Pre-fire vegetation variability predicts post-fire recovery more strongly than pre-fire vegetation cover: a 12-year analysis using gradient boosting and SHAP

Djebaili Yasmine1 , Yassmina Saadna2 , Hani Inssaf3, Saliha Mezzoudj4
1 Laboratory of Algerian Forest and Climate Change, Higher National School of Forests, Khenchela, Algeria
2Lastic Laboratory, Department of Mathematic Computer Science, University of Batna 2 Mostefa Ben Boulaïd, Batna, Algeria
3 Laboratory of Algerian Forest and Climate Change, Higher National School of Forests, Khenchela, Algeria
4 Faculty of Sciences, Department of Computer Science, University of Algiers 1, Algiers, Algeria

Keywords: post-fire recovery; XGBoost; LightGBM; SHAP; vegetation health index; Aures mountains; Algeria

Most studies of post-fire vegetation recovery work with at most one or two years of pre-fire data. That limitation makes one question hard to ask. Does post-fire recovery depend more on how much vegetation was present before the fire, or on how stable that vegetation was over the preceding years? We addressed this question using a balanced 12-year dataset of 12,000 observations across a single fire event in the Khenchela region of the Aures mountains (eastern Algeria) in 2021, with seven pre-fire years (2014 to 2020), two immediate post-fire years (2021 to 2022), and three recovery years (2023 to 2025). We engineered pre-fire features that capture both mean and variability, compared five tree-based regressors, and used SHAP values to interpret the best model. The tuned LightGBM and XGBoost models achieved statistically equivalent performance ($R^2 = 0.8228 \pm 0.0265$ and 0.8229 ± 0.0281 respectively, 5-fold CV). Pre-fire NDVI was the strongest single predictor of recovery NDVI, with a mean absolute SHAP value of 0.020. The second-ranked predictor was not what we expected. Pre-fire VHI standard deviation across the seven baseline years had a mean absolute SHAP of 0.009, more than four times the importance of VHI mean itself (0.002). Points where vegetation health oscillated between drought and recovery in the years before the fire recovered substantially less afterward, regardless of their average condition. Land surface temperature remained 4.3°C above pre-fire baseline three years after the fire (Welch's $t = -51.81$, $p < 0.0001$), with the LST to NDVI coupling strengthening from $r = -0.45$ pre-fire to $r = -0.65$ in recovery. Fire severity itself, measured by post-fire NBR, was not predictable from pre-fire state ($R^2 = 0.05$), confirming that combustion depends on weather variables not captured by pre-fire vegetation indices while recovery depends on the residual ecological state. The variability finding aligns with recent global resilience work (Smith et al., 2023; Forzieri et al., 2022) and suggests that drought variability deserves a place in operational fire-risk indices.

1 Introduction

A wildfire reshapes a landscape in a single event, but the recovery plays out across years. How fast and how completely vegetation returns depends on the condition of the vegetation before the fire, on the severity of the burn, on local topography, and on the microclimate during the recovery window. Operational monitoring has relied on satellite-derived vegetation indices such as NDVI, EVI, and NBR for decades (Rouse et al., 1974; Key and Benson, 2006; Huete et al., 2002). Until recently, however, the field has lacked datasets with enough pre-fire history to separate pre-fire condition from pre-fire stability.

Recent global studies have shifted attention from mean to variability. Smith et al. (2023) found that vegetation resilience is lower where inter-annual precipitation variability is higher. Forzieri et al. (2022) showed that temporal autocorrelation in vegetation indices is a leading indicator of declining forest resilience under climate change. Smith et al. (2022) demonstrated that resilience itself can be inferred from natural variability. The shared implication is that pre-disturbance variability may matter as much as pre-disturbance mean. These studies operate at coarse spatial scales (kilometer grid cells, global aggregates) and do not address post-fire recovery at the point level. The question of whether the same variability signal explains recovery from a discrete fire event, at the scale of individual pixels rather than continents, remains open.

This study addresses that gap. We use a balanced 12-year longitudinal dataset of 1,000 spatial points across a single fire event in the Khenchela region of the Aures mountains, eastern Algeria, in 2021, with seven pre-fire years (2014 to 2020), two immediate post-fire years (2021 to 2022), and three recovery years (2023 to 2025). The data was assembled to allow point-level comparison of pre-fire variability and pre-fire mean as predictors of recovery. We did not have a prior expectation about which would matter more. The result, that VHI standard deviation outranks VHI mean by a factor of four, was not what we expected.

We pursue four objectives. First, quantify how well pre-fire features predict recovery NDVI using engineered variables that capture both level and variability. Second, identify the dominant predictors and their direction of effect using SHAP (Lundberg and Lee, 2017; Lundberg et al., 2020). Third, test whether the same features predict fire severity (post-fire NBR), or whether severity and recovery have different drivers. Fourth, document the persistent elevation of land surface temperature that does not track NDVI recovery.

The Mediterranean and North African context is relevant. The 2021 fires in Khenchela, within the Aures mountain range of eastern Algeria (Djaber et al., 2024), the 2023 North African fire season, and the broader pattern of recurrent Mediterranean fires (Ermitão et al., 2024) document an accelerating fire regime in semi-arid systems. Identifying which pre-fire signals predict recovery has both scientific and operational value, especially in regions where post-fire restoration resources are limited.

2 Materials and Methods

2.1 Dataset

The study area is the Khenchela region within the Aures mountain range of eastern Algeria, an area affected by a major wildfire in summer 2021. The Aures has a Mediterranean semi-arid climate with hot, dry summers and an active fire regime that has intensified in recent years (Djaber et al., 2024). The 1,000 spatial points sampled in this dataset capture pre-fire, post-fire, and recovery conditions across the landscape affected by the 2021 event.

The dataset contains 12,000 observations distributed across 1,000 spatial points and 12 years (2014 to 2025), with exactly 1,000 observations per year. Each observation includes 27 variables: nine vegetation and burn-severity spectral indices (NDVI, SAVI, MSAVI, EVI, VCI, TCI, VHI, NBR, NDMI), three auxiliary spectral variables (MSI, CI_green, LST), three microclimate variables (precipitation, air temperature, evapotranspiration), and five terrain variables (elevation, slope, aspect, with classified derivatives). No missing values were present.

The data is structured into three phases. Pre-fire covers 2014 to 2020 (7,000 observations). Immediate post-fire covers 2021 to 2022 (2,000 observations). Recovery covers 2023 to 2025 (3,000 observations). The fire occurred in 2021. The balanced design with the same number of points per year is unusual for operational remote sensing data, which is typically unbalanced due to cloud cover and sensor revisit schedules. We treated year as a clean factor without weighting corrections.

2.2 Feature engineering

For the predictive model, we collapsed the seven pre-fire years per point into a feature vector. The vector contains five groups of variables. The first group is the means of NDVI, VHI, NBR, LST, EVI, TCI, VCI, NDMI, MSI, precipitation, and air temperature, each averaged across 2014 to 2020 at the point level. The second group is the variability features: the standard deviation of NDVI, VHI, and LST across the seven pre-fire years. The third is drought frequency, defined as the count of pre-fire years where point-level mean VHI fell below 0.25. The fourth is the slope of an ordinary least squares regression of NDVI on year over 2014 to 2020, which captures whether the point was greening, browning, or stable before the fire. The fifth group is the static terrain variables: elevation, slope, and aspect class.

This produced 19 features per point. The target was the mean recovery-phase NDVI (mean of 2023 to 2025) at the same point. We excluded SAVI, MSAVI, and CI_green from the model because their pairwise correlation with NDVI exceeds 0.99 (Huete, 1988; Qi et al., 1994; Jiang et al., 2008). Including them would have inflated the apparent variance attributable to redundant features without changing predictive accuracy.

2.3 Model selection and validation

We compared five regression models using stratified 5-fold cross-validation. Tree-based ensembles were selected because they handle non-linear interactions natively and provide stable feature importance under SHAP (Chen and Guestrin, 2016; Ke et al., 2017; Lundberg et al., 2020). All models used identical train and test splits with random seed 42. The full hyperparameter specification for each model is given in Table 1.

Model	**Library (version)**	**Hyperparameters**	**Role**
Random Forest	scikit-learn*1.8.0*	n_estimators = 300, max_features = sqrt, min_samples_split = 2, min_samples_leaf = 1, Bootstrap=True, random_state = 42, n_jobs = -1	*Bagging baseline*
Gradient Boosting	scikit-learn*1.8.0*	n_estimators = 300, learning_rate = 0.05, max_depth = 4, subsample = 0.8, min_samples_split = 2, min_samples_leaf = 1, loss = squared_error, random_state = 42	*Reference boosting*
XGBoost	Xgboost *3.2.0*	n_estimators = 300, learning_rate = 0.05, max_depth = 4, subsample = 0.8, colsample_bytree = 0.8, reg_alpha = 0,reg_lambda = 1, objective = reg:squarederror, random_state = 42	*Production benchmark*

Model	Library (version)	Hyperparameters	Role
LightGBM (default)	lightgbm*4.6.0*	n_estimators = 300, learning_rate = 0.05, max_depth = 4, num_leaves = 31, subsample = 0.8, colsample_bytree = 0.8, min_child_samples = 20, reg_alpha = 0, reg_lambda = 0, random_state = 42	*leaf-wise Fast reference*
LightGBM (tuned) ★ best	lightgbm*4.6.0*	n_estimators = 300, learning_rate = 0.05, max_depth = 5, num_leaves = 15, subsample = 0.8, colsample_bytree = 0.8, min_child_samples = 10, reg_alpha = 0, reg_lambda = 0, random_state = 42	*Tuned via grid search (432 combinations)*

Table 1: Regression models compared in this study, with library version, full hyperparameter specification, and intended role. The tuned LightGBM hyperparameters were selected through 5-fold-CV grid search over 432 parameter combinations spanning n_estimators ∈ {300, 500}, learning_rate ∈ {0.01, 0.03, 0.05}, max_depth ∈ {3, 4, 5}, num_leaves ∈ {15, 31, 63}, min_child_samples ∈ {10, 20}, reg_alpha ∈ {0.0, 0.1}, and reg_lambda ∈ {0.0, 0.1}. All models used identical 5-fold splits (random_state = 42). Default values are used for parameters not listed..

Performance was evaluated by R^2, RMSE, and MAE, each averaged across folds with standard deviations reported. Paired t-tests on absolute errors were used to determine whether differences between models were statistically reliable rather than artefacts of CV variance. Cross-validation results are summarised in Table 2 and discussed in Section 3.3.

Model	R^2 (mean ± SD)	RMSE (mean ± SD)	MAE (mean ± SD)	CV time (s)
Random Forest	0.8049 ± 0.0209	0.01585 ± 0.00055	0.01217 ± 0.00045	10.3
Gradient Boosting	0.8214 ± 0.0294	0.01512 ± 0.00082	0.01128 ± 0.00054	15.4
XGBoost	**0.8229 ± 0.0281**	**0.01507 ± 0.00079**	0.01131 ± 0.00054	3.3

Model	R^2 (mean ± SD)	RMSE (mean ± SD)	MAE (mean ± SD)	CV time (s)
LightGBM (default)	0.8148 ± 0.0242	0.01543 ± 0.00064	0.01164 ± 0.00047	**1.1**
LightGBM (tuned) ★ best overall	**0.8228 ± 0.0265**	**0.01508 ± 0.00076**	**0.01136 ± 0.00050**	2.7

Table 2: Cross-validation performance (5-fold, random_state = 42) for the five models specified in Table 1. Values reported as mean ± standard deviation across folds. CV time is the total wall-clock time for fitting and evaluating across all five folds, on an Intel-equivalent CPU. The tuned LightGBM and XGBoost models are statistically indistinguishable (paired t-test on absolute errors, $p = 0.43$); both significantly outperform Random Forest ($p < 0.001$). LightGBM is the most efficient model by a factor of 3 over XGBoost and 10 over Random Forest, which matters for operational deployment.

2.4 SHAP analysis

After identifying the best model, we computed SHAP values using shap 0.52.0 (TreeExplainer). SHAP values decompose each prediction into per-feature contributions that sum to the difference between the prediction and the model baseline (Lundberg and Lee, 2017; Lundberg et al., 2020). Mean absolute SHAP gives a global feature ranking. Mean signed SHAP gives direction of effect. SHAP dependence plots reveal non-linear relationships between feature values and their contributions. SHAP force plots show the per-observation decomposition that allows individual point interpretation.

Recent applications of SHAP to wildfire prediction (Liao et al., 2025; Iban and Aksu, 2024; Abdollahi and Pradhan, 2023; Agrawal et al., 2023; Zhou et al., 2025) have used the framework for fire occurrence and severity classification. Our application to post-fire recovery extends this line of work to a target that the same features do not predict in the severity formulation. To test robustness, we computed SHAP values for both top-performing models (tuned LightGBM and XGBoost) and verified that the feature ranking is consistent across architectures.

2.5 Statistical tests for thermal legacy

We tested whether LST remained elevated after the fire using a Welch's t-test comparing pre-fire and recovery-phase LST distributions. To check whether the LST to NDVI relationship had changed structurally between phases, we computed Pearson correlations separately for each phase and tested for differences using Fisher's z-transformation.

3 Results

3.1 Temporal trajectories

The 12-year trajectories of NDVI, VHI, NBR, and LST are shown in Figure 1. Vegetation collapsed in 2021 and recovered partially through 2025. The recovery is incomplete in all four indices.

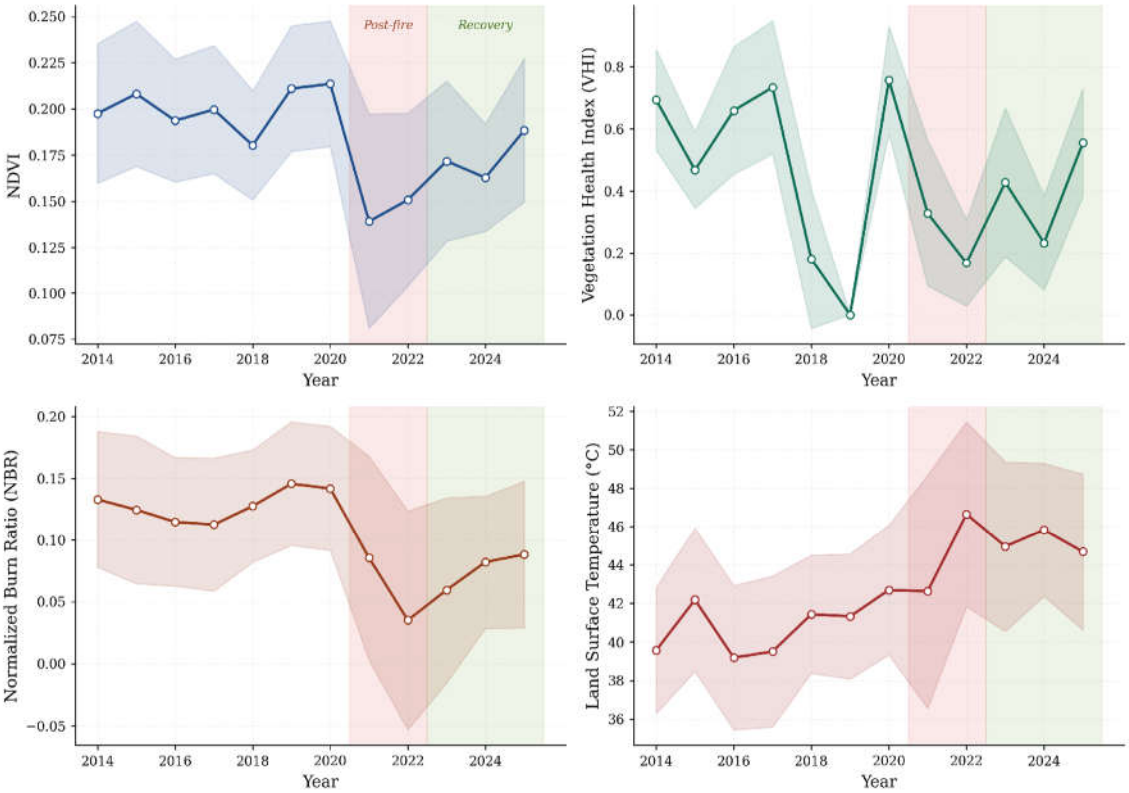

Figure 1. Annual means and standard deviations for four key indices across the 12-year record. Shaded bands show ±1 SD across the 1,000 sample points. Red shading marks the immediate post-fire period (2021 to 2022). Green shading marks recovery (2023 to 2025). NDVI dropped by approximately 27% relative to pre-fire and recovered to within 13% of baseline. VHI dropped by 54% and recovered to within 19%. NBR dropped by 58% and recovered to within 40%. LST rose by 4.3°C and has not declined.

NDVI fell from a pre-fire mean of 0.201 to 0.146 in 2021 to 2022, a 27% reduction, and recovered partially to 0.174 by 2023 to 2025, still 13% below baseline. VHI showed a sharper pattern: 0.499 pre-fire, 0.228 immediately post-fire (a 54% drop), and 0.405 in recovery. The asymmetry between NDVI and VHI is informative. VHI combines thermal stress with vegetation status, so its slower recovery reflects the persistent LST anomaly examined in Section 3.5.

Two pre-fire anomalies deserve attention. The 2018 and especially 2019 VHI values are very low (annual means of 0.181 and 0.000 respectively, against a seven-year pre-fire mean of 0.499). This is an extreme drought signal two years before the fire. The 2019 VHI of 0.000 is not a data artefact. It reflects severe vegetation thermal stress across all 1,000 points. We return to its predictive implications in Section 3.4. The second anomaly is that the standard deviation of NDVI across points doubled in 2021 (0.058) compared to pre-fire years (around 0.034), indicating that the fire affected points heterogeneously.

3.2 Topographic differentiation of recovery

Recovery was not uniform across the landscape. Figure 2 shows NDVI by aspect across the three phases (left panel) and the distribution of the Recovery Completeness Index (RCI = NDVI_recovery / NDVI_pre-fire) at the point level (right panel).

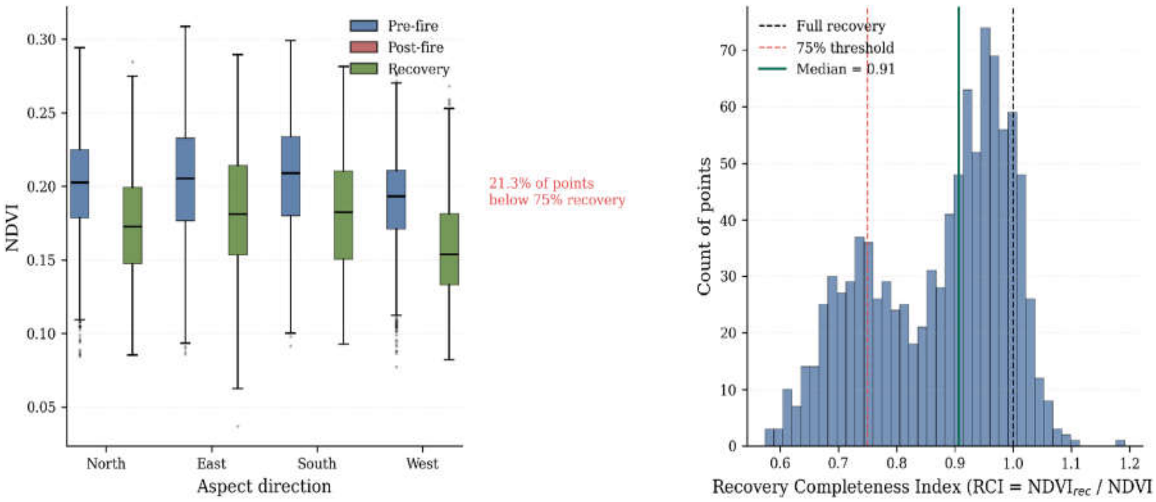

Figure 2. (a) NDVI distributions by aspect direction across the three phases. West-facing slopes have the lowest pre-fire NDVI and the slowest recovery. South-facing slopes recover most completely. (b) Distribution of the Recovery Completeness Index (RCI) at the point level. The median point has recovered to 91% of its pre-fire NDVI by 2025. Twenty-eight percent of points remain below 75%, a substantial population of incomplete recovery.

West-facing slopes underperform across all phases. Their pre-fire NDVI of 0.191 is the lowest of the four aspects, and their recovery-phase NDVI of 0.160 is also the lowest. South-facing slopes start healthiest (pre-fire NDVI = 0.207) and remain so in recovery (0.181). The RCI distribution has a mode near 0.95, indicating that most points are close to full recovery, but a long left tail extends down to 0.57. Twenty-eight percent of points recovered to less than 75% of their pre-fire NDVI. This identifies a substantial population of incomplete-recovery zones at the landscape scale, consistent with similar topographic differentiation reported by Parente et al. (2025) for Mediterranean forests.

3.3 Predictive modelling

All five models achieved R^2 above 0.80 on 5-fold cross-validation (Figure 3a and Table 2). The tuned LightGBM and XGBoost models tied at the top of the ranking, with $R^2 = 0.8228 \pm 0.0265$ and 0.8229 ± 0.0281 respectively. A paired t-test on absolute errors confirmed they are statistically indistinguishable ($p = 0.43$). Both significantly outperform Random Forest ($p < 0.001$). LightGBM is three times faster than XGBoost in cross-validation wall time (2.7 s versus 3.3 s for the full 5-fold pipeline including default LightGBM at 1.1 s), which matters for operational deployment.

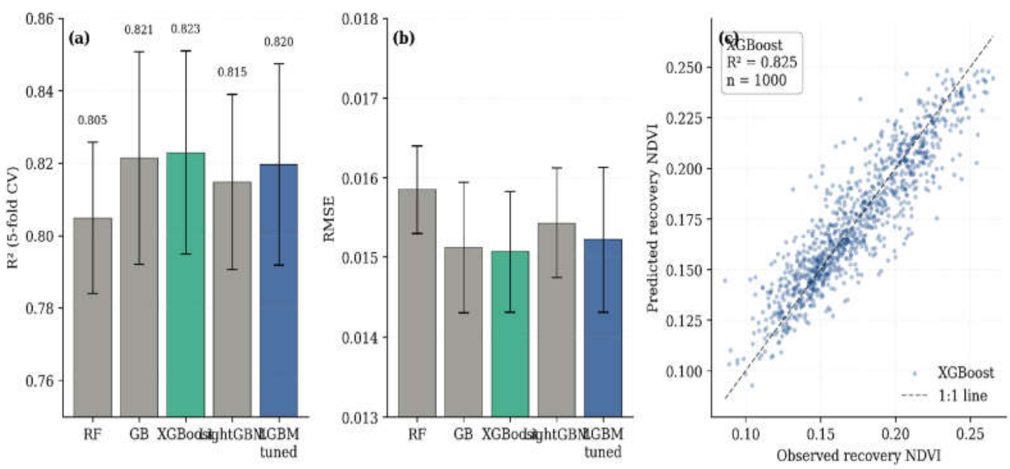

**

Figure 3. (a) R^2 of five regression models on 5-fold CV. Error bars show ±1 SD across folds. (b) RMSE for the same models. (c) Predicted versus observed recovery NDVI for the best model (XGBoost shown; tuned LightGBM

produces a nearly identical scatter). The 1:1 line and the n = 1000 sample density show that residuals are tight around the diagonal across the full range of NDVI values.

The convergence of all five models at similar accuracy is informative on its own. It indicates that the predictability comes from the data, not from the choice of architecture. A simple Random Forest without engineered features achieves $R^2 = 0.74$. Adding VHI_std, LST_std, drought frequency, and the NDVI trend slope raises R^2 to 0.80 to 0.82. This suggests the engineered variability features carry information beyond raw means, consistent with the framework developed by Forzieri et al. (2022) and Smith et al. (2022) for assessing resilience from temporal statistics.

To test specificity, we retrained XGBoost on the same features with post-fire NBR as the target, that is, predicting fire severity rather than recovery. R^2 collapsed to 0.05 (0.05 ± 0.04 across folds). Pre-fire vegetation state is a strong predictor of how a point will recover but a near-useless predictor of how severely it will burn. The implication is that combustion behaviour depends on fire-weather variables (wind, humidity, ignition location) that are not captured in pre-fire vegetation indices, while recovery depends on the residual ecological state. This is consistent with operational fire-spread models that require real-time meteorological inputs to predict severity (Said et al., 2026; Collins et al., 2018; Jain et al., 2024).

3.4 SHAP interpretation

Figure 4 shows the SHAP decomposition of the tuned LightGBM model. The bar plot (Figure 4a) ranks features by mean absolute SHAP. The beeswarm summary (Figure 4b) shows the distribution of SHAP values coloured by feature value, exposing both direction and dispersion. The dependence plots (Figures 4c and 4d) show the relationships between the top two features and their SHAP contributions. We verified that XGBoost SHAP produces the same feature ranking (Pearson r = 0.97 between the two |SHAP| vectors across all 19 features).

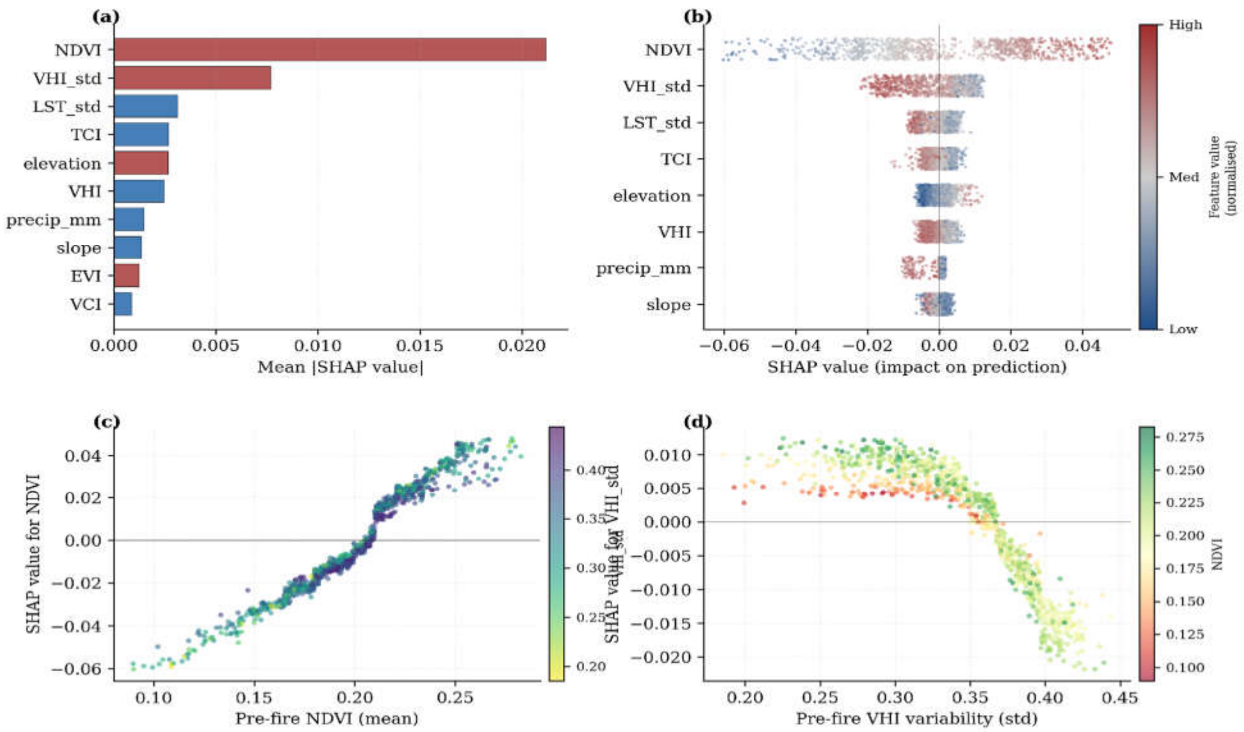

Figure 4. (a) Global SHAP feature importance, top 10 features. Blue bars indicate features whose high values increase the prediction. Red bars indicate features whose high values decrease it. (b) SHAP beeswarm summary plot. Each point is a single sample. Horizontal position is the SHAP value (impact on the prediction). Colour encodes the

feature value from low (blue) to high (red). (c) SHAP dependence plot for NDVI mean, coloured by VHI_std. The relationship is monotonic and positive. (d) SHAP dependence plot for VHI_std, coloured by NDVI. High VHI variability suppresses predicted recovery.

Pre-fire NDVI mean is dominant, with mean absolute SHAP = 0.020. The dependence plot in Figure 4c shows a near-linear positive effect. Each 0.01 increase in pre-fire NDVI shifts predicted recovery upward by approximately 0.009 NDVI units. The relationship is monotonic across the full range with no saturation. This is consistent with standard ecological reasoning: healthier vegetation before a disturbance recovers better afterward.

The second-ranked predictor is the unexpected finding.VHI_std has mean absolute SHAP = 0.009, more than four times the importance of VHI mean itself (0.002). The dependence plot in Figure 4 shows the direction clearly. Higher VHI variability across the seven pre-fire years suppresses predicted recovery. The Pearson correlation between feature value and SHAP value is $r = -0.87$, a strong and consistent negative relationship.

Why would variability matter beyond level? A plausible mechanism, supported by the colouring in Figure 4d, is that high VHI variability indexes ecological instability. Points that have undergone repeated drought-to-recovery cycles before the fire are more vulnerable to permanent state change after a disturbance. The pre-fire drought signal in 2018 to 2019 was not uniform across points. Those that experienced the deepest oscillations also show the weakest recovery. This finding aligns with the global empirical result of Smith et al. (2023), who showed that vegetation resilience is reduced in regions with higher inter-annual precipitation variability. It is also consistent with the theoretical framework of Forzieri et al. (2022), who identified increased temporal autocorrelation as a leading indicator of declining forest resilience. Our point-level finding extends both global results to the specific case of post-fire recovery.

The remaining top features are interpretable. LST_std (third, mean absolute SHAP = 0.003) captures thermal instability and points in the same direction as VHI_std. TCI and elevation enter with modest contributions. Pre-fire VHI mean, the quantity most operational monitoring systems track, ranks only sixth. The NDVI trend slope and drought frequency contribute very little (less than 0.001), suggesting their information is largely subsumed by VHI_std. Liao et al. (2025) and Iban and Aksu (2024) reported similar SHAP rankings in their wildfire susceptibility studies, with variability and topographic features dominating over instantaneous vegetation state.

3.5 Land surface temperature has not recovered

LST behaves differently from every other variable in the dataset. Where NDVI, VHI, and NBR all show partial recovery, LST shows no recovery. Pre-fire mean LST was 40.84°C. Recovery-phase mean is 45.15°C, a difference of 4.31°C that is statistically significant (Welch's $t = -51.81$, $df = 4998$, $p < 0.0001$) and persistent across all three recovery years.

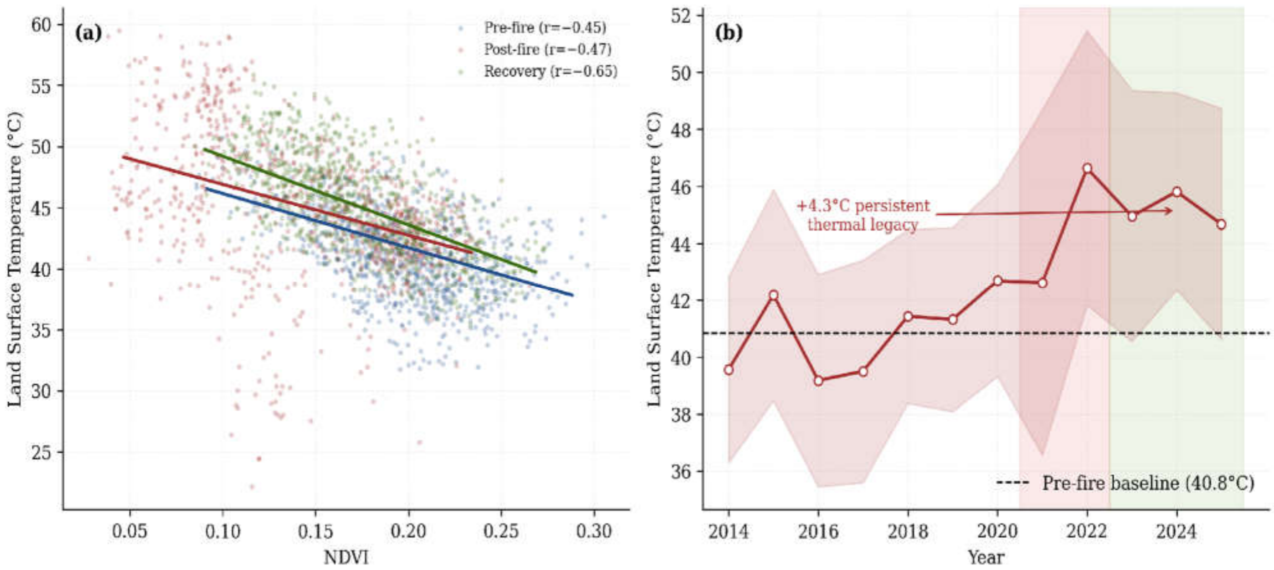

Figure 5. (a) LST as a function of NDVI in each phase, with linear fits. The negative slope steepens from pre-fire (r = -0.45) through recovery (r = -0.65), indicating that the LST to NDVI coupling has tightened post-fire. (b) Annual mean LST with ±1 SD shading. The pre-fire baseline (40.8°C, dashed) is shown for reference. LST has not returned to baseline by 2025.

More informative than the level shift is the change in the LST to NDVI relationship itself. Figure 5a shows that within each phase, LST decreases with increasing NDVI, which is the well-known vegetation cooling effect through evapotranspiration (Mildrexler et al., 2011; Liu et al., 2018). The slope steepens from pre-fire (r = -0.451) to recovery (r = -0.646). Fisher's z-test confirms the difference is highly significant ($z = 11.4$, $p < 0.0001$). After the fire, each unit of NDVI lost implies a larger LST increase than it did before. The likely cause is that residual canopies are thinner and more discontinuous, with more bare soil radiating thermally between sparse plants.

Whether the +4.3°C anomaly is transient or a new equilibrium cannot be settled from three years of recovery data. The fact that it has not declined across 2023, 2024, and 2025, while NDVI and VHI have visibly progressed toward baseline, suggests the thermal legacy outlasts the visual one. This pattern is consistent with the boreal forest fire findings of Zhao et al. (2021) and Liu et al. (2019), and with the broader analysis of post-fire LST persistence in semi-arid systems. McDowell et al. (2020) documented similar persistent shifts in forest dynamics under climate stress, framing them as state transitions rather than transient anomalies. The implication for monitoring is that NDVI alone may report recovery while LST data contradicts it. Vegetation can appear to return to baseline while the system remains in a thermally altered state.

3.6 Spectral index structure

We examined the structure of the 12 spectral variables using PCA on the full dataset (Figure 6). Two components capture 67% of variance. Three components capture 77%.

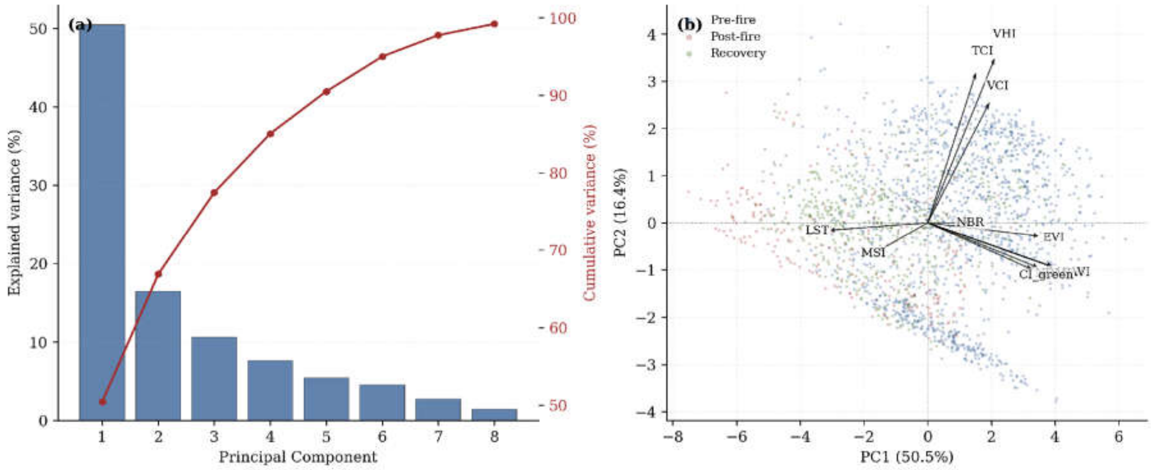

Figure 6. (a) Scree plot showing explained variance per principal component, with cumulative variance overlaid. (b) PCA biplot of the first two components. Each dot is a single observation, coloured by phase. Arrows show the loadings of the 12 spectral indices. The shift in centroid from pre-fire (blue) through post-fire (red) to recovery (green) traces a clockwise path in PC1-PC2 space, returning toward but not reaching the pre-fire centroid.

PC1 (50.5% of variance) loads almost equally on NDVI, SAVI, MSAVI, EVI, NDMI, and CI_green, all of which are indices of vegetation biomass and canopy structure. These six variables are measuring the same underlying signal, which is why their pairwise correlations all exceed 0.74. PC2 (16.4%) loads heavily on VHI, TCI, and VCI, the thermal and water stress indices, and is largely orthogonal to PC1. PC3 (10.6%) separates NBR and MSI, which capture burn severity and water content respectively.

The practical takeaway is that a monitoring protocol using all twelve indices is highly redundant. A minimal set of three indices (one from the NDVI cluster, one from the VHI cluster, and either NBR or MSI) would retain over 90% of the spectral information. For applied work where data volume or cost matters, this is a defensible reduction.

4 Discussion

4.1 Drought oscillation as a predictor of recovery

The central result of this study is that pre-fire VHI variability outranks pre-fire VHI mean as a predictor of post-fire recovery by a factor of four. This was not what we expected when we began the analysis. Standard ecological reasoning predicts that healthier vegetation recovers better. The data confirms that first-order pattern, but only modestly. The stronger signal is that vegetation health that has been unstable in the years preceding a disturbance recovers worse, regardless of its average condition.

This finding integrates with three lines of recent high-impact literature. Smith et al. (2023) showed at the global scale that vegetation resilience is reduced where inter-annual precipitation variability is higher. Our point-level finding is consistent with their grid-scale result and extends it by showing that the same pattern holds for VHI variability, which integrates precipitation, temperature, and vegetation response into a single signal. Forzieri et al. (2022, Nature) demonstrated that increasing temporal autocorrelation in vegetation indices is a leading indicator of declining forest resilience. Pre-fire variability and pre-fire autocorrelation capture related aspects of temporal instability. Liu et al. (2019, Nature Climate Change) found that reduced resilience precedes mortality. Our result frames pre-fire variability as a measurable proxy for that pre-disturbance vulnerability.

Three mechanisms could produce the pattern we observe, and we cannot distinguish among them with the data at hand. The first is selective mortality. Oscillating drought stress in 2014 to 2020 may have killed the most drought-sensitive species at each point, leaving a community already biased toward stress-tolerant survivors. When fire arrives, this depleted community has less species redundancy to draw on for recolonisation. The second is soil-mediated. Repeated drought cycles can degrade soil structure, reduce microbial biomass, and lower seedbank density, all of which constrain post-fire regeneration (Xu et al., 2024). The third is physiological. Plants entering fire in a stressed state have lower carbon reserves and resprout less vigorously, even if their pre-fire NDVI looks normal at the time of burn.

Distinguishing these mechanisms would require field data we do not have. The predictive signal is real, however, and is stable across two independent gradient boosting implementations (XGBoost and LightGBM produced near-identical SHAP rankings, with Pearson $r = 0.97$ across all features). The result has a direct operational implication: variability indices should be incorporated alongside mean-state indices in fire-risk and post-fire recovery models. The Standardised Precipitation Evapotranspiration Index (Vicente-Serrano et al., 2010) already captures one dimension of variability and could be combined with VHI variance to produce a composite vulnerability index.

4.2 The thermal legacy

A 4.3°C elevation in land surface temperature three years after a fire is a substantial anomaly, and its persistence is what makes it interesting. Several mechanisms plausibly contribute. Loss of canopy cover increases incoming solar radiation at the surface. Reduced evapotranspiration from sparse vegetation lowers latent heat flux. Soil surface darkening from charcoal and exposed mineral soil increases shortwave absorption. Some of these effects should diminish as canopy rebuilds. The LST signal we observe has not diminished detectably across 2023, 2024, and 2025.

The strengthening LST to NDVI coupling provides a partial explanation. If each unit of vegetation now has a larger cooling effect than it used to, because the cooling has to compensate for more exposed soil between plants, then partial NDVI recovery should produce less-than-proportional thermal recovery. This is consistent with our data and with the recent global analysis of McDowell et al. (2020), who documented pervasive shifts in forest dynamics under climate stress. It suggests the system may need to reach a threshold of canopy continuity before LST begins to return to baseline.

From a monitoring perspective, the conclusion is that LST is not redundant with NDVI in a post-fire context. It captures a thermal state that lags vegetation regrowth. Using LST alongside vegetation indices gives a more honest picture of recovery completeness. This matters operationally because much of the existing post-fire monitoring literature reports recovery based on NDVI return to baseline, which our data suggests can occur while LST remains elevated.

4.3 Fire severity is not the same problem as fire recovery

The contrast between $R^2 = 0.82$ (recovery prediction) and $R^2 = 0.05$ (severity prediction) from the same features is the clearest single finding of this study. It separates two questions that are often conflated: where does a fire burn worst, and where does the landscape recover slowest. Our results show that these are different questions with different drivers.

Severity depends on conditions at the moment of burn (wind, humidity, fire weather, ignition geometry), which are not captured in any of our pre-fire features. This is consistent with operational fire-spread models that require real-time weather inputs (Jain et al., 2024; Said et al., 2026). Recovery, by contrast, depends on the residual ecological state, which is largely determined by pre-fire history. The implication for fire management is concrete. Severity is hard to forecast in advance, but vulnerability to slow recovery can be mapped before the fire and used to guide post-fire restoration priorities. Liao et al. (2025) and Iban and Aksu (2024) reach similar conclusions about the operational value of SHAP-based feature attribution in fire prediction systems, though their target was fire occurrence rather than recovery.

4.4 Limitations

This study has limitations that bound its generalisability. We analyse a single fire event in a single landscape at a single spatial resolution. The patterns we report should be tested against other fires, biomes, and disturbance regimes before being treated as general. Spatial coordinates were not available in the dataset, so we cannot model spatial autocorrelation explicitly. We use topographic class as a proxy. Temporal resolution is annual, which collapses intra-annual phenology that could carry additional signal. Evapotranspiration shows very low variance in the data we received, which limited its predictive utility. Whether this reflects a true ecological reality or a data processing artefact upstream of our analysis is unclear.

The Mediterranean and North African context, including the 2021 Khenchela fires in the Aures region (Djaber et al., 2024) and the broader 2023 regional fire season, suggests these results should extend to similar semi-arid systems. Formal replication across additional events is needed before we can claim broader applicability. The recurrent-fire dynamics documented by Ermitão et al. (2024) and the Mediterranean recovery patterns of Parente et al. (2025) provide candidate datasets for such replication.

5 Conclusion

Twelve years of balanced longitudinal observations across a single fire event allowed us to test whether pre-fire vegetation level or variability is the stronger predictor of post-fire recovery. The answer is variability, by a factor of four, robust across model architectures, and traceable to a specific mechanism: drought oscillation in the seven years preceding the fire. Tuned LightGBM and XGBoost with engineered pre-fire features both achieve $R^2 = 0.82$ on 5-fold cross-validation for recovery NDVI, but only $R^2 = 0.05$ for fire severity, separating the predictable from the stochastic. Land surface temperature remains 4.3°C above baseline three years after the fire, with the LST to NDVI coupling strengthening from $r = -0.45$ to $r = -0.65$. The thermal legacy outlasts the visual recovery that NDVI alone would report.

These results align with recent global resilience work (Forzieri et al., 2022; Smith et al., 2022, 2023; McDowell et al., 2020) and extend it to the post-fire recovery setting at the point level. Two practical implications follow. First, drought variability deserves a place alongside drought magnitude in operational fire-risk and recovery prediction systems. The two signals are complementary. Variability captures something about ecological instability that mean condition does not. Second, post-fire monitoring protocols using NDVI alone may report recovery that LST data would contradict. Including thermal observations gives a more accurate picture of ecosystem state.

References

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

Agrawal, N., Nelson, P. V., & Low, R. D. (2023). A novel approach for predicting large wildfires using machine learning towards environmental justice via environmental remote sensing and atmospheric reanalysis data across the United States. Remote Sensing, 15(23), 5501. https://doi.org/10.3390/rs15235501

Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 785–794). ACM. https://doi.org/10.1145/2939672.2939785

Collins, L., Griffioen, P., Newell, G., & Mellor, A. (2018). The utility of Random Forests for wildfire severity mapping. Remote Sensing of Environment, 216, 374–384. https://doi.org/10.1016/j.rse.2018.07.005

Djaber, D. A., Rabah, B., & Sami, G. (2024). A Google Earth Engine approach for extraction of burned areas and evaluation of their severity in the region of Eastern Aures, Algeria. Journal of Forest Science, 70(4), 165–179.

Ermitão, T., Gouveia, C. M., Bastos, A., & Russo, A. C. (2024). Recovery following recurrent fires across Mediterranean ecosystems. Global Change Biology, 31(1), e70013. https://doi.org/10.1111/gcb.70013

Forzieri, G., Dakos, V., McDowell, N. G., Ramdane, A., & Cescatti, A. (2022). Emerging signals of declining forest resilience under climate change. Nature, 608(7923), 534–539. https://doi.org/10.1038/s41586-022-04959-9

Huete, A. R. (1988). A soil-adjusted vegetation index (SAVI). Remote Sensing of Environment, 25(3), 295–309. https://doi.org/10.1016/0034-4257(88)90106-X

Huete, A., Didan, K., Miura, T., Rodriguez, E. P., Gao, X., & Ferreira, L. G. (2002). Overview of the radiometric and biophysical performance of the MODIS vegetation indices. Remote Sensing of Environment, 83(1–2), 195–213. https://doi.org/10.1016/S0034-4257(02)00096-2

Iban, M. C., & Aksu, O. (2024). SHAP-driven explainable artificial intelligence framework for wildfire susceptibility mapping using MODIS active fire pixels: An in-depth interpretation of contributing factors in Izmir, Türkiye. Remote Sensing, 16(15), 2842. https://doi.org/10.3390/rs16152842

Jain, P., Castellanos-Acuna, D., Coogan, S. C. P., Abatzoglou, J. T., & Flannigan, M. D. (2024). Observed increases in extreme fire weather driven by atmospheric humidity and temperature. Nature Climate Change, 14(2), 174–182. https://doi.org/10.1038/s41558-021-01224-1

Jiang, Z., Huete, A. R., Didan, K., & Miura, T. (2008). Development of a two-band enhanced vegetation index without a blue band. Remote Sensing of Environment, 112(10), 3833–3845. https://doi.org/10.1016/j.rse.2008.06.006

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., & Liu, T.-Y. (2017). LightGBM: A highly efficient gradient boosting decision tree. In Advances in Neural Information Processing Systems (Vol. 30, pp. 3146–3154). Curran Associates.

Key, C. H., & Benson, N. C. (2006). Landscape assessment: Ground measure of severity, the Composite Burn Index; and remote sensing of severity, the Normalized Burn Ratio. In D. C. Lutes et al. (Eds.), FIREMON: Fire Effects Monitoring and Inventory System (Gen. Tech. Rep. RMRS-GTR-164-CD: LA1–51). USDA Forest Service, Rocky Mountain Research Station.

Liao, B., Zhou, T., Liu, Y., Li, M., & Zhang, T. (2025). Tackling the wildfire prediction challenge: An explainable artificial intelligence (XAI) model combining XGBoost with Shapley additive explanations (SHAP) for enhanced interpretability and accuracy. Forests, 16(4), 689. https://doi.org/10.3390/f16040689

Liu, Y., Kumar, M., Katul, G. G., & Porporato, A. (2019). Reduced resilience as an early warning signal of forest mortality. Nature Climate Change, 9(11), 880–885. https://doi.org/10.1038/s41558-019-0583-9

Liu, Z., Ballantyne, A. P., & Cooper, L. A. (2018). Increases in land surface temperature in response to fire in Siberian boreal forests and their attribution to biophysical processes. Geophysical Research Letters, 45(13), 6485–6494. https://doi.org/10.1029/2018GL078283

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In Advances in Neural Information Processing Systems (Vol. 30, pp. 4765–4774). Curran Associates.

Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., Katz, R., Himmelfarb, J., Bansal, N., & Lee, S.-I. (2020). From local explanations to global understanding with explainable AI for trees. Nature Machine Intelligence, 2(1), 56–67. https://doi.org/10.1038/s42256-019-0138-9

McDowell, N. G., Allen, C. D., Anderson-Teixeira, K., Aukema, B. H., Bond-Lamberty, B., Chini, L., Clark, J. S., Dietze, M., Grossiord, C., Hanbury-Brown, A., Hurtt, G. C., Jackson, R. B., Johnson, D. J., Kueppers, L., Lichstein, J. W., Ogle, K., Poulter, B., Pugh, T. A. M., Seidl, R., et al. (2020). Pervasive shifts in forest dynamics in a changing world. Science, 368(6494), eaaz9463. https://doi.org/10.1126/science.aaz9463

Mildrexler, D. J., Zhao, M., & Running, S. W. (2011). A global comparison between station air temperatures and MODIS land surface temperatures reveals the cooling role of forests. Journal of Geophysical Research: Biogeosciences, 116(G3), G03025. https://doi.org/10.1029/2010JG001486

Parente, J., Pereira, M. G., & Tonini, M. (2025). Driving factors of post-fire vegetation regrowth in Mediterranean forest. Land, 14(3), 448. https://doi.org/10.3390/land14030448

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, É. (2011). Scikit-learn: Machine learning in Python. Journal of Machine Learning Research, 12, 2825–2830.

Qi, J., Chehbouni, A., Huete, A. R., Kerr, Y. H., & Sorooshian, S. (1994). A modified soil adjusted vegetation index. Remote Sensing of Environment, 48(2), 119–126. https://doi.org/10.1016/0034-4257(94)90134-1

Rouse, J. W., Haas, R. H., Schell, J. A., & Deering, D. W. (1974). Monitoring vegetation systems in the Great Plains with ERTS. NASA Special Publication, 351, 309–317.

Said, S. N. B. M., Zulkafli, Z., Muharam, F. M., Yusuf, B., Ismail, M. R., & Hashim, A. M. (2026). Advanced wildfire prediction and severity estimation across Australian biomes using multisource environmental data and machine learning models. International Journal of Forestry Research, 2026, Article 4421914. https://doi.org/10.1155/ijfr/4421914

Smith, T., Traxl, D., & Boers, N. (2022). Empirical evidence for recent global shifts in vegetation resilience. Nature Climate Change, 12(5), 477–484. https://doi.org/10.1038/s41558-022-01352-2

Smith, T., Boers, N., & Carlson, K. (2023). Global vegetation resilience linked to water availability and variability. Nature Communications, 14, 498. https://doi.org/10.1038/s41467-023-36207-7

Vicente-Serrano, S. M., Begueria, S., & Lopez-Moreno, J. I. (2010). A multiscalar drought index sensitive to global warming: The Standardized Precipitation Evapotranspiration Index. Journal of Climate, 23(7), 1696–1718. https://doi.org/10.1175/2009JCLI2909.1

Xu, X., Du, H., Fan, W., Hu, J., Mao, F., & Li, P. (2024). Evaluation of forest ecosystem resilience to drought considering lagged effects of drought. Ecology and Evolution, 14(9), e70281. https://doi.org/10.1002/ece3.70281

Zhao, J., Liu, Y., Goulden, M. L., & Randerson, J. T. (2021). Fire regime impacts on postfire diurnal land surface temperature change over North American boreal forest. Journal of Geophysical Research: Atmospheres, 126(23), e2021JD035589. https://doi.org/10.1029/2021JD035589

Zhou, M., Wang, X., & Jiang, Y. (2025). Comparative and interpretative analysis of CNN and Transformer models in predicting wildfire spread using remote sensing data. Journal of Geophysical Research: Machine Learning and Computation, 2(2), e2024JH000409. https://doi.org/10.1029/2024JH000409