Ecological Informatics 92 (2025) 103455

Contents lists available at ScienceDirect

Ecological Informatics

journal homepage: www.elsevier.com/locate/ecolinf

Predicting wildfire occurrences in Portugal using machine learning classification models

Jorge Caiado [a,b,*], Mariana Marques [a]

[a] ISEG Lisbon School of Managament and Economics, Universidade de Lisboa, Portugal
[b] ISEG Research/CEMAPRE, Universidade de Lisboa, Portugal

ARTICLE INFO

Keywords:
Wildfire prediction
Machine learning
Fire risk assessment
Climate and fire modeling
Land-use and wildfires
SMOTE and imbalanced data

ABSTRACT

Wildfires pose significant environmental and socio-economic challenges, particularly in fire-prone regions such as Portugal. The ability to predict wildfire occurrences is essential for improving preparedness and mitigation strategies. This study evaluates the effectiveness of three machine learning classification models (Logistic Regression, Random Forest and XGBoost) in forecasting wildfire occurrences across four Portuguese districts: Lisbon, Porto, Setúbal and Viseu. Using historical fire occurrence data and meteorological variables, the models were trained and tested on different land-use categories, including settlements, brush and agriculture. The results indicate that brush fires are the most predictable due to strong climatic influences, with models achieving F1-scores above 0.93. Settlement fires, in contrast, were more challenging to predict, likely due to human-driven variability, whereas agricultural fires exhibited intermediate predictability. To address dataset imbalances, the Synthetic Minority Oversampling Technique (SMOTE) was applied, leading to improvements in recall but a trade-off in precision. Feature importance analysis highlights the influence of long-term temporal trends, meteorological conditions and human activity on wildfire risk. These findings demonstrate the potential of machine learning models in wildfire forecasting and provide valuable insights for policymakers and fire management authorities in designing targeted prevention strategies.

1. Introduction

Wildfire management is an escalating concern in Mediterranean and southern European contexts, where regions such as Portugal confront rising fire frequency and intensity linked to climate extremes, rural depopulation, land abandonment and urban encroachment into vegetated landscapes (Costa et al., 2011; San-Miguel-Ayanz et al., 2022). In Portugal, the confluence of hotter, drier summers and shifts in land use have amplified the risk and impact of wildfires dramatically in the past two decades. The 2017 fire season, for example, led to catastrophic losses and underlined the urgent need for data-driven forecasting systems that can underpin proactive preparedness and targeted mitigation (Johnston et al., 2020; Parente et al., 2016).

Effective wildfire forecasting not only supports the allocation of suppression resources and the activation of early warning systems but is also critical for informing land-use policy in rural and peri-urban zones (Jain et al., 2020). While wildfires are conventionally defined as uncontrolled fires in natural vegetation, the growing overlap of built and natural environments (the so-called wildland–urban interface, WUI) means that fires in brush, agricultural fields, and settlement areas now pose substantial combined ecological and social risks (Phelps and Woolford, 2021; San-Miguel-Ayanz et al., 2022). Consequently, embracing a broad operational definition of wildfire, which includes fires in brush, agricultural, and peri-urban settlement areas, is increasingly adopted in the literature and essential for comprehensive risk assessment.

Historically, wildfire prediction models were built on statistical or empirical approaches, relying on meteorological thresholds, historical fire occurrence rates or simple spatial risk indicators (Martínez et al., 2009). However, these approaches struggle to capture the non-linear, high-dimensional relationships characteristic of wildfire ignition and propagation in complex, human-modified landscapes (Alkhatib et al., 2023; Phelps and Woolford, 2021). The rapid development of data science and machine learning (ML) technologies in the last five years has ushered in a new era for wildfire modeling. ML models, including ensemble and deep learning methods, can leverage large-scale environmental, temporal and social datasets, thereby delivering marked improvements in real-time risk mapping, anomaly detection and

* Corresponding author at: ISEG Lisbon School of Managament and Economics, Universidade de Lisboa, Portugal.
E-mail address: jcaiado@iseg.ulisboa.pt (J. Caiado).

https://doi.org/10.1016/j.ecoinf.2025.103455
Received 19 February 2025; Received in revised form 2 October 2025; Accepted 2 October 2025
Available online 6 October 2025

potential spread forecasting (Bot and Borges, 2022; Xu et al., 2025).

Recent research has expanded to hybrid and deep learning approaches that integrate remote sensing, climatic and socio-economic data to predict fire occurrence and spread with increasing accuracy (Abdollahi and Pradhan, 2023; Marjani and Mesgari, 2023; Masoudian et al., 2025). Furthermore, explainable AI tools are being deployed to interpret the relative influence of climate, land management and anthropogenic drivers in ML-based models (Abdollahi and Pradhan, 2023). However, most studies either focus on a single land-cover type (usually forests/brushland) or do not disaggregate predictions by land use or region, thus limiting their operational usefulness.

Despite the rapid growth of machine learning and ensemble methods in wildfire prediction, most existing studies adopt global or national perspectives, often treating wildfire occurrence as a single, undifferentiated event. This approach overlooks critical heterogeneity in fire risk across land-use types and regions. While recent studies such as Jones et al. (2024) have identified global pyromes with distinct fire drivers, and Varela et al. (2018) and Shmuel et al. (2025) have emphasized the need for region-specific fire danger assessments, few have combined district-level, land-use-disaggregated analysis with machine learning for Mediterranean Europe. Our study addresses this gap by conducting a district-level, land-use disaggregated analysis for Portugal, systematically comparing predictive performance across settlement, brush and agricultural fires. This fine-grained perspective provides novel insights into how drivers and predictability vary spatially and across land-use categories, offering practical implications for targeted fire management and prevention strategies.

Building on this motivation, we evaluate and compare three widely used ML classifiers (Logistic Regression, Random Forest and XGBoost) for predicting wildfire occurrence across three land-use types (settlements, brush and agriculture) in four diverse Portuguese districts. Using an extensive dataset spanning from 2001 to 2023, this work aims to: (1) benchmark ML model performance across land-use classes and regions; (2) assess the impact of class imbalance and SMOTE on predictive results; and (3) extract the most important features driving fire occurrence risk, leveraging model explainability frameworks.

This work advances the literature by providing spatially and thematically disaggregated wildfire risk insights, relevant both for Portuguese policymakers and the wider Mediterranean context, supporting climate adaptation and land-use planning with state-of-the-art ML methods.

The structure of this paper is as follows: Section 2 reviews relevant literature on wildfire prediction and risk assessment, highlighting previous findings. Section 3 describes the methodology, including the machine learning models used and training-test configurations. Section 4 presents the data used. Section 5 discusses model results across different Portuguese districts. Section 6 provides a detailed discussion of the key empirical findings. Finally, Section 7 concludes with key findings and recommendations for future research.

2. Literature review

2.1. Machine learning for wildfire prediction

Machine learning has become ubiquitous in wildfire science, with a sharp rise in applications since 2020 as computational power and data streams from remote sensing, climate models and citizen reporting have expanded (Bot and Borges, 2022; Jain et al., 2020; Yang et al., 2021). Early work, such as that by Cortez and Morais (2007) in Portugal, established the superiority of ML models like neural networks and decision trees over traditional statistical predictors. The field has since evolved towards more complex architectures, with studies consistently finding that gradient-boosted trees (e.g., XGBoost) outperform simpler models, particularly on large datasets (Shmuel and Heifetz, 2022; Xie and Peng, 2019).

Recent advances are dominated by deep and hybrid learning approaches. Deep learning models incorporating convolutional architectures and spatiotemporal attention mechanisms now achieve high F1-scores in fire ignition and spread prediction (Marjani and Mesgari, 2023; Masrur et al., 2024; Radke et al., 2019). Furthermore, hybrid models that fuse remote sensing metrics, meteorological time series, and social indices are increasingly prominent, pushing the boundaries of predictive accuracy (Masoudian et al., 2025; Xu et al., 2025). Parallel to these methodological innovations, the interpretability challenge has spurred the integration of explainable AI (XAI) frameworks to unpack model decisions and identify key drivers (Abdollahi and Pradhan, 2023).

2.2. The critical need for regional and land-use specificity

A persistent limitation in the literature, noted by systematic reviews, is the tendency to prioritize brush/forest fires and to treat wildfire as a monolithic event, often from a global or national perspective (Bot and Borges, 2022). This overlooks critical heterogeneity, as fire drivers and predictability are known to vary significantly by region and land-use context (Alkhatib et al., 2023; Oliveira et al., 2012).

A growing body of work underscores the importance of this regional disaggregation. Jones et al. (2024) demonstrated at a global scale that forest ecoregions cluster into distinct "pyromes" with unique sensitivities to climatic, human, and vegetation controls, revealing that emissions trends are dominated by shifts in specific extratropical pyromes. Complementing this macro-view, studies such as Shmuel et al. (2025) and Varela et al. (2018) have shown that tailoring fire danger assessments to regional patterns significantly outperforms one-size-fits-all models. Shmuel et al. (2025) specifically developed regionally-tailored fire weather indices, while Varela et al. (2018) designed a localized Fire Weather Index (FWI) classification for Greece that proved more operationally useful than standard classifications.

However, while these studies compellingly argue for regional specificity, few have combined this principle with a fine-grained, land-use-disaggregated analysis using machine learning, particularly within the high-risk context of Mediterranean Europe. Most regional models still aggregate fire events or focus on a single land-cover type. Our work bridges this gap by conducting a district-level analysis for Portugal that systematically stratifies predictions and model performance across settlement, brush, and agricultural land-use categories.

2.3. Wildfire risk drivers and data challenges

The drivers of wildfire risk are a complex interplay of environmental and human factors. While meteorological conditions remain central, human activity (e.g., population density, agricultural practices, infrastructure) is equally critical, especially in peri-urban and agricultural landscapes (Carrasco et al., 2021; Johnston et al., 2020; Masoudian et al., 2025). This complexity is further compounded by the universal challenge of class imbalance in fire occurrence datasets, where "fire" events are rare compared to "no-fire" days. To address this, research has focused on techniques like the Synthetic Minority Oversampling Technique (SMOTE) and cost-sensitive learning, which improve recall for fire detection but often involve a trade-off with precision (Andrianarivony and Akhloufi, 2024; Phelps and Woolford, 2021). Consequently, the field has moved towards multi-metric evaluation (e.g., F1-score, ROC-AUC) over accuracy alone (Bot and Borges, 2022).

3. Methodology

This study applies three established machine learning (ML) algorithms – Logistic Regression, Random Forest and XGBoost - for the prediction of wildfire occurrence across distinct land-use types and administrative regions in Portugal. These models were selected to reflect a diversity of algorithmic complexity, interpretability and performance when applied to structured environmental, climatic and anthropogenic datasets. The following outlines the principles, training procedure and evaluation strategies used in the research.

3.1. Machine learning models

Logistic Regression provides a baseline linear approach for binary classification of fire occurrence, offering high interpretability and computational efficiency (Garcia et al., 1995; Hosmer et al., 2013). Random Forest enhances predictive capacity by leveraging an ensemble of decision trees, which are trained on bootstrapped samples and include random feature selection at each node to capture complex, non-linear interactions and reduce overfitting (Breiman, 2001; Oliveira et al., 2012). XGBoost (Extreme Gradient Boosting) further advances predictive performance through sequential boosting and regularization, making it particularly effective for large, imbalanced, and high-dimensional wildfire datasets (Chen and Guestrin, 2016; Shmuel and Heifetz, 2022).

Each model outputs a class probability or label for fire occurrence (yes/no) for each observation. Model selection and hyperparameter tuning were conducted based on cross-validation performance and standard best practices for each algorithm.

These models complement each other in their ability to explore different aspects of wildfire prediction. Logistic Regression offers interpretability, Random Forest captures feature interactions and XGBoost maximizes predictive performance (Alkhatib et al., 2023; Bot and Borges, 2022).

For readers interested in the technical and mathematical formulation of each algorithm, including probability estimation in Logistic Regression, decision tree construction and aggregation in Random Forest, and the boosting, loss minimization, and regularization components in XGBoost, a detailed supplement is provided as an attachment to this manuscript.

3.2. Model performance metrics

Model performance was assessed using five widely adopted classification metrics: accuracy, precision, recall, F1-score, and the area under the receiver operating characteristic curve (ROC AUC). Together, these indicators provide a comprehensive picture of predictive power and reliability. In the wildfire context, high recall is crucial for detecting actual fire occurrences, while high precision prevents unnecessary false alarms and resource expenditure. The F1-score is emphasized as the primary criterion for model comparison, as it balances the trade-off between recall and precision under class imbalance.

Given the imbalanced nature of wildfire occurrence data, the ROC AUC was also computed, since it is a standard metric for evaluating classifiers under skewed distributions (Hoo et al., 2017; Huang and Ling, 2005). ROC AUC provides a threshold-independent measure of a model's ability to discriminate between fire and no-fire events, making it particularly informative in imbalanced settings.

Standard approach for balancing imbalanced datasets was applied, employing the Synthetic Minority Oversampling Technique (SMOTE) to augment minority class samples and mitigate bias towards the majority class, particularly for settlement and agricultural fire predictions (Chawla et al., 2002). These evaluation metrics are widely used in fire prediction research (Andrianarivony and Akhloufi, 2024) and are considered appropriate for imbalanced classification problems such as wildfire occurrence.

4. Data

4.1. Wildfire dataset: Sources, land-use categories and temporal features

This study compiled historical wildfire occurrence data from the Instituto da Conservação da Natureza e das Florestas (ICNF), covering fire events in mainland Portugal from 2001 to 2023. Due to resource constraints, the analysis focused on four districts with high policy relevance: Lisbon, Porto, Setúbal, and Viseu. Fire occurrence data was integrated with meteorological variables — temperature, precipitation, and humidity — obtained from the Instituto Português do Mar e da Atmosfera (IPMA). While meteorological data for Lisbon was readily available online, data for the other districts required direct requests to IPMA.

In line with official ICNF statistics, we classify wildfire occurrences into three categories: settlements, brush and agriculture. This classification is also consistent with international research on the wildland–urban interface (WUI), where fire risks at the boundaries of human habitation and agricultural landscapes are particularly critical (Johnston et al., 2020; Phelps and Woolford, 2021). Including settlements and agriculture as distinct categories allows us to capture human-related fire risk, a growing concern in Portugal, especially in peri-urban areas and regions where traditional agricultural practices intersect with fire-prone ecosystems. This categorization therefore reflects both national reporting standards and current scientific perspectives on human–fire interactions.

Fire risk estimation was conducted using historical values of temperature (minimum, average, and maximum), rainfall, and humidity for each day, generating daily occurrence probabilities across the three land-use types. Brush fires represent traditional vegetation-driven wildfires, agricultural fires include those occurring in croplands, pastures, and managed fields, and settlement fires reflect ignitions near or within the WUI.

Temporal features were represented using one-hot encoding for both years and months. Specifically, variables such as year_2001, year_2002, etc. indicate the calendar year in which each fire or non-fire observation was recorded. Similarly, variables such as month_1, month_2, …, month_12 correspond to the month of the year (from January to December). This encoding allows the models to identify year-specific trends (e.g., changes in fire policy or climate anomalies) and seasonal patterns (e.g., peak fire months) that influence wildfire occurrence probabilities across land-use types and regions.

4.2. Geographic context and fire occurrence distribution

Fig. 1 highlights the four Portuguese districts selected for this study (Lisbon, Porto, Setúbal and Viseu), each representing distinct geographic, climatic and socio-economic profiles relevant to wildfire risk analysis.

Lisbon is located in western Portugal along the Atlantic coast and includes the nation's capital city. Covering approximately 2761 km^2, the Lisbon district is relatively small compared to others such as Setúbal and Viseu. However, it is the most densely populated district in Portugal, with a population exceeding 2.8 million, primarily concentrated in the metropolitan area. As the country's political, economic and cultural center, Lisbon consists of highly urbanized areas alongside suburban and rural zones. The district has a Mediterranean climate, characterized by hot, dry summers and mild, rainy winters, with the Atlantic Ocean moderating extreme temperature fluctuations. While Lisbon enjoys a moderate climate, wildfire risk persists, particularly in brush areas and rural outskirts during the summer due to dry conditions.

The dataset for Lisbon consisted of 5419 daily observations, distributed across the three land-use categories. Fire occurrences were highly imbalanced, particularly in brush, where most observations corresponded to fire events (4833 instances), while only 536 observations recorded no fire. In the agriculture category, fire occurrences (3290) outnumbered no-fire cases (2129), whereas settlements exhibited a more balanced but still skewed distribution, with 1929 fire occurrences and 3490 no-fire cases.

Unlike the usual imbalance in wildfire prediction, where non-fire days vastly outnumber fire days, this inversion in Lisbon (and also observed in Viseu) reflects the ICNF's reporting practice: data are compiled at the daily level, and when multiple brush fires occur in a district on the same day, the day is classified as a “fire event.” Consequently, brush categories show more fire-labeled days than no-fire days.

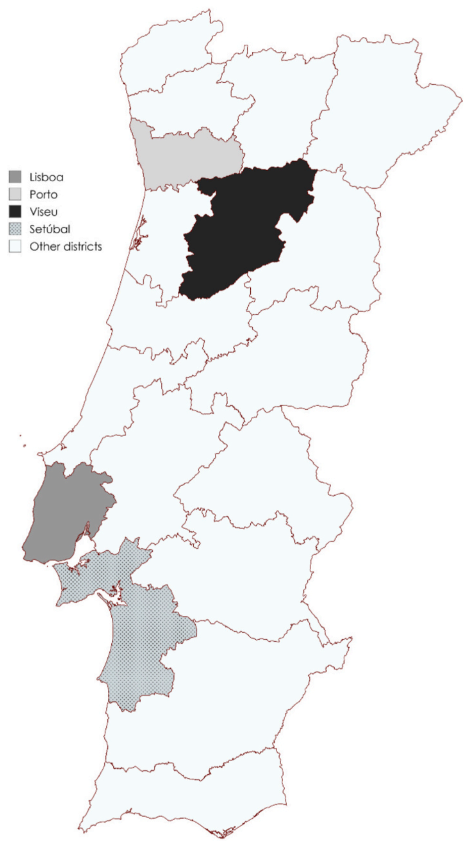

Fig. 1. Highlighted districts in Portugal: Lisbon, Porto, Setúbal and Viseu.

This local reporting convention underscores the need to contextualize results within the national fire monitoring system and highlights that model outcomes should be interpreted with these data structures in mind.

Porto, located in northern coastal Portugal, is the second most populated district after Lisbon, with around 1.8 million residents concentrated in urban and suburban areas. The district spans approximately 2400 km^2 and has a mild Mediterranean climate characterized by humid winters and dry summers. These climatic conditions, combined with dense human settlement, shape the region's wildfire risk profile. The wildfire dataset for Porto reflects a strong imbalance across land-use categories. Fire occurrences were more frequent than non-fire events in both settlements (3701 versus 1460) and brush (4904 versus 266), while agriculture showed a more balanced distribution (2120 fires versus 3050 no-fire cases).

Setúbal, located south of Lisbon along Portugal's southwestern coast, spans approximately 5064 km^2 and hosts around 900,000 residents. With a Mediterranean climate and diverse landscapes, ranging from urbanized areas to agricultural zones and protected brushland, Setúbal is particularly vulnerable to wildfires during hot, dry summers. Fire occurrence patterns in the district reflect this vulnerability. Brushland recorded the highest number of fire events (4024), followed by agriculture (2673), while settlements had fewer fire instances (1514), though still substantial. This distribution shapes the predictive challenge for each land-use category.

Viseu, an inland and rural district in central Portugal, spans approximately 5007 km^2 and is characterized by extensive agricultural and forested lands. With a population of around 370,000 and a landscape shaped by vineyards, forestry and mountainous terrain, the district is particularly susceptible to brush fires. Its Mediterranean climate, combined with continental influence, leads to dry, hot summers and cold winters, conditions favorable to wildfire ignition. Fire occurrence patterns confirm this vulnerability. Brushland recorded 4640 fire events, vastly outnumbering no-fire instances (301), followed by settlements (3255 fires versus 1686 no-fires), while agriculture showed the opposite pattern, with 3989 no-fire instances and only 952 fires.

5. Results

The predictive modeling employed three classification algorithms: Logistic Regression, Random Forest Classifier, and XGBoost Classifier. These models were evaluated based on classification accuracy, precision, recall, and F1-score to ensure a comprehensive assessment of predictive performance.

For reproducibility, Table 3 (Appendix 1) summarizes the main hyperparameter settings adopted for each model. Logistic Regression was implemented with the *liblinear* solver, L2 regularization and an increased maximum iteration limit (1000) to ensure convergence. Random Forest models were initialized with scikit-learn defaults and tuned by grid search, with the most consistent configurations involving 200–500 estimators and a maximum depth between 15 and 20. XGBoost classifiers were trained with the *logloss* evaluation metric, a learning rate of 0.1, maximum depth of 6–8, and up to 200 boosting rounds, with early stopping used to mitigate overfitting. To address class imbalance, particularly acute in settlement and agricultural fire classes, the Synthetic Minority Oversampling Technique (SMOTE) was applied to the training sets before fitting the models. These design choices follow common practices in wildfire prediction using ML (e.g., Patil et al., 2024; Shmuel and Heifetz, 2022), ensuring comparability while maintaining computational efficiency.

The following sections present detailed results for Lisbon, followed by similar analyses for Porto, Setúbal and Viseu.

5.1. Wildfire occurrence prediction in Lisbon

Table 1 presents the wildfire occurrence prediction results for Lisbon across the three land-use types, using Random Forest, XGBoost, and Logistic Regression classifiers. These models were tested with different data splits (40 %, 30 %, and 20 %) to evaluate their stability and generalizability.

For wildfire occurrence in settlements, Logistic Regression consistently demonstrated the highest accuracy across all test sizes, ranging between 0.706 and 0.708. However, its recall values remained low (between 0.380 and 0.400), indicating that it failed to identify a significant number of fire occurrences. XGBoost achieved a more balanced performance, with slightly higher recall values (between 0.428 and 0.457) compared to Logistic Regression. In contrast, Random Forest exhibited the lowest recall, suggesting lower sensitivity to fire occurrences, but it performed slightly better in terms of precision (between 0.573 and 0.552), reducing the number of false positives. Across all models, ROC AUC values for settlements remained around 0.68–0.71, reflecting only moderate discriminative ability and confirming the difficulty of distinguishing between fire and no-fire events in highly heterogeneous urban contexts.

For brush fires, model performance remained consistently high across all test sizes. Logistic Regression attained the highest recall (between 0.997 and 0.999) but at the cost of slightly lower precision (between 0.907 and 0.908). Both Random Forest and XGBoost maintained robust performance, with F1-scores ranging from 0.936 to 0.946, indicating strong overall predictive power. ROC AUC values ranged between 0.73 and 0.79, suggesting good but not perfect separation between fire and non-fire classes; this reflects the fact that while brush fire events are generally easier to detect, some overlap in predictor signals remains.

For agriculture fires, all models exhibited similar predictive performance, with accuracy values between 0.697 and 0.725 and F1-scores ranging from 0.766 to 0.783. While both XGBoost and Random Forest demonstrated slightly higher precision, Logistic Regression tended to favor recall, meaning it was more effective at identifying fire occurrences but also generated more false positives. The ROC AUC values

Table 1
Model performance metrics for wildfire occurrence prediction in Lisbon.

		Test size = 40 %				
Model	Target	Accuracy	Precision	Recall	F1-Score	ROC AUC
Random Forest	Settlements	0.687	0.573	0.434	0.494	0.691
XGBoost	Settlements	0.694	0.584	0.445	0.505	0.690
Logistic Regression	Settlements	0.706	0.628	0.400	0.489	0.703
Random Forest	Brush	0.897	0.911	0.982	0.945	0.749
XGBoost	Brush	0.881	0.914	0.959	0.936	0.725
Logistic Regression	Brush	0.906	0.907	0.999	0.951	0.762
Random Forest	Agriculture	0.704	0.737	0.803	0.769	0.767
XGBoost	Agriculture	0.708	0.745	0.794	0.769	0.768
Logistic Regression	Agriculture	0.702	0.732	0.811	0.769	0.772

		Test size = 30 %				
Model	Target	Accuracy	Precision	Recall	F1-Score	ROC AUC
Random Forest	Settlements	0.675	0.540	0.417	0.471	0.680
XGBoost	Settlements	0.707	0.603	0.457	0.520	0.702
Logistic Regression	Settlements	0.706	0.626	0.379	0.472	0.704
Random Forest	Brush	0.897	0.911	0.983	0.946	0.739
XGBoost	Brush	0.887	0.912	0.969	0.940	0.734
Logistic Regression	Brush	0.906	0.908	0.998	0.951	0.767
Random Forest	Agriculture	0.715	0.749	0.806	0.776	0.777
XGBoost	Agriculture	0.725	0.760	0.807	0.783	0.777
Logistic Regression	Agriculture	0.702	0.731	0.815	0.771	0.773

		Test size = 20 %				
Model	Target	Accuracy	Precision	Recall	F1-Score	ROC AUC
Random Forest	Settlements	0.683	0.552	0.422	0.478	0.694
XGBoost	Settlements	0.690	0.567	0.428	0.488	0.699
Logistic Regression	Settlements	0.708	0.628	0.380	0.473	0.710
Random Forest	Brush	0.896	0.908	0.985	0.945	0.737
XGBoost	Brush	0.891	0.909	0.978	0.942	0.733
Logistic Regression	Brush	0.903	0.906	0.997	0.949	0.792
Random Forest	Agriculture	0.708	0.741	0.802	0.770	0.775
XGBoost	Agriculture	0.709	0.741	0.805	0.772	0.785
Logistic Regression	Agriculture	0.697	0.724	0.814	0.766	0.780

(0.77–0.79) indicate solid discriminative performance across all models, reinforcing that agricultural fires are comparatively well captured by the feature set.

To address class imbalance, an oversampling technique known as the Synthetic Minority Oversampling Technique or SMOTE (Chawla et al., 2002) was applied. This method generated synthetic examples of the minority classes, improving dataset balance.

Table 2 presents model performance after SMOTE was implemented. The primary effect of SMOTE was an improvement in recall for settlements, with Logistic Regression increasing from 0.379 to 0.626, though at the cost of reduced precision, leading to more false positives. For brush fires, accuracy of Logistic Regression showed a sharply decrease, from a pre-SMOTE range of 0.903 to 0.906 to a post-SMOTE range of 0.686 to 0.696. However, accuracy and recall values remained high for both Random Forest and XGBoost models, demonstrating that these models retained strong sensitivity to fire occurrences. Notably, ROC AUC values remained stable before and after SMOTE, confirming that oversampling primarily affects the balance between recall and precision rather than the overall discriminative power of the classifiers. In the case of agriculture, performance remained stable, with F1-scores ranging between 0.734 and 0.776, suggesting minimal impact of SMOTE on balancing sensitivity and precision.

Fig. 2 presents the most important factors influencing wildfire occurrence in Lisbon, as identified by the XGBoost model. The results vary across land-use types, reflecting different fire drivers. For brush fires, the key factors are climatic, especially high temperatures and seasonal conditions. Fires tend to peak during the driest months, when vegetation is more flammable. Certain years also stand out, likely reflecting prolonged droughts or changing weather patterns. In settlements, fire risk is more closely tied to human activity. Days of the week matter, suggesting more fires happen during weekends or periods of higher outdoor activity. Specific years also play a role, possibly due to changes in urban development or fire prevention policies. Seasonal effects like heat and low humidity also increase fire likelihood in populated areas. For agricultural fires, historical land-use patterns are most important. Years in the early 2000s show strong influence, which may be linked to common practices like field burning. Seasonality matters too, especially after harvests when crop residues dry out. Weekdays appear relevant, suggesting that certain farming operations might increase fire risk on specific days.

5.2. *Wildfire occurrence prediction in Porto, Setúbal and Viseu*

While Lisbon served as the reference case for detailed model performance analysis, results from Porto, Setúbal, and Viseu largely confirmed similar overall patterns, albeit with notable regional variations in predictive accuracy and feature influence across land-use types.

In Porto, all models performed exceptionally well for brush fires, with F1-scores exceeding 0.96, and Logistic Regression achieving top recall (F1 = 0.977). Settlement fires also showed good model performance, with Logistic Regression attaining high recall (0.919) in populated zones, while agricultural fire prediction remained more modest. ROC AUC values support these findings, with brush fires reaching 0.73–0.79 (indicating good separability), settlements showing only moderate discrimination (0.69–0.71), and agriculture yielding intermediate performance (0.70–0.76). The application of SMOTE led to

Table 2
Model performance metrics after SMOTE for wildfire occurrence prediction in Lisbon.

Model	Target	Test size = 40 %				
		Accuracy	Precision	Recall	F1-Score	ROC AUC
Random Forest	Settlements	0.678	0.548	0.469	0.505	0.687
XGBoost	Settlements	0.690	0.575	0.457	0.509	0.696
Logistic Regression	Settlements	0.658	0.511	0.623	0.561	0.703
Random Forest	Brush	0.872	0.920	0.940	0.930	0.748
XGBoost	Brush	0.884	0.916	0.960	0.937	0.723
Logistic Regression	Brush	0.688	0.956	0.687	0.799	0.760
Random Forest	Agriculture	0.707	0.749	0.784	0.766	0.769
XGBoost	Agriculture	0.721	0.762	0.790	0.776	0.770
Logistic Regression	Agriculture	0.689	0.769	0.702	0.734	0.773

Model	Target	Test size = 30 %				
		Accuracy	Precision	Recall	F1-Score	ROC AUC
Random Forest	Settlements	0.666	0.522	0.445	0.480	0.681
XGBoost	Settlements	0.686	0.557	0.463	0.506	0.697
Logistic Regression	Settlements	0.657	0.504	0.615	0.554	0.703
Random Forest	Brush	0.875	0.919	0.944	0.931	0.750
XGBoost	Brush	0.883	0.915	0.961	0.937	0.751
Logistic Regression	Brush	0.686	0.962	0.681	0.797	0.767
Random Forest	Agriculture	0.710	0.754	0.785	0.769	0.775
XGBoost	Agriculture	0.713	0.757	0.785	0.771	0.777
Logistic Regression	Agriculture	0.693	0.772	0.711	0.740	0.772

Model	Target	Test size = 20 %				
		Accuracy	Precision	Recall	F1-Score	ROC AUC
Random Forest	Settlements	0.675	0.535	0.449	0.488	0.696
XGBoost	Settlements	0.695	0.571	0.463	0.511	0.711
Logistic Regression	Settlements	0.666	0.513	0.626	0.564	0.712
Random Forest	Brush	0.871	0.916	0.944	0.930	0.749
XGBoost	Brush	0.882	0.907	0.969	0.937	0.743
Logistic Regression	Brush	0.696	0.964	0.689	0.804	0.787
Random Forest	Agriculture	0.703	0.746	0.778	0.762	0.776
XGBoost	Agriculture	0.714	0.750	0.797	0.773	0.781
Logistic Regression	Agriculture	0.698	0.777	0.708	0.741	0.781

improvements in recall for settlements and agriculture, enhancing sensitivity to minority classes without severely compromising precision. Importantly, ROC AUC values remained stable before and after SMOTE, suggesting that the oversampling technique improved balance between recall and precision without altering the underlying discriminative power of the classifiers.

In Setúbal, the models continued to excel in predicting brush fires, with pre-SMOTE recall values near or above 0.99 for Logistic Regression. Agricultural fires showed intermediate predictability; here, Logistic Regression favored higher recall, whereas XGBoost and Random Forest allowed better balance between recall and precision. As in other districts, settlement fires proved more challenging. Although SMOTE improved recall for settlements (up to 0.683 for Logistic Regression), it also resulted in increased false positives, especially with more complex models. ROC AUC values reinforce this pattern: settlements exhibited the lowest values (0.66–0.76), indicating moderate separability, brush fires showed slightly higher values (0.68–0.72), and agriculture achieved stronger discrimination (0.73–0.77). The stability of ROC AUC across SMOTE and non-SMOTE conditions indicates that oversampling mostly reshaped recall/precision trade-offs without substantially improving separability.

For Viseu, brush fire prediction accuracy was consistently high across all algorithms, with recall often exceeding 0.95 and, for Logistic Regression, reaching perfect recall pre-SMOTE. Settlement fires were moderately well predicted, with Logistic Regression showing high recall pre-SMOTE and XGBoost maintaining robust F1-scores after SMOTE. Agricultural fire prediction was the least reliable in this district, displaying marked improvements in recall only after SMOTE application, albeit with a decrease in precision. ROC AUC values in Viseu highlight these contrasts: brush fires achieved modest separability (0.64–0.72), settlement fires were similar (0.62–0.68), while agriculture achieved the highest values overall (0.86–0.88), reflecting a stronger capacity to distinguish agricultural fire events despite weaker precision-recall tradeoffs.

To synthesize patterns without repetition, we analyzed feature importance derived from the XGBoost models for all districts and land-use types. The analysis reveals several district-specific and category-specific trends:

- Brush Fires: In all districts, the primary features driving predictions were climatic and seasonal variables—such as maximum and average temperature, humidity, and key summer or drought months. This indicates a strong and consistent link between brush fire risk and environmental conditions.
- Settlement Fires: The top predictors were often temporal features (notably certain years or months), human activity indicators (weekday effects), and occasionally meteorological variables. The temporal and human-element dominance suggests urban and peri-urban fires are highly influenced by anthropogenic drivers and policy or population changes, adding stochasticity and complexity to their prediction.
- Agricultural Fires: Importance rankings highlight a mix of historical patterns (influential years, possibly reflecting policy or land-use changes), seasonality, and farming schedule indicators (such as specific months or weekdays), implicating both climatic and management practices in agricultural fire risk.

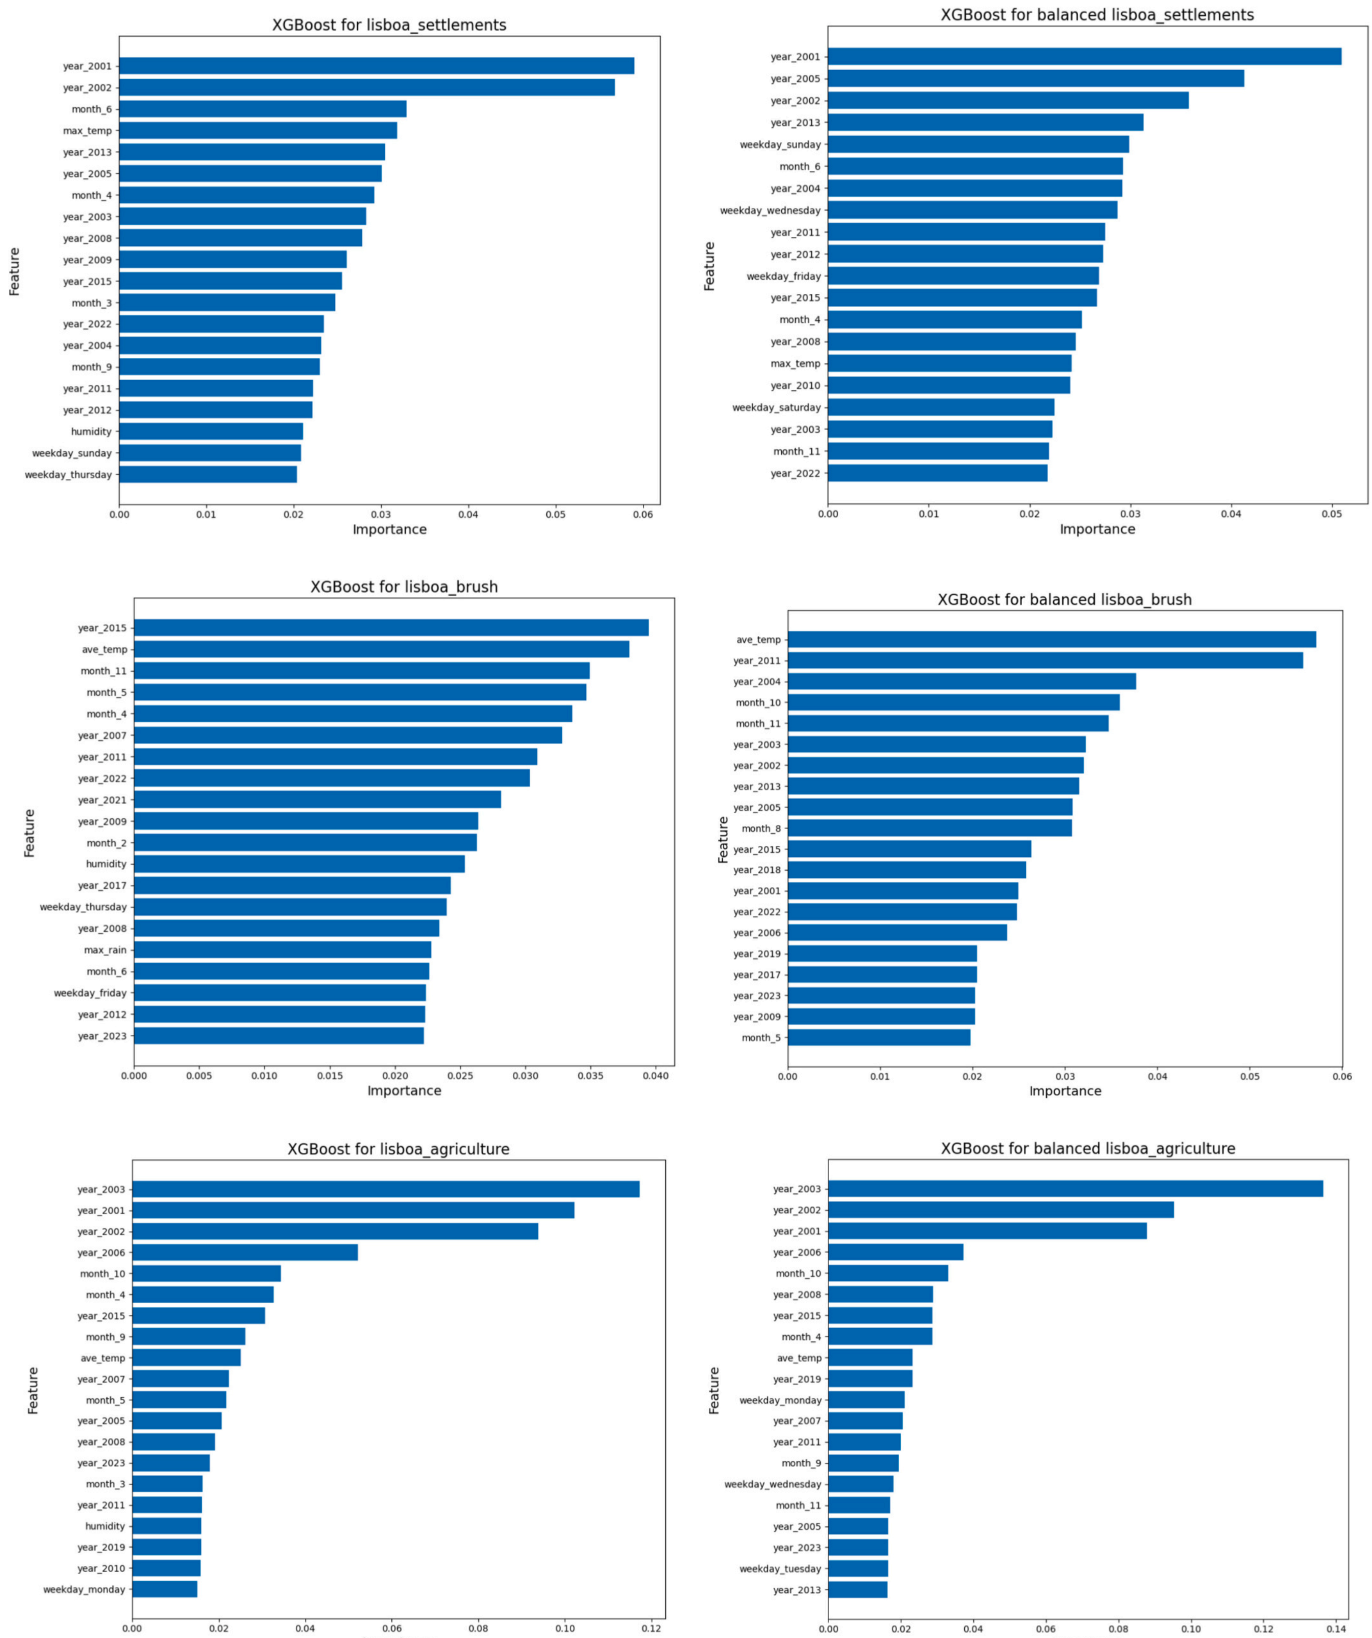

Fig. 2. Feature importance rankings of the XGBoost model in Lisbon, by land-use type.

A comparative summary across all three districts shows that brush fires are generally the most predictable due to their regular environmental and seasonal patterns. Settlement fires present greater modeling challenges, given their sensitivity to human activity and infrastructural factors. Agricultural fires display the most variability in predictability, reflecting the complexity of agricultural land management and its interaction with climatic variables. ROC AUC analyses across districts further underscore these patterns: agriculture consistently produced the highest separability values ($\geq$0.75 in Porto and Setúbal, and $\geq$ 0.85 in Viseu), brush fires maintained good though less consistent discrimination, while settlements remained the most difficult to separate, with values rarely exceeding 0.73.

Detailed performance metrics (Tables 4, 5 and 6) and graphical feature importance summaries (Figs. 3, 4 and 5) for Porto, Setúbal and Viseu are provided in the Appendix 2. These materials offer comprehensive model comparisons and illustrate key predictive factors across districts and land-use categories for readers seeking additional detail.

6. Discussion

The results of this study demonstrate considerable variation in the effectiveness of machine learning models for predicting wildfire occurrences across land-use types and Portuguese districts. These differences reinforce the complex interplay between environmental factors, human activity, regional characteristics and temporal dynamics, trends that have been emphasized in recent literature (Alkhatib et al., 2023; Costa et al., 2011; Johnston et al., 2020).

Consistent with prior studies (Oliveira et al., 2012; Vilar et al., 2010), our findings show that brush fires are the most predictable, with models achieving F1-scores above 0.93. The strong association between brush fire risk and climatic variables (maximum temperature, average temperature, humidity and seasonality), as revealed by feature importance rankings, aligns with the established literature on the central role of weather in fire ignition and spread (Turco et al., 2015; Van Beusekom et al., 2018; Xi et al., 2019). Multiple recent ML studies report similarly high predictive performance for wildland fire occurrence when environmental drivers dominate, particularly when using tree-based ensemble models or hybrid deep learning techniques (Masrur et al., 2024; Shmuel and Heifetz, 2022).

In contrast, settlement fires present greater modeling challenges, as reflected in their lower recall scores and higher variability across models. Our finding that temporal variables (years, months) and human activity indicators (weekday effects) are more important than climatic variables in settlement fire prediction echoes the conclusions of Johnston et al. (2020) and Rodrigues and de la Riva (2014), who identified anthropogenic factors and policy changes as critical but difficult-to-model elements in the wildland–urban interface (WUI). Masoudian et al. (2025) similarly found that in peri-urban and settlement areas, human behaviors, land management and population density played a more pronounced role in driving fire occurrence than local weather.

Agricultural fires were found to have intermediate predictability, influenced by both meteorological conditions and structural temporal features. The relevance of specific years and months suggests links to recurring agricultural practices (e.g., field burning, residue management), which previous works have highlighted as a dominant risk source in Mediterranean and southern European agro-ecosystems (Carrasco et al., 2021; Martínez et al., 2009). The relatively lower precision and recall for agricultural fires in some districts suggest, as Bot and Borges (2022) point out, that the integration of additional contextual data, such as crop cycles, land management and even remote sensing indicators of vegetation status, may be necessary for further improvements.

Temporal features (notably specific years) were top-ranked for all land-use types. This result is in concordance with Alkhatib et al. (2023) and Parente et al. (2016), who argue that time variables can encode not only inter-annual climate variability (e.g., drought episodes) but also shifts in policy, land use, and reporting systems. The importance of annual or seasonal anomalies is likewise confirmed in recent deep learning wildfire forecasting research (Masrur et al., 2024; Xu et al., 2025).

To address dataset imbalance, the application of SMOTE substantially improved recall for minority fire instances, especially for settlements and agriculture. This supports the findings of Phelps and Woolford (2021) and Chawla et al. (2002), who demonstrated that balancing techniques can be critical for enhancing detection of rare but high-consequence events in fire occurrence data. The observed trade-off between recall and precision reflects a broader issue noted in the literature on rare event classification, where sensitivity gains often increase false positives (Andrianarivony and Akhloufi, 2024).

The models used in this study, especially XGBoost, consistently outperformed more traditional linear models, an observation echoing the conclusions of Shmuel and Heifetz (2022), Xie and Peng (2019), and Chen and Guestrin (2016). Recent review articles (Alkhatib et al., 2023; Bot and Borges, 2022) have emphasized that tree-based ensemble and boosting algorithms provide state-of-the-art accuracy, particularly when tackling imbalanced, non-linear, and high-dimensional wildfire datasets.

Our multi-district, multi-land-use approach moves beyond the majority of prior ML-based wildfire studies, which typically focus on forest or brushland fire prediction alone (Bot and Borges, 2022; Jain et al., 2020). Our results align with Jones et al. (2024), who found that extratropical forests are increasingly vulnerable to climate-driven fires, and with Shmuel et al. (2025), who demonstrated that regionally tailored models outperform global ones. Similarly, Varela et al. (2018) showed that localized FWI classification improves operational utility, a principle we extend through ML-based, land-use-specific prediction in Portugal. By including settlement and agricultural contexts, this study provides new evidence for the unique risk profiles and modeling challenges of the WUI and agricultural areas, supporting recent calls in the literature for more granular, context-specific analyses (Masoudian et al., 2025; San-Miguel-Ayanz et al., 2022).

Recent advances increasingly emphasize the integration of deep learning, spatiotemporal models, and remote sensing data for improved wildfire occurrence and spread prediction (Marjani and Mesgari, 2023; Radke et al., 2019; Xu et al., 2025). Explainable AI (XAI) frameworks are also gaining traction, providing insights into model decision processes and highlighting the relative influence of climate, land management and anthropogenic drivers (Abdollahi and Pradhan, 2023). Our own feature importance analysis is in line with these advancements, revealing the centrality of both environmental and human factors across different land-use contexts.

Despite strong performance in brush fire prediction, the limited predictability in settlements and agriculture confirms that further inclusion of spatially explicit variables (e.g., NDVI, land cover), detailed human activity proxies and real-time environmental data is necessary for advancing fire risk modeling (Bot and Borges, 2022; Masoudian et al., 2025). Additionally, generalization across regions with differing fire regimes remains a major challenge highlighted by Xu et al. (2025) and others.

Overall, our results reinforce and expand upon the growing consensus that integrating historical fire records, climate variables and human activity indicators with advanced machine learning tools constitutes the most promising direction for regional wildfire forecasting and risk management (Abdollahi and Pradhan, 2023; Jain et al., 2020). Our findings also support the view that model transparency and explainability, through techniques such as feature importance ranking and XAI, are essential for translating predictions into actionable policy and management insights.

At the same time, we emphasize that this study's contribution is primarily applied and policy-oriented. By systematically comparing prediction performance across districts and land-use types, we provide an evidence base for tailoring fire management strategies to distinct socio-ecological contexts. This complements, rather than replaces, ongoing methodological innovations in the field. In this sense, our results bridge the gap between cutting-edge ML developments and their practical translation into decision-making, highlighting the value of disaggregated, context-specific applications for local and regional wildfire management.

7. Conclusions

This study evaluated three machine learning models (Logistic Regression, Random Forest and XGBoost) for predicting wildfire occurrence across four Portuguese districts and three land-use types. Our land-use-disaggregated, district-level analysis addresses a key gap in wildfire prediction literature.

We found that brush fires are highly predictable (F1-scores >0.93) due to strong climatic controls, with XGBoost consistently outperforming other models. In contrast, settlement and agricultural fires proved more challenging, reflecting the greater influence of anthropogenic factors and complex land-management practices. The application of Synthetic Minority Oversampling Technique (SMOTE) improved recall for these minority classes, albeit with a precision trade-off, highlighting a limitation that future work could address with more advanced imbalance-handling techniques.

Feature importance analysis underscored the combined role of climatic, temporal, and human-related variables, reinforcing the need for integrated risk models. Future research should leverage artificial intelligence approaches such as SHapley Additive exPlanations (SHAP) and Local Interpretable Model-agnostic Explanations (LIME) for deeper model interpretability.

Our primary contribution lies in demonstrating the critical importance of land-use and regional disaggregation for operational wildfire forecasting. The heterogeneity in drivers and predictability across contexts provides actionable insights for targeted fire management in Portugal and similar Mediterranean regions.

To build on this work, future studies should incorporate richer predictor sets, including remote sensing indices, topography and socio-economic data, and explicitly test model transferability across regions through cross-validation. Such efforts, coupled with stakeholder collaboration, are essential for developing robust, context-aware strategies for wildfire prevention and mitigation.

CRediT authorship contribution statement

Jorge Caiado: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Mariana Marques:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Investigation, Data curation.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors were supported by national funds through FCT – Fundação para a Ciência e a Tecnologia, I.P., in the framework of the unit UID/06522/2023.

Appendix A. Appendix 1: Machine Learning Models

This appendix provides methodological descriptions of the machine learning models (Logistic Regression, Random Forest and XGBoost) employed for wildfire occurrence prediction.

Logistic Regression.

Logistic Regression is a fundamental statistical model widely used for binary classification problems such as predicting fire occurrence (Hosmer et al., 2013). The model assumes a linear relationship between the predictor variables X and the log-odds of the dependent variable $Y \in \{0, 1\}$ (fire occurrence or no fire occurrence). The logistic function (sigmoid) maps any real-valued number to a probability in the range $(0, 1)$, ensuring a probabilistic interpretation of the predictions:

$$P(Y_i = 1|X) = P(Y_i = 1|X_{1i}, X_{2i}, \ldots, X_{ki}) = \frac{1}{1 + e^{-(b_0 + b_1X_{1i} + b_2X_{2i} + \cdots + b_kX_{ki})}},$$

where: $P(Y_i = 1|X)$ represents the probability of fire occurrence; $b_1, b_2, \cdots, b_k$ are the coefficients corresponding to the predictor variables $X_{1i}, X_{2i}, \ldots, X_{ki}$, b_0 is the intercept term, and e is the base of the natural logarithm.

The model is trained using Maximum Likelihood Estimation (MLE), where the likelihood function is maximized to obtain the best-fitting coefficients. Logistic Regression is computationally efficient and provides interpretable results, making it a useful baseline for comparison with more complex models.

Random Forest.

Random Forest is an ensemble learning method that constructs multiple decision trees during training and aggregates their predictions to improve accuracy and robustness (Breiman, 2001). Each tree is built using a subset of the training data (bootstrapping) and a random selection of features, reducing overfitting and increasing generalization.

Given a dataset D with predictor variables X and target variable Y, Random Forest generates multiple decision trees $T_1, T_2, \ldots, X_m$. The final classification output is obtained via majority voting:

$$\widehat{Y} = mode(T_1(X), T_2(X), \ldots, X_m(X))$$

where $\widehat{Y}$ is the predicted class (fire occurrence or no fire occurrence). Each decision tree follows a recursive partitioning process based on an impurity measure such as Gini Index or Entropy.

Random Forest is highly effective in handling non-linearity and interactions between variables, making it a robust choice for fire occurrence prediction. Additionally, it reduces overfitting by averaging multiple tree predictions, improving generalization.

XGBoost.

Extreme Gradient Boosting (XGBoost) is an optimized implementation of gradient boosting that enhances speed and predictive performance (Chen and Guestrin, 2016). Unlike Random Forest, which builds trees independently, XGBoost sequentially constructs trees, where each new tree corrects errors from the previous iterations.

The model minimizes a regularized objective function consisting of a loss function L and a complexity penalty $\Omega(f)$ to prevent overfitting:

$$O = \sum_{i=1}^{N} L(Y_i, \widehat{Y}_i) + \sum_{k=1}^{K} \Omega(f_k),$$

where: $L(Y_i, \widehat{Y}_i)$ represents the logistic loss function; $\Omega(f_k)$ is the regularization term that controls tree complexity; k is the number of trees; and f_k and is a function that maps input X to output Y.

Gradient boosting updates predictions iteratively by fitting new trees to the negative gradient of the loss function, adjusting weights to minimize errors:

$$\widehat{Y}_i^{(t+1)} = \widehat{Y}_i^{(t)} + \eta f_t(X_i),$$

where η is the learning rate controlling step size.

The main hyperparameter settings for each algorithm are provided in Table 3.

Appendix Table 3
Main hyperparameter settings for the machine learning models.

Model	Key Hyperparameters	Notes
Logistic Regression	Solver = *liblinear*; Penalty = L2; max_iter = 1000	Increased iteration limit to ensure convergence in multi-class setting.
Random Forest	Estimators = 200–500; Maximum depth = 15–20; Default bootstrap = True	Best configurations identified via grid search; scikit-learn defaults otherwise retained.
XGBoost	Learning rate = 0.1; Maximum depth = 6–8; Boosting rounds = up to 200; eval_metric = *logloss*; Early stopping applied	Parameters reflect standard practice in wildfire ML tasks; tuned to balance performance and overfitting.
All Models	Features normalized with Min–Max scaling; SMOTE applied to balance minority classes in training data	SMOTE particularly important for settlement and agricultural fires.

Training and test sets.

To ensure robust assessment of generalization and minimize overfitting, we experimented with several training–test split ratios: 60:40, 70:30, and 80:20. Through comparative evaluation, the 70:30 split was selected as optimal for balancing training data sufficiency and test sample robustness. All models were trained and evaluated on the same data splits, stratified to preserve class distribution across fire and no-fire cases.

Appendix B. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.ecoinf.2025.103455.

Data availability

All data, code and supplementary materials used in this study are publicly available at: https://github.com/jorgecaiado/fire-forecasting-portugal. This ensures full reproducibility of the analysis, in accordance with the journal's guidelines (Huettmann and Arhonditsis, 2023).

References

Abdollahi, A., Pradhan, B., 2023. Explainable artificial intelligence (XAI) for interpreting the contributing factors feed into the wildfire susceptibility prediction model. Sci. Total Environ. 879, 163004. https://doi.org/10.1016/j.scitotenv.2023.163004.

Alkhatib, R., Sahwan, W., Alkhatieb, A., Schutt, B., 2023. A brief review of machine learning algorithms in forest fires science. Appl. Sci. 13, 8275. https://doi.org/10.3390/app13148275.

Andrianarivony, H.S., Akhloufi, M.A., 2024. Machine learning and deep learning for wildfire spread prediction: A review. Fire 7, 482. https://doi.org/10.3390/fire7120482.

Bot, K., Borges, J.G., 2022. A systematic review of applications of machine learning techniques for wildfire management decision support. Inventions 7 (1), 15.

Breiman, L., 2001. Random forests. Mach. Learn. 45, 5–32. https://doi.org/10.1023/A:1010933404324.

Carrasco, J., Acuna, M., Miranda, A., Alfaro, G., Pais, C., Weintraub, A., 2021. Exploring the multidimensional effects of human activity and land cover on fire occurrence for territorial planning. J. Environ. Manag. 297, 113428.

Chawla, N.V., Bowyer, K.W., Hall, L.O., Kegelmeyer, W.P., 2002. SMOTE: synthetic minority over-sampling technique. J. Artif. Intell. Res. 16, 321–357.

Chen, T., Guestrin, C., 2016. XGBoost: A scalable tree boosting system. In: In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16). Association for Computing Machinery, New York, NY, USA, pp. 785–794. https://doi.org/10.1145/2939672.2939785.

Cortez, P., Morais, A., 2007. A data mining approach to predict forest fires using meteorological data. In: Neves, J., Santos, M.F., Machado, J. (Eds.), 13th Portuguese Conference on Artificial Intelligence, New Trends in Artificial Intelligence, pp. 512–523.

Costa, L., Thonicke, K., Poulter, B., Badeck, F.W., 2011. Sensitivity of Portuguese forest fires to climatic, human, and landscape variables: subnational differences between fire drivers in extreme fire years and decadal averages. Reg. Environ. Chang. 11, 543–551. https://doi.org/10.1007/s10113-010-0169-6.

Garcia, C.V., Woodard, P.M., Titus, S.J., Adamowicz, W.L., Lee, B.S., 1995. A logit model for predicting the daily occurrence of human caused forest-fires. Int. J. Wildland Fire 5, 101–111.

Hoo, Z.H., Candlish, J., Teare, D., 2017. What is an ROC curve? Emerg. Med. J. 34, 357–359.

Hosmer, D.W., Lemeshow, S., Sturdivant, R.X., 2013. Applied Logistic Regression, 3rd edition. John Wiley & Sons. https://onlinelibrary.wiley.com/doi/book/10.1002/9781118548387.

Huang, J., Ling, C.X., 2005. Using AUC and accuracy in evaluating learning algorithms. IEEE Trans. Knowl. Data Eng. 17 (3), 299–310.

Huettmann, F., Arhonditsis, G., 2023. Towards an ecological informatics scholarship that is reflective, repeatable, transparent, and sharable! Eco. Inform. 76, 102132. https://doi.org/10.1016/j.ecoinf.2023.102132.

Jain, P., Coogan, S.C., Subramanian, S.G., Crowley, M., Taylor, S., Flannigan, M.D., 2020. A review of machine learning applications in wildfire science and management. Environ. Rev. 28, 478–505.

Johnston, L.M., Wang, X., Erni, S., Taylor, S.W., McFayden, C.B., Oliver, J.A., Stockdale, C., Christianson, A., Boulanger, Y., Gauthier, S., Arseneault, D., Wotton, M., Parisien, M., Flannigan, M.D., 2020. Wildland fire risk research in Canada. Environ. Rev. 28, 164–186. https://doi.org/10.1139/er-2019-0046.

Jones, M.W., Veraverbeke, S., Andela, N., Doerr, S.H., Kolden, C., Mataveli, G., Pettinari, M.L., Le Quéré, C., Rosan, T.M., van der Werf, G.R., van Wees, D., Abatzoglou, J.T., 2024. Global rise in forest fire emissions linked to climate change

in the extratropics. Science 386, 6719. https://www.science.org/doi/10.1126/science.adl5889.
Marjani, M., Mesgari, M.S., 2023. The large-scale wildfire spread prediction using a multi-kernel convolutional neural network. ISPRS Ann. Photogramm. Remote Sens. Spatial Inform. Sci. X-4/W1-2022, 483–488. https://doi.org/10.5194/isprs-annals-X-4-W1-2022-483-2023.
Martínez, J., Vega-García, C., Chuvieco, E., 2009. Human-caused wildfire risk rating for prevention planning in Spain. J. Environ. Manag. 90, 1241–1252.
Masoudian, E., Mirzaei, A., Bagheri, H., 2025. Assessing wildfire susceptibility in Iran: leveraging machine learning for geospatial analysis of climatic and anthropogenic factors. Trees, Forests People 19, 100774. https://doi.org/10.1016/j.tfp.2025.100774.
Masrur, A., Yu, M., Taylor, A., 2024. Capturing and interpreting wildfire spread dynamics: attention-based spatiotemporal models using ConvLSTM networks. Eco. Inform. 82, 102760. https://doi.org/10.1016/j.ecoinf.2024.102760.
Oliveira, S., Oehler, F., San-Miguel-Ayanz, J., Camia, A., Pereira, J.M., 2012. Modeling spatial patterns of fire occurrence in Mediterranean Europe using multiple regression and random Forest. For. Ecol. Manag. 275, 117–129. https://doi.org/10.1016/j.foreco.2012.03.003.
Parente, P., Pereira, M.G., Tonini, M., 2016. Space-time clustering analysis of wildfires: the influence of dataset characteristics, fire prevention policy decisions, weather and climate. Sci. Total Environ. 559, 151–165. https://doi.org/10.1016/j.scitotenv.2016.03.129.
Patil, R., Pawar, J., Shah, K., Shetty, D., Ajith, A., Jadhav, S., 2024. Machine learning based Forest fire prediction: A comparative approach. Int. Res. J. Multidiscip. Technovat. https://doi.org/10.54392/irjmt2413.
Phelps, N., Woolford, D.G., 2021. Comparing calibrated statistical and machine learning methods for wildland fire occurrence prediction: a case study of human-caused fires in lac La Biche, Alberta, Canada. Int. J. Wildland Fire 30, 850–870. https://doi.org/10.1071/WF20139.
Radke, D., Hessler, A., Ellsworth, D., 2019. FireCast: leveraging deep learning to predict wildfires spread. In: Proceedings of the 28th international joint conference on artificial intelligence, pp. 4575–4581. https://doi.org/10.24963/ijcai.2019/636.
Rodrigues, M., de la Riva, J., 2014. An insight into machine-learning algorithms to model human caused wildfire occurrence. Environ. Model. Softw. 57, 192–201. https://doi.org/10.1016/j.envsoft.2014.03.003.
San-Miguel-Ayanz, J., Durrant, T., Boca, R., Maianti, P., Liberta`, G., Artes Vivancos, T., Jacome Felix Oom, D., Branco, A., De Rigo, D., Ferrari, D., Pfeiffer, H., Grecchi, R., Onida, M., Loffler, P., 2022. Forest Fires in Europe, Middle East and North Africa 2021. Publications Office of the European Union. https://data.europa.eu/doi/10.2760/34094.
Shmuel, A., Heifetz, E., 2022. Global wildfire susceptibility mapping based on machine learning models. Forests 13, 1050. https://doi.org/10.3390/f13071050.
Shmuel, A., Lazebnik, T., Heifetz, E., Glickman, O., e Price, C., 2025. Fire weather indices tailored to regional patterns outperform global models. npj Nat Hazards 2, 74. https://doi.org/10.1038/s44304-025-00126-y.
Turco, M., Palazzi, E., von Hardenberg, J., Provenzale, A., 2015. Observed climate change hotspots. Geophys. Res. Lett. 42, 3521–3528. https://doi.org/10.1002/2015GL063891.
Van Beusekom, A.E., Gould, W.A., Monmany, A.C., Khalyani, A.H., Quiñones, M., Fain, S. J., González, G., 2018. Fire weather and likelihood: characterizing climate space for fire occurrence and extent in Puerto Rico. Clim. Chang. 146, 117–131.
Varela, V., Sfetsos, A., Vlachogiannis, D., Gounaris, N., 2018. Fire Weather Index (FWI) classification for fire danger assessment applied in Greece. Tethys 15, 31–40. https://doi.org/10.3369/tethys.2018.15.03.
Vilar, L., Woolford, D.G., Martell, D.L., Martín, M.P., 2010. A model for predicting human-caused wildfire occurrence in the region of Madrid, Spain. Int. J. Wildland Fire 19, 325–337. https://doi.org/10.1071/WF09030.
Xi, D.D., Taylor, S.W., Woolford, D.G., Dean, C.B., 2019. Statistical models of key components of wildfire risk. Annu. Rev. Stat. Appl. 6, 197–222.
Xie, Y., Peng, M., 2019. Forest fire forecasting using ensemble learning approaches. Neural Comput. & Applic. 31, 4541–4550. https://doi.org/10.1007/s00521-018-3515-0.
Xu, Z., Li, J., Cheng, S., Rui, X., Zhao, Y., He, H., Guan, H., Sharma, A., Erxleben, M., Chang, R., Xu, L.L., 2025. Deep learning for wildfire risk prediction: integrating remote sensing and environmental data. ISPRS J. Photogramm. Remote Sens. 227, 632–677. https://doi.org/10.1016/j.isprsjprs.2025.06.002.
Yang, S., Lupascu, M., Meel, K.S., 2021. Predicting Forest fire using remote sensing data and machine learning. Proc. AAAI Conf. Artif. Intell. 35 (17), 14983–14990. https://doi.org/10.1609/aaai.v35i17.17758.