
Mesogeos: A multi-purpose dataset for data-driven wildfire modeling in the Mediterranean

Spyros Kondylatos^{1,2}, Ioannis Prapas^{1,2}, Gustau Camps-Valls², and Ioannis Papoutsis¹

¹Orion Lab, Institute for Astronomy, Astrophysics, Space Applications, and Remote Sensing,
National Observatory of Athens

²Image Processing Laboratory (IPL), Universitat de València
{skondylatos, iprapas, ipapoutsis}@noa.gr
{gustau.camps}@uv.es

Abstract

We introduce Mesogeos¹, a large-scale multi-purpose dataset for wildfire modeling in the Mediterranean. Mesogeos integrates variables representing wildfire drivers (meteorology, vegetation, human activity) and historical records of wildfire ignitions and burned areas for 17 years (2006-2022). It is designed as a cloud-friendly spatio-temporal dataset, namely a datacube, harmonizing all variables in a grid of 1km x 1km x 1-day resolution. The datacube structure offers opportunities to assess machine learning (ML) usage in various wildfire modeling tasks. We extract two ML-ready datasets that establish distinct tracks to demonstrate this potential: (1) short-term wildfire danger forecasting and (2) final burned area estimation given the point of ignition. We define appropriate metrics and baselines to evaluate the performance of models in each track. By publishing the datacube, along with the code to create the ML datasets and models, we encourage the community to foster the implementation of additional tracks for mitigating the increasing threat of wildfires in the Mediterranean.

1 Introduction

Wildfires play a key role in the ecosystem [1–4], yet they present risks to both humans and the environment [5]. The threat is inflated by climate change, which aggravates the frequency and extremity of wildfire events [6], particularly in Mediterranean-type climate regions [7, 8]. The changes are expected to be more and more prevalent in the following years [9]; thus, there is a pressing need for innovative solutions to enhance wildfire preparedness and management, enabling adaptation to evolving conditions. The development of such solutions is hampered by the complexity to model wildfires, resulting from the dynamic interactions between several fire drivers such as climate, vegetation, and human activity [10], operating across different spatial and temporal scales.

Traditional models [11–13] ignore these intricate interactions. In contrast, Machine Learning (ML) offers the potential to capture them in a data-centric manner. Nevertheless, the application of ML in the context of wildfires requires careful consideration [14]. The wildfire occurrence is stochastic, which means that the same environmental conditions may lead or not to a fire ignition. Moreover, wildfires are rare events that can lead to imbalanced or sparse datasets. Despite the challenges, ML has been employed successfully in several applications [15]. Particularly, Deep Learning (DL) has been suggested as a method for modeling Earth System problems, including wildfires [16, 17].

¹Inspired from the Greek word Μεσόγειος (Mesógeios), widely adopted to refer to the Mediterranean Sea.

Although the potential of DL in wildfire modeling appears promising, its adoption is still not widespread. One major obstacle is the limited availability of extensive datasets necessary to support its utilization. The vast amount of data required to model wildfires at a larger scale presents difficulties in the collection and curation of the data. The data sources are often scattered across different platforms and become available in diverse formats and resolutions. Thus, the community lacks a large-scale dataset suitable for various ML tasks in the context of wildfires.

In this work, we introduce *Mesogeos*, an extensive multi-purpose dataset designed to support the development of ML models for various wildfire applications in the Mediterranean. It contains a complete set of variables associated with fire drivers, i.e. meteorological conditions, vegetation characteristics, and anthropogenic factors. It also encompasses past burned areas, fire ignition points, and burned area sizes that can serve as predictands for diverse ML tasks. *Mesogeos* is harmonized in a standard spatiotemporal grid format, namely a datacube [18], with a daily temporal resolution and a spatial resolution of $1km \times 1km$, containing data from 2006 to 2022. The datacube structure facilitates the extraction of ML-ready datasets for numerous applications. To the best of our knowledge, *Mesogeos* is the largest harmonized, multi-purpose dataset for data-driven wildfire modeling.

To demonstrate the datacube’s potential applications, we extract from it two ready-to-consume ML datasets: one tailored for the next day’s wildfire danger forecasting and one for burned area size prediction, given the ignition. We employ DL models to establish benchmarks for the two datasets. Furthermore, we propose several additional directions for utilizing the dataset, suggesting its capabilities for addressing other wildfire-related applications. To encourage further research and facilitate the development of similar datasets, we openly publish the *Mesogeos* datacube, the derived datasets and models, and the code used to generate them [19]. We also provide a github repository: <https://github.com/Orion-AI-Lab/mesogeos> and a website for the project: <https://orion-ai-lab.github.io/mesogeos/> with information on how to use the data and code. These resources can be a valuable reference for future implementations and extractions of similar datasets.

2 Related Work

DL has demonstrated successful applications in various tasks related to wildfires. Huot et al. [20] have built segmentation models for predicting fire danger with U-Net-type architectures. Radke et al. [21] developed FireCast, a fire spread model leveraging Convolutional Neural Networks (CNNs) that demonstrated superior performance compared to physics-based models. Similarly, Hodges and Lattimer [22] and Burge et al. [23] employed DL techniques to predict fire evolution by training on fire simulations. Lastly, Ba et al. [24] addressed the fire detection task, by developing SmokeNet, a CNN-based model that was trained to predict hotspots, as provided by the Moderate Resolution Imaging Spectroradiometer (MODIS) Active Fire (AF) data product [25]. Although these studies showcased the potential of DL in various wildfire applications, the datasets used in each work remain unpublished, thus it is impossible for the community to reproduce or improve the results.

When it comes to modeling wildfires, many studies rely on satellite-derived data, such as AF products that detect thermal anomalies or burned area products that locate rapid reflectance changes. MODIS and VIIRS satellites offer such openly accessible products and are commonly used due to their high temporal resolution, offering daily global coverage. The MODIS AF product exhibits a spatial resolution of $1km \times 1km$ and has been generating data since 2002. It operates by employing thermal sensors to identify anomalous thermal signatures associated with ongoing fires [25]. The VIIRS satellite, introduced in 2012, follows a similar AF detection methodology for fires, holding an enhanced spatial resolution of $375m \times 375m$, which leads to a better response to relatively small fires and possesses an improved nighttime performance [26, 27]. Several studies have been undertaken to assess the quality of these products by comparing their outcomes against human-collected fire databases. These analyses have brought to light certain limitations associated with their utilization for the assessment of wildfires. In the United States and China, MODIS demonstrated a moderate level of concurrence with actual fire data [28, 29]. Moreover, both MODIS and VIIRS products exhibited disagreements when evaluated against real fire occurrences in Turkey, a Mediterranean-type region, with more favorable results observed for larger fires [30].

Alternative datasets sourced from MODIS include MOD14A1 [31] and MCD64A1 [32]. The former is an open-source product, containing a collection of daily fire mask composites at a spatial resolution

of $1km \times 1km$. The latter, becomes available at a spatial resolution of $500m \times 500m$, mapping the spatial extent and approximate date of biomass burning worldwide. Several validation studies have been undertaken to assess the accuracy of these datasets, revealing instances of disagreement between their outputs and reliable fire records [33–36]. Apart from MODIS products, there are several other publicly available fire datasets derived from Earth Observation satellites. These include global datasets such as FRY [37] and Fire Atlas [38] which are designated to deliver the total burned area of fires and GlobFire [39] which provides daily fire perimeters. In Europe, the European Forest Fire Information System (EFFIS) [40] provides accurate burned area estimates following a semi-supervised approach that uses different satellite sensors, estimating about 95% of the total area that burns in Europe every year [41]. The EFFIS burned area product is used in this work because of its improved accuracy in the Mediterranean region.

While these datasets provide only fire data, a comprehensive wildfire analysis and modeling needs to incorporate variables related to fire drivers, such as vegetation, weather, drought, and topography. In this direction, several studies have published datasets that integrate fire targets with variables related to fire ignition and spread. These datasets often have limitations, such as focusing on specific small-scale regions or exhibiting coarse spatial and temporal resolutions. Furthermore, they are typically designed for specific tasks tailored to a single ML objective. Kondylatos et al. [42] have published a dataset covering Greece, which is specifically designed for forecasting the next day’s wildfire danger. Though they also publish a datacube, its applications are limited by its small size, only covering a part of the eastern Mediterranean. Huot et al. [43], Singla et al. [44], Diao et al. [45] have introduced datasets designated for wildfire spread prediction in the continental US. The former relies on the MODIS AF product as the target variable, while the others utilize the VIIRS AF product. Moreover, Sayad et al. [46] have shared a dataset tailored to wildfire modeling in a small region of Canada, recording a limited number of fire events. Finally, Prapas et al. [47] presented a global dataset for seasonal fire danger forecasting, but in a coarse spatial and temporal resolution.

Comparison with Existing Datasets and Contributions. Mesogeos is a large-scale, versatile, multi-purpose dataset designed to cater to a multitude of ML tasks related to wildfire modeling. It sets itself apart from datasets [42–46], which focus solely on single ML tasks. It offers a broad scope, by encompassing a wide range of daily inputs covering many relevant fire drivers across the entire area of interest. It is provided in a cloud-optimized datacube structure, offering spatio-temporal metadata, that uniquely associates data points with specific date, longitude, and latitude values. This structure empowers researchers to select subsets in any dimension, retrieve variables, calculate new ones, or even augment the datacube with other data. Such inherent flexibility simplifies data access, facilitates the extraction of diverse ML datasets, and enables the expansion of the dataset. Thus, it allows its adaptation to a wide range of ML applications, based on individual research needs. Moreover, by leveraging the EFFIS burned area product while refining the provided ignition dates of fires through cross-comparison with the MODIS AF product (as illustrated in Section 3), we achieve a more reliable representation of burned areas and fire ignitions compared to existing datasets [25, 26, 31, 32, 37–39, 48]. Finally, it is worth noting that Mesogeos is the first dataset of this resolution and scale tailored for wildfire modeling in the Mediterranean region, a fire-prone area that has lacked dedicated datasets of this nature.

3 Mesogeos Datacube

General information. The Mesogeos dataset is structured as a spatio-temporal datacube with three dimensions: longitude, latitude, and time. The datacube encompasses 27 variables related to meteorology, vegetation, land cover, and human activity. All these data are well-known fire drivers and can be used as predictors in wildfire-related applications. Mesogeos also includes historical burned areas, ignitions, and burned area sizes as separate variables. It has $1km \times 1km \times daily$ resolution and contains the values of the variables covering the period from 2006 to 2022. It incorporates data from the wide Mediterranean area and spans a total area of $4714km \times 1753km$ and 6026 days.

Data Sources. Hantson et al. [49] study the complex interactions between the variables that control fire. They divide fire controllers into three main categories: human, weather, and vegetation. Weather conditions and vegetation play a crucial role in determining the rates of fuel drying and therefore affecting fire occurrence and spread. Topography also influences fire behavior as fire fronts travel faster uphill because of the upward convection of heat. While natural factors such as weather,

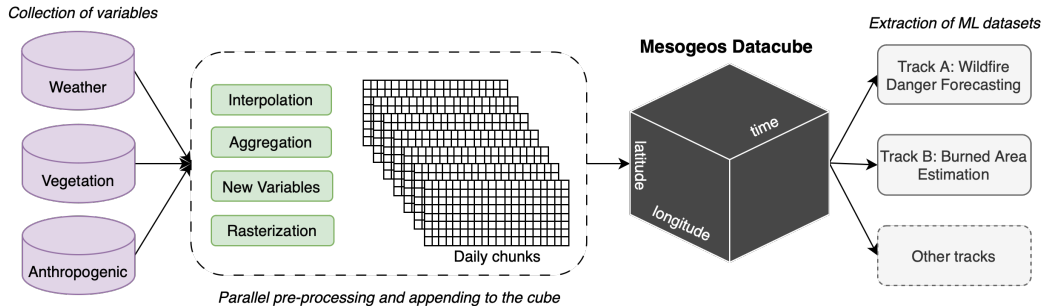


Figure 1: Pipeline of the datacube construction. The data are collected from various sources. The data inputs are pre-processed using interpolation, aggregation, calculation of new variables, and rasterization and the final daily chunks are appended in the datacube on the corresponding date. This process runs in parallel for multiple days to enhance the efficiency of the process. After the creation of the datacube, the ML datasets are extracted from it.

vegetation, and fuel load impact fire occurrence, human activity-related outcomes such as intentional or accidental fire ignition, land conversion, and population density also significantly shape fire regimes. In this study, in an attempt to cover all the factors influencing fire occurrence, we collect data sources containing information about all the aforementioned fire drivers.

The meteorological data (temperature, wind speed, wind direction, dewpoint temperature, surface pressure, relative humidity, total precipitation, surface solar radiation downwards) are collected from the ERA5-Land database [50], which contains historical hourly land weather measurements from 1950 to today. We use day’s and night’s land surface temperature [51], Normalized Difference Vegetation Index (NDVI) [52], and Leaf Area Index (LAI) [53] from MODIS and soil moisture index from the European Drought Observatory (EDO) [54]. These data are used as proxies for the vegetation status and drought. Distance from roads and population are downloaded from Worldpop [55] and are used as indicators of human activity. At the same time, topography data, i.e. elevation, slope, aspect, and curvature are gathered from the Copernicus DEM - Global Digital Elevation Model (COP-DEM) [56]. The land cover classes are collected from the Copernicus Climate Change Service [57]. The burned areas come from EFFIS. Finally, MODIS AF product [25] is used to estimate ignition cells and the ignition date.

Datacube Creation. Creating a unified dataset that stores all wildfire-related information in a standard format will later permit an easy extraction of ML-ready datasets for different tasks. Consequently, we have opted to gather and harmonize all the data into a spatio-temporal datacube format to leverage this structure’s various capabilities regarding easy access, manipulation, and extraction of data. Nevertheless, the creation of such a datacube poses significant challenges. The substantial volume of data (TBs of unprocessed data) necessitates significant downloading and storage capabilities. Additionally, the data are sourced from different providers, each with their own access points and formats (e.g. vector or raster), making data acquisition a challenging task. Furthermore, the original resolution of each variable varies, requiring harmonization to match the expected resolution of the datacube. Consequently, the construction of the datacube demands careful and efficient development for minimizing time and resource requirements.

To address the challenges above, we create the pipeline illustrated in Figure 1 for the creation of the datacube. We follow these steps: Firstly, we collect and store all relevant variables from the various input sources. Subsequently, we construct the datacube’s structure by generating a grid of dimensions $1km \times 1km \times 1 - day$ and defining daily chunks. The daily chunking applies independently to each variable, which means that values are stored in different files for each day. Finally, we append each input source to the datacube on a day-to-day basis to prevent memory errors. For each day, we perform all the necessary pre-processing steps, such as converting data into raster format, conducting temporal or spatial interpolation/aggregation, changing coordinate systems, and doing variable calculations. Then, we store the values in the chunk that refers to the corresponding date. As chunks are stored independently, we make this process totally in parallel. We use the xarray [58] python package for development. Using the default Zarr [59] compression, the datacube occupies a storage space of 648 GB, while the memory needed to load the datacube is much larger, at around 3.2 TB, assuming 32-bit

floats for the dynamic variables. The code for creating this datacube is made available and can be consulted to further enhance the existing dataset. Notably, this pipeline can be adapted with minor adjustments for generating similar datacubes applicable to various Earth science domains. For a more comprehensive understanding of the pre-processing procedures undertaken for each variable, please refer to the Supplementary Material.

Burned Areas Dataset. The burned areas dataset provided by EFFIS is an improvement over the MODIS products, offering a more reliable and credible resource. EFFIS enhances the burned area data obtained from MODIS by employing a semi-supervised processing of imagery from various satellites, i.e. Sentinel-2, and VIIRS. This process involves semi-automatic procedures aimed at enhancing the quality of fire maps [41]. Despite its improved quality, it is important to note that the dataset does not offer the same coverage of burned areas across all Mediterranean countries throughout the specified timeframe, resulting in the absence of data for certain countries in specific years. This is further analyzed in the Supplementary Material.

Ignition Date Calculation. The start dates of the fires provided by EFFIS may not always correspond to the date of ignition of the fire [41]. However, the accurate calculation of the first detection of a wildfire is extremely important in order to avoid data leakage and enhance the transparency and precision of the training process. For this, we implement a method that involves intersecting burned areas from EFFIS with AF obtained from MODIS. We use each burned area identified by EFFIS as a representative instance of a distinct fire event. We then select hotspots from a $1km$ spatial buffer zone surrounding the burned area and a temporal buffer of 7 days around the date of the ignition as provided by EFFIS. From this selection, we identify the hotspot with the oldest date within the buffer as the ignition point of the fire and its date as the ignition date of the fire. We discard the specific wildfire from the dataset if no hotspots are detected within the designated buffer zone. Notably, MODIS AF are used solely for ignition date refinement and ignition point identification and are not employed as primary anchors for fire events. Despite the improvements achieved through our approach, it is essential to recognize that some misalignments in the ignition dates may persist.

4 Machine Learning tracks

4.1 Track A: Wildfire danger forecasting

Task formulation. For a given cell and a given day t , we define fire danger as the probability of a fire occurring on the day t and becoming large, given the values of the different fire drivers x_t in the preceding days. We assume that a wildfire exceeding 30 hectares indicates high wildfire danger. To measure this danger, we treat the ignition point of the fire as a representative point of the event and the final burned area size resulting from it as an indicator of the corresponding danger level. Conversely, low wildfire danger is associated with the absence of any wildfire within a specified buffer zone surrounding a given pixel. In alignment with prior research in the field of data-driven wildfire danger forecasting [20, 42, 60], the task is formulated as a binary ML classification problem. One class signifies increased danger, while the other represents low-danger instances. However, in this work, we slightly modify the standard classification loss, involving a weighting scheme that considers the burned area size of each distinct event. This modification aims to interpret a more significant fire expansion as an indication of higher danger. The resulting softmax probabilities of the trained classification model serve as indicators of the level of fire danger.

Dataset extraction. We extract a time-series dataset, consisting of days $t - 1, t - 2, \dots, t - 30$ of the dynamic input observations and the static features repeated in time. Positive class examples consist of the ignition points of the fires that started on the day t . For the negative class examples, we select cells outside a buffer of $62km$ from any fire that started on this day to mitigate the risk of choosing cells that imply great danger but did not burn. Moreover, we follow the sampling strategy as in [42] and sample i) two times more negatives than positives, ii) the negatives following the land cover distribution of the positives.

Experimental Setup. For the experiments, we use a Long Short-Term Memory (LSTM) architecture [61] and the encoder of a Transformer model [62] as standard models for time-series data. Moreover, we employ a Gated Transformer Network (GTN) [63] that uses the attention mechanism both in time

and in variables, which could aid in modeling the complex interactions of the variables in the current task. The models are optimized using the cross-entropy (CE) loss. To let the ML models learn to assign greater danger to larger fires, we weigh the loss based on the size of the wildfire’s burned area. For this, we multiply the standard CE loss value of a given sample by the corresponding burned area size resulting from it. In practice, in order to prevent the larger fires from totally dominating the learning process, we apply a logarithmic transformation to the burned area size multiplier, in an attempt to narrow the penalization gap between small and large fires. Negative samples are assigned weights equal to the minimum burned area size among the positive samples. This ensures that the negatives, representing low-danger instances, receive adequate attention during the training process, but not more than any high-danger instance. The hyperparameters for each model are tuned separately using the validation set. The years 2006 – 2019 are used for the training set, 2020 is used as the validation set, and the years 2021 – 2022 are used as a test set. The final dataset consists of 25722 samples (8574 positives and 17148 negatives), from which 19353 (6451 positives and 12902 negatives) are in the training set, 2262 (754 positives and 1508 negatives) are in the validation set and 4107 (1369 positives and 2738 negatives) are in the test set. All the available input variables from the datacube are used in the experiments. They are all normalized before passing into the model. Precision, Recall, and Area Under Precision-Recall Curve (AUPRC) are used as metrics for the evaluation of the performance of the models. The details about the architectures of the models and the hyperparameters are provided in the Supplementary Material.

Table 1: Results of the fire danger forecasting track

Model	Precision	Recall	F1	AUPRC
LSTM [61]	0,763	0,812	0,786	0,853
Transformer [62]	0,802	0,759	0,780	0,856
GTN [63]	0,781	0,790	0,786	0,858




Figure 2: A wildfire danger map of the Mediterranean, produced by the Transformer model. The fire danger is indicated by the softmax probabilities of the trained model.

Results. The results of the models are shown in Table 1. The promising results from all the models show that they can distinguish between high and low fire danger instances. It should be noted that the optimal performance model varies for each metric, thus making it difficult to define the overall best-performing model. In addition to the quantitative metrics, we also provide a visualization map generated by the Transformer model, presented in Figure 2. The map displays the model’s softmax probabilities, indicating fire danger levels. This provides an example of how a daily fire danger map can be generated for operational scenarios in the Mediterranean region. The high spatial resolution of the dataset employed in this study could enable improved fire management strategies and demonstrate potential operational benefits. In this example, concerning fire danger for a day in the summer of 2022 notable variations in fire danger are observed across different areas of the Mediterranean, as well as within individual regions of each country.

4.2 Track B: Final Burned Area Prediction

Task formulation. This track focuses on predicting the likely extent of a wildfire’s final burned area, given the ignition point and a set of variables available at the time of ignition, representing the fire drivers inside a neighborhood around the ignition point. These fire drivers encompass factors that influence fire behavior and spread. Thus, the objective of the ML task is to estimate the likelihood of the neighboring pixels surrounding the ignition point, to be eventually contained within the final burned area of the wildfire. The final burned area prediction is treated as a segmentation task, with two classes, indicating whether a pixel will experience burning or remain unaffected by the ignited fire. The resulting softmax probabilities of the trained model serve as indicators of the likelihood of a pixel being burnt.

Dataset extraction. For every fire event, we extract $64km \times 64km$ patches, that are centered around the fire’s ignition point, usually containing the whole burned area of a given fire event. The extracted samples include all the values of the variables of the datacube for the date of the fire’s occurrence.

Experimental Setup. The 64×64 patches are randomly cropped to 32×32 during the training process. This approach ensures that the ignition point remains within the patch while preventing the model from generating a bias towards fire expansion solely from the central cell. As this is a segmentation task, we use the U-Net architecture [64] with an EfficientNet-B1 [65] encoder. Different input variables are stacked as separate channels. The cross-entropy loss is used to train the models’ parameters. To define a baseline for the task, we train an additional model that uses as input only the ignition points. We do a temporal split to avoid leaking data from fire events happening close in time, using 2006 – 2019 for training (12550 samples), 2020 for validation (1781 samples), and 2021 – 2022 for testing (3527 samples). The loss in the validation is used for early stopping. Input variables are scaled with the minimum, and maximum values in the range $[0, 1]$ before being served as inputs to the model. As evaluation metrics, we report the CE loss and the AUPRC. The exact architecture and the values of the hyperparameters are provided in the Supplementary Material.

Table 2: Results of the final burned area prediction track

Model	CE Loss	AUPRC
U-Net (only ignitions)	0.0177	0.394
U-Net (all variables)	0.0166	0.418




Figure 3: Two examples of predictions (softmax values for the positive class) from the U-Net using *all variables*, and the U-Net using *only ignition* points. The predictions are presented together with binary maps representing the initial *ignition points* and the *ground truth* burned areas.

Results. The experimental results are presented in Table 2, which shows that the model incorporating all variables slightly outperforms the baseline model relying only on the ignition as input. In Figure 3 we see two examples from the models’ predictions in the test set. It is notable that models predict higher spread danger around the ignition point, with the model that uses input from all variables showing some enhanced skill compared to the no-skill baseline. To improve on the task, it might be important to include variables before and after the fire starts and not only from the date of ignition. Notably, Coffield et al. [66] find that the performance can be improved using weather data forecasts for 1 – 5 days after the start of a fire. The dataset extraction script that accompanies the track includes the capabilities to extract datasets that include time series of arbitrary length before and after the start of each fire.

5 Potential other tracks to explore

Mesogeos incorporates burned areas and their sizes, as well as ignition points as predictands, and thus can be used for various ML applications. Thus, users can address additional tracks beyond those presented in the current study. The following ideas serve as suggestions for potential paths of research.

Fire size prediction. This application involves predicting the final size of fires rather than the presence of burned areas, which was addressed in this work. This can be framed as either a regression or a multi-class classification task. As a regression task, the target variable would be the burned area in hectares and therefore predictive models could be developed to estimate the final size of fires. As a classification task, a model could be employed to categorize fires into small, medium, or large. Notably, the dataset initially created for the wildfire danger forecasting track can be used as is for these two tasks.

Extreme events forecasting. Forecasting extreme events holds significant importance for effective fire management strategies. Concerning the Mediterranean, most of the damage caused by wildfires is the result of only a few large fires [67]. Therefore, anticipating these wildfires is crucial for ecosystem preservation and optimization of fire management. Mesogeos includes extreme events, such as the 5 massive fires in Greece in the summer of 2021, that burned nearly 94,000 hectares [68], the extreme events in Portugal in October 2017, with 5 events burning more than 18,000 hectares each [69] and the fire in Valencia, Spain in 2012 that burned over 50,000 hectares [70]. Extracting a dedicated dataset for this track offers an opportunity to develop models explicitly predicting these extreme events. As these extreme events are rare, one potential approach to tackle this task would be treating it as an outlier detection problem.

Wildfire susceptibility mapping. Wildfire susceptibility is defined as the static probability of wildfires in a certain area, depending on the characteristics of the terrain and prevailing meteorological conditions [71]. This task can be framed as a binary classification task. In this context, the dataset extraction process would involve identifying positive samples as the number of all pixels affected by fire over multiple years. Negative samples would be obtained from pixels that have never experienced burning, i.e. not belonging to any burned area of the dataset. An ML model could then be employed to distinguish between the fire-susceptible and non-susceptible samples.

Self-supervised learning. The vast amount of data in the datacube remain untapped when extracting task-specific datasets. Self-supervised learning (SSL) [72] offers a promising approach to take advantage of the full capabilities of these data. SSL allows for acquiring a representation that can be utilized across various downstream tasks, including those mentioned earlier. Concerning the SSL track, extracting specific datasets is unnecessary, as the training samples can be directly extracted from the datacube during the data loading process. Careful engineering is essential when selecting and extracting samples to minimize the time required for the model to be trained.

Modeling at different spatio-temporal scales. Mesogeos has a resolution of $1km \times 1km \times daily$, enabling the examination of problems at that specific temporal and spatial scale. However, the flexibility of the datacube format allows for resampling in various temporal and spatial dimensions through appropriate aggregations. This would enable the treatment of other tasks in coarser temporal or spatial scales, like seasonal or sub-seasonal fire modeling.

Beyond traditional ML. When doing ML-based wildfire modeling for decision support it is many times important to dive deeper into understanding the underlying processes that drive the models’ predictions. In that respect, Mesogeos can foster the development of explainable AI techniques [73] toward a better understanding of models and subsequently the interactions of the fire drivers that result in wildfires. Additionally, causal inference methods [74] could be used to assess the effects of human controls, such as agricultural practices or land use on wildfire regimes. Moreover, considering the stochastic nature of fire processes, noise commonly appears on the labels. Methods that take into account the noisy labels [75] and especially those estimating the inherent aleatoric uncertainty [76], could enhance the reliability of the models and support the decision-making. When existing layers in the datacube are not enough, the datacube can be easily enhanced with extra information such as socio-economic factors, settlement, and infrastructure.

6 Limitations

Despite the advantages of Mesogeos, we would like to acknowledge certain limitations. Firstly, the datasets inherit inaccuracies of the original data sources. Factors such as the satellites’ spatial resolution and missing data resulting from cloud cover can influence the precise determination of the fires’ location and size. Additionally, the acquisition of the fire ignition date is challenging and prone to deviations, as discussed in Section 3. Another limitation arises from the types of fires included in the burned areas’ products. It is possible that there are fires resulting from prescribed burning or agricultural burning [41], which cannot be modeled using just the variables within Mesogeos. Moreover, while the daily temporal and $1km \times 1km$ spatial resolution of the datacube is appropriate for the applications suggested in this research, it cannot be used to address other, nowcasting-type problems related to other cycles of fire management such as fire spread, fire detection, or fire recovery related applications. For the target variable, Mesogeos considers the highly reliable final burned areas from EFFIS, but ignores intermediate temporal information of the wildfire evolution, compared to work explicitly targeting wildfire spread [20, 45, 44]. Furthermore, the dataset lacks information regarding fire suppression efforts. The interventions of firefighters and responders have the potential to influence fire dynamics both at the time of ignition, achieved through water application to weaken fire spread, as well as during winter months by means of fuel cleaning or controlled burning. It should be acknowledged that the absence of such data could influence the modeling of some of the tracks. For example, when training the model to predict fire danger, it remains unknown whether the fire would have grown larger (indicating higher danger) in the absence of any wildfire suppression measures, or the opposite. Finally, it is important to highlight that while ML can assist in wildfire modeling, an operational application in wildfire management necessitates a thorough evaluation across fire seasons and against operational baselines, including domain experts and wildfire responders in the process.

7 Availability and Maintenance

The Mesogeos datacube and the datasets utilized in this study are made publicly available. The project’s website <https://orion-ai-lab.github.io/mesogeos/> will hold updated links to the data and code repository, as well as a leaderboard for the ML tracks. The repository contains code for generating Mesogeos, extracting datasets for the tracks, and running the models, enabling the reproduction of the results presented in this work. We encourage the community to further contribute with more ML tracks and models and advance data-driven wildfire modeling using the Mesogeos datacube.

8 Conclusion

In conclusion, this work introduces Mesogeos, a valuable resource for data-driven wildfire modeling. By leveraging the structure of a datacube and incorporating variables that represent various fire drivers and historical wildfires, Mesogeos facilitates the extraction of diverse datasets, empowering researchers to model various fire-related tasks. In this work, we demonstrate two tracks focusing on fire danger forecasting and burned area prediction to showcase the effectiveness and potential of the dataset. Lastly, we present several alternative tracks that address a wide range of challenges and tasks associated with anticipating and understanding wildfires, thereby paving the way for new avenues of research and advancement in wildfire modeling.

Acknowledgments and Disclosure of Funding

This work has received funding from the European Union’s Horizon 2020 Research and Innovation Projects DeepCube and TREEADS, under Grant Agreement Numbers 101004188 and 101036926 respectively.

References

- [1] Juli G. Pausas, Jon E. Keeley, and Dylan W. Schwilk. Flammability as an ecological and evolutionary driver. *Journal of Ecology*, 105(2):289–297, 2017. ISSN 1365-2745. doi: 10.1111/1365-2745.12691. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/1365-2745.12691>. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/1365-2745.12691>.
- [2] William J. Bond and Jon E. Keeley. Fire as a global ‘herbivore’: the ecology and evolution of flammable ecosystems. *Trends in Ecology & Evolution*, 20(7):387–394, July 2005. ISSN 0169-5347. doi: 10.1016/j.tree.2005.04.025. URL <https://www.sciencedirect.com/science/article/pii/S0169534705001321>.
- [3] Andrew C. Scott and Ian J. Glasspool. The diversification of paleozoic fire systems and fluctuations in atmospheric oxygen concentration. *Proceedings of the National Academy of Sciences*, 103(29):10861–10865, 2006. doi: 10.1073/pnas.0604090103. URL <https://www.pnas.org/doi/abs/10.1073/pnas.0604090103>.
- [4] David M. J. S. Bowman, Jennifer K. Balch, Paulo Artaxo, William J. Bond, Jean M. Carlson, Mark A. Cochrane, Carla M. D’Antonio, Ruth S. DeFries, John C. Doyle, Sandy P. Harrison, Fay H. Johnston, Jon E. Keeley, Meg A. Krawchuk, Christian A. Kull, J. Brad Marston, Max A. Moritz, I. Colin Prentice, Christopher I. Roos, Andrew C. Scott, Thomas W. Swetnam, Guido R. van der Werf, and Stephen J. Pyne. Fire in the earth system. *Science*, 324(5926):481–484, 2009. doi: 10.1126/science.1163886. URL <https://www.science.org/doi/abs/10.1126/science.1163886>.
- [5] M. Lucrecia Pettinari and Emilio Chuvieco. Fire Danger Observed from Space. *Surveys in Geophysics*, 41(6):1437–1459, November 2020. ISSN 1573-0956. doi: 10.1007/s10712-020-09610-8. URL <https://doi.org/10.1007/s10712-020-09610-8>.
- [6] Juli G Pausas and Jon E Keeley. Wildfires and global change. *Frontiers in Ecology and the Environment*, 19(7):387–395, 2021. ISSN 1540-9309. doi: 10.1002/fee.2359. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/fee.2359>. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/fee.2359>.
- [7] Francisco Moreira, Davide Ascoli, Hugh Safford, Mark A Adams, José M Moreno, José M C Pereira, Filipe X Catry, Juan Armesto, William Bond, Mauro E González, Thomas Curt, Nikos Koutsias, Lachlan McCaw, Owen Price, Juli G Pausas, Eric Rigolot, Scott Stephens, Cagatay Tavsanoglu, V Ramon Vallejo, Brian W Van Wilgen, Gavriil Xanthopoulos, and Paulo M Fernandes. Wildfire management in Mediterranean-type regions: paradigm change needed. *Environmental Research Letters*, 15(1):011001, January 2020. ISSN 1748-9326. doi: 10.1088/1748-9326/ab541e. URL <https://iopscience.iop.org/article/10.1088/1748-9326/ab541e>.
- [8] Enric Batllori, Marc-André Parisien, Meg A. Krawchuk, and Max A. Moritz. Climate change-induced shifts in fire for Mediterranean ecosystems. *Global Ecology and Biogeography*, 22(10):1118–1129, 2013. ISSN 1466-8238. doi: 10.1111/geb.12065. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/geb.12065>. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/geb.12065>.
- [9] Kirk R. Klausmeyer and M. Rebecca Shaw. Climate Change, Habitat Loss, Protected Areas and the Climate Adaptation Potential of Species in Mediterranean Ecosystems Worldwide. *PLoS ONE*, 4(7):e6392, July 2009. ISSN 1932-6203. doi: 10.1371/journal.pone.0006392. URL <https://dx.plos.org/10.1371/journal.pone.0006392>.

- [10] Sally Archibald, David P. Roy, Brian W. van WILGEN, and Robert J. Scholes. What limits fire? An examination of drivers of burnt area in Southern Africa. *Global Change Biology*, 15(3):613–630, March 2009. ISSN 13541013, 13652486. doi: 10.1111/j.1365-2486.2008.01754.x. URL <https://onlinelibrary.wiley.com/doi/10.1111/j.1365-2486.2008.01754.x>.
- [11] CE Van Wagner et al. *Structure of the Canadian forest fire weather index*, volume 1333. Environment Canada, Forestry Service Ontario, 1974.
- [12] Patricia L. Andrews. BEHAVE: fire behavior prediction and fuel modeling system-BURN subsystem, part I. Technical report, 1986. URL <https://doi.org/10.2737%2Fint-gtr-194>.
- [13] Mark A. Finney. FARSITE: Fire area simulator-model development and evaluation. Technical report, 1998. URL <https://doi.org/10.2737%2Frmrs-rp-4>.
- [14] Ioannis Prapas, Spyros Kondylatos, Ioannis Papoutsis, Gustau Camps-Valls, Michele Ronco, Miguel-Ángel Fernández-Torres, Maria Piles Guillem, and Nuno Carvalhais. Deep Learning Methods for Daily Wildfire Danger Forecasting, November 2021. URL <http://arxiv.org/abs/2111.02736>. arXiv:2111.02736 [cs].
- [15] Piyush Jain, Sean C.P. Coogan, Sriram Ganapathi Subramanian, Mark Crowley, Steve Taylor, and Mike D. Flannigan. A review of machine learning applications in wildfire science and management. *Environmental Reviews*, 28(4):478–505, 2020. doi: 10.1139/er-2020-0019. URL <https://doi.org/10.1139/er-2020-0019>.
- [16] Markus Reichstein, Gustau Camps-Valls, Bjorn Stevens, Martin Jung, Joachim Denzler, and Nuno Carvalhais. Deep learning and process understanding for data-driven Earth system science. *Nature*, 566(7743):195–204, 2019. Publisher: Nature Publishing Group.
- [17] Gustau Camps-Valls, Devis Tuia, Xiao Xiang Zhu, and Markus Reichstein. *Deep learning for the Earth Sciences: A comprehensive approach to remote sensing, climate science and geosciences*. John Wiley & Sons, 2021. URL <https://github.com/DL4ES>.
- [18] Miguel D. Mahecha, Fabian Gans, Gunnar Brandt, Rune Christiansen, Sarah E. Cornell, Normann Fomferra, Guido Kraemer, Jonas Peters, Paul Bodesheim, Gustau Camps-Valls, Jonathan F. Donges, Wouter Dorigo, Lina M. Estupinan-Suarez, Victor H. Gutierrez-Velez, Martin Gutwin, Martin Jung, Maria C. Londoño, Diego G. Miralles, Phillip Papastefanou, and Markus Reichstein. Earth system data cubes unravel global multivariate dynamics. *Earth System Dynamics*, 11(1):201–234, February 2020. ISSN 2190-4987. doi: 10.5194/esd-11-201-2020. URL <https://esd.copernicus.org/articles/11/201/2020/>.
- [19] Spyros Kondylatos, Ioannis Prapas, Gustau Camps-Valls, and Ioannis Papoutsis. mesogeos: A multi-purpose dataset for data-driven wildfire modeling in the Mediterranean, March 2023. URL <https://doi.org/10.5281/zenodo.7473331>.
- [20] Fantine Huot, R. Lily Hu, Matthias Ihme, Qing Wang, John Burge, Tianjian Lu, Jason Hickey, Yi-Fan Chen, and John Anderson. Deep Learning Models for Predicting Wildfires from Historical Remote-Sensing Data. *arXiv:2010.07445 [cs]*, November 2020. URL <http://arxiv.org/abs/2010.07445>. arXiv: 2010.07445.
- [21] David Radke, Anna Hessler, and Dan Ellsworth. FireCast: Leveraging Deep Learning to Predict Wildfire Spread. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, pages 4575–4581, Macao, China, August 2019. International Joint Conferences on Artificial Intelligence Organization. ISBN 978-0-9992411-4-1. doi: 10.24963/ijcai.2019/636. URL <https://www.ijcai.org/proceedings/2019/636>.
- [22] Jonathan L. Hodges and Brian Y. Lattimer. Wildland Fire Spread Modeling Using Convolutional Neural Networks. *Fire Technology*, 55(6):2115–2142, November 2019. ISSN 1572-8099. doi: 10.1007/s10694-019-00846-4. URL <https://doi.org/10.1007/s10694-019-00846-4>.
- [23] John Burge, Matthew Bonanni, Matthias Ihme, and Lily Hu. Convolutional LSTM Neural Networks for Modeling Wildland Fire Dynamics, April 2021. URL <http://arxiv.org/abs/2012.06679>. arXiv:2012.06679 [cs].

- [24] Rui Ba, Chen Chen, Jing Yuan, Weiguo Song, and Siuming Lo. SmokeNet: Satellite Smoke Scene Detection Using Convolutional Neural Network with Spatial and Channel-Wise Attention. *Remote Sensing*, 11(14):1702, January 2019. doi: 10.3390/rs11141702. URL <https://www.mdpi.com/2072-4292/11/14/1702>. Number: 14 Publisher: Multidisciplinary Digital Publishing Institute.
- [25] Louis Giglio, Wilfrid Schroeder, and Christopher O. Justice. The collection 6 MODIS active fire detection algorithm and fire products. *Remote Sensing of Environment*, 178:31–41, June 2016. ISSN 0034-4257. doi: 10.1016/j.rse.2016.02.054. URL <https://www.sciencedirect.com/science/article/pii/S0034425716300827>.
- [26] Wilfrid Schroeder, Patricia Oliva, Louis Giglio, and Ivan A. Csiszar. The New VIIRS 375 m active fire detection data product: Algorithm description and initial assessment. *Remote Sensing of Environment*, 143:85–96, March 2014. ISSN 00344257. doi: 10.1016/j.rse.2013.12.008. URL <https://linkinghub.elsevier.com/retrieve/pii/S0034425713004483>.
- [27] Ivan Csiszar, Wilfrid Schroeder, Louis Giglio, Evan Ellicott, Krishna P. Vadrevu, Christopher O. Justice, and Brad Wind. Active fires from the Suomi NPP Visible Infrared Imaging Radiometer Suite: Product status and first evaluation results. *Journal of Geophysical Research: Atmospheres*, 119(2):803–816, 2014. ISSN 2169-8996. doi: 10.1002/2013JD020453. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/2013JD020453>. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/2013JD020453>.
- [28] Emily J. Fusco, John T. Finn, John T. Abatzoglou, Jennifer K. Balch, Sepideh Dadashi, and Bethany A. Bradley. Detection rates and biases of fire observations from MODIS and agency reports in the conterminous United States. *Remote Sensing of Environment*, 220:30–40, January 2019. ISSN 0034-4257. doi: 10.1016/j.rse.2018.10.028. URL <https://www.sciencedirect.com/science/article/pii/S0034425718304826>.
- [29] Davide Fornacca, Guopeng Ren, and Wen Xiao. Performance of Three MODIS Fire Products (MCD45A1, MCD64A1, MCD14ML), and ESA Fire_cci in a Mountainous Area of Northwest Yunnan, China, Characterized by Frequent Small Fires. *Remote Sensing*, 9(11):1131, November 2017. ISSN 2072-4292. doi: 10.3390/rs9111131. URL <https://www.mdpi.com/2072-4292/9/11/1131>. Number: 11 Publisher: Multidisciplinary Digital Publishing Institute.
- [30] Ka 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*, 15(2): 85–94, April 2022. ISSN 19717458. doi: 10.3832/ifor3754-015. URL <https://iforest.sisef.org/?doi=ifor3754-015>.
- [31] Louis Giglio and Christopher Justice. MOD14A1 MODIS/Terra Thermal Anomalies/Fire Daily L3 Global 1km SIN Grid V006, 2015. URL <https://lpdaac.usgs.gov/products/mod14a1v006/>.
- [32] Louis Giglio, Luigi Boschetti, David P. Roy, Michael L. Humber, and Christopher O. Justice. The Collection 6 MODIS burned area mapping algorithm and product. *Remote Sensing of Environment*, 217:72–85, November 2018. ISSN 0034-4257. doi: 10.1016/j.rse.2018.08.005. URL <https://www.sciencedirect.com/science/article/pii/S0034425718303705>.
- [33] Mhosisi Masocha, Timothy Dube, Ndumezulu T. Mpofo, and Silvester Chimunhu. Accuracy assessment of MODIS active fire products in southern African savannah woodlands. *African Journal of Ecology*, 56(3):563–571, 2018. ISSN 1365-2028. doi: 10.1111/aje.12494. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/aje.12494>. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/aje.12494>.
- [34] Joanne V. Hall, Fernanda Argueta, and Louis Giglio. Validation of MCD64A1 and FireCCI51 cropland burned area mapping in Ukraine. *International Journal of Applied Earth Observation and Geoinformation*, 102:102443, October 2021. ISSN 1569-8432. doi: 10.1016/j.jag.2021.102443. URL <https://www.sciencedirect.com/science/article/pii/S0303243421001501>.

- [35] Julia A. Rodrigues, Renata Libonati, Allan A. Pereira, Joana M. P. Nogueira, Filipe L. M. Santos, Leonardo F. Peres, Ananda Santa Rosa, Wilfrid Schroeder, José M. C. Pereira, Louis Giglio, Isabel F. Trigo, and Alberto W. Setzer. How well do global burned area products represent fire patterns in the Brazilian Savannas biome? An accuracy assessment of the MCD64 collections. *International Journal of Applied Earth Observation and Geoinformation*, 78: 318–331, June 2019. ISSN 1569-8432. doi: 10.1016/j.jag.2019.02.010. URL <https://www.sciencedirect.com/science/article/pii/S0303243419300248>.
- [36] Chunmao Zhu, Hideki Kobayashi, Yugo Kanaya, and Masahiko Saito. Size-dependent validation of MODIS MCD64A1 burned area over six vegetation types in boreal Eurasia: Large underestimation in croplands. *Scientific Reports*, 7(1):4181, July 2017. ISSN 2045-2322. doi: 10.1038/s41598-017-03739-0. URL <https://www.nature.com/articles/s41598-017-03739-0>.
- [37] Pierre Laurent, Florent Mouillot, Chao Yue, Philippe Ciais, M. Vanesa Moreno, and Joana M. P. Nogueira. FRY, a global database of fire patch functional traits derived from space-borne burned area products. *Scientific Data*, 5(1):180132, July 2018. ISSN 2052-4463. doi: 10.1038/sdata.2018.132. URL <https://www.nature.com/articles/sdata2018132>. Number: 1 Publisher: Nature Publishing Group.
- [38] Niels Andela, Douglas C. Morton, Louis Giglio, Ronan Paugam, Yang Chen, Stijn Hantson, Guido R. van der Werf, and James T. Randerson. The Global Fire Atlas of individual fire size, duration, speed and direction. *Earth System Science Data*, 11(2):529–552, April 2019. ISSN 1866-3508. doi: 10.5194/essd-11-529-2019. URL <https://essd.copernicus.org/articles/11/529/2019/>. Publisher: Copernicus GmbH.
- [39] Tomàs Artés, Duarte Oom, Daniele de Rigo, Tracy Houston Durrant, Pieralberto Maianti, Giorgio Libertà, and Jesús San-Miguel-Ayanz. A global wildfire dataset for the analysis of fire regimes and fire behaviour. *Scientific Data*, 6(1):296, November 2019. ISSN 2052-4463. doi: 10.1038/s41597-019-0312-2. URL <https://www.nature.com/articles/s41597-019-0312-2>. Number: 1 Publisher: Nature Publishing Group.
- [40] Jesus San-Miguel-Ayanz, PM Barbosa, G Schmuck, G Libertà, and J Meyer-Roux. The european forest fire information system (effis). In *Proceedings of 6th AGILE conference on Geographic Information Science*, pages 24–26, 2003.
- [41] European Forest Fire Information System. Effis - rapid damage assessment. <https://effis.jrc.ec.europa.eu/about-effis/technical-background/rapid-damage-assessment>, 2023. [Online; accessed 7-June-2023].
- [42] Spyros Kondylatos, Ioannis Prapas, Michele Ronco, Ioannis Papoutsis, Gustau Camps-Valls, María Piles, Miguel-Ángel Fernández-Torres, and Nuno Carvalhais. Wildfire Danger Prediction and Understanding With Deep Learning. *Geophysical Research Letters*, 49(17):e2022GL099368, 2022. ISSN 1944-8007. doi: 10.1029/2022GL099368. URL <https://onlinelibrary.wiley.com/doi/abs/10.1029/2022GL099368>. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2022GL099368>.
- [43] Fantine Huot, R. Lily Hu, Nita Goyal, Tharun Sankar, Matthias Ihme, and Yi-Fan Chen. Next Day Wildfire Spread: A Machine Learning Dataset to Predict Wildfire Spreading From Remote-Sensing Data. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–13, 2022. ISSN 1558-0644. doi: 10.1109/TGRS.2022.3192974. Conference Name: IEEE Transactions on Geoscience and Remote Sensing.
- [44] Samridhhi Singla, Ayan Mukhopadhyay, Michael Wilbur, Tina Diao, Vinayak Gajjewar, Ahmed Eldawy, Mykel Kochenderfer, Ross Shachter, and Abhishek Dubey. WildfireDB: An Open-Source Dataset Connecting Wildfire Spread with Relevant Determinants, November 2021. URL <https://zenodo.org/record/5636429>.
- [45] Tina Diao, Samridhhi Singla, Ayan Mukhopadhyay, Ahmed Eldawy, Ross Shachter, and Mykel Kochenderfer. Uncertainty Aware Wildfire Management, October 2020. URL <http://arxiv.org/abs/2010.07915>. arXiv:2010.07915 [cs, stat].

- [46] Younes Oulad Sayad, Hajar Mousannif, and Hassan Al Moatassime. Predictive modeling of wildfires: A new dataset and machine learning approach. *Fire Safety Journal*, 104:130–146, March 2019. ISSN 0379-7112. doi: 10.1016/j.firesaf.2019.01.006. URL <https://www.sciencedirect.com/science/article/pii/S0379711218303941>.
- [47] Ioannis Prapas, Akanksha Ahuja, Spyros Kondylatos, Ilektra Karasante, Eleanna Panagiotou, Lazaro Alonso, Charalampos Davalas, Dimitrios Michail, Nuno Carvalhais, and Ioannis Papoutsis. Deep Learning for Global Wildfire Forecasting, November 2022. URL <http://arxiv.org/abs/2211.00534>. arXiv:2211.00534 [cs].
- [48] Tom Toulouse, Lucile Rossi, Antoine Campana, Turgay Celik, and Moulay A. Akhloufi. Computer vision for wildfire research: An evolving image dataset for processing and analysis. *Fire Safety Journal*, 92:188–194, September 2017. ISSN 0379-7112. doi: 10.1016/j.firesaf.2017.06.012. URL <https://www.sciencedirect.com/science/article/pii/S0379711217302114>.
- [49] Stijn Hantson, Almut Arneth, Sandy P. Harrison, Doug I. Kelley, I. Colin Prentice, Sam S. Rabin, Sally Archibald, Florent Mouillot, Steve R. Arnold, and Paulo Artaxo. The status and challenge of global fire modelling. *Biogeosciences*, 13(11):3359–3375, 2016. Publisher: Copernicus Publications.
- [50] Joaquín Muñoz-Sabater, Emanuel Dutra, Anna Agustí-Panareda, Clément Albergel, Gabriele Arduini, Gianpaolo Balsamo, Souhail Boussetta, Margarita Choulga, Shaun Harrigan, Hans Hersbach, et al. Era5-land: A state-of-the-art global reanalysis dataset for land applications. *Earth System Science Data Discussions*, pages 1–50, 2021.
- [51] Z. Wan, S. Hook, and G. Hulley. MOD11a1 MODIS/terra land surface temperature and the emissivity daily l3 global 1km SIN grid, 2018.
- [52] K. Didan. Mod13a2 modis/terra vegetation indices 16-day l3 global 1km sin grid v006 [data set]. *NASA EOSDIS LP DAAC* <https://doi.org/10.5067/MODIS/MOD13A2>, 6, 2015.
- [53] R. B. Myneni, N. V. Shabanov, Y. Knyazikhin, W. Yang, H. Dong, and B. Tan. MOD15a2: Global LAI and FPAR. In *AGU Fall Meeting Abstracts*, volume 2002, pages B61B–0719, 2002.
- [54] Carmelo Cammalleri, Jürgen V. Vogt, Bernard Bisselink, and Ad de Roo. Comparing soil moisture anomalies from multiple independent sources over different regions across the globe. *Hydrology and Earth System Sciences*, 21(12):6329–6343, December 2017. ISSN 1027-5606. doi: 10.5194/hess-21-6329-2017. URL <https://hess.copernicus.org/articles/21/6329/2017/hess-21-6329-2017.html>. Publisher: Copernicus GmbH.
- [55] Andrew J Tatem. Worldpop, open data for spatial demography. *Scientific data*, 4(1):1–4, 2017.
- [56] European Space Agency and Airbus. Copernicus DEM, 2022. Title of the publication associated with this dataset: Copernicus DEM.
- [57] Copernicus Climate Change Service. Land cover classification gridded maps from 1992 to present derived from satellite observations, 2019.
- [58] S. Hoyer and J. Hamman. xarray: N-D labeled arrays and datasets in Python. *Journal of Open Research Software*, 5(1), 2017. doi: 10.5334/jors.148. URL <https://doi.org/10.5334/jors.148>.
- [59] Alistair Miles, John Kirkham, Martin Durant, James Bourbeau, Tarik Onalan, Joe Hamman, Zain Patel, shikharsg, Matthew Rocklin, Raphael Dussin, Vincent Schut, Elliott Sales de Andrade, Ryan Abernathey, Charles Noyes, sbalmer, Pyup io Bot, Tommy Tran, Stephan Saalfeld, Justin Swaney, Josh Moore, Joe Jevnik, Jerome Kelleher, Jan Funke, George Sakkis, Chris Barnes, and Anderson Banihirwe. zarr-developers/zarr-python: v2.4.0, 2020.
- [60] Guoli Zhang, Ming Wang, and Kai Liu. Forest Fire Susceptibility Modeling Using a Convolutional Neural Network for Yunnan Province of China. *International Journal of Disaster Risk Science*, 10(3):386–403, September 2019. ISSN 2095-0055, 2192-6395. doi: 10.1007/s13753-019-00233-1. URL <http://link.springer.com/10.1007/s13753-019-00233-1>.

- [61] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8): 1735–1780, 1997.
- [62] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention Is All You Need, December 2017. URL <http://arxiv.org/abs/1706.03762>. arXiv:1706.03762 [cs].
- [63] Minghao Liu, Shengqi Ren, Siyuan Ma, Jiahui Jiao, Yizhou Chen, Zhiguang Wang, and Wei Song. Gated Transformer Networks for Multivariate Time Series Classification, March 2021. URL <http://arxiv.org/abs/2103.14438>. arXiv:2103.14438 [cs].
- [64] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation, May 2015. URL <http://arxiv.org/abs/1505.04597>. arXiv:1505.04597 [cs].
- [65] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019.
- [66] Shane R. Coffield, Casey A. Graff, Yang Chen, Padhraic Smyth, Efi Foufoula-Georgiou, and James T. Randerson. Machine learning to predict final fire size at the time of ignition. *International Journal of Wildland Fire*, 28(11):861, 2019. ISSN 1049-8001. doi: 10.1071/WF19023. URL <http://www.publish.csiro.au/?paper=WF19023>.
- [67] Jesús San-Miguel-Ayanz, Jose Manuel Moreno, and Andrea Camia. Analysis of large fires in European Mediterranean landscapes: Lessons learned and perspectives. *Forest Ecology and Management*, 294:11–22, April 2013. ISSN 0378-1127. doi: 10.1016/j.foreco.2012.10.050. URL <https://www.sciencedirect.com/science/article/pii/S0378112712006561>.
- [68] Theodore M. Giannaros, Georgios Papavasileiou, Konstantinos Lagouvardos, Vassiliki Kotroni, Stavros Dafis, Athanasios Karagiannidis, and Eleni Dragozi. Meteorological Analysis of the 2021 Extreme Wildfires in Greece: Lessons Learned and Implications for Early Warning of the Potential for Pyroconvection. *Atmosphere*, 13(3):475, March 2022. ISSN 2073-4433. doi: 10.3390/atmos13030475. URL <https://www.mdpi.com/2073-4433/13/3/475>. Number: 3 Publisher: Multidisciplinary Digital Publishing Institute.
- [69] Marc Castellnou, Nuno Guiomar, Francisco Rego, and Paulo Fernandes. *Fire growth patterns in the 2017 mega fire episode of October 15, central Portugal*. November 2018. ISBN 978-989-26-1650-6. doi: 10.14195/978-989-26-16-506_48.
- [70] Ana M. Vicedo-Cabrera, Ana Esplugues, Carmen Iníguez, Marisa Estarlich, and Ferran Ballester. Health effects of the 2012 Valencia (Spain) wildfires on children in a cohort study. *Environmental Geochemistry and Health*, 38(3):703–712, June 2016. ISSN 1573-2983. doi: 10.1007/s10653-015-9753-5. URL <https://doi.org/10.1007/s10653-015-9753-5>.
- [71] Andrea Trucchia, Giorgio Meschi, Paolo Fiorucci, Antonello Provenzale, Marj Tonini, and Umberto Pernice. Wildfire hazard mapping in the eastern Mediterranean landscape. *International Journal of Wildland Fire*, 32(3):417–434, March 2023. ISSN 1049-8001, 1448-5516. doi: 10.1071/WF22138. URL <https://www.publish.csiro.au/WF/WF22138>.
- [72] Randall Balestriero, Mark Ibrahim, Vlad Sobal, Ari Morcos, Shashank Shekhar, Tom Goldstein, Florian Bordes, Adrien Bardes, Gregoire Mialon, Yuandong Tian, Avi Schwarzschild, Andrew Gordon Wilson, Jonas Geiping, Quentin Garrido, Pierre Fernandez, Amir Bar, Hamed Pirsiavash, Yann LeCun, and Micah Goldblum. A Cookbook of Self-Supervised Learning, April 2023. URL <http://arxiv.org/abs/2304.12210>. arXiv:2304.12210 [cs].
- [73] Giulia Vilone and Luca Longo. Explainable Artificial Intelligence: a Systematic Review. *arXiv:2006.00093 [cs]*, October 2020. URL <http://arxiv.org/abs/2006.00093>. arXiv: 2006.00093.
- [74] Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. Towards Causal Representation Learning. *arXiv:2102.11107 [cs]*, February 2021. URL <http://arxiv.org/abs/2102.11107>. arXiv: 2102.11107.

- [75] Filipe R. Cordeiro and Gustavo Carneiro. A Survey on Deep Learning with Noisy Labels: How to train your model when you cannot trust on the annotations? In *2020 33rd SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI)*, pages 9–16, Recife/Porto de Galinhas, Brazil, November 2020. IEEE. ISBN 978-1-72819-274-1. doi: 10.1109/SIBGRAPI51738.2020.00010. URL <https://ieeexplore.ieee.org/document/9266015/>.
- [76] Mark Collier, Basil Mustafa, Efi Kokiopoulou, Rodolphe Jenatton, and Jesse Berent. A Simple Probabilistic Method for Deep Classification under Input-Dependent Label Noise, November 2020. URL <http://arxiv.org/abs/2003.06778>. arXiv:2003.06778 [cs, stat].
- [77] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- [78] William Falcon et al. Pytorch lightning. *GitHub. Note:* <https://github.com/PyTorchLightning/pytorch-lightning>, 3, 2019.
- [79] Pavel Iakubovskii. Segmentation models pytorch. https://github.com/qubvel/segmentation_models.pytorch, 2019.

A Supplementary Material

In the Supplementary Material we:

- Provide a detailed description of the dataset.
- Provide the details for the deep learning models that were used in the paper.

A.1 Dataset Description

This subsection provides a detailed description of the dataset.

A.1.1 Motivation

Mesogeos is created to enable the development of data-driven wildfire modeling in the Mediterranean. Wildfires are one of the most threatening natural phenomena and their impact is expected to aggravate even more due to climate change. While several physical models are used in various fire-related applications, they do not have the capabilities to model the complex interactions between the different fire drivers, like the data-driven models. However, the lack of comprehensive datasets hinders the widespread adoption of data-driven modeling in wildfire-related applications. Besides, even the existing datasets in the domain are restricted to a specific ML application. Thus, Mesogeos is introduced as a multi-purpose dataset that can be used for the development of several applications. Moreover, it covers the Mediterranean, one of the most fire-prone regions on Earth, where no other wildfire-related datasets are available.

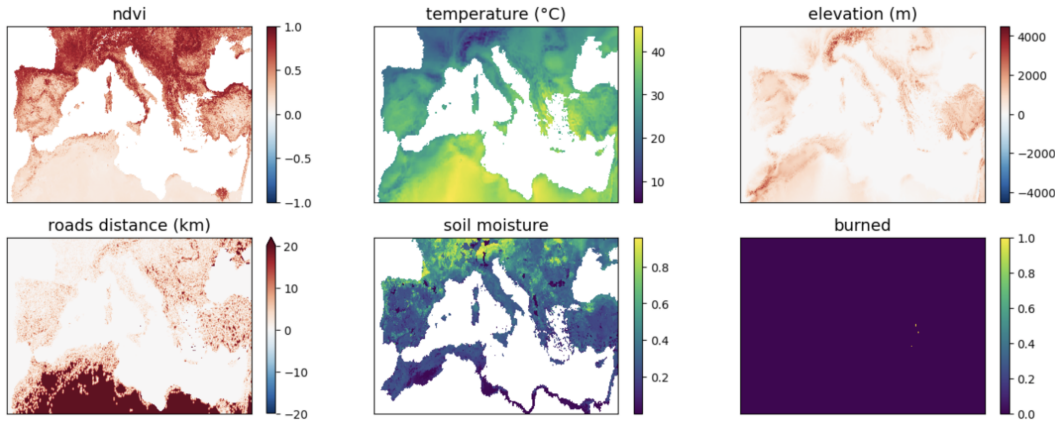


Figure 4: Visualization of some of the variables in Mesogeos datacube for a day.

Mesogeos was developed by the Orion Lab of the Institute for Astronomy, Astrophysics, Space Applications, and Remote Sensing of the National Observatory of Athens in collaboration with the Image Processing Laboratory (IPL) of the University of Valencia. This work has received funding from the European Union’s Horizon 2020 Research and Innovation Projects DeepCube and TREEADS, under Grant Agreement Numbers 101004188 and 101036926 respectively.

A.1.2 Composition

The Mesogeos dataset is structured in a spatio-temporal grid format, namely a datacube with three dimensions: longitude, latitude, and time. The datacube encompasses 27 variables related to known fire drivers such as meteorology, vegetation, land cover, and human activity. Mesogeos also includes historical burned areas, ignitions, and burned area sizes as separate variables. It has $1km \times 1km \times daily$ resolution and contains the values of the variables covering the period from 2006 to 2022. It incorporates data from the wide Mediterranean area and spans a total area of $4714km \times 1753km$ and 6026 days, for a total of 47.796.706.692 different data points for each dynamic variable.

The detailed list of the variables with their input sources, original temporal and spatial resolution, and their units are provided in Table 3.

Table 3: Table of all the variables in Mesogeos.

Variable	Source	Sp. Res.	Temp. Res.	Units
<i>Dynamic variables</i>				
Max Temperature	ERA5-Land	9km	hourly	K
Max Wind Speed	ERA5-Land	9km	hourly	m/s
Max Wind Direction	ERA5-Land	9km	hourly	°
Max Dewpoint Temperature	ERA5-Land	9km	hourly	K
Max Surface Pressure	ERA5-Land	9km	hourly	Pa
Min Relative Humidity	ERA5-Land	9km	hourly	%/100
Total Precipitation	ERA5-Land	9km	hourly	m
Mean Surface Solar Radiation Downwards	ERA5-Land	9km	hourly	J/m^2
Day Land Surface Temperature	MODIS	1km	daily	K
Night Land Surface Temperature	MODIS	1km	daily	K
Normalized Difference Vegetation Index (NDVI)	MODIS	500m	16-days	-
Leaf Area Index (LAI)	MODIS	500m	8-days	-
Soil moisture	EDO	5km	10-days	-
Burned Areas	EFFIS	1km	vector	{0, 1}
Ignition Points	MODIS	1km	vector	hectares
<i>Semi-static variables</i>				
Population	Worldpop	1km	yearly	$people/km^2$
Fraction of agriculture	Copernicus CCS	300m	yearly	%/100
Fraction of forest	Copernicus CCS	300m	yearly	%/100
Fraction of grassland	Copernicus CCS	300m	yearly	%/100
Fraction of settlements	Copernicus CCS	300m	yearly	%/100
Fraction of shrubland	Copernicus CCS	300m	yearly	%/100
Fraction of sparse vegetation	Copernicus CCS	300m	yearly	%/100
Fraction of water bodies	Copernicus CCS	300m	yearly	%/100
Fraction of wetland	Copernicus CCS	300m	yearly	%/100
<i>Static variables</i>				
Roads distance	Worldpop	1km	static	km
Elevation	COP-DEM	30m	static	m
Slope	COP-DEM	30m	static	rad
Aspect	COP-DEM	30m	static	°
Curvature	COP-DEM	30m	static	rad

We present a visualization of some of the variables in the datacube for a day in Figure 4.

Mesogeos inherits inaccuracies of the original data sources. Many of the data come from satellites (e.g. burned areas and fire ignitions, land surface temperature (LST), NDVI, LAI). Thus satellites' spatial resolution and missing data from cloud cover cascade to the datacube. Moreover, due to satellites' unavailability for specific days, there are some data missing for the LST MODIS product that was replaced by the values of the previous days (172nd of 2006, 190th of 2015, 265th of 2021, 344th of 2019, 85th of 2017, 169th of 2019, the dates are in Julian date format). The same holds for the 233rd day of 2019 for the LAI product. Moreover, the burned areas dataset provided by EFFIS does not offer the same coverage of burned areas across all Mediterranean countries throughout the specified timeframe, resulting in the absence of data for certain countries in specific years. The distribution in the years from the EFFIS product is provided in Figure 5. However, EFFIS is continuously post-processing fires from the countries spanning in the past, thus when they are made available, our intention is to append them in the datacube. In the meanwhile, if as more accurate products that refer to fires are produced, they can easily be appended to the datacube as separate variables.

As also noted in the main text according to the EFFIS website, the start dates of the fires provided by EFFIS may not always correspond to the date of ignition of the fire. For this, we make a specific post-processing to acquire the first detection of a wildfire, as this is very important in wildfire



Figure 5: The distribution of the number of fires per country in the burned areas product provided by EFFIS.

modeling. Despite the improvements achieved through our approach, it is essential to recognize that some misalignments in the ignition dates may persist in the datacube.

From the multi-purpose dataset Mesogeos, in this work, we extract two ML-ready datasets which are created by sampling the original datasets. One tailored to fire danger forecasting and one to final burned area prediction. The sampling scheme for the two datasets is analytically presented in the main paper. Moreover, the training, validation, test splits are also provided in the main text.

A.1.3 Collection Process

All the data in Mesogeos were collected directly from the data sources in the format originally provided by each data source, and a recent crawling was used close to the final date of the datacube, September 29, 2022. The collection process for each of the data sources used for Mesogeos was different for each product:

- **ERA5-Land:** The hourly data from ERA5-Land were collected from <https://cds.climate.copernicus.eu/cdsapp#!/dataset/reanalysis-era5-land?tab=form> using the API request tool provided by the platform. Access to data is based on a principle of full, open, and free access (<https://cds.climate.copernicus.eu/api/v2/terms/static/licence-to-use-copernicus-products.pdf>). The days and the timespan used was decided to match the shape of the datacube.
- **MODIS:** The data from MODIS (day's/night's land surface temperature, NDVI, LAI) were downloaded from NASA's portal <https://modis.gsfc.nasa.gov/data/> using the official python script provided for downloading that needs to declare the days, tiles and products that need to be downloaded. There are no restrictions on the use of data (<https://lpdaac.usgs.gov/data/data-citation-and-policies/>). We downloaded the data covering the temporal and spatial dimen-

sions of the datacube. Specifically, we downloaded tiles h17v04, h18v04, h19v04, h20v04, h17v05, h18v05, h19v05, and h20v05 for all the years from 2006 to 2022.

- **European Drought Observatory (EDO):** The soil moisture index was downloaded by a simple wget command from the EDO website <https://edo.jrc.ec.europa.eu/gdo/php/index.php?id=2112>. The data are distributed under a free open access (<https://edo.jrc.ec.europa.eu/gdo/php/index.php?id=2044>). We downloaded the whole data covering the timespan of the datacube.
- **European Forest Fire Information System (EFFIS):** The EFFIS data were downloaded from the real-time updated burnt areas database of the Data and Services tab of the EFFIS website <https://effis.jrc.ec.europa.eu/applications/data-and-services>. The data are distributed under the Creative Commons Attribution 4.0 International License. From all the data we used the data from 2006 til 2022, as these are the data with corrected dates. The previous years have not been processed yet from the EFFIS, as we noticed that all dates of all the fires were dated as 01/01/year.
- **Ignition Points:** The ignition points were downloaded from NASA’s FIRMS download archive <https://firms.modaps.eosdis.nasa.gov/download/create.php> portal by a manual request for the specific region of interest and timespan. There are no restrictions on the use of data (<https://lpdaac.usgs.gov/data/data-citation-and-policies/>).
- **Land Cover:** The land cover data were downloaded from <https://cds.climate.copernicus.eu/cdsapp#!/dataset/satellite-land-cover?tab=form> for the years 2006-2022 using the API request tool provided by the platform. Access to data is based on a principle of full, open and free access <https://cds.climate.copernicus.eu/api/v2/terms/static/licence-to-use-copernicus-products.pdf>.
- **Worldpop data:** Roads distance and population data were downloaded via wget from <https://www.worldpop.org> for each country in the datacube’s spatial span separately. Then the data from the separate countries were merged in a single grid using the GDAL (<https://gdal.org/>) merge command. The data are distributed under the Creative Commons Attribution 4.0 International License. The population data were provided yearly, thus the data that were downloaded were from 2006 to 2022, while for the road distance product, there is one single instance.
- **Digital Elevation Models:** The elevation product was downloaded from COP-DEM <https://spacedata.copernicus.eu/en/web/guest/collections/copernicus-digital-elevation-model> using a wget script. The product is available under a free license <https://spacedata.copernicus.eu/collections/copernicus-digital-elevation-model>. Then a post-process was used to generate slope, aspect, and curvature as described in the next section. This product is static.

It is unknown to the authors of the dataset if any ethical review processes were conducted by the providers of the datasets.

A.1.4 Pre-processing, Cleaning, Labeling

We store the collected data in an internal machine and post-process them in order to append them in the $1 \text{ km} \times 1 \text{ km} \times 1 \text{ day}$ resolution datacube.

- **ERA5-Land Data.** We combine 2 m dewpoint temperature (DT) and 2 m temperature (T), to calculate relative humidity using the following equation:

$$100 * \frac{\exp \frac{17.625 * DT}{243.04 + DT}}{\exp \frac{17.625 * T}{243.04 + T}} \quad (1)$$

We also combine 10 m wind u -component and 10 m wind v -component to compute wind speed and direction. Then, we calculate the maximum daily value of 2 m temperature, 2 m dewpoint temperature, and surface pressure, the minimum daily value of relative humidity, the mean value of the surface solar radiation downwards, and the day’s total precipitation. All these are calculated based on the hourly values of the day. For the wind, we calculate daily values of maximum wind speed. The daily wind speed direction is calculated at the

hour that the maximum wind speed occurs. Finally, we use linear interpolation to map the 9 km spatial resolution to the 1 km spatial resolution of the datacube.

- **LAI, NDVI.** We concatenate the separate tiles downloaded from MODIS and we use the nearest interpolation to map these variables to 1 km spatial resolution. Moreover, as the time resolution of the products is 8 and 16 days, we forward-fill the values in time, to fill the temporal gaps and let the variables have a daily temporal resolution.
- **Day and Night Land Surface Temperature.** No pre-processing was conducted for these products.
- **Soil Moisture Index.** We reproject this variable in EPSG:4326 and we use the nearest interpolation to map it to 1 km spatial resolution. As the temporal resolution of the variables is 10 days, we also forward-fill it in time, to fill the temporal gaps.
- **Roads Distance** is a static variable that has no time dimension. No pre-processing is needed.
- **Yearly population density.** We collect each year's population density covering the years 2006-2020. As the years 2021, and 2022 are not available, we use the year's 2020 value as a proxy for them.
- **Elevation, Slope, Aspect, Curvature.** After gathering the elevation, we upscale it to a 1 km spatial resolution by using mean aggregation. After having the 1 km spatially resolved elevation at hand, we calculate slope, aspect, and curvature at 1 km spatial resolution, using GDAL <https://gdal.org/>.
- **Land Cover.** The land cover product comes in 300 m spatial resolution. In order to take a 1 km spatial resolution product, we create 8 variables, each one related to the fraction of one out of 8 classes of interest, which are based on the following subclasses: agriculture, forest, grassland, settlements, shrubland, sparse vegetation, water bodies, wetland. For this, we calculate the percentage of each subclass in the 1km pixel of interest. Thus, in each 1 km spatially resolved cell, each of these 8 variables has a value between 0 and 1, which represents the fraction of the class presence in the pixel. We repeat this process for all the years.
- **Burned Areas & Ignition Points.** The gathered shapefiles representing final burned areas, produced by different ignitions need some pre-processing to identify the date of the ignition of the fire and to combine several burned areas, that are produced by the same fire event, into one. To this end, we associate burned area products with the product of active fires. In order to identify the date of the ignition, we create a 1 km buffer around the burned area shapefiles and we look for an active fire inside this area that has been produced within a week's temporal buffer from the date that EFFIS provides for the burned area. Moreover, to combine several burned areas into one, we also use a buffer of 1 km and we do the following: We use as an anchor a burned area and we calculate its ignition point following the procedure above. If this ignition point is also lying inside buffered burned areas other than the anchor one and these burned areas have a date within a week before the anchor one, we consider that they belong to the same event as the anchor one. Having the ignition points and burned areas at hand, we rasterize them in the same 1 km spatial grid of all the input variables. Each ignition point also holds the burned area size of the fire.

The final datacube has 29 variables.

We shift backward the variables related to fire events (burned areas, ignition points). Moreover, we also shift backward the meteorological data, as we assume that will be available through forecasts for the day of interest.

A.2 Models Details

A.2.1 Track A: Next Day's Wildfire Danger Forecasting

Long Short-term Memory (LSTM), Transformer, and Gated Transformer are the models used for Track A. All the hyperparameters included in this section were defined using the best-performing model in the validation set. All the models were trained for 30 epochs with the binary cross-entropy loss with ℓ_2 -norm regularization, the Adam optimizer, and a batch size of 256. The data were normalized before serving as input. We filled with the temporal aggregate of the time series the nulls

existing in the inputs of the models. When the value was still null, we were filling it with 0, which is the mean of each variable, after being normalized.

Regarding LSTM’s architecture, a normalization layer is followed by an LSTM layer with 128 neurons, which is followed by two linear hidden layers with 128 and 64 neurons, respectively, and an output 2-class softmax layer. The final model is trained with a 0.0063 weight decay and 0.004 learning rate, which is divided by 10 every 15 epochs. All linear layers, but the last, are followed by a dropout with probability $p = 0.5$ and the ReLU activation function.

Regarding Transformer’s architecture, the time-series input is passed through a linear layer with 256 neurons and then through a standard positional encoding layer. Then, it is followed by 2 standard Transformer Encoder Layers which are made up of self-attention and feed-forward layers. The number of heads in each layer is 2 and the neurons in each feed-forward layer are 512. The final layer is a 2-class softmax layer. The final model is trained with a 0.0018 weight decay and 0.00029 learning rate, which is divided by 10 every 15 epochs.

Regarding Gated Transformer’s architecture, for the time-Transformer block, the original time-series input is passed through a linear layer with 256 neurons and then through a standard positional encoding layer. This is followed by 4 standard Transformer Encoder Layers which are made up of self-attention and feed-forward layers. The number of heads in each layer is 4 and the neurons in each feed-forward layer are 512. Regarding the channel-Transformer block, following the original paper, the time-series input is transposed in order for the attention to act in the dimension of the variables. The input is then passed through 4 Transformer Encoder Layers with the same parameters as the time one. The outputs of the two Transformers blocks are then concatenated and passed through a softmax layer. The first logit of the softmax is multiplied with the outputs of time-Transformer and the second by the output of the channel-Transformer. The final outputs are concatenated and pass through the final 2-class softmax layer. The final model is trained with a 0.0045 weight decay and 0.00012 learning rate, which is divided by 10 every 15 epochs.

All the models were trained in a system with 2 GPUs (NVIDIA GeForce RTX 3080), each one having a memory of 10 GB. The total RAM of the system is 128GB and it also has 48 CPUs. The models were trained using one of the GPUs.

All models were developed using the PyTorch library [77]

A.2.2 Track B: Final Burned Area Prediction

For track B, the U-Net was employed.

A standard U-Net model was used with an EfficientNet-B1 Encoder, using Pytorch [77], Pytorch Lightning [78] and the segmentations_model_pytorch library [79]. The model was trained for a maximum of 50 epochs, using early stopping, with the binary cross-entropy loss with ℓ_2 -norm regularization, the Adam optimizer, and a batch size of 128. The standard U-Net architecture was used with no modifications. The final model is trained with a 0.000001 weight decay and 0.005 learning rate. The data are scaled with the minimum and maximum values in the range [0, 1] before being served as inputs to the model. We fill the nulls with -1. The model was trained in a system with 2 GPUs (NVIDIA GeForce RTX 3090), each one having a memory of 24 GB. The total RAM of the system is 128GB and it also has 32 CPUs. The models were trained using one of the GPUs.