBoost (XGB). These algorithms are trained using all the computed features associated with the hotspot points. The algorithm yielding the best result acts as baseline for the subsequent experiments.

RN-Van: Vanilla ResNet. We want to utilize the complete information available from Sentinel-3 and LULC data, which covers an area of $90km^2$ centered around the hotspot point. For this purpose, we exploit standard computer vision models to efficiently encode the information into denser, semantically meaningful, vector representation. Specifically, in this work we employ a ResNet-18 [9] on the S3 inputs.

RN-E2E: End-To-End ResNet. To assess the potential contribution of additional data sources, we evaluate the performance of training the ResNet-18 with a downstream Multi Layer Perceptron. The latter takes as input the concatenation of Sentinel-3 and LULC features embedding provided by ResNet-18 and the ones from MODIS, VIIRS, Time and NPH, as shown in Figure 2.

5. EXPERIMENTS

5.1. Implementation details

In our experiments, we define six Feature Sets (FS) as outlined in Table 1 in order to analyze the performance of different scenarios with varying data availability. For the experiments with ResNet, we use maps with dimension of 32×32 pixels and 33 channels (32 from Sentinel-3 and 1 from LULC).

	FS1	FS2	FS3	FS4	FS5	FS6
MODIS	✓	✓	✓	✓		
VIIRS	✓	✓	✓	✓		
Time	✓	✓	✓	✓	✓	
Land Cover		✓	✓	✓	✓	✓
Sentinel-3			✓	✓	✓	✓
NPH				✓	✓	

Table 1. Feature Sets configuration.

We divide the dataset into 50 splits following the methodology described in Section 3. We then select 28 splits for train, 14 for validation and 8 for tests, resulting in the following ratios: 56%, 28%, 16%, respectively.

Using a grid search approach, we select the optimal hyperparameters for the XGBoost model, specifically a learning

rate of $\lambda = 0.1$, a tree maximum depth of 12, a feature subsample ratio of 0.8, and a positive class weight equal to 10.

For experiments with RN-Van and RN-E2E, we only consider the FS6 and FS4 configurations, namely the feature sets with every source included in the first case, and with only S3 and LULC in the second case. All the training processes are conducted for 20 epochs, using a batch size of 128, a learning rate of $1e^{-3}$, using AdamW as optimizer, and Binary Cross Entropy (BCE) Loss function. While baselines are trained on CPU only, the ResNet experiments require around 20 hours on a single NVIDIA A100 GPU.

5.2. Results

	LR	MLP	XGB (baseline)
FS1	6.80	34.51	46.88

Table 2. Comparison of baseline models in terms of F1 score

Features Set	FS1	FS2	FS3	FS4	FS5	FS6
Model	XGB	XGB	XGB	XGB RN-E2E	XGB	XGB RN-Van
F1	46.88	56.21	78.99	80.53 76.47	80.02	71.40 76.16

Table 3. Results of XGBoost (XGB) and ResNet architectures in terms of F1 score.

In every experiment, we adopt the *F1 score* as performance metric. Table 2 highlights the results obtained by the baseline approaches. We observe that simple solutions such as the linear regression are not sufficient to tackle the task at hand, while the even a MLP is heavily outperformed by XGBoost. For this reason, we select the latter as baseline for the following experiments.

As evidenced by the results presented in Table 3, expanding the feature set with additional data sources proves to be beneficial for the task. Specifically, XGBoost manages to reach an F1 score of 80.53 in the FS4 configuration, resulting in a +33.65 performance improvement with respect to the baseline. Furthermore, the scores achieved with the FS5 and FS6 configurations indicate that the proposed system exhibits robust performance even in the absence of MODIS and VIIRS data sources. This capability enables the system to be deployed in near real-time scenarios, conditioned by the availability of Sentinel-3 data. Finally, we observe that RN-Van outperforms XGBoost on FS6 with a +4.76 performance increase, as it likely leverages the data from the entire area more effectively, unlike XGBoost which only considers individual pixels. On the other hand, using the RN-E2E architecture incorporating all the point features with FS4, leads to a negligible improvement compared to RN-Van (i.e., +0.31). Similarly, this can probably be attributed to the underlying architecture that cannot exploit the external features as well as the visual ones. XGBoost still demonstrates superior results in this configuration.

6. CONCLUSIONS AND FUTURE WORKS

We proposed a novel multimodal supervised machine learning approach to enable a better automated near real-time wildfire detection system. The proposed approach process data from multiple satellite sources, showing that by adding more characterizing data sources, the wildfire detection performance increases. Moreover, promising results can also be obtained using data from Sentinel-3 and LULC only.

Future research endeavours could encompass the extension of the applicability of this work to address diverse fire types, including small ones, or focus on a more precise real-time fire detection system by leveraging time series analysis on geostationary satellites such as the upcoming Meteosat generation three.

7. REFERENCES

- [1] The Intergovernmental Panel on Climate Change (IPCC), "Climate Change 2022: Impacts, Adaptation and Vulnerability," 2022, [Online; accessed 17-May-2023].
- [2] Louis Giglio, Jacques Descloitres, Christopher O. Justice, and Yoram J. Kaufman, "An enhanced contextual fire detection algorithm for modis," *Remote Sensing of Environment*, vol. 87, pp. 273–282, 2003.
- [3] Wilfrid Schroeder, Patricia Oliva, Louis Giglio, and Ivan A. Csiszar, "The new viirs 375m active fire detection data product: Algorithm description and initial assessment," *Remote Sensing of Environment*, vol. 143, pp. 85–96, 2014.
- [4] W Schroeder, I Csiszar, L Giglio, E Ellicott, and C Justice, "Satellite active fire product validation using high spatial resolution reference data," .
- [5] Chermelle B. Engel, Simon D. Jones, and Karin Reinke, "A seasonal-window ensemble-based thresholding technique used to detect active fires in geostationary remotely sensed data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 59, no. 6, pp. 4947–4956, 2021.
- [6] Yuyun Fu, Rui Li, Xuewen Wang, Yves Bergeron, Osvaldo Valeria, Raphaël D. Chavardès, Yipu Wang, and Jiheng Hu, "Fire detection and fire radiative power in forests and low-biomass lands in northeast asia: Modis versus viirs fire products," *Remote Sensing*, vol. 12, no. 18, 2020.
- [7] Kadir Alperen Coskuner, "Assessing the performance of modis and viirs active fire products in the monitoring of wildfires: a case study in turkey," *iForest - Biogeosciences and Forestry*, , no. 2, pp. 85–94, 2022.
- [8] Krishna Karra, Caitlin Kontgis, Zoe Statman-Weil, Joseph C. Mazariello, Mark Mathis, and Steven P. Brumby, "Global land use / land cover with sentinel 2 and deep learning," in *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*, 2021, pp. 4704–4707.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," *CoRR*, vol. abs/1512.03385, 2015.