EcoGuard: Acoustical Forest Threat & Fire Risk Predictor

Prof. Harshal D. Misalkar[1], Devayani V. Korade[2], Shrawani A. Gedam[3],
Rutuja M. Pande[4], Dhananjay Khajuria[5]
[1,2,3,4,5]*Department of Information Technology, PRMITR, Badnera, India*

***Abstract*—Forests are threatened by illegal logging, poaching, and wildfires, but detecting these threats can be difficult in large, isolated areas, so it is important to monitor them constantly and ensure the preservation of forest resources. EcoGuard is an AI-based forest monitoring system capable of detecting potential threats and analyzing the risk of wildfires at any given moment. The system processes environmental sounds and weather data, presenting the information on a web dashboard for real-time monitoring. To classify the audio data, EcoGuard uses ResNet-18 Convolutional Neural Networks to transform raw sound files into Mel-spectrogram images and classify them as chainsaw, gunshots, animal sounds, rainfall, and other sounds. These files are processed to identify potential threats such as illegal logging or poaching activities. In addition, to predict the possibility of wildfires, EcoGuard employs a Random Forest Regressor that uses weather data, including temperature, humidity, rain, and wind, to analyze the potential spread of a wildfire. The system then compiles the information into a web dashboard that displays all of the gathered data, including audio, weather, detected threats with their locations on a map, alerts, and more. By combining the processing of audio files to detect threats and analyze the weather, EcoGuard increases the efficiency of forest monitoring systems, helping to reduce environmental damage and ensuring forest conservation.**

***Index Terms*—Forest monitoring, Audio classification, ResNet-18, Fire Risk Prediction, Random Forest Regressor**

I. INTRODUCTION

Forests are essential ecosystems that sustain biodiversity, conserve natural resources and the ecological balance, and host various wildlife species. Unfortunately, they are subject to illegal deforestation, poaching, and other anthropogenic activities, as well as natural disasters such as wildfires. These activities can severely impact the forests and their flora and fauna, which often requires immediate intervention to prevent large scale damage. Such threats call for constant monitoring, which is challenging in remote regions and often requires people to put themselves at risk. Monitoring chainsaw sounds or even gunshots can alert park rangers about illegal logging or poaching activities. Similarly, temperature, humidity, rainfall, and wind speed can be critical to assessing a forest's state and predicting wildfire risk. Constant vigilance and activity analysis can help protect the forests and their ecological importance, such as carbon sequestration, water resources, and biodiversity.

Methods such as human monitoring, watchtowers, sensors, and patrols can be insufficient, as they might require continuous human operations, have limited scopes, or be too slow to prevent large-scale damage. Moreover, such approaches do not always allow proactive measures, such as predicting wildfires, and may only help on smaller scales, such as localized threat detection. Therefore, a system that can constantly monitor the forest and its perimeters, analyze the data to identify and classify potential threats, and warn the authorities in real-time is needed to address the research problem. EcoGuard is proposed as an AI-based forest monitoring and wildfire prediction system that utilizes threat detection based on sound identification, as well as weather data analysis to predict potential wildfires. The system will feature a dashboard that can display the classification results and detection data, as well as warn the operators in the case of threats detected. Specifically, the acoustic data will be converted to Mel-spectrograms and processed with a ResNet-18 CNN to identify potential sounds, such as gunshots, chainsaws, animal calls, rainfall, and others. Then, temperature, humidity, wind speed, and rainfall data will be analyzed with a Random Forest Regressor to predict potential wildfires. The dashboard will then provide relevant information based on the detected sounds and

the wildfire prediction model. The research will focus on creating a system that monitors forests with artificial intelligence to detect potential threats such as illegal logging or poaching activities and predict wildfires. EcoGuard will utilize an AI-based approach to analyze sounds and weather data to identify threats and provide information for proactive intervention. By combining two methods, such as audio analysis and predictive weather modeling, the project will propose a unique solution that will address the research problem and provide a reliable method to ensure forest conservation and prevent large-scale damage.

II. RELATED WORK

The existing literature on forest threat and wildfire monitoring covers a range of topics such as acoustic-based wildfire detection, deep learning environmental sound classification, audio feature extraction, machine-learning based wildfire prediction, remote-sensing and environment-based wildfire prediction, among others. It has been established that multiple approaches and data sources can be used in detecting or predicting forest-fire events. However, the different studies cover varying aspects, such as input data source, target prediction, model architecture, and deployment environment. The following review will break down the existing research into six major themes and compare the previous studies with the proposed EcoGuard system.

A. Acoustic-Based Forest Threat and Wildfire Detection

Acoustic-based forest-threat and wildfire detection is a recent area of research, as sound can carry information about events occurring in forest environments. Huang et al. [1] propose an audio-based wildfire detection pipeline for embedded systems. They developed a three-stage system, consisting of audio preprocessing, feature engineering, and classification. They showed that it is possible to detect wildfire sounds in the presence of forest background noise. The system processes a one-second audio clip in 66ms on a Raspberry Pi 4 and reported precision up to 99% in their experiments. The study further discusses challenges with visual and infrared wildfire detection, such as obstacles in the field of view. Acoustic sensing provides another means of obtaining information about forest events. Damasevicius et al. [2] propose a hybrid machine learning model for forest wildfire detection, using sound as the input data. They developed a hybrid model that uses spectrogram-based audio representation with a convolutional processing stage and an LSTM layer to extract spatial and temporal information about the wildfire sounds. They also used Particle Swarm Optimization to optimize the neural-network weights. Their model achieved 94. 7% accuracy, 93. 2% precision, 96. 1% recall, 94. 6% F1-score, and AUC of 0. 987. The two studies demonstrate that it is possible to use acoustic information for wildfire detection, though Huang et al. [1] focused on embedded system deployment while Damasevicius et al. [2] classified only wildfire and non-wildfire sounds, leaving the question of using acoustic data for multiple forest threats unclear.

Relevance to EcoGuard: The EcoGuard system expands the idea of using acoustic information to a more general forest-surveillance setting. According to the documentation, EcoGuard uses a ResNet-18 CNN for acoustic threat classification that recognizes 27 classes of audio, including chainsaws, gunshots, fire, and fauna. The system further processes the audio using Mel-spectrogram and MFCC representations before classification. Unlike the approach by Huang et al. [1], which specifically targets embedded system deployment on a Raspberry Pi, EcoGuard is considered in this project as a software-based system. Its documented live-monitoring interface uses the browser Web Audio API to fetch the microphone audio, which is then processed by the acoustic classification pipeline.

B. Deep Learning for Environmental Sound Classification

Deep learning methods for environmental sound classification have been studied, as audio data contains complex temporal and frequency information. Damasevicius et al. [2] used a hybrid convolutional and recurrent neural network to classify wildfire and non-wildfire sounds from spectrogram-based audio representations. The model's convolutional layers processed the time-frequency spectrogram information while the LSTM layer captured the temporal information of the audio samples. They reported an accuracy of 94. 7% and a recall of 96. 1%, indicating that the model was able to identify most of the wildfire samples in their dataset. Huang et al. [1], on the other hand, used a machine

learning pipeline with audio preprocessing, feature engineering, and classification, instead of a large deep neural network. They specifically discussed computational and memory limitations of embedded systems, which makes deep neural network inference challenging. The limitations of the deep learning model used by Damasevicius et al. [2], beyond the information discussed in the provided paper, are Not clearly stated in the provided paper.

Relevance to EcoGuard: EcoGuard uses the PyTorch ResNet-18 CNN for acoustic threat classification, which takes in a single-channel 128 × 128 input representation and classifies the audio into 27 sound classes. The README also documents a custom dropout-based classification head. Therefore, EcoGuard is different from the studies that only perform binary classification of wildfire vs non-wildfire, as its documented acoustic component is designed to perform multi-class forest-threat classification. The README indicates 27 audio classes instead of only two wildfire classes.

C. Audio Feature Extraction and Mel-Spectrogram Analysis

Audio feature extraction is a crucial step in many audio classification pipelines, as raw audio data is difficult to be used directly by conventional machine learning models. Huang et al. [1] transformed audio signals into spectrogram-based representations for classification. The paper discusses the use of Mel-spectrograms, which provides a time-frequency representation of the sound while reducing the dimensionality, compared to a conventional short-time Fourier transform. Damasevicius et al. [2] also used spectrogram representation as the input to their hybrid deep learning model. The spectrogram provides both the spectral and temporal information about the sound, as the convolutional layers of the model extract the patterns in the spectrogram while the LSTM component models the temporal information. The study also discussed audio features such as MFCCs, which were evaluated as part of their feature analysis. The transformation of audio into a time-frequency representation provides useful information for environmental sound classification, though the limitations of the feature extraction step identified by the reviewed papers are Not clearly stated in the provided papers.

Relevance to EcoGuard: EcoGuard's documented acoustic pipeline uses the Mel-spectrogram and MFCC processing. The README also documents Librosa for Mel-spectrogram, MFCC, RMS-energy and zero-crossing-rate processing, with SciPy and NumPy used for signal processing operations. The documented monitoring process transforms the incoming microphone audio into an acoustic analysis representation before passing it to the ResNet-18 classification model. Thus, the approach used in EcoGuard aligns with the general approach of previous acoustic research while applying the extracted audio representation to a more general multi-class forest-threat classification task.

D. Machine Learning-Based Wildfire Risk Prediction

Apart from detecting an active fire event, machine learning methods have been used to estimate the probability or risk of wildfire occurrence from environmental factors. Gao et al. [3] investigated forest-fire risk prediction in the Heihe area of Heilongjiang Province, China, using Random Forest (RF) and Backpropagation Neural Network (BPNN) models. They used historical forest-fire data from 1995–2015 together with meteorological, topographic, and geographic information. Random Forest feature importance evaluation and logistic stepwise regression were used to identify important fire-driving factors prior to developing the prediction models. They identified several important factors, including the number of days after the last rainfall, relative humidity, maximum temperature, water-vapor pressure, minimum relative humidity, and the distance to settlements.

They reported that RF and BPNN had similar overall prediction accuracy and goodness of fit and concluded that both were suitable for forest-fire occurrence prediction in the studied region. Hiremath and Kannan [4] proposed a Random Forest-based classifier for predicting fire occurrence one day in advance in the Odisha region. They considered skin temperature, near-surface air temperature, near-surface humidity, and soil moisture. Further parameters were derived from these variables, and mutual information was used to identify the most relevant parameters. The top nine parameters were then supplied to the Random Forest model. Weighted sampling and SMOTE were used to handle class imbalance. The authors specifically identified that much of the previous research had

focused on diagnostic analysis of forest fires, while fewer studies attempted proactive forecasting.

Relevance to EcoGuard: EcoGuard has a separate Wildfire Risk Predictor component. According to the README, this component uses a Scikit-Learn Random Forest Regressor, nine climatological features, and Canadian Fire Weather Index (FWI) synthesis to produce a continuous wildfire ignition and spread risk value. This is different from Hiremath and Kannan [4], where Random Forest is used as a classifier for fire/no-fire prediction one day in advance. EcoGuard's documented wildfire component uses a Random Forest Regressor to generate continuous risk information. Therefore, the two approaches should not be considered identical even though both use Random Forest.

E. Remote-Sensing and Environmental Data-Based Wildfire Prediction

Remote sensing and environmental data have also been widely used for wildfire prediction. Yang, Lupascu, and Meel [5] proposed a machine learning approach using remote-sensing data to predict forest fires in Indonesia. They discussed limitations of existing systems based on handcrafted features and expensive ground instruments and proposed a cost-effective machine learning approach. Their model achieved an area under the ROC curve greater than 0. 81 while the baseline approach did not exceed 0.70. The performance remained above 0.81 when evaluated using reduced data. Apostolakis et al. [6] investigated next-day forest-fire risk prediction using information up to the previous day. Their methodology involved topographic, meteorological, Earth Observation, and historical fire-proneness features. The authors highlighted the difficulty caused by high spatial requirements and extreme class imbalance. They evaluated multiple machine learning algorithms using different cross-validation schemes and reported sensitivity above 90% and specificity above 65%. The study included algorithms such as Random Forest, Extra Trees, XGBoost, and shallow neural networks. The researchers also emphasized the importance of appropriate validation and evaluation methods for assessing model generalizability. Huot et al. [7] investigated deep learning models for wildfire prediction using historical remote sensing information. Their dataset combined historical wildfire observations with terrain, vegetation, and weather information. The researchers formulated wildfire likelihood estimation as multiple image and sequence segmentation problems and evaluated four models: a convolutional autoencoder, residual U-Net, convolutional autoencoder with ConvLSTM, and residual U-Net with ConvLSTM. Hernandez and Hoskins [8] compared Decision Trees, Random Forest, Naïve Bayes, and Neural Networks for wildfire prediction in selected areas of California's Central Valley. Their dataset included MERRA-2 data, Landsat 8 data, and fire history from 2012–2023. The best reported model in their experiment was a Decision Tree with an F1-score of 0. 689. They also found a substantial difference between random and chronological data splits, with the chronological split producing zero true positives for all tested methods except Naïve Bayes. These studies demonstrate the importance of environmental, meteorological, geographical, and historical information for wildfire prediction. They also show that model performance can depend strongly on geographic region, data characteristics, class imbalance, and validation strategy.

Relevance to EcoGuard: EcoGuard does not document a satellite-based wildfire prediction pipeline. Instead, its documented wildfire risk component uses nine climatological features with a Random Forest Regressor and Canadian FWI synthesis. Therefore, remote sensing approaches such as those used by Yang et al. [5], Apostolakis et al. [6], and Huot et al. [7] should be treated as related environmental prediction approaches rather than technologies implemented in EcoGuard.

III. RESEARCH GAP

The reviewed studies suggest that using machine learning, deep learning, acoustic sensing, remote sensing, and environmental data can be effective for forest-fire detection and wildfire-risk prediction. However, the reviewed literature largely focuses on a single task: detection or prediction.

A. Separate Treatment of Acoustic Detection and Wildfire-Risk Prediction

Previous studies have investigated acoustic-based wildfire detection and environmental wildfire-risk

prediction in separate problem formulations. For instance, Huang et al. [1] developed an audio-based wildfire detection approach that can distinguish wildfire sounds from forest background sounds. Damasevicius et al. [2] have also investigated wildfire detection based on acoustic features and a hybrid CNN-LSTM model. Meanwhile, Gao et al. [3] and Hiremath and Kannan [4] focused on wildfire prediction using meteorological and environmental parameters. Other works used remote sensing, geographical information, vegetation data, and historical fire information for wildfire prediction [5] [8]. As a result, from the perspective of the reviewed literature, it can be observed that the field of forest monitoring has not fully integrated the two information sources of acoustic event detection and environmental wildfire risk. The two sources of information are relevant but investigated in different research directions.

B. Single-Purpose Nature of Existing Acoustic Detection Studies

The reviewed acoustic-based studies primarily focus on wildfire detection. For instance, Huang et al. [1] classify audio signals into fire and forest categories. Damasevicius et al. [2] investigated the use of acoustic information for wildfire/non-wildfire classification. While the reviewed studies effectively demonstrate the applicability of acoustic information for wildfire detection, the objective of these studies is primarily limited to identifying wildfire-related audio information. As such, the literature provides a research opportunity to shift the focus from single-purpose acoustic detection to a more generalized approach of acoustic classification. A generalized approach would allow for classification of various types of forest-related sounds, which can provide information about different events, rather than focusing only on the task of wildfire detection.

C. Lack of a Common Framework Combining Acoustic and Environmental Information

The reviewed papers use different types of input and different target problems. Acoustic-based studies process audio signals and spectrogram-based representations [1], [2], whereas wildfire risk studies use meteorological, geographical, remote sensing, vegetation, and historical-fire information [3][8]. The reviewed literature does not demonstrate a unified framework of forest monitoring that integrates the two information sources. Acoustic analysis can provide information about the events detected using sound, whereas environmental analysis can provide information about the likelihood of a wildfire based on the current environmental conditions. A research opportunity exists to investigate a more generalized approach of forest monitoring that uses both sources of information.

D. Regional and Dataset-Specific Limitations

A number of the reviewed studies are developed and evaluated using data from specific geographical regions. Gao et al. [3] investigated wildfire risk in the Heihe area of Heilongjiang Province, China. Hiremath and Kannan [4] focused on the Odisha region, whereas Yang et al. [5] studied forest-fire prediction using remote sensing data in Indonesia. Sharma and Khanal [8] investigated spatial wildfire prediction using data from a specific study region, and Hernandez and Hoskins [9] focused on California's Central Valley. Such studies demonstrate the effectiveness of machine learning in specific geographical areas. However, the effectiveness of these approaches can be dependent on the specific characteristics of the region, the dataset, the environmental factors, and the evaluation strategy. As such, there is a need to further investigate the application of forest-monitoring approaches in different environmental conditions and with different datasets to assess their effectiveness and generalizability.

E. Challenges in Real-World Wildfire Prediction and Evaluation

The reviewed literature also suggests some of the challenges related to wildfire prediction in the real world. For instance, Apostolakis et al. [7] discussed the challenges of extreme class imbalance and spatial granularity in next-day wildfire prediction. Hernandez and Hoskins [9] demonstrated the differences between random and chronological data splitting, showing the importance of the evaluation strategy in wildfire prediction. The literature suggests that the performance of a given approach in one experimental setup does not necessarily reflect its performance in a different geographical area and a different temporal setting. As such, there is a need for wildfire-monitoring approaches to take these aspects into account when evaluating their performance.

F. Overall Research Gap

The main research gap identified from the review is the separate treatment of acoustic forest monitoring and environmental wildfire-risk prediction. Studies [1] and [2] demonstrate the effectiveness of acoustic signals for wildfire detection, and studies [3][9] illustrate the use of meteorological, geographical, remote sensing, historical, and machine learning approaches for wildfire prediction. However, these research directions are predominantly investigated separately, and acoustic-based wildfire detection studies generally focus on relatively narrow fire detection objectives, whereas environmental wildfire prediction studies do not primarily provide information about acoustic events. As such, a research opportunity exists to develop a more generalized approach of forest monitoring that integrates both sources of information with consideration of multiple types of acoustic forest-related events. This would capitalize on the strengths of the individual studies rather than asserting that acoustic wildfire detection or machine-learning-based wildfire prediction has not been investigated previously.

IV. PROPOSED SYSTEM

A. Problem Addressed

EcoGuard is proposed as an integrated forest-monitoring and wildfire-risk-intelligence system for addressing the major forest threats such as illegal deforestation, wildlife poaching, and wildfires. The system combines acoustic analysis and environmental wildfire-risk prediction with the output information on forest-related sound events and prevailing fire-risk conditions. Acoustic events such as chainsaw activity, gunshots, fire, vehicle sounds, woodchopping, fauna, and other environmental sounds can be classified, while meteorological and fire-weather parameters are used to estimate wildfire ignition and spread probability. The resulting information is presented through a centralized web-based monitoring platform that supports visualization, alerts, historical monitoring, geographical analysis, and report generation.

B. Overall Architecture

EcoGuard follows a multi-layer architecture consisting of frontend interface, backend services, machine-learning components, data-persistence layer, and external data services. The core of the system contains two complementary modules: the Acoustic Threat Detection Module and the Wildfire Risk Prediction Module. The acoustic module receives audio and processes it through signal-processing techniques before classification using a ResNet-18 convolutional neural network. The wildfire module processes meteorological and Fire Weather Index parameters and uses a Random Forest Regressor to estimate wildfire probability. The backend provides communication between the frontend, machine-learning models, database, analytics, and reporting components. MongoDB provides persistent storage, while an in-memory state buffer supports rapid access to recent information.

C. Acoustic Threat Detection Module

The Acoustic Threat Detection Module is responsible for identifying forest-related acoustic events from microphone input or uploaded audio files. The system accepts common audio formats and processes the incoming signal before classification. The acoustic pipeline uses digital signal processing to generate a Mel-spectrogram representation and extract additional acoustic characteristics. The resulting representation is supplied to a ResNet-18 CNN, which classifies the sound into one of the supported sound categories and produces a confidence score. The detected class is subsequently associated with a threat-severity level such as Critical, High, Medium, or Low.

D. Audio Preprocessing

The audio preprocessing stage converts the raw audio signal into a representation suitable for deep-learning classification. The incoming audio is first validated and converted into a standard 22,050 Hz mono signal. Fast Fourier Transform (FFT) is then applied using 2,048-point windows to obtain frequency-domain information. The frequency bins are mapped to the Mel scale to generate a 128 × 128 two-dimensional Mel-spectrogram. In addition, MFCC, RMS energy, and Zero Crossing Rate (ZCR) are extracted as acoustic features. The processed spectrogram is subsequently converted into the input tensor required by the acoustic classification model.

E. ResNet-18 Model

The acoustic classification component uses a ResNet-18 convolutional neural network implemented with

PyTorch. The generated 128 × 128 Mel-spectrogram is converted into a four-dimensional tensor with the shape [1, 1, 128, 128] and provided to the network. ResNet-18 learns acoustic patterns from the spectrogram representation and classifies the input into the supported sound categories. The model produces the predicted sound class and corresponding confidence value, after which the result is assigned a threat-severity level. The reported system performance for the acoustic model is 98. 8% validation accuracy with sub-50ms inference latency.

F. Wildfire Risk Prediction Module

The Wildfire Risk Prediction Module estimates wildfire ignition and spread probability from meteorological and fire-weather information. The module uses temperature, humidity, wind, rainfall, and derived Fire Weather Index parameters as input features. These variables are combined into a nine-feature input vector and supplied to a Random Forest Regressor. The model produces a continuous wildfire probability ranging from 0% to 100%, which is subsequently converted into a risk category. The resulting risk level can be displayed through the system's wildfire-risk interface and used for monitoring and reporting.

G. Random Forest Model and Fire-Weather Features

The wildfire-risk component uses a Scikit-Learn Random Forest Regressor trained using the UCI Algerian Forest Fires Dataset (ID 547). The model uses nine features consisting of temperature, humidity, wind, rain, Fine Fuel Moisture Code (FFMC), Duff Moisture Code (DMC), Drought Code (DC), Initial Spread Index (ISI), and Build-Up Index (BUI). The FWI-derived variables represent different aspects of fuel moisture, drought conditions, and potential fire spread. The predicted wildfire probability is classified into four risk levels: Low for probabilities below 35%, Medium for 35% to below 55%, High for 55% to below 75%, and Extreme for probabilities of 75% or higher. The reported system performance for the Random Forest model is an R^2 score of at least 0. 97 across the specified drought-season benchmarks.

H. Datasets

Two datasets are used for the machine-learning components of the proposed system. The FSC22 Forest Sound Dataset is used for the acoustic ResNet-18 model, while the UCI Algerian Forest Fires Dataset (ID 547) is used for the wildfire-risk Random Forest model. The two datasets support the separate acoustic-classification and wildfire-risk-prediction tasks within the proposed system.

I. Backend and APIs

The backend is implemented using FastAPI and acts as the communication layer between the frontend, machine-learning models, database, analytics services, and report-generation components. It contains controllers for acoustic classification, wildfire-risk prediction, analytics, report generation, and machine-learning management. RESTful APIs are provided for submitting audio for classification, requesting wildfire-risk predictions, retrieving historical acoustic and wildfire-risk records, obtaining dashboard and map data, generating reports, and managing model-training records. The backend also provides a health-check endpoint for monitoring service status.

I. Frontend

The frontend provides the user-facing interface for interacting with EcoGuard. It is developed using HTML5, CSS3, and Vanilla JavaScript, with Chart. js used for data visualization and Leaflet. js used for geographical mapping. The Web Audio API supports browser-based microphone capture and audio visualization. The interface provides pages for live monitoring, wildfire-risk prediction, weather information, dashboards, maps, alerts, reports, historical records, machine-learning information, settings, and user profiles. The live-monitoring interface provides waveform and frequency visualization while allowing audio analysis and display of classification results.

K. Database and Data Management

MongoDB is used as the primary persistent database for storing system information. The database maintains acoustic detection records, wildfire-risk predictions, user and role information, forest-zone and sensor-related information, and machine-learning training records. An in-memory state buffer is additionally used for rapid access to recent system information. This combination provides both persistent storage and fast access to operational monitoring data.

L. Report Generation

The system provides a reporting facility for presenting forest monitoring information in a structured form. Reports can include time-based trends, sound-category distributions, zone-wise risk comparisons, and historical information. The system supports CSV data export and PDF report generation. The backend uses ReportLab for programmatic generation of PDF reports, allowing monitoring information to be converted into downloadable forest-intelligence reports.

M. Complete System Workflow

The complete EcoGuard workflow consists of two complementary processing paths. In the acoustic path, audio is captured or uploaded, validated, resampled, transformed using FFT, converted into a Mel-spectrogram, and processed to obtain MFCC, RMS, and ZCR features. The resulting spectrogram is passed to the ResNet-18 CNN, which produces the sound classification and confidence score. The result is then assigned a severity level and made available for storage, dashboard visualization, and alerts. In the wildfire-risk path, meteorological parameters are processed together with Fire Weather Index components to form a nine-feature input vector. This vector is provided to the Random Forest Regressor, which generates wildfire probability and the corresponding risk tier. The resulting information is stored and presented through the dashboard and reporting components.

V. DATASET, METHODOLOGY AND EXPERIMENTAL RESULTS

A. Dataset Description

EcoGuard utilizes two distinct datasets related to its two machine-learning components: acoustic threat classification and wildfire-risk prediction. Acoustic processing uses FSC22 Forest Sound Dataset, while the wildfire-risk component uses the UCI Algerian Forest Fires Dataset (ID 547). The datasets are used independently for the two different learning tasks and provide the input data necessary for acoustic-event classification and wildfire-risk estimation. The FSC22 Forest Sound Dataset is associated with the ResNet-18 convolutional neural network that powers EcoGuard's acoustic classification feature. The dataset provides forest-related audio data representing different environmental and potentially threatening sound events.

The supported sound categories include chainsaw activity, gunshots, fire, vehicle sounds, woodchopping, fauna, and other environmental sounds. These categories make the dataset relevant to forest surveillance, since acoustic events can provide indications of human activities, environmental conditions, or potential threats within a monitored forest environment. The project documentation does not specify the number of samples, class distribution, or separate train/validation/test counts for this dataset. For wildfire-risk prediction, EcoGuard uses the UCI Algerian Forest Fires Dataset (ID 547). The dataset provides environmental and fire-weather information used to estimate wildfire probability. The model utilizes temperature, humidity, wind, rainfall, and Fire Weather Index-related parameters. The final input consists of nine features: Temperature, Humidity, Wind, Rain, FFMC, DMC, DC, ISI, and BUI. The features represent environmental conditions and different aspects of fuel moisture, drought conditions, and potential fire spread. The resulting model output is a continuous wildfire probability ranging from 0% to 100%, which is subsequently converted into predefined risk categories.

Table I: Description of the Datasets Used in EcoGuard

Data's et	Purpose	Input Type	Main Features/Cl asses	Model
FSC2 2 Forest Soun d Data's et	Acoustic threat classification	Audio	Forest and environment al sound categories	ResNe t-18
UCI Alger ian Forest Fires Data's et (ID 547)	Wildfire-risk prediction	Environmental/ weather data	Temperature, Humidity, Wind, Rain, FFMC, DMC, DC, ISI, BUI	Rando m Forest Regres sor

B. Data Preprocessing

The two machine-learning pipelines apply different preprocessing procedures due to the differing characteristics of their input data. The acoustic pipeline converts raw audio signals into a two-dimensional spectral representation, while the wildfire pipeline organizes environmental and fire-weather variables into a numerical feature vector. For acoustic

classification, the input can come from microphone capture or uploaded audio files. The incoming audio is validated and converted into a standard 22,050 Hz mono signal. Fast Fourier Transform (FFT) is then applied using 2,048-point windows to obtain frequency-domain information. The frequency information is mapped onto the Mel scale to produce a 128 × 128 two-dimensional Mel-spectrogram. Additional acoustic characteristics, namely Mel-Frequency Cepstral Coefficients (MFCC), Root Mean Square (RMS) energy, and Zero Crossing Rate (ZCR) are extracted during preprocessing. The processed spectrogram is then converted into the tensor representation required by the classification model. The Mel-spectrogram provides a compact representation of how the frequency content of an audio signal varies across time. This representation allows the convolutional neural network to process the acoustic information as a two-dimensional form and learn patterns associated with different forest sound categories. In EcoGuard, the generated 128 × 128 Mel-spectrogram is converted into a four-dimensional input tensor with the shape [1, 1, 128, 128] before being supplied to the ResNet-18 model. For wildfire-risk prediction, the environmental input is represented using nine features: [Temperature, Humidity, Wind, Rain, FFMC, DMC, DC, ISI, BUI]. The first four features describe meteorological conditions, while FFMC, DMC, DC, ISI, and BUI are Fire Weather Index components used to represent fuel moisture, drought conditions, and potential fire-spread characteristics. The nine variables are combined into a feature vector and supplied to the Random Forest Regressor for probability estimation.

C. Model Development

Acoustic Threat Classification using ResNet-18 The acoustic threat classification component uses a ResNet-18 convolutional neural network implemented with PyTorch. The model receives the processed Mel-spectrogram as its input representation and learns acoustic patterns corresponding to the supported forest sound categories. The use of a convolutional architecture enables the model to operate on the spatial structure of the spectrogram and identify relevant patterns within the time-frequency representation. In the EcoGuard pipeline, the preprocessed 128 × 128 Mel-spectrogram is converted into a four-dimensional tensor of shape [1, 1, 128, 128]. This tensor is supplied to ResNet-18, which produces the predicted sound class and its corresponding confidence value. The predicted class is then associated with a threat-severity level such as Critical, High, Medium, or Low. Thus, the model output is not confined to identifying an acoustic event but is also incorporated into the threat-monitoring process of the system. The project documentation reports a validation accuracy of 98. 8% for the acoustic classification model and an inference latency below 50ms. These values represent the documented validation and inference performance of the implemented acoustic model. Training hyperparameters such as the optimizer, learning rate, batch size, number of epochs, and loss-function configuration are not specified in the available project documentation and are therefore not included here.

Wildfire Risk Prediction using Random Forest

The wildfire-risk component uses a Random Forest Regressor implemented with Scikit-Learn. The model is intended to estimate a continuous wildfire probability from environmental and Fire Weather Index features. Random Forest is used in the implementation to process the combination of meteorological and fire-weather variables and produce the wildfire probability used by the monitoring system. The model receives the nine-feature vector consisting of temperature, humidity, wind, rain, FFMC, DMC, DC, ISI, and BUI. The Random Forest Regressor produces a wildfire probability between 0% and 100%. This numerical probability is subsequently mapped to four risk levels: Low for probabilities below 35%, Medium for probabilities from 35% to below 55%, High for probabilities from 55% to below 75%, and Extreme for probabilities of 75% or higher. The documented performance of the wildfire model is an R^2 score of at least 0. 97 across the specified drought-season benchmarks. The available documentation does not provide additional training hyperparameters or a detailed train-test split, and therefore these details are not assumed in the methodology.

D. Overall Methodology and Workflow

The overall EcoGuard system comprises two distinct but interconnected machine learning pipelines. The first pipeline is used to detect and classify forest-related sound events, and the second one is used to predict the likelihood of wildfires based on environmental and weather-related variables. The pipelines also have specific applications that allow

users to monitor the activity, set up alerts, store the data, and create reports. The first pipeline's input consists of microphone input or uploaded files in WAV or MP3 format. These files are then converted into a Mel-spectrogram, and the features, such as MFCCs, RMS Energy, and Zero-Crossing Rate, are extracted. Then the ResNet-18 CNN model is used to classify the sound events and extract threat-related information from the sound. The second machine learning pipeline uses nine weather- and fire-related features, including temperature, wind speed, humidity, and others, as input variables. The Random Forest Regressor is then used to predict the probability of wildfires and classify them by severity. In turn, this information is used to produce alerts, reports, and dashboards in the application.

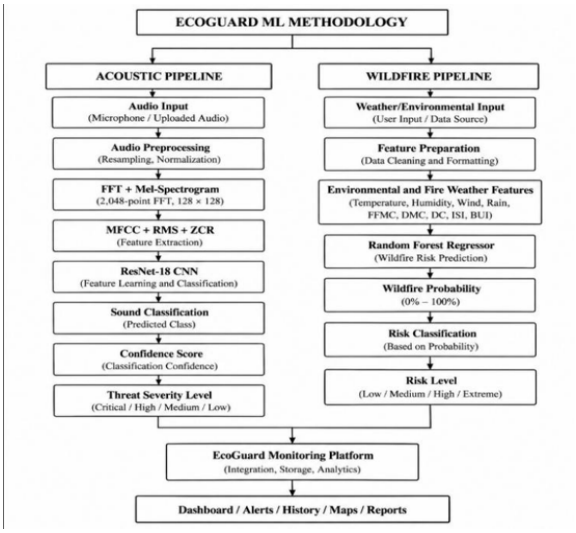


Fig: overall EcoGuard ML methodology showing acoustic threat classification and wildfire-risk prediction pipelines.

E. System Implementation

EcoGuard is implemented as a web-based software platform that integrates the machine-learning models with frontend, backend, database, visualization, and reporting components. The backend is implemented using FastAPI and provides the communication layer between the frontend, machine-learning models, database, analytics services, and reporting components. The backend contains functionality for acoustic classification, wildfire-risk prediction, analytics, report generation, and machine-learning management. The frontend is developed using HTML5, CSS3, and Vanilla JavaScript. Chart.js is used for analytical visualization, while Leaflet.js provides geographical mapping. The Web Audio API supports browser-based microphone capture and audio visualization. The interface includes live monitoring, wildfire-risk prediction, dashboards, maps, alerts, reports, historical records, and other system-management pages. The backend exposes RESTful APIs for submitting audio for classification, requesting wildfire-risk predictions, retrieving historical acoustic and wildfire-risk records, obtaining dashboard and map information, generating reports, and managing machine-learning records.

This API-based structure allows the frontend to communicate with the machine-learning components without directly executing the models within the user interface. MongoDB is used as the primary persistent database. It stores acoustic detection records, wildfire-risk predictions, user and role information, forest-zone and sensor-related information, and machine-learning training records. An in-memory state buffer is additionally used to provide rapid access to recent system information. The prediction outputs are subsequently incorporated into the monitoring interface. Acoustic results can include the detected sound class, confidence score, threat severity, and associated action information, while wildfire predictions provide the estimated probability and corresponding risk category. These results can be retained for historical analysis and presented through dashboards, maps, alerts, and reports.

F. Experimental Setup

The experimental evaluation takes into account the documented performance of the two machine-learning components. For acoustic threat classification, the reported metric is validation accuracy together with inference latency. The ResNet-18 model achieves a documented validation accuracy of 98. 8%, with reported inference latency below 50ms. Since the available documentation identifies this as validation accuracy, it is not treated as test accuracy. For wildfire-risk prediction, the Random Forest Regressor is evaluated using the R^2 metric. The documented performance is an R^2 score of at least 0. 97 across the specified drought-season benchmarks. The model generates a continuous probability value that is subsequently converted into Low, Medium, High, or Extreme risk levels according to the defined probability thresholds.

G. Results and Discussion

Table II: Results and Discussion of the EcoGuard System

Model	Metric	Result	Output
ResNet-18	Validation Accuracy	98. 8%	Sound class + confidence + severity
ResNet-18	Inference Latency	< 50ms	Acoustic prediction
Random Forest Regressor	R2	≥ 0. 97	Wildfire probability
Random Forest Regressor	Risk Levels	4 levels	Low, Medium, High, Extreme

The documented results of the two machine-learning components are summarized in Table II. The acoustic model demonstrates high validation performance while maintaining low reported inference latency. The wildfire model provides a high documented R^2 value for its regression task and converts its continuous output into interpretable risk categories. For acoustic threat classification, the 98. 8% validation accuracy indicates that the implemented ResNet-18 model performs strongly on the documented validation evaluation. The sub-50ms inference latency also indicates that individual acoustic predictions can be produced promptly enough to support the application's monitoring workflow. The classification output is further enhanced by the assignment of threat-severity levels, allowing the result to be used not only as a sound label but also as monitoring information. For wildfire-risk prediction, the reported R^2 value of at least 0. 97 indicates strong agreement between the model's regression output and the target used in the documented evaluation benchmarks. Rather than producing only a binary fire/no-fire decision, the model produces a continuous probability that is converted into four risk categories.

This allows the monitoring interface to represent different levels of wildfire risk and supports the presentation of risk information through dashboards and reports. The two models address different dimensions of forest monitoring and therefore complement one another. The acoustic model focuses on identifying sound-based events associated with forest activities and threats, whereas the Random Forest model focuses on environmental conditions associated with wildfire risk.

Their outputs are integrated at the application level, allowing EcoGuard to present acoustic threats and wildfire-risk information through a common monitoring platform. The results should be interpreted within the scope of the documented experiments. The reported model metrics demonstrate the performance of the implemented machine-learning components under their respective evaluation conditions; they do not by themselves establish performance under every possible real-world forest environment. The project therefore uses the model outputs as components of a broader software-based monitoring and intelligence platform rather than claiming universal field performance.

VI. CONCLUSION

The chosen design concept for forest monitoring refers to an AI-assisted system, which utilizes the ResNet-18 CNN for addressing the acoustic threat detection problem and trains the Random Forest Regressor for wildfire risk prediction. In developing the proposed solution, specific attention was paid to the use of the ResNet-18 CNN architecture with the application of Fast Fourier Transform and Mel-spectrogram audio preprocessing techniques for transforming raw audio data into spectrograms. In turn, the forest monitoring system was trained on a set of 27 classes of sounds with a reported 98. 8% accuracy on the validation set. Concerning the wildfire risk prediction algorithm, the proposed Random Forest Regressor leverages the set of meteorological parameters and Canadian Fire Weather Index to predict the probability of a wildfire occurrence with report R2 score of of ≥ 0.97.

Altogether, by integrating the AI-driven solutions for acoustic threat detection and wildfire-risk prediction, along with the dashboard, alerts, reports, and historical data, the proposed forest monitoring system offers a unified software interface for addressing the problem. Nevertheless, the accuracy of the proposed solution may depend on the quality and coverage of the training set, audio-specific characteristics, and the relevance of the wildfire data set, which requires further validation in the actual forest environment.

VII. FUTURE SCOPE

- Expansion of Acoustic Training Data: The acoustic

dataset can be expanded with additional forest sounds recorded under different weather, distances, background-noise levels, and environmental conditions. This could improve the robustness of classification in diverse monitoring environments.

- Detection of Unseen and Rare Threats: Future versions can investigate methods for identifying rare or previously unseen acoustic events that are not sufficiently represented in the existing classification categories. This can make the monitoring system more adaptable to new forest threats.
- Larger and More Diverse Wildfire Datasets: The wildfire-risk model can be trained and evaluated using larger datasets covering multiple geographical regions and longer time periods. This would help assess and potentially improve the model's generalizability across different environmental conditions.
- Advanced Multimodal Risk Assessment: Future development can incorporate additional environmental information and investigate more advanced methods for jointly analyzing acoustic and environmental observations. This could provide a richer representation of forest conditions and wildfire risk.
- Real-World Field Validation: The system can be evaluated through controlled field experiments and long-term monitoring in actual forest environments. Such validation would help measure acoustic classification and wildfire-risk prediction performance under real environmental conditions.
- Explainable AI: Explainability techniques can be incorporated to provide users with more understandable information about why a particular acoustic class or wildfire-risk prediction was generated. This can improve transparency and support better interpretation of AI-generated results.
- Scalability for Larger Forest Regions: Future versions can improve the system's scalability to support larger numbers of monitoring zones, audio sources, and environmental data streams. This would allow the platform to be extended from individual monitoring areas to larger forest regions.

REFERANCES

[1].H.-T. Huang, A. R. J. Downey, and J. D. Bakos, “Audio-based wildfire detection on embedded systems,” Electronics, vol. 11, no. 9, Art. no. 1417, 2022.

[2].R. Damasevicius et al., “A hybrid machine learning model for forest wildfire detection using sounds,” in Proc. Federated Conference on Computer Science and Information Systems (FedCSIS), pp. 99–106, 2024.

[3].C. Gao, H. Lin, and H. Hu, “Forest-fire-risk prediction based on Random Forest and backpropagation neural network of Heihe area in Heilongjiang Province, China,” Forests, vol. 14, no. 1, Art. no. 170, 2023.

[4].H. Hiremath and S. R. Kannan, “Integrated anomaly detection and early warning system for forest fires in the Odisha region,” Atmosphere, vol. 15, Art. no. 1284, 2024.

[5].S. Yang, M. Lupascu, and K. S. Meel, “Predicting Forest fire using remote sensing data and machine learning,” 2021.

[6].A. Apostolakis, S. Girtsou, G. Giannopoulos, N. S. Bartsotas, and C. Kontoes, “Estimating next day’s forest fire risk via a complete machine learning methodology,” Remote Sensing, vol. 14, Art. no. 1222, 2022.

[7].F. Huot et al., “Deep learning models for predicting wildfires from historical remote-sensing data,” in Proc. NeurIPS 2020 AI for Humanitarian Assistance and Disaster Response Workshop, 2020.

[8].K. Hernandez and A. B. Hoskins, “Machine learning algorithms applied to wildfire data in California’s Central Valley,” Trees, Forests and People, vol. 15, Art. no. 100516, 2024.

[9].S. Sharma and S. Khanal, “Forest fire prediction: A spatial machine learning and neural network approach,” Fire, vol. 7, no. 6, Art. no. 205, 2024.