

# MVeLMA: Multimodal Vegetation Loss Modeling Architecture for Predicting Post-fire Vegetation Loss

Meenu Ravi  
ravim@vt.edu  
Virginia Tech  
Alexandria, Virginia, USA

Shailik Sarkar  
shailik@vt.edu  
Virginia Tech  
Alexandria, Virginia, USA

Yanshen Sun  
yansh93@vt.edu  
Virginia Tech  
Alexandria, Virginia, USA

Vaishnavi Singh  
vs699@georgetown.edu  
Georgetown University  
Washington D.C., USA

Chang-Tien Lu  
clu@vt.edu  
Virginia Tech  
Alexandria, Virginia, USA

## Abstract

Understanding post-wildfire vegetation loss is critical for developing effective ecological recovery strategies and is often challenging due to the extended time and effort required to capture the evolving ecosystem features. Recent works in this area have not fully explored all the contributing factors, their modalities, and interactions with each other. Furthermore, most research in this domain is limited by a lack of interpretability in predictive modeling, making it less useful in real-world settings. In this work, we propose a novel end-to-end ML pipeline called MVeLMA (Multimodal **V**egetation **L**oss **M**odeling **A**rchitecture) to predict county-wise vegetation loss from fire events. MVeLMA uses a multimodal feature integration pipeline and a stacked ensemble-based architecture to capture different modalities while also incorporating uncertainty estimation through probabilistic modeling. Through comprehensive experiments, we show that our model outperforms several state-of-the-art (SOTA) and baseline models in predicting post-wildfire vegetation loss. Furthermore, we generate vegetation loss confidence maps to identify high-risk counties, thereby helping targeted recovery efforts. The findings of this work have the potential to inform future disaster relief planning, ecological policy development, and wildlife recovery management.

## CCS Concepts

• **Computing methodologies** → **Boosting**; • **Applied computing** → *Environmental sciences*.

## Keywords

Wildfire, Gaussian Process, Spatiotemporal Forecasting, Uncertainty Estimation, Multimodal, Vegetation Loss, Ecological Modeling, Regression

## ACM Reference Format:

Meenu Ravi, Shailik Sarkar, Yanshen Sun, Vaishnavi Singh, and Chang-Tien Lu. 2025. MVeLMA: Multimodal Vegetation Loss Modeling Architecture for Predicting Post-fire Vegetation Loss. In *The 33rd ACM International*

*Conference on Advances in Geographic Information Systems (SIGSPATIAL '25)*, November 3–6, 2025, Minneapolis, MN, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3748636.3764183>

## 1 Introduction

The prevalence of wildfires has increased significantly in recent years. Driven by climate change and other environmental factors, these natural disasters can disrupt the ecological balance, destroy wildlife habitats, devastate agricultural areas, and contribute to global warming by releasing harmful gases into the atmosphere [3, 4, 6, 17, 35]. Identifying affected areas and monitoring land coverage are essential for ecological management and informing environmental and land-use policy decisions at various scales. To establish effective ecological recovery strategies and manage land area more effectively, it is critical to understand how fire affects vegetation and to identify spatiotemporal patterns in post-fire vegetation loss.

In recent years, there has been an increasing interest in predicting wildfire-related damage, particularly burned areas, wildlife loss, and infrastructure destruction [51]. The use of remote sensing data for vegetation health monitoring, specifically through indicators such as vegetation indices have become prevalent. Particularly, the Normalized Difference Vegetation Index (NDVI) is frequently used to evaluate vegetation health, study agricultural productivity, and monitor vegetation responses to climate-related events [9, 12, 14, 32, 36, 42, 51, 52, 59–61, 68].

Existing literature primarily focuses on time series forecasting approaches and machine learning (ML) models applied to remote sensing indices such as NDVI. Vegetation forecasting research has predominantly focused on leveraging meteorological variables (e.g., precipitation, temperature, and climate data), along with soil properties, crop characteristics, and remote sensing indices such as NDVI [51]. A wide range of frameworks have been developed to address spatial and temporal complexities involved in vegetation modeling. Particularly, deep learning algorithms such as Deep Neural Networks (DNN), Convolutional Neural Networks (CNN), Recurrent Neural Networks (RNN), and CNNLSTM have been used due to their ability to extract dependencies from large datasets. These methods have been applied to various use cases, including drought prediction, landslide detection, crop yield estimation, drought damage assessment, plant disease detection, and wildfire spread modeling




This work is licensed under a Creative Commons Attribution 4.0 International License. *SIGSPATIAL '25*, Minneapolis, MN, USA

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2086-4/2025/11

<https://doi.org/10.1145/3748636.3764183>



**Figure 1: Overview of the proposed pipeline architecture for vegetation loss prediction, combining multimodal data sources, temporal encoding, uncertainty estimation, and predictive modeling.**

[2, 5, 14, 29, 36, 41, 45, 52, 65]. However, a clear research gap remains in forecasting vegetation loss specifically caused by wildfires. Additionally, limited work exists that leverages land cover features and variables beyond traditional meteorological inputs, such as solar radiation, fire-specific temporal data, and vapor pressure.

Traditional time-series forecasting models, while efficient for stationary time series, struggle to handle non-linear relationships present in multivariable and remote-sensing data. Similarly, traditional ML approaches can find it challenging to capture the spatiotemporal dynamics and high-dimensionality of such data. Furthermore, these works largely ignore interpretability through uncertainty quantification, which remains significant in real-world applications of such predictive tools.

To address these limitations, we propose an end-to-end, efficient, uncertainty-quantifying and scalable multimodal framework for predicting county-level vegetation loss following wildfire events, called MVeLMA.

Our approach focuses on modeling nine dynamic meteorological features along with twenty-seven enriched and static topographical features collected over multiple years. This framework combines deep neural recurrent networks for capturing temporal patterns, Gaussian Process Regression (GPR) for modeling spatial correlations that produce uncertainty-aware predictions, and an ensemble decision tree-based residual network to refine these predictions by incorporating static topographic and enriched features. To the best of our knowledge, no previous studies have implemented such a

multimodal approach for wildfire-related vegetation loss prediction. Therefore, this study serves as a novel solution to the primary objective of our study, which is to predict vegetation loss following wildfire events. The key contributions of our work can be summarized as follows:

- **Proposing a multimodal ensemble framework for probabilistic vegetation loss prediction.** Existing approaches to wildfire impact prediction often rely on fewer than fifteen meteorological or static environmental variables. The proposed MVeLMA framework integrates heterogeneous data that captures both temporal weather patterns and static geographic variables into a multimodal architecture. It leverages temporal encoding, uncertainty-aware regression, and residual learning to provide vegetation loss forecasts with confidence intervals.
- **Integrating a diverse set of feature variables:** MVeLMA combines nearly thirty-six dynamic and static variables, including meteorological trends, vegetation indices, and topographical features. This cross-modal design allows the model to capture granular regional differences and their influence on post-fire vegetation response.
- **Extensively conducting empirical performance evaluations:** The proposed model was tested across more than fifty counties in a wildfire-prone state, demonstrating consistent improvement over baseline methods. It achieved an approximate 10% increase in  $R^2$  and an 11% reduction in NRMSE,

highlighting the advantages of the multimodal architecture and uncertainty modeling.

- **Generating post-wildfire risk maps:** This study produces county-level vegetation loss predictions with uncertainty estimates. These maps identify high-risk counties, enabling targeted post-fire recovery, resource allocation, and policy planning.

## 2 Related Works

The application of statistical modeling, machine learning (ML), and deep learning (DL) techniques in wildfire prediction and management has seen significant advancements. This section reviews recent literature, categorizing studies based on their methodological approaches.

### 2.1 Statistical and ML Methods

Traditional statistical and ML methods remain prevalent in wildfire prediction [1, 31, 37]. Joseph et al. [34] apply Bayesian models to estimate extreme wildfire sizes by modeling upper quantiles from spatiotemporal predictors. Malik et al. [43] utilize random forest models to predict wildfire risks by analyzing comprehensive datasets, including powerlines, terrain, and vegetation, to improve spatial and temporal accuracy. Gizatullin et al. [20] construct linear models between fire risk and multiple factors. Castel-Clavera et al. [4] design region-specific logistic regression model. Zhang et al. [69] present a two-part model combining Integrated Nested Laplace Approximation (INLA) and Stochastic Partial Differential Equations (SPDE) to jointly model wildfire counts and burnt areas. Koh et al. [40] model wildfire ignitions using marked point processes. Illarionova et al. [27] combine random forest and XGBoost. Schmitt et al. [56] emphasize the role of ecological factors in fire susceptibility. Xu et al. [64] employ spatiotemporal prediction and a non-parametric dynamic threshold. Ji et al. [32] infer wildfire events from satellite images with a simple change-point analysis procedure.

### 2.2 DL Methods

**Deep Learning Approaches for Wildfire Prediction** Several studies have leveraged deep learning architectures to enhance wildfire prediction accuracy [14]. Bhowmik et al. [2] developed a multimodal wildfire prediction and early-warning system using a U-Convolutional-LSTM neural network, achieving over 97% accuracy by integrating environmental, meteorological, and geological data. Similarly, Burnt-Net [57] utilized post-fire Sentinel-2 imagery with a deep learning morphological neural network for accurate burned area mapping. Deng et al. [5] employed an attention mechanism and a graph convolution network for spatiotemporal daily wildfire risk prediction by mining global and local dependencies. CNNs are commonly used to infer fire impact from satellite images [51]. Long Short-Term Memory (LSTM) networks are widely used for fire forecasting with time series data [41, 52, 59]. Conv-LSTM-based methods [36, 42, 45] are also popular methods in spatiotemporal data modeling. Yao et al. [65] design a 3D-CNN with LSTM. CNN-BiLSTM [44] performs daily spread prediction with remote sensing and weather data. Wang et al. [60] add attention mechanism (AM) on top of CNN-BiLSTM. Jadouli et al. [28] introduce a modular

transformer-based deep learning architecture with pretrained internal “world model” layers. Yu et al. [67] develop a spatiotemporal transformer model to predict PM2.5 levels in wildfire-prone regions. Röscher et al. [53] introduce a GNN-based framework to model wildfire spread across Europe using spatial node embedding and dynamic temporal graphs, allowing generalization across locations. Parellada et al. [49] leverage a BiLSTM with a GNN.

### 2.3 Mixed Models and Data Sources

**Mixture of Models and Multi-Modal Data Integration** Integrating deep learning and machine learning models, as well as multi-modal spatiotemporal data [63] has been a focus in recent studies. For instance, Jadouli et al. [30] propose a modular forecasting pipeline combining deep learning, ensemble methods, and transfer learning. Huot et al. [26] provide a dataset for wild fire spread prediction with remote sensing data while the dataset WildfireSpreadTS [17] includes multi-modal temporal sequences. Kabir [35] utilizes both ARIMA and LSTM. Jadouli et al. [28] fuse CNNs, LSTMs, ensembles, and transfer learning into a unified wildfire forecasting framework. Zhu et al. [70] carefully process and identify important factors for defoliation forecasting. Dong et al. [10] emphasize localized modeling using topographic and meteorological factors. Khanmohammadi [37] incorporate vegetation type, slope, and meteorological data. Ghorbanzadeh et al. [18] involve field GPS survey data. Zhang et al. [68] identify climate and human activity as key spatiotemporal drivers. Deng et al. [6] incorporate rainfall, slope, vegetation, and historical ignitions. Michail [46] performs seasonal fire prediction with vegetation dryness, past fires, and climate data. Parellada et al. [49] introduce terrain and wind direction dynamics. Moghim et al. [47] highlight regional variability and the need for localized modeling. CREDS [33] is designed to handle wildfires with Unmanned Aerial Vehicles (UAVs).

While prior research demonstrates advancements in spatiotemporal modeling, most approaches either focus on deterministic predictions or rely on isolated data sources. In contrast, MVeLMA integrates deep sequential encoding, uncertainty quantification, and multimodal inputs to produce region-specific and confidence-aware predictions of post-fire vegetation loss.

## 3 Methodology

### 3.1 Preliminaries

In this work, we aim to build a predictive model that not only focuses on predicting the binary event of wildfire but also its effects. To make the application more suitable for real-world use cases, we also focus on incorporating spatially coherent uncertainty estimation to guide policy-level interventions and optimize resource allocation during emergencies. Therefore, we use vegetation loss as an indicator of the impact of the fire and formulate it as a spatiotemporal multivariate time series forecasting problem. The problem is formally defined as follows:

**Input:**

- A set of county-level wildfire events  $F = \{f_1, f_2, f_3, \dots, f_n\}$  where each  $f_i$  represents a wildfire event within the respective county with spatial and temporal information.
- A 30-day multivariate time series prior to each fire event  $f_i$  represented by  $X^T$ . This 3-D array is of shape  $N \times T \times W$

**Table 1: Summary of data variables used in the study**

| Data Features          | Units            | Sources                      | Time Range | Usage  | Importance (%) |
|------------------------|------------------|------------------------------|------------|--------|----------------|
| Average Temperature    | °C               | Univ. of Idaho (EPSCOR) [39] | 2018–2024  | Input  | 5.30           |
| Precipitation          | mm               | Univ. of Idaho (EPSCOR)      | 2018–2024  | Input  | 8.48           |
| Relative Min Humidity  | %                | Univ. of Idaho (EPSCOR)      | 2018–2024  | Input  | 14.75          |
| Relative Max Humidity  | %                | Univ. of Idaho (EPSCOR)      | 2018–2024  | Input  | 4.87           |
| Solar Radiation        | W/m <sup>2</sup> | Univ. of Idaho (EPSCOR)      | 2018–2024  | Input  | 3.68           |
| Minimum Temperature    | °C               | Univ. of Idaho (EPSCOR)      | 2018–2024  | Input  | 8.19           |
| Maximum Temperature    | °C               | Univ. of Idaho (EPSCOR)      | 2018–2024  | Input  | 4.09           |
| Vapor Pressure Deficit | kPa              | Univ. of Idaho (EPSCOR)      | 2018–2024  | Input  | 3.88           |
| Wind Speed             | m/s              | Univ. of Idaho (EPSCOR)      | 2018–2024  | Input  | 5.09           |
| Elevation              | m                | USGS SRTMGL1 v003 [13]       | 2018       | Input  | 6.16           |
| Land coverage          | 0–1 ratio        | MODIS MCD12Q1 v6.1 [15]      | 2018       | Input  | 25.68          |
| Fire Events            | dates            | MODIS C6.1 Fire Archive [19] | 2018–2024  | Input  | 14.75          |
| NDVI                   | scaled           | MODIS MOD13Q1 v6.1 [8]       | 2018–2024  | Target | –              |

**Table 2: Data sources and respective spatial/temporal resolutions.**

| Source                                   | Resolution              |
|------------------------------------------|-------------------------|
| GRIDMET (Abatzoglou, 2013)               | 4 km, daily             |
| SRTMGL1 v003 (Farr et al., 2007)         | 30 m                    |
| MODIS MCD12Q1 v6.1 (Friedl et al., 2019) | 500 m, annual           |
| MODIS MOD13Q1 v6.1 (Didan, 2015)         | 250 m, 16-day composite |
| MODIS MCD14ML (FIRMS)                    | 1 km, daily             |

where  $N$  represents the number of fire events,  $T = 30$ , the 30-day sequence of meteorological history, and  $W = 9$ , which represents the number of daily meteorological features.

- Enriched features represented by  $X^S$ . This 2-D array is of shape  $N \times S$  where  $N$  represents the number of fire events and  $S = 27$ , which represents the number of land coverage features and fixed features.

Given  $X = \{x_1, x_2, \dots, x_N\}$ , Let  $\mathbf{x}_i^T \in \mathbb{R}^{T \times W}$  denote the multivariate time series for  $T = 30$  days and  $W$  meteorological features, and  $\mathbf{x}_i^S \in \mathbb{R}^S$  represent  $S$  enriched spatial features (e.g., land cover, elevation). The complete input is given by:

$$\mathbf{x}_i = \left( \mathbf{x}_i^T, \mathbf{x}_i^S \right)$$

**Target Variable:** Vegetation loss is measured using NDVI values. NDVI loss values for each fire event  $f \in F$ , measured as  $y = V_{abf} - V_{maf}$  where  $V_{abf}$  refers to the mean NDVI for 30 days before fire event  $f$  and  $V_{maf}$  refers to the minimum NDVI within the 30 days after the fire event.

**Problem formulation:** We formulate post-fire vegetation loss prediction as a supervised regression problem over a multiscale feature space that captures both temporal dynamics and spatial heterogeneity. Let  $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$  denote our dataset, where each  $\mathbf{x}_i$  encodes spatial location, 30-day pre-fire meteorological and

topographical time series, and enriched county-level geophysical descriptors for wildfire event  $f_i$ , and  $y_i \in \mathbb{R}$  denotes the observed NDVI loss.

#### Objective:

- Maximize the accuracy and precision of the stacked time series model for estimating post-fire vegetation loss
- Calculate and aggregate coherent uncertainty measures at the county-level to support identification of at-risk and vulnerable regions.

## 3.2 Model development

We adopt a multimodal data integration pipeline that couples with a stacked modeling approach in which model integration is based on the modality of the data and the latent feature space. This stacked architecture leverages the strength of each component for the various data modalities (i.e., temporal, spatial, and static).

The overall workflow is illustrated in Figure 1. The model is a stacked model comprising an encoding layer, GPR, and a Random Forest (RF) regression layer.

**3.2.1 Integrating Temporal Weather Data through Recurrent Sequential Modeling.** Long Short-Term Memory (LSTM) is a type of Recurrent Neural Network (RNN), known for its efficacy in capturing long-term temporal dependencies [23]. Like LSTM, BiLSTM is effective at modeling long-term dependencies in time series data [24], making it effective for incorporating 30-day weather sequences preceding fire events. However, BiLSTM extends the traditional LSTM architecture, allowing the model to capture temporal dependencies by processing input sequences in both forward and backward directions. This allows the model to identify vegetation patterns that may not be strictly causal or unidirectional. This bidirectional context improves the model’s ability to capture the temporal relationship between fire and climate features, considering the influence of both preceding and succeeding weather patterns on fire conditions.

This layer processes a 3-D tensor input with dimensions  $N \times T \times W$ , where  $N$  is the number of fire events in the training set,  $T$  is the

30-day sequence length, and  $W$  is the number of weather features (10) at each time step. To extract temporal signatures from  $\mathbf{x}_i^T$ , the Bidirectional LSTM (BiLSTM) network computes a sequence of hidden states  $\{\mathbf{h}_{it}\}_{t=1}^T$  in both forward and backward directions. The attention mechanism helps the model focus on the most important features in the data and the most informative time steps. Therefore, a context vector  $\mathbf{c}_i \in \mathbb{R}^{d_h}$  is computed via attention weights  $\alpha_{it}$ :

$$\mathbf{c}_i = \sum_{t=1}^T \alpha_{it} \mathbf{h}_{it}, \quad \alpha_{it} = \frac{\exp(e_{it})}{\sum_{k=1}^T \exp(e_{ik})}, \quad e_{it} = \text{score}(\mathbf{h}_{it}) \quad (1)$$

This context vector is then projected to a latent representation  $\mathbf{z}_i \in \mathbb{R}^D$  via a dense layer:

$$\mathbf{z}_i = \phi(\mathbf{W}\mathbf{c}_i + \mathbf{b}) \quad (2)$$

where  $D = 20$  is the latent dimension, and  $\phi$  is a non-linear activation. This latent vector is an encoding of the weather sequence, structured to capture temporal trends involved in vegetation loss.

**3.2.2 Modeling Spatial Dependency through Multivariate Gaussian Process.** GPR is a type of non-parametric probabilistic method for regression and time-series forecasting problems due to its flexibility in quantifying predictive uncertainty [21]. It defines prior probability distributions over latent functions and uses Bayes' Theorem to calculate the posterior distribution given observed data. The prior distribution is characterized by the regression mean function and a covariance matrix. The covariance matrix is developed by a kernel function that encodes information about the shape, periodicity, and structure that the function is expected to have. For our pipeline, we take the output of the previous module, i.e., the latent vector  $\mathbf{z}_i$ , and concatenate it with enriched features  $\mathbf{x}_i^S$  to form  $\tilde{\mathbf{x}}_i = [\mathbf{z}_i; \mathbf{x}_i^S] \in \mathbb{R}^{D+S}$ . The GPR models ground truth vegetation loss values ( $y_i$ ) as a distribution over functions in the following way:

$$f(\cdot) \sim \mathcal{GP}(m(\cdot), k(\cdot, \cdot)) \quad (3)$$

where  $m(\cdot)$  is the mean function (set to constant) and  $k(\cdot, \cdot)$  is the kernel function. In our experiments, we explore multiple kernel choices and their combination:

- **RBF kernel (baseline):**

$$k_{\text{RBF}}(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j) = \sigma^2 \exp\left(-\frac{\|\tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_j\|^2}{2\ell^2}\right) \quad (4)$$

The RBF kernel is widely used in prediction and regression problems because it effectively maps the data onto an infinite-dimensional feature space, making it flexible. This distance function-based kernel is a more compactly supported kernel function, thereby reducing the computational complexity of the training process [58].

- **Matérn kernel:**

$$k_{\text{Mat}}(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j) = \sigma^2 \cdot \frac{2^{1-\nu}}{\Gamma(\nu)} \left( \sqrt{2\nu} \frac{\|\tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_j\|}{\ell} \right)^\nu K_\nu \left( \sqrt{2\nu} \frac{\|\tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_j\|}{\ell} \right) \quad (5)$$

with  $\nu = 2.5$ , allowing for a balance between smoothness and flexibility.

- **Composite kernel (Matérn + Periodic):**

$$k_{\text{comp}} = k_{\text{Mat}} + k_{\text{Per}}, \quad k_{\text{Per}}(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j) = \exp\left(-\frac{2 \sin^2(\pi \|\tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_j\|/p)}{\ell^2}\right) \quad (6)$$

The composite kernel was chosen to augment the Matérn kernel with the periodic kernel's ability to explicitly model data that exhibits seasonality.

Primarily, the GPR module produces two outputs for each input  $\tilde{\mathbf{x}}_i$ . First, the posterior predictive mean  $\mu_i$ , and second, the predictive variance  $\sigma_i^2$ . The mean  $\mu_i$  represents the model's expected value of NDVI loss for the fire event  $i$ , which incorporates both temporal dynamics (via  $\mathbf{z}_i$ ) and spatial features (via  $\mathbf{x}_i^S$ ). The predictive variance represents epistemic uncertainty. These outputs are computed using standard Gaussian Process posterior inference equations conditioned on the training data  $\mathbf{X}$  for a test location  $\tilde{\mathbf{x}}$ :

$$\mu^* = k(\tilde{\mathbf{x}}^*, \mathbf{X})^\top [K + \sigma_n^2 \mathbf{I}]^{-1} \mathbf{y} \quad (7)$$

$$\sigma^{*2} = k(\tilde{\mathbf{x}}^*, \tilde{\mathbf{x}}^*) - k(\tilde{\mathbf{x}}^*, \mathbf{X})^\top [K + \sigma_n^2 \mathbf{I}]^{-1} k(\tilde{\mathbf{x}}^*, \mathbf{X}) \quad (8)$$

where,  $\mu^*$  is the Predictive mean of NDVI loss for test input  $\tilde{\mathbf{x}}^*$ ;  $\sigma^{*2}$  denotes the predictive variance (uncertainty) for test input  $\tilde{\mathbf{x}}_i$ ;  $k(\tilde{\mathbf{x}}^*, \mathbf{X})$  is the Kernel vector between test input  $\tilde{\mathbf{x}}_i$  and all training inputs  $\mathbf{X}$ ;  $k(\tilde{\mathbf{x}}^*, \tilde{\mathbf{x}}^*)$  is the Kernel evaluated at the test input with itself (i.e., prior variance at  $\tilde{\mathbf{x}}^*$ );  $K = k(\mathbf{X}, \mathbf{X})$  represents the covariance Kernel matrix (Gram matrix) of all training inputs;  $\sigma_n^2$ : Noise variance hyperparameter, learned during training;  $\mathbf{I}$ : Identity matrix of shape  $N \times N$  and  $\mathbf{y}$ : Vector of ground truth NDVI loss values for the training data.

**3.2.3 Random Forest Regression for Residual Connection and Robustness:**

This module addresses the limitation of deep neural networks' deterministic predictions and limited capacity of a probabilistic learner to model high-frequency interaction in low-data regions. We concatenate the posterior mean from the GPR with the latent vector derived from the Bi-LSTM layer to construct the input for an RF module.

$$\mathbf{x}_i^{\text{RF}} = [\mathbf{z}_i; \mu_i^*] \in \mathbb{R}^{D+1} \quad (9)$$

RF is an ensemble model consisting of multiple decision trees that learn the patterns from the combined input and output a final vegetation loss prediction for each test sample by calculating the collective average of the decision trees. The incorporation of RF enables the model to be generalizable to spatially heterogeneous regions where Gaussian Process may underperform. Furthermore, it makes the model robust to overfitting and capable of performing implicit feature selection. Additionally, RF has lower sensitivity to hyperparameters, requiring minimal fine-tuning to achieve strong performance compared to architectures such as multi-layer perceptrons (MLP) and support vector machines (SVM) [66]. Lastly, it serves as a computationally efficient refinement for final prediction by training it as a residual learner.

$$\hat{y}_i = \text{RF}(\mathbf{x}_i^{\text{RF}}) \quad (10)$$

The final vegetation loss predictions from the RF layer are validated against the observed vegetation loss. The results of the validation will indicate how accurately the hybrid stacked model performs.

**Algorithm 1** MVeLMA Joint Optimization Model

---

```

1: for each fire event  $f_i$  do
2:   Get temporal sequence  $\mathbf{x}_i^T$ , enriched features  $\mathbf{x}_i^S$ , and latent
   features vector  $\mathbf{z}_i$ 
3:   Compute hidden states:  $\text{BiLSTM}(\mathbf{x}_i^T) \rightarrow \{\mathbf{h}_{it}\}_{t=1}^T$ 
4:   Compute attention weights  $e_{it} = \text{Linear}(\mathbf{h}_{it})$ 
5:   Normalize attention weights:  $\text{softmax}(e_{it}) \rightarrow \alpha_{it}$ 
6:   Compute context vector  $\mathbf{c}_i = \sum_{t=1}^T \alpha_{it} * \mathbf{h}_{it}$   $\triangleright$  Weighted
   sum of hidden states
7:   Encode latent vector:  $\text{Linear}(\mathbf{c}_i) \in \mathbb{R}^D \rightarrow \mathbf{z}_i$ 
8:   Train GPR on  $\mathbf{z}_i$ 
9:   for epoch  $e < 500$  do
10:    Compute GP posterior mean  $\mu_i$ 
11:    Backpropagate and optimize:  $\mathcal{L}_e = -\text{MLL}(\mu_i, y_i)$ 
12:    if early stopping criteria is met then  $\triangleright$  uses patience
    and minimum loss as criteria
13:      break
14:    end if
15:  end for
16:  Predict GPR posterior mean:  $\mu_i$  and variance  $\sigma_i^2$ 
17:  Stack features:  $\mathbf{x}_i^{RF} \leftarrow [\mathbf{z}_i, \mathbf{x}_i^S, \mu_i]$ 
18:  for each RF iteration do  $\triangleright$  Train RF to correct GP
    prediction errors
19:    Train RF on  $\mathbf{x}_i \rightarrow \text{residual}$ 
20:    Predict residual
21:    Final prediction:  $\hat{y}_i = \text{RF}(\mathbf{x}_i^{RF})$ 
22:  end for
23:  Evaluate: MAE,  $R^2$ , MAPE, NRMSE
24: end for

```

---

**3.2.4 Joint Optimization Strategy.** While each component in our architecture incorporates different modalities (temporal sequences, spatial metadata, etc.), we ensure coherent integration of these modules through a sequential stacking strategy. Algorithm 1 presents pseudocode summarizing the three phases of the training pipeline.

*Loss Function:* We optimize the GPR layer on the latent features produced by the BiLSTM encoder. The loss function is modeled by the negative log marginal likelihood (MLL) of the GPR posterior mean.

$$L_{\text{GPR}, y_{\text{BiLSTM}}} = -\frac{1}{2} y^T ( [K + \sigma_n^2 \mathbf{I}]^{-1} y ) + \frac{1}{2} \log |K + \sigma_n^2 \mathbf{I}| + \frac{n}{2} \log(2\pi) \quad (11)$$

This loss function is minimized when the BiLSTM is trained with the GPR, as its parameters are updated through backpropagation based on the GPR’s marginal likelihood loss using an optimizer. The GPR posterior mean, BiLSTM-encoded latent vectors, and enriched features are concatenated and passed to the RF model. This model is trained to minimize the residual error between the GPR posterior mean and the ground truth value, using a mean squared error loss function.

$$L_{\text{RF}} = \frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y_i)^2 \quad (12)$$

**Table 3: Enriched features used in the RF model alongside learned latent representations.**

| Feature Name       | Description                                              |
|--------------------|----------------------------------------------------------|
| fire_duration_days | Length the fire lasted in days                           |
| fire_month         | Month the fire started                                   |
| fire_dayofyear     | Day of the year the fire started                         |
| ndvi_7d_slope      | NDVI slope over the 7 days before the fire               |
| ndvi_7d_std        | NDVI standard deviation over the 7 days before the fire  |
| ndvi_14d_slope     | NDVI slope over the 14 days before the fire              |
| ndvi_14d_std       | NDVI standard deviation over the 14 days before the fire |
| ndvi_before_slope  | NDVI slope over the 30 days before the fire              |
| ndvi_before_std    | NDVI standard deviation over the 30 days before the fire |
| elevation          | Average elevation of the county where the fire occurred  |
| LC_00 [7]          | Water                                                    |
| LC_01              | Evergreen needleleaf forest                              |
| LC_02              | Evergreen broadleaf forest                               |
| LC_03              | Deciduous needleleaf forest                              |
| LC_04              | Deciduous broadleaf forest                               |
| LC_05              | Mixed forests                                            |
| LC_06              | Closed shrublands                                        |
| LC_07              | Open shrublands                                          |
| LC_08              | Woody savannas                                           |
| LC_09              | Savannas                                                 |
| LC_10              | Grasslands                                               |
| LC_11              | Permanent wetlands                                       |
| LC_12              | Croplands                                                |
| LC_13              | Urban and built-up                                       |
| LC_14              | Cropland/natural vegetation mosaic                       |
| LC_15              | Snow and ice                                             |
| LC_16              | Barren or sparsely vegetated                             |

**3.3 Experiments**

In this section, we present the experimentation process we have conducted using wildfires in the state of California to illustrate the accuracy of our algorithms and compare qualitatively the vegetation loss prediction using a curated set of candidate models. We have implemented the stacked deep learning model using multivariate time series, as described in Section 3.2, and provide a comprehensive analysis of the performance of the model along with the role of each layer.

**3.3.1 Data collection and processing.** California was selected as the study area due to its recurrent, large-scale wildfires and the availability of high-resolution environmental and remote sensing data, making it a representative and data-rich region for modeling vegetation loss. Effective modeling of vegetation loss requires integrating core fire event data with relevant meteorological and topographical features that influence a wildfire’s impact on vegetation. This study incorporates diverse datasets covering wildfire behavior and environmental features using Google Earth Engine (GEE) scripts. The components are described below, with Table 1 summarizing key features and Table 2 detailing their spatial and temporal resolutions. Feature importance analysis was conducted using an RF

model trained on aggregate feature representations to assess the contribution of individual meteorological and topographical features to model performance. Specifically, importance was calculated using Gini-based measures [11]. As shown in Table 1, aggregate land coverage features and historical fire events greatly influence predictors (23.68% and 14.75% respectively), while meteorological features show varying contributions to model performance.

**NDVI.** Vegetation productivity, measured by the Normalized Difference Vegetation Index (NDVI), is derived from remote sensing data that assesses vegetation density and health based on light reflectance. NDVI values range from -1 to 1, where higher values indicate dense and healthy vegetation, and lower values correspond to sparse, stressed, or barren surfaces [71]. NDVI reflects vegetation dynamics influenced by seasonal patterns, and its values can fluctuate depending on the time of year, regardless of fire activity. Additionally, the 16-day sampling interval of the satellite introduces temporal misalignment with fire events and daily weather data. To address both the coarse temporal resolution and the underlying seasonality, NDVI values from the 30 days before and after each fire were aggregated using the mean and minimum, respectively, without interpolation.

**Fire events.** Fire data derived from satellite observations were filtered to retain only high-confidence detections (i.e., confidence  $\geq 60$ ), ensuring the reliability of identified fire pixels. These detections were then mapped to historical wildfire records in California using the California Open Data Portal [48]. For each fire event, the acquisition date and geographic coordinates were extracted and mapped to their corresponding counties. To enhance forecasting performance, additional temporal features were derived, as summarized in Table 3.

**Meteorological data.** Nine climate and weather variables were selected, each available as daily measurements at the county level. To handle partially missing values within the aggregation window (i.e., 7 days before to 14 days after the fire), the mean of each weather variable was computed using the available daily values for the respective fire event and county. Fire events with completely missing weather data across any one of the selected variables were excluded from the final dataset.

**Topographical data.** County-level mean elevation and seventeen land cover fractions for the year 2018 were incorporated as enriched static features. Missing elevation values were imputed using the overall mean elevation across all counties. For land cover fractions, missing values were interpreted as the absence of the corresponding class and, therefore, replaced with zeros.

**3.3.2 Baselines.** The performance of MVeLMA in comparison with both SOTA deterministic and stochastic models, which adopt various strategies for handling spatiotemporal data and integrating feature sets, is summarized in Table 4. The details about the baselines are listed as follows:

- **Deterministic:**
  - *Dokhan et al.* [9]: A feed-forward neural network model (FFNN) comprising fully connected hidden layers with

**Table 4: Comparison of MVeLMA with Baseline Methods**

| Model         | MAE           | R <sup>2</sup> | MAPE (%)    | NRMSE         |
|---------------|---------------|----------------|-------------|---------------|
| <b>MVeLMA</b> | <b>0.0109</b> | <b>0.7232</b>  | <b>4.77</b> | <b>0.5261</b> |
| FFNN          | 0.0126        | 0.6175         | 5.45        | 0.6185        |
| UFSTP         | 0.0151        | 0.4815         | 9.77        | 0.7201        |
| GRU           | 0.0133        | 0.4850         | 5.15        | 0.7176        |
| 1D-CNN        | 0.0138        | 0.5382         | 6.87        | 0.6795        |
| LSTM-MC       | 0.0159        | 0.4410         | 10.77       | 0.7477        |

nonlinear activation functions that was developed to forecast the burned area of forest fires using meteorological and spatiotemporal inputs.

- *Gated Recurrent Units (GRU)* [50]: A type of RNN composed of gated recurrent units, well-known for capturing dynamic temporal dependencies.
- *One-dimensional CNN (1D-CNN)* [55]: A neural network model that leverages temporal convolutional layers to identify temporal patterns in multivariate time series and extract spatial features between time steps, which has been applied extensively for time series forecasting.
- **Stochastic:**
  - *Urban Fire Spatial-Temporal Prediction (UFSTP)* [62]: A SOTA spatiotemporal and multi-source data enhanced predictive model. UFSTP uses a multi-layer perceptron (MLP) to encode enriched spatial features and Gated Recurrent Units (GRU) to encode dynamic temporal sequences. It integrates a spatiotemporal propagation mechanism to model interactions across regions over time, concatenating both enriched and temporal representations to forecast fire occurrence.
  - *Long Short-Term Memory-Monte Carlo (LSTM-MC)* [54]: A SOTA probabilistic framework for one-step-ahead daily streamflow forecasting that uses a Monte Carlo-based deep learning model based on LSTM to account for model uncertainty.

**3.3.3 Model Configuration.** The proposed model was implemented in PyTorch and trained on a CPU. The first component is a BiLSTM that encodes 30-day weather sequences prior to each fire event. The hidden size is set to 64 ( $d_h = 64$ ), resulting in a 128-dimensional output from the bidirectional LSTM. An attention mechanism applies a linear layer to each time step, followed by a softmax operation across time to compute attention weights. These weights are used to obtain a weighted sum of the BiLSTM outputs, which is passed through a fully connected layer to produce a 20-dimensional latent representation ( $l = 20$ ).

The second component, the GPR layer, is implemented using GPyTorch [16]. To model the target variable, vegetation loss, we used a Matérn kernel with a smoothness parameter of  $\nu = 2.5$ , chosen for its flexibility and ability to capture spatial and temporal variability. The kernel is wrapped in a ScaleKernel to capture the overall magnitude of variation in the target function. The GPR model was trained using the Adam optimizer [38] for up to 80 epochs with a learning rate of 0.1, minimizing the exact negative

**Table 5: Comparison of time cost (in seconds) of MVeLMA and baseline models**

| Model                              | Avg./epoch (s) | Training (s) | Eval (s)                                |
|------------------------------------|----------------|--------------|-----------------------------------------|
| <b>MVeLMA<br/>(GPR Layer only)</b> | <b>0.29</b>    | <b>23.23</b> | <b><math>2.88 \times 10^{-5}</math></b> |
| FFNN                               | 0.002          | 1.260        | $1 \times 10^{-3}$                      |
| UFSTP                              | 0.002          | 0.030        | $5 \times 10^{-3}$                      |
| GRU                                | 0.610          | 84.33        | $9 \times 10^{-3}$                      |
| 1D-CNN                             | 0.250          | 23.37        | $5 \times 10^{-3}$                      |
| LSTM-MC                            | 0.650          | 83.58        | $9 \times 10^{-3}$                      |

MLL to optimize the kernel parameters and the posterior mean function.

For each fire event, the input to train the RF was assembled by concatenating three features to form a 48-dimensional feature vector as follows: First, a 20-dimensional latent vector produced by the BiLSTM with an attention layer; second, a 27-dimensional vector of enriched covariates of shape  $(F, 27)$ , where  $F$  represents the total number of fire events in the dataset shown in Table 3; and finally, scalar value representing the posterior mean prediction derived from the GPR layer. The RF was fine-tuned to balance model complexity and ensure stable predictions, using 500 trees. Parameters were selected to prevent under- and overfitting while also capturing dependencies between latent and enriched features.

**3.3.4 Model Performance Comparison.** We compare our model with the five aforementioned SOTA models and traditional multivariate time series models typically applied to natural disaster forecasting (i.e., FFNN, UFSTP, GRU, 1D-CNN, LSTM-MC). Model performance was evaluated using four metrics: mean absolute error (MAE), coefficient of determination ( $R^2$ ), mean absolute percentage error (MAPE), and normalized root mean squared error (NRMSE). Based on the results, we have described several observations as follows.

First, as summarized in Table 4, MVeLMA achieves a consistent lead in all four metrics. The improvement can largely be attributed to the model’s ability to integrate multimodal feature sources and capture their dependencies through distinct modeling components. Conversely, most existing time-series-based approaches for vegetation estimation rely on a limited set of features, typically 10 to 13 variables including spectral bands, meteorological indicators, and vegetation indices, and rarely explore how different model types align with different feature modalities [22].

The proposed model outperforms existing approaches by addressing key limitations in both probabilistic and deterministic baselines. Probabilistic methods often struggle to generalize in regions with limited data availability. Deterministic models, while effective in capturing temporal patterns from meteorological and topographical variables, fail to model spatial dependencies and lack the ability to quantify predictive uncertainty.

**3.3.5 Scalability Studies.** Beyond the performance of the predictions, the computational efficiency is a vital component of real-world use cases. In this stage, we compare the runtime of each of the three layers, the time stratified by kernel choice, and the train

**Table 6: Ablation Study of Model Components on Vegetation Loss Prediction**


| Model Variant                     | MAE           | $R^2$         | MAPE (%)    | NRMSE         |
|-----------------------------------|---------------|---------------|-------------|---------------|
| <b>MVeLMA</b>                     | <b>0.0109</b> | <b>0.7232</b> | <b>4.77</b> | <b>0.5261</b> |
| w/o BiLSTM                        | 0.0111        | 0.6567        | 4.8244      | 0.5859        |
| w/o GPR                           | 0.0113        | 0.6832        | 5.1196      | 0.5629        |
| w/o RF                            | 0.0121        | 0.6250        | 6.4640      | 0.6124        |
| w/o BiLSTM and GPR                | 0.0116        | 0.6677        | 5.2043      | 0.5764        |
| w/o BiLSTM and RF                 | 0.0127        | 0.6118        | 5.8763      | 0.6230        |
| w/o GPR and RF                    | 0.0137        | 0.5446        | 5.5768      | 0.6748        |
| <b>MVeLMA (Oregon Validation)</b> | <b>0.0247</b> | <b>0.7390</b> | <b>1.51</b> | <b>0.5109</b> |

and test times of MVeLMA with five baselines. As shown in Table 5, MVeLMA’s training time per epoch is comparable to that of the candidate models. However, its inference time is over 90% faster than the fastest among them (i.e., the FFNN), demonstrating its efficiency at prediction time.

**3.3.6 Ablation Studies and Cross-regional Validation Study.** In this subsection, we present ablated studies to evaluate the effectiveness of each component of the proposed model. We design six variants of MVeLMA and compare them with the complete version of MVeLMA. The differences between these seven models are described as follows:

- *w/o BiLSTM*: removing the BiLSTM component i.e., no longer leveraging the pre- and post-fire temporal sequence
- *w/o GPR*: removing the GPR component
- *w/o RF*: removing the RF component i.e., no longer probabilistically modeling how latent features relate to vegetation loss
- *w/o BiLSTM and GPR*: removing both BiLSTM and GPR i.e., using only RF prediction
- *w/o BiLSTM and RF*: removing BiLSTM and RF i.e., using only GPR
- *w/o GPR and RF*: removing GPR and RF i.e., using only BiLSTM
- *MVeLMA*: the final version of the proposed model

Table 6 summarizes the results. We delineate our observations as follows: (i) Replacing the BiLSTM with a 30-day mean-aggregated meteorological sequence reduces complexity but limits the model’s ability to capture temporal dynamics such as seasonal shifts or weather changes around fire events. This simplification leads to a modest performance drop (MAPE: 4.82%,  $R^2$ : 0.6567 compared to 4.77% and 0.7232, respectively). (ii) By removing the GPR layer, the model fails to capture uncertainty, reducing its adaptability to less frequently observed or spatially distinct conditions. Without the GPR posterior mean as an input to RF, the overall feature richness is also reduced, further lowering performance. (iii) Excluding the RF layer removes both the final refinement step and the contribution of enriched static features. This variant exhibits a significant performance decline (MAPE: 6.46%,  $R^2$ : 0.6250), highlighting the importance of integrating contextual information through post-GPR correction. (iv) Finally, a purely deterministic RF model, using only



**Figure 2: Performance analysis across different architectural variations across  $R^2$  and MAE metrics. Results highlight the effectiveness of BiLSTM over unidirectional LSTM for capturing temporal fire event and climate dependencies, and RF over MLP for ensemble prediction.**

aggregated meteorological and enriched features, performs better than the ablated variants but lacks uncertainty estimation and fails to capture temporal dependencies. While reasonably competitive (MAPE: 5.26%,  $R^2$ : 0.6677), it underperforms compared with the MVeLMA model, reiterating the advantage of a stacked ensemble model.

The predictive accuracy of MVeLMA was compared against architectural variants, specifically using unidirectional LSTM and MLP instead of BiLSTM and RF, respectively. As shown in Figure 2, replacing BiLSTM with LSTM resulted in a performance decrease in  $R^2$  and MAE scores, highlighting the effectiveness of the bidirectional approach in capturing temporal dependencies between fire events and climate patterns. Additionally, using MLP instead of RF significantly lowered performance, with prediction accuracy decreasing by nearly 50%. This underscores the robustness of RF and its reduced sensitivity to hyperparameter tuning.

To evaluate the model’s cross-regional adaptability, MVeLMA was tested on vegetation loss in Oregon, a state with a wetter, more temperate climate and different elevation and vegetation patterns compared to California [25]. The Oregon validation used 64 wildfire events from May to July 2023, with metrics shown in Table 6. The  $R^2$  value for this region was comparable to that of California at 0.7390, suggesting MVeLMA’s architecture effectively captures the relationship between topographical and climate features and wildfire events. Specifically, the model achieved lower MAPE (1.51 vs. 4.77), indicating more consistent predictions across the different regional fire patterns. However, the higher MAE (0.0247 vs. 0.0109) suggests that while the model captures the relationship between the covariates and wildfire events, additional fine-tuning would likely be required to better account for region-specific landscape variances.

Additionally, to validate temporal generalization, MVeLMA was tested on Oregon data from an earlier seasonal period (May to July) compared to the California training period (July to September). This tests the model’s adaptability to not only geographical variations but also seasonal variations in fire conditions, as fire seasonality

patterns vary between Oregon’s wetter climate and California’s drier Mediterranean climate.

### 3.4 Case Study

To evaluate model performance and interpretability, we conducted a spatial and temporal analysis of post-fire vegetation loss across selected California counties using observed post-fire vegetation loss (OPFVL), predicted post-fire vegetation loss (PPVL), and predictive confidence (APC).


Across the test counties, OPFVL ranged from 0.0006 (Alameda County) to 0.0873 (Colusa County), with higher values concentrated in counties experiencing larger, more severe fires. The model’s PPVL ranged from 0.0092 to 0.0684, while APC remained consistently high, between 0.87 and 0.98. Figure 3 illustrates this spatial distribution, showing that counties with larger, more frequent fire events exhibited higher predictive confidence, likely due to more training data. For temporal analysis, we selected Butte County due to its high frequency of fire events and seasonal vegetation variability. Fire events in Butte County were grouped into overlapping three-week periods, with each period starting one week after the previous period, based on fire start dates. During the initial period (July 17–August 6), Butte County exhibited an average OPFVL of 0.0123, with a predicted PPVL of 0.0091 and a corresponding APC of 0.984, as shown in Figure 4. In the subsequent period (August 7–August 27), both OPFVL and PPVL increased to 0.0265 and 0.0248 respectively, while APC remained above 0.97 throughout the season.

The results of this case study demonstrate MVeLMA’s ability to capture both spatial and temporal dynamics of post-fire vegetation loss, while maintaining consistently high predictive confidence across diverse fire conditions.


## 4 Conclusion and Future Work

In this paper, we propose a novel end-to-end probabilistic model called MVeLMA for forecasting county-level vegetation loss from fire events using a multimodal feature integration pipeline and a stacked ensemble-based architecture. To the best of our knowledge, this is the first work to combine the spatiotemporal learning capabilities of BiLSTM, the uncertainty quantification of GPR, and the residual correction strength of RF. MVeLMA model can be used for spatiotemporal prediction of vegetation loss in states with 30-day lead times. Extensive experiments and comparisons to common spatiotemporal models, as well as SOTA models, demonstrate the effectiveness and spatiotemporal predictive capabilities of our proposed model. Furthermore, ablation studies showcase the effectiveness of each individual component of the model. The integration of prediction uncertainty makes it applicable in real-world scenarios as demonstrated by the case study done on wildfire-affected counties of California.

This work can be further extended in the future by considering the impact of point of interest data, such as landmarks and human-centered activities. Future work should also explore calibration techniques to better align predicted uncertainty with actual prediction errors to handle potential overconfidence. Lastly, incorporating such anthropogenic data into the model can improve the



**Figure 3:** This figure displays the spatial distribution of MVeLMA's observed, predicted, and predictive confidence estimates across all California counties, where darker colors represent higher vegetation loss and predictive confidences respectively.



**Figure 4:** Spatial distribution of MVeLMA's observed, predicted, and predictive-confidence estimates across Butte, California, for four overlapping three-week intervals from July 17 to August 27, 2023. The inset on shows the county's location within California. The darker colors represent higher vegetation loss and predictive confidence, respectively.

model's accuracy and robustness and provide additional insight for agricultural planning and ecosystem management.

## 5 Acknowledgements

Approved for Public Release; Distribution Unlimited. Public Release Case Number 25-1835. The author's affiliation with The MITRE Corporation is provided for identification purposes only and is not intended to convey or imply MITRE's concurrence with, or support for, the positions, opinions, or viewpoints expressed by the author.

## References

- [1] Gözde Bayat and Kazım Yıldız. 2022. Comparison of the Machine Learning Methods to Predict Wildfire Areas. *Turkish Journal of Science and Technology* 17, 2 (Sept. 2022), 241–250. <https://doi.org/10.55525/tjst.1063284>
- [2] Rohan T. Bhowmik, Youn Soo Jung, Juan A. Aguilera, Mary Prunicki, and Kari Nadeau. 2023. A multi-modal wildfire prediction and early-warning system based on a novel machine learning framework. *Journal of Environmental Management* 341 (Sept. 2023), 117908. <https://doi.org/10.1016/j.jenvman.2023.117908>
- [3] Shiv Bolan, Shailja Sharma, Santanu Mukherjee, Daniel F. Gomez Isaza, Essie M. Rodgers, Pingfan Zhou, Deyi Hou, Facundo Scordo, Sudeep Chandra, Kadambot H. M. Siddique, and Nanthi Bolan. 2025. Wildfires under changing climate, and their environmental and health impacts. *Journal of Soils and Sediments* (April 2025). <https://doi.org/10.1007/s11368-025-04020-y>
- [4] Jorge Castel-Clavera, François Pimont, Thomas Opitz, Julien Ruffault, Miguel Rivière, and Jean-Luc Dupuy. 2022. Disentangling the factors of spatio-temporal patterns of wildfire activity in south-eastern France. *International Journal of Wildland Fire* 32, 1 (Dec. 2022), 15–28. <https://doi.org/10.1071/WF22086>
- [5] Jiehang Deng, Bin Hong, Weiming Wang, and Guosheng Gu. 2025. Daily Wildfire Risk Prediction by Mining Global and local spatio-temporal dependency. *Earth Science Informatics* 18, 3 (March 2025), 316. <https://doi.org/10.1007/s12145-024-01652-5>
- [6] Jiehang Deng, Weiming Wang, Guosheng Gu, Zhiqiang Chen, Jing Liu, Guobo Xie, Shaowei Weng, Lei Ding, and Chuan Li. 2023. Wildfire susceptibility prediction using a multisource and spatiotemporal cooperative approach. *Earth Science Informatics* 16, 4 (Dec. 2023), 3511–3529. <https://doi.org/10.1007/s12145-023-01104-6>
- [7] Sadashiva Devadiga. [n. d.]. [landweb.modaps.eosdis.nasa.gov/tsdata/C6\\_scheme.html](https://landweb.modaps.eosdis.nasa.gov/tsdata/C6_scheme.html)
- [8] Kamel Didan. 2021. MODIS/Terra Vegetation Indices 16-Day L3 Global 250m SIN Grid V061. <https://doi.org/10.5067/MODIS/MOD13Q1.061>
- [9] Sabri Nasser Hussein Murshed Dokhan and Mohammed Fuad Mohammed Ahmed Saif. 2024. A Neural Network Approach to Forest Fire Burned Area Prediction. *Journal of Applied Science, Technology and Computing* 1, 1 (April 2024), 52–62. <https://publisher.uthm.edu.my/ojs/index.php/jastec/article/view/16481>
- [10] Hao Dong, Han Wu, Pengfei Sun, and Yunhong Ding. 2022. Wildfire Prediction Model Based on Spatial and Temporal Characteristics: A Case Study of a Wildfire

- in Portugal's Montesinho Natural Park. *Sustainability* 14, 16 (Jan. 2022), 10107. <https://doi.org/10.3390/su141610107>
- [11] A. Druckman and T. Jackson. 2008. Measuring resource inequalities: The concepts and methodology for an area-based Gini coefficient. *Ecological Economics* 65, 2 (April 2008), 242–252. <https://doi.org/10.1016/j.ecolecon.2007.12.013>
  - [12] Behnam Ebadati, Reza Attarzadeh, Mohammad Alikhani, Fahimeh Youssefi, and Saied Pirasteh. 2024. Rapid Post-Wildfire Burned Vegetation Assessment with Google Earth Engine (Case Study: 2023 Canada Wildfires). *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences XLVIII-3-W3-2024* (Nov. 2024), 45–53. <https://doi.org/10.5194/isprs-archives-XLVIII-3-W3-2024-45-2024>
  - [13] Tom G. Farr, Paul A. Rosen, Edward Caro, Robert Crippen, Riley Duren, Scott Hensley, Michael Kobrick, Mimi Paller, Ernesto Rodriguez, Ladislav Roth, David Seal, Scott Shaffer, Joanne Shimada, Jeffrey Umland, Marian Werner, Michael Oskin, Douglas Burbank, and Douglas Alsdorf. 2007. The Shuttle Radar Topography Mission. *Reviews of Geophysics* 45, 2 (June 2007), 2005RG000183. <https://doi.org/10.1029/2005RG000183>
  - [14] Aya Ferchichi, Ali Ben Abbes, Vincent Barra, and Imed Riadh Farah. 2022. Forecasting vegetation indices from spatio-temporal remotely sensed data using deep learning-based approaches: A systematic literature review. *Ecological Informatics* 68 (May 2022), 101552. <https://doi.org/10.1016/j.ecoinf.2022.101552>
  - [15] Mark Friedl and Damien Sulla-Menashe. 2022. MODIS/Terra+Aqua Land Cover Type Yearly L3 Global 500m SIN Grid V061. <https://doi.org/10.5067/MODIS/MCD12Q1.061>
  - [16] Jacob R. Gardner, Geoff Pleiss, David Bindel, Kilian Q. Weinberger, and Andrew Gordon Wilson. 2021. GPYtorch: Blackbox Matrix-Matrix Gaussian Process Inference with GPU Acceleration. <https://doi.org/10.48550/arXiv.1809.11165>
  - [17] Sebastian Gerard, Yu Zhao, and Josephine Sullivan. 2023. Wildfire-SpreadTS: A dataset of multi-modal time series for wildfire spread prediction. *Advances in Neural Information Processing Systems* 36 (Dec. 2023), 74515–74529. [https://proceedings.neurips.cc/paper\\_files/paper/2023/hash/ebd545176bdaa9cd5d45954947bd74b7-Abstract-Datasets\\_and\\_Benchmarks.html](https://proceedings.neurips.cc/paper_files/paper/2023/hash/ebd545176bdaa9cd5d45954947bd74b7-Abstract-Datasets_and_Benchmarks.html)
  - [18] Omid Ghorbanzadeh, Khalil Valizadeh Kamran, Thomas Blaschke, Jagannath Aryal, Amin Naboureh, Jamshid Einali, and Jinhu Bian. 2019. Spatial Prediction of Wildfire Susceptibility Using Field Survey GPS Data and Machine Learning Approaches. *Fire* 2, 3 (Sept. 2019), 43. <https://doi.org/10.3390/fire2030043>
  - [19] Louis Giglio, Wilfrid Schroeder, and Joanne V Hall. [n. d.]. MODIS Collection 6 and Collection 6.1 Active Fire Product User's Guide. ([n. d.]).
  - [20] Almaz T. Gizatullin and N. A. Alekseenko. 2022. Prediction of Wildfires Based on the Spatio-Temporal Variability of Fire Danger Factors. *GEOGRAPHY, ENVIRONMENT, SUSTAINABILITY* 15, 2 (June 2022), 31–37. <https://doi.org/10.24057/2071-9388-2021-139>
  - [21] Yangming Guo, Xiaolei Li, Guanghan Bai, and Jiezhong Ma. 2012. Time Series Prediction Method Based on LS-SVR with Modified Gaussian RBF. In *Neural Information Processing*, Tingwen Huang, Zhigang Zeng, Chuandong Li, and Chi Sing Leung (Eds.). Springer, Berlin, Heidelberg, 9–17. [https://doi.org/10.1007/978-3-642-34481-7\\_2](https://doi.org/10.1007/978-3-642-34481-7_2)
  - [22] Trung H. Nguyen, Simon Jones, Mariela Soto-Berelov, Andrew Haywood, and Samuel Hislop. 2020. Landsat Time-Series for Estimating Forest Aboveground Biomass and Its Dynamics across Space and Time: A Review. *Remote Sensing* 12, 1 (Jan. 2020), 98. <https://doi.org/10.3390/rs12010098>
  - [23] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory [textbook Neural Computation. <https://dl.acm-org.ezproxy.lib.vt.edu/doi/10.1162/neco.1997.9.8.1735>
  - [24] Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional LSTM-CRF Models for Sequence Tagging. <https://doi.org/10.48550/arXiv.1508.01991> arXiv:1508.01991 [cs].
  - [25] USDA Climate Hubs. [n. d.]. Climate Change and Wildfire in Idaho, Oregon, and Washington. 2023 ([n. d.]). <https://www.climatehubs.usda.gov/hubs/northwest/topic/climate-change-and-wildfire-idaho-oregon-and-washington>
  - [26] Fantine Huot, R. Lily Hu, Nita Goyal, Tharun Sankar, Matthias Ihme, and Yi-Fan Chen. 2022. Next Day Wildfire Spread: A Machine Learning Data Set to Predict Wildfire Spreading from Remote-Sensing Data. *IEEE Transactions on Geoscience and Remote Sensing* 60 (2022), 1–13. <https://doi.org/10.1109/TGRS.2022.3192974> arXiv:2112.02447 [cs].
  - [27] Svetlana Illarionova, Dmitrii Shadrin, Fedor Gubanov, Mikhail Shutov, Usman Tasuev, Ksenia Evtseva, Maksim Mironenko, and Evgeny Burnaev. 2025. Exploration of geo-spatial data and machine learning algorithms for robust wildfire occurrence prediction. *Scientific Reports* 15, 1 (March 2025), 10712. <https://doi.org/10.1038/s41598-025-94002-4>
  - [28] Ayoub Jadouli and Chaker El Amrani. 2024. Enhancing Wildfire Forecasting Through Multisource Spatio-Temporal Data, Deep Learning, Ensemble Models and Transfer Learning. *Advances in Artificial Intelligence and Machine Learning* 04, 03 (2024), 2614–2628. <https://doi.org/10.54364/AIAML.2024.43152>
  - [29] Ayoub Jadouli and Chaker El Amrani. 2025. Deep Learning with Pretrained 'Internal World' Layers: A Gemma 3-Based Modular Architecture for Wildfire Prediction. <https://doi.org/10.48550/arXiv.2504.18562>
  - [30] Ayoub Jadouli and Chaker El Amrani. 2024. Enhancing Wildfire Forecasting Through Multisource Spatio-Temporal Data, Deep Learning, Ensemble Models and Transfer Learning. *Advances in Artificial Intelligence and Machine Learning* 04, 03 (2024), 2614–2628. <https://doi.org/10.54364/AIAML.2024.43152>
  - [31] Piyush Jain, Sean C.P. Coogan, Sriram Ganapathi Subramanian, Mark Crowley, Steve Taylor, and Mike D. Flannigan. 2020. A review of machine learning applications in wildfire science and management. *Environmental Reviews* 28, 4 (Dec. 2020), 478–505. <https://doi.org/10.1139/er-2020-0019>
  - [32] Cuicui Ji, Changbin Wu, Xiaosong Li, Fuyang Sun, and Bin Sun. 2024. Methods for spatial and temporal detection of forest wildfire disturbance based on time series Eco-environment indicators. *Ecological Indicators* 168 (Nov. 2024), 112765. <https://doi.org/10.1016/j.ecolind.2024.112765>
  - [33] Josy John, Shridhar Velhal, and Suresh Sundaram. 2025. A Resource-Efficient Decentralized Sequential Planner for Spatiotemporal Wildfire Mitigation. *IEEE Transactions on Automation Science and Engineering* (2025), 1–1. <https://doi.org/10.1109/TASE.2025.3536356>
  - [34] Maxwell B. Joseph, Matthew W. Rossi, Nathan P. Mietkiewicz, Adam L. Mahood, Megan E. Cattau, Lise Ann St. Denis, R. Chelsea Nagy, Virginia Iglesias, John T. Abatzoglou, and Jennifer K. Balch. 2019. Spatiotemporal prediction of wildfire size extremes with Bayesian finite sample maxima. *Ecological Applications* 29, 6 (2019), e01898. <https://doi.org/10.1002/eap.1898>
  - [35] Muhammad Khubayeb Kabir, Kawshik Kumar Ghosh, Md Fahim Ul Islam, and Jia Uddin. 2024. Wildfire Prediction in the United States Using Time Series Forecasting Models. *Annals of Emerging Technologies in Computing (AETiC)* 8, 2 (April 2024), 32–42. <https://doi.org/10.33166/AETiC.2024.02.003>
  - [36] Serkan KARTAL, Muzafer Can IBAN, and Alihsan SEKERTEKIN. 2024. Next-level vegetation health index forecasting: A ConvLSTM study using MODIS Time Series. *Environmental Science and Pollution Research* 31, 12 (March 2024), 18932–18948. <https://doi.org/10.1007/s11356-024-32430-x>
  - [37] Sadegh Khammohammadi, Mehrdad Arashpour, Emadaldin Mohammadi Golafshani, Miguel G. Cruz, Abbas Rajabifard, and Yu Bai. 2022. Prediction of wildfire rate of spread in grasslands using machine learning methods. *Environmental Modelling & Software* 156 (Oct. 2022), 105507. <https://doi.org/10.1016/j.envsoft.2022.105507>
  - [38] Diederik P. Kingma and Jimmy Ba. 2017. Adam: A Method for Stochastic Optimization. <https://doi.org/10.48550/arXiv.1412.6980>
  - [39] Andrew Kliskey. [n. d.]. Idaho EPSCoR. <https://www.idahoepscor.org/>
  - [40] Jonathan Koh, François Pimont, Jean-Luc Dupuy, and Thomas Opitz. 2023. Spatiotemporal wildfire modeling through point processes with moderate and extreme marks. *The Annals of Applied Statistics* 17, 1 (March 2023), 560–582. <https://doi.org/10.1214/22-AOAS1642>
  - [41] Xufeng Lin, Zhongyuan Li, Wenjing Chen, Xueying Sun, and Demin Gao. 2023. Forest Fire Prediction Based on Long- and Short-Term Time-Series Network. *Forests* 14, 4 (April 2023), 778. <https://doi.org/10.3390/f14040778>
  - [42] Jiahe Liu and Xiaodi Wang. 2023. Forecasting Post-Wildfire Vegetation Recovery in California using a Convolutional Long Short-Term Memory Tensor Regression Network. <https://doi.org/10.48550/arXiv.2311.02492>
  - [43] Ashima Malik, Megha Rajam Rao, Nandini Puppala, Prathusha Koouri, Venkata Anil Kumar Thota, Qiao Liu, Sen Chiao, and Jerry Gao. 2021. Data-Driven Wildfire Risk Prediction in Northern California. *Atmosphere* 12, 1 (Jan. 2021), 109. <https://doi.org/10.3390/atmos12010109>
  - [44] Mohammad Marjani, Masoud Mahdianpari, and Fariba Mohammadimanesh. 2024. CNN-BiLSTM: A Novel Deep Learning Model for Near-Real-Time Daily Wildfire Spread Prediction. *Remote Sensing* 16, 8 (Jan. 2024), 1467. <https://doi.org/10.3390/rs16081467>
  - [45] Arif Masrur, Manzhou Yu, and Alan Taylor. 2024. Capturing and interpreting wildfire spread dynamics: attention-based spatiotemporal models using ConvLSTM networks. *Ecological Informatics* 82 (Sept. 2024), 102760. <https://doi.org/10.1016/j.ecoinf.2024.102760>
  - [46] Dimitrios Michail, Lefki-Ioanna Panagiotou, Charalampos Davalas, Ioannis Prapas, Spyros Kondylatos, Nikolaos Ioannis Bountos, and Ioannis Papoutsis. 2024. Seasonal Fire Prediction using Spatio-Temporal Deep Neural Networks. <https://doi.org/10.48550/arXiv.2404.06437>
  - [47] Sanaz Moghim and Majid Mehrabi. 2024. Wildfire assessment using machine learning algorithms in different regions. *Fire Ecology* 20, 1 (Dec. 2024), 104. <https://doi.org/10.1186/s42408-024-00335-2>
  - [48] California Department of Forestry and Fire Protection (CAL FIRE) and its Fire and Resource Assessment Program (FRAP). [n. d.]. California Historical Fire Perimeters - California Open Data. <https://data.ca.gov/dataset/california-historical-fire-perimeters>
  - [49] Benjami Parellada I Calderer. 2024. *Predicting California wildfire spread with time-series and spatial-temporal models*. Master's thesis. Universitat Politècnica de Catalunya. <https://upcommons.upc.edu/handle/2117/420301>
  - [50] Rylan Perumal and Terence L van Zyl. 2020. Comparison of Recurrent Neural Network Architectures for Wildfire Spread Modelling. In *2020 International SAUPEC/RobMech/PRASA Conference*. 1–6. <https://doi.org/10.1109/SAUPEC/RobMech/PRASA48453.2020.9078028>

- [51] R. Shanmuga Priya and K. Vani. 2024. Vegetation change detection and recovery assessment based on post-fire satellite imagery using deep learning. *Scientific Reports* 14, 1 (June 2024), 12611. <https://doi.org/10.1038/s41598-024-63047-2>
- [52] D. Sushma Reddy and P. Rama Chandra Prasad. 2018. Prediction of vegetation dynamics using NDVI time series data and LSTM. *Modeling Earth Systems and Environment* 4, 1 (April 2018), 409–419. <https://doi.org/10.1007/s40808-018-0431-3>
- [53] Moritz Rösch, Michael Nolde, Tobias Ullmann, and Torsten Riedlinger. 2024. Data-Driven Wildfire Spread Modeling of European Wildfires Using a Spatiotemporal Graph Neural Network. *Fire* 7, 6 (June 2024), 207. <https://doi.org/10.3390/fire7060207>
- [54] Ali Sattari, Ehsan Foroumandi, Keyhan Gavahi, and Hamid Moradkhani. 2025. A probabilistic machine learning framework for daily extreme events forecasting. *Expert Systems with Applications* 265 (March 2025), 126004. <https://doi.org/10.1016/j.eswa.2024.126004>
- [55] Tobias Schmieg and Carsten Lanquillon. 2024. Time Series Representation Learning: A Survey on Deep Learning Techniques for Time Series Forecasting. In *Artificial Intelligence in HCI*, Helmut Degen and Stavroula Ntoa (Eds.). Springer Nature Switzerland, Cham, 422–435. [https://doi.org/10.1007/978-3-031-60606-9\\_25](https://doi.org/10.1007/978-3-031-60606-9_25)
- [56] Eric Alexander Schmitt, Evan Zaremba, Neha Ananthavaram, Li Liu, Mario Giraldo, and Xunfei Jiang. 2024. Ecosystem-Based Wildfire Risk Prediction with Machine Learning. In *2024 IEEE International Conference on Big Data (BigData)*. 7540–7545. <https://doi.org/10.1109/BigData62323.2024.10825794>
- [57] Seyd Teymoor Seydi, Mahdi Hasanlou, and Jocelyn Chanussot. 2022. Burnt-Net: Wildfire burned area mapping with single post-fire Sentinel-2 data and deep learning morphological neural network. *Ecological Indicators* 140 (July 2022), 108999. <https://doi.org/10.1016/j.ecolind.2022.108999>
- [58] Parviz Sohrabi, Behshad Jodeiri Shokri, and Hesam Dehghani. 2023. Predicting coal price using time series methods and combination of radial basis function (RBF) neural network with time series. *Mineral Economics* 36, 2 (June 2023), 207–216. <https://doi.org/10.1007/s13563-021-00286-z>
- [59] Christos Vasilakos, George E. Tsekouras, and Dimitris Kavroudis. 2022. LSTM-Based Prediction of Mediterranean Vegetation Dynamics Using NDVI Time-Series Data. *Land* 11, 6 (June 2022), 923. <https://doi.org/10.3390/land11060923>
- [60] Yin Wang, Nan Zhang, Mingjie Chen, Yabing Zhao, Famiao Guo, Jingxian Huang, Daoli Peng, and Xiaohui Wang. 2025. Prediction and Spatiotemporal Dynamics of Vegetation Index Based on Deep Learning and Environmental Factors in the Yangtze River Basin. *Forests* 16, 3 (March 2025), 460. <https://doi.org/10.3390/f16030460>
- [61] Harry West, Nevil Quinn, and Michael Horswell. 2019. Remote sensing for drought monitoring & impact assessment: Progress, past challenges and future opportunities. *Remote Sensing of Environment* 232 (Oct. 2019), 111291. <https://doi.org/10.1016/j.rse.2019.111291>
- [62] Haiyu Xiang, Lizhi Wu, Zidong Guo, and Shaoyun Ren. 2025. Urban Fire Spatial–Temporal Prediction Based on Multi-Source Data Fusion. *Fire* 8, 5 (May 2025), 177. <https://doi.org/10.3390/fire8050177>
- [63] Chen Xu, Yao Xie, Daniel A. Zuniga Vazquez, Rui Yao, and Feng Qiu. 2023. Spatio-Temporal Wildfire Prediction Using Multi-Modal Data. *IEEE Journal on Selected Areas in Information Theory* 4 (2023), 302–313. <https://doi.org/10.1109/JSAIT.2023.3276054>
- [64] Zekun Xu, Zhaoming Zhang, Guojin He, Shuaizhang Zhang, Tengfei Long, and Guizhou Wang. 2025. Adaptive Early Wildfire Monitoring Based on Spatiotemporal Prediction and Himawari 8/9. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 18 (2025), 9396–9408. <https://doi.org/10.1109/JSTARS.2025.3554892>
- [65] Jiaqi Yao, Fan Mo, Haoran Zhai, Shiyi Sun, Karl-Heinz Feger, Lulu Zhang, Xinming Tang, Guoyuan Li, and Hong Zhu. 2022. A spatio-temporal prediction model theory based on deep learning to evaluate the ecological changes of the largest reservoir in North China from 1985 to 2021. *Ecological Indicators* 145 (Dec. 2022), 109618. <https://doi.org/10.1016/j.ecolind.2022.109618>
- [66] Haihui You, Zengyi Ma, Yijun Tang, Yuelan Wang, Jianhua Yan, Mingjiang Ni, Kefa Cen, and Qunxing Huang. 2017. Comparison of ANN (MLP), ANFIS, SVM, and RF models for the online classification of heating value of burning municipal solid waste in circulating fluidized bed incinerators. *Waste Management* 68 (Oct. 2017), 186–197. <https://doi.org/10.1016/j.wasman.2017.03.044>
- [67] Manzhu Yu, Arif Masrur, and Christopher Blaszcak-Boxe. 2023. Predicting hourly PM<sub>2.5</sub> concentrations in wildfire-prone areas using a SpatioTemporal Transformer model. *Science of The Total Environment* 860 (Feb. 2023), 160446. <https://doi.org/10.1016/j.scitotenv.2022.160446>
- [68] Meng Zhang, Hui Lin, Xiangren Long, and Yaotong Cai. 2021. Analyzing the spatiotemporal pattern and driving factors of wetland vegetation changes using 2000–2019 time-series Landsat data. *Science of The Total Environment* 780 (Aug. 2021), 146615. <https://doi.org/10.1016/j.scitotenv.2021.146615>
- [69] Zhongwei Zhang, Elias Krainski, Peng Zhong, Harvard Rue, and Raphaël Huser. 2023. Joint modeling and prediction of massive spatio-temporal wildfire count and burnt area data with the INLA-SPDE approach. *Extremes* 26, 2 (June 2023), 339–351. <https://doi.org/10.1007/s10687-023-00463-z>
- [70] Chenghao Zhu, Xiaoli Zhang, Ning Zhang, Mohammed Abdelmanan Hassan, and Lin Zhao. 2018. Assessing the Defoliation of Pine Forests in a Long Time-Series and Spatiotemporal Prediction of the Defoliation Using Landsat Data. *Remote Sensing* 10, 3 (March 2018), 360. <https://doi.org/10.3390/rs10030360>
- [71] Murat Özyavuz, Cemil Bilgili, and Aylin Salıcı. 2015. Determination of vegetation changes with NDVI method. *Journal of environmental protection and ecology* 16 (Jan. 2015), 264–273.