TerraHeal: An Automated Framework for Identifying Post-Wildfire Vegetation Recovery Anomalies Using Sentinel-2 Imagery and Machine Learning

Atharv Khare
Independent Researcher
Bhopal, India

Abstract

Monitoring vegetation recovery after wildfire is essential for ecosystem management, but many operational products emphasize immediate burn severity rather than longer-term recovery dynamics. This paper presents TerraHeal, an automated satellite-based framework that detects persistent post-wildfire recovery anomalies, which we refer to as *recovery cold spots*. Using Google Earth Engine (GEE) and 10 m Sentinel-2 imagery, we derive multi-temporal Normalized Difference Vegetation Index (NDVI) composites at fixed intervals (3, 6, 12, 18, and 24 months after fire) across multiple wildfire events. We implement two complementary approaches to anomaly detection. The first uses heuristic, rule-based thresholds that compare post-fire recovery to pre-fire baselines. The second applies an unsupervised Isolation Forest model to identify pixel-level deviations in recovery trajectories. We then verify whether detected cold spots persist through 24 months to distinguish transient setbacks from sustained recovery failures. Applied to major California wildfires, including the Camp Fire and Dixie Fire, TerraHeal highlights localized areas where vegetation regrowth remains substantially delayed relative to expected recovery patterns. By moving beyond burn-severity mapping to operational detection of recovery failures, the framework supports targeted restoration planning. The GEE-based workflow is reproducible and designed to scale to broader wildfire monitoring applications.

Keywords: Wildfire recovery, NDVI, Sentinel-2, Google Earth Engine, Isolation Forest, anomaly detection, remote sensing

1 Introduction

Wildfires are increasing in frequency and severity globally, driven by climate change, fuel accumulation, and expanding wildland-urban interfaces [1, 2]. While immediate fire impacts receive substantial attention through burn severity mapping, the long-term ecological trajectory of burned landscapes remains understudied. Traditional wildfire assessment tools focus on quantifying initial damage using indices such as the differenced Normalized Burn Ratio (dNBR) [3, 4], yet provide limited insight into post-fire vegetation recovery dynamics.

Post-fire ecosystem recovery is a complex, non-linear process influenced by burn severity, pre-fire vegetation composition, soil degradation, topography, and climate variability [5, 6]. While most burned areas exhibit natural regeneration, certain locations experience anomalously slow or stalled vegetation regrowth—what we term "recovery cold spots." These areas may indicate severe soil degradation, loss of native seed banks, or increased vulnerability to invasive species colonization [7]. Identifying such locations is critical for prioritizing limited restoration resources and preventing long-term ecosystem state transitions.

Recent advances in satellite remote sensing and cloud computing have enabled systematic monitoring of post-fire landscapes at unprecedented spatial and temporal scales [8]. The Normalized Difference Vegetation Index (NDVI), derived from multispectral imagery, is a widely used proxy for vegetation condition and photosynthetic activity [9]. Time-series analyses of NDVI have been used to characterize post-fire recovery trajectories [10, 11], although many studies rely on averaged recovery curves that can obscure localized anomalies.

This research addresses a critical gap in operational wildfire recovery monitoring by developing an automated framework for pixel-level anomaly detection in post-fire vegetation recovery. Our approach combines high-resolution satellite imagery (Sentinel-2), cloud-based geospatial processing (Google Earth Engine), and unsupervised machine learning (Isolation Forest) to systematically identify and characterize recovery cold spots. Unlike conventional approaches that average recovery patterns, our method explicitly targets outlier pixels whose recovery trajectories deviate significantly from expected baselines.

1.1 Research Objectives

This study pursues five objectives. First, we develop an automated, reproducible framework to compute multi-temporal NDVI profiles at standardized post-fire intervals using Sentinel-2 Surface Reflectance data. Second, we implement two complementary anomaly-detection strategies: heuristic, rule-based thresholds and an unsupervised Isolation Forest model. Third, we evaluate whether detected cold spots persist through 24 months to separate transient effects from sustained recovery failures. Fourth, we demonstrate operational applicability across multiple large wildfires in California. Finally, we provide reusable code and a documented workflow intended to scale to other regions.

1.2 Contributions

This work contributes an operational framework that shifts emphasis from immediate burn severity to longer-horizon recovery outcomes. Methodologically, it analyzes recovery trajectories at the pixel level rather than relying on spatial averages that can mask localized failures. It also introduces an explicit persistence check at 12 and 24 months so that short-lived anomalies are not conflated with sustained cold spots. Finally, it integrates an unsupervised Isolation Forest model within a reproducible, Google Earth Engine pipeline designed for scalable processing.

Nomenclature

Symbol/Term	Definition
NDVI	Normalized Difference Vegetation Index
NBR	Normalized Burn Ratio
dNBR	Differenced Normalized Burn Ratio
GEE	Google Earth Engine
IF	Isolation Forest
SR	Surface Reflectance
NIR	Near-Infrared
SWIR	Shortwave Infrared
t_0	Immediate post-fire period (0–3 months)
t_3	3-month post-fire interval
t_6	6-month post-fire interval
t_{12}	12-month post-fire interval
t_{18}	18-month post-fire interval
t_{24}	24-month post-fire interval
R_{3-12}	Recovery rate from 3 to 12 months
R_{12-24}	Recovery rate from 12 to 24 months
P_{25}	25th percentile

2 Related Work

2.1 Burn Severity Assessment

The Normalized Burn Ratio (NBR) and its differenced variant (dNBR) have become standard tools for wildfire severity mapping [3]. Developed by Key and Benson [3], dNBR compares pre-fire and immediate post-fire spectral signatures in the near-infrared (NIR) and shortwave infrared (SWIR) bands:

$$\text{NBR} = \frac{\text{NIR} - \text{SWIR}}{\text{NIR} + \text{SWIR}} \tag{1}$$

$$\text{dNBR} = \text{NBR}_{\text{pre-fire}} - \text{NBR}_{\text{post-fire}} \tag{2}$$

While dNBR effectively quantifies immediate burn severity [4, 24], it provides limited information about subsequent recovery dynamics. The Monitoring Trends in Burn Severity (MTBS) program has mapped large fires across the United States since 1984 using dNBR [25], establishing a valuable baseline but not tracking recovery outcomes.

2.2 Vegetation Recovery Monitoring

Several studies have characterized post-fire vegetation recovery using NDVI time-series. White et al. [11] analyzed NBR recovery in western North American forests, finding that recovery averaged 30-44% at 5 years and 47-72% at 10 years post-fire, with substantial variability by forest type and severity. Meng et al. [5] demonstrated that post-fire climate, particularly precipitation and minimum temperatures, significantly influences NDVI recovery rates in Sierra Nevada mixed conifer forests.

Recent work has begun applying Google Earth Engine for post-fire monitoring. Zahabnazouri et al. [12] used Sentinel-2 imagery and dNDVI to compare recovery across different fire

history zones in Mediterranean ecosystems, finding that repeated burns led to slower and incomplete recovery. Similarly, Wassner et al. [13] demonstrated that low-severity zones recover more rapidly than high-severity areas, with NDVI increasing most rapidly in the first 2 years post-fire.

However, these studies typically analyze mean recovery trends rather than identifying specific anomalous pixels. Our approach explicitly targets outlier detection to guide targeted management interventions.

2.3 Machine Learning in Fire Science

Machine learning has been increasingly applied to wildfire prediction and severity mapping. Random Forest and Support Vector Machines have been used for burn severity classification [14], while more recent work has explored XGBoost and neural networks for fire risk prediction [15].

The Isolation Forest algorithm [16, 17] is an unsupervised anomaly detection method particularly well-suited to high-dimensional data where anomalies are rare and difficult to detect through simple thresholds. Originally developed for fraud detection and network security, it has recently been applied to remote sensing applications including hyperspectral anomaly detection [18, 19] and land cover change detection [20]. To our knowledge, this is the first application of Isolation Forest to post-fire vegetation recovery trajectory analysis.

2.4 Google Earth Engine Applications

Google Earth Engine (GEE) [8] has revolutionized large-scale remote sensing analysis by providing cloud-based access to multi-petabyte archives of satellite imagery with integrated processing capabilities. Recent wildfire applications include automated burned area mapping [21], multi-index burn severity assessment [22], and wildfire susceptibility prediction [23].

Our framework leverages GEE's capabilities for automated Sentinel-2 time-series processing, enabling reproducible analysis at scale without local computational infrastructure requirements.

3 Study Area and Data

3.1 Wildfire Events

We selected three major California wildfires as case studies that span different ecosystems and severities. The Camp Fire ignited on November 8, 2018 and burned 153,336 acres in Butte County; it was among California's most destructive events, causing 85 fatalities and destroying much of Paradise, California. The Dixie Fire ignited on July 13, 2021 and burned 963,309 acres across five counties, becoming California's largest single wildfire at the time; it affected mixed conifer, ponderosa pine, and montane hardwood ecosystems. The August Complex ignited on August 17, 2020 and burned 1,032,648 acres in Northern California; it comprised multiple ignitions and primarily affected the Mendocino National Forest.

Fire perimeter data were obtained from the California Department of Forestry and Fire Protection (CAL FIRE) Fire and Resource Assessment Program (FRAP) in GeoJSON format.

3.2 Satellite Data

3.2.1 Sentinel-2 Surface Reflectance

We used the Sentinel-2 Multi-Spectral Instrument (MSI) Level-2A Surface Reflectance product (COPERNICUS/S2_SR_HARMONIZED) accessed through Google Earth Engine. Sentinel-2 provides 10 m spatial resolution for the visible and near-infrared bands (B2, B3, B4, and B8) and an effective 5-day revisit time when combining Sentinel-2A and Sentinel-2B. The instrument

includes 13 spectral bands spanning the visible through SWIR, and Level-2A products are atmospherically corrected using the Sen2Cor processor. For quality control, we used the scene classification (SCL) band to mask clouds and related artifacts.

Sentinel-2's higher spatial resolution (10 m vs. 30 m for Landsat) enables detection of finer-scale recovery heterogeneity, particularly important in fragmented landscapes.

3.2.2 Temporal Windows

For each wildfire, we defined the following temporal windows:

The pre-fire baseline window spans the two years prior to the ignition date (for example, 2016–2018 for the Camp Fire). Post-fire recovery is summarized using six three-month windows: t_0 (0–3 months), t_3 (3–6 months), t_6 (6–9 months), t_{12} (12–15 months), t_{18} (18–21 months), and t_{24} (24–27 months) after ignition.

Three-month windows allowed sufficient image acquisition while maintaining temporal specificity.

4 Methodology

4.1 Preprocessing Pipeline

4.1.1 Cloud and Shadow Masking

Sentinel-2 Level-2A products include a Scene Classification Layer (SCL) that labels clouds, cloud shadows, snow, and related features. We applied strict quality filtering by masking cloud shadows (class 3), medium-probability clouds (class 8), high-probability clouds (class 9), thin cirrus (class 10), and snow/ice (class 11).

This masking approach ensures that computed NDVI values represent actual vegetation rather than atmospheric or snow contamination.

4.1.2 NDVI Computation

The Normalized Difference Vegetation Index (NDVI) quantifies vegetation greenness and photosynthetic activity:

$$\text{NDVI} = \frac{\text{NIR} - \text{Red}}{\text{NIR} + \text{Red}} = \frac{B8 - B4}{B8 + B4} \tag{3}$$

where B8 (NIR) and B4 (Red) are Sentinel-2 bands. NDVI values range from -1 to $+1$, with healthy vegetation typically > 0.6.

Within each temporal window, we filtered the Sentinel-2 collection by date range and fire perimeter, applied the cloud and shadow masks, computed NDVI for each image, and then produced a median composite to reduce residual cloud contamination and temporal noise.

4.2 Feature Engineering

We constructed a multi-band feature stack for each fire perimeter. It includes the median pre-fire baseline NDVI (two-year window) and the median NDVI composites for $t_0, t_3, t_6, t_{12}, t_{18}$, and t_{24}. We also compute interval-specific NDVI drops, $\Delta\text{NDVI} = \text{NDVI}_{\text{baseline}} - \text{NDVI}_{t_i}$, and two recovery-rate features:

$$R_{3-12} = \frac{\text{NDVI}_{t_{12}} - \text{NDVI}_{t_3}}{9 \text{ months}} \tag{4}$$

$$R_{12-24} = \frac{\text{NDVI}_{t_{24}} - \text{NDVI}_{t_{12}}}{12 \text{ months}} \tag{5}$$

This resulted in a 15-band feature stack per pixel, capturing both absolute values and temporal dynamics.

4.3 Anomaly Detection Approaches

We implemented two complementary methodologies for identifying recovery cold spots:

4.3.1 Heuristic Rule-Based Detection

We defined cold spots using two criteria:

Rule 1 - Absolute threshold: Pixels where NDVI at 12 months remains below 50% of baseline:

$$\text{Cold Spot}_{\text{threshold}} = \text{NDVI}_{t_{12}} < 0.5 \times \text{NDVI}_{\text{baseline}} \tag{6}$$

This threshold is ecologically meaningful, as vegetation recovering to less than half its pre-fire vigor after one year indicates severe regeneration limitation.

Rule 2 - Relative recovery rate: Pixels with recovery rates in the bottom 25th percentile of the study area:

$$\text{Cold Spot}_{\text{percentile}} = R_{3-12} < P_{25}(R_{3-12}) \tag{7}$$

where P_{25} denotes the 25th percentile. A pixel was classified as a cold spot if either criterion was met:

$$\text{Cold Spot}_{\text{heuristic}} = \text{Cold Spot}_{\text{threshold}} \vee \text{Cold Spot}_{\text{percentile}} \tag{8}$$

4.3.2 Isolation Forest Machine Learning

The Isolation Forest (IF) algorithm [16] is an unsupervised outlier detection method based on the principle that anomalies are "few and different" and can therefore be isolated more quickly by recursive partitioning.

At a high level, the method repeatedly selects a feature and a split value at random to partition the data into binary trees. Points that can be isolated in fewer splits are assigned higher anomaly scores. The anomaly score for a sample x in a dataset of size n is

$$s(x, n) = 2^{-\frac{E(h(x))}{c(n)}}, \tag{9}$$

where $E(h(x))$ is the expected path length for x across the forest and $c(n)$ is the average path length of an unsuccessful search in a binary tree of size n.

Shorter average path lengths (easier isolation) indicate higher anomaly probability.

Implementation:

We exported the 15-band feature stack from GEE as a GeoTIFF and processed it in Python. The raster was read with `rasterio`, reshaped into a pixel-by-feature array ($n_{\text{pixels}} \times 15$), and masked to remove NoData values. We then fit an Isolation Forest model using `scikit-learn` with a contamination parameter of 0.1, corresponding to the assumption that roughly 10% of pixels within a fire perimeter exhibit anomalous recovery trajectories.

4.4 Persistence Verification

To distinguish transient recovery delays from sustained failures, we used a two-stage verification scheme. Cold spots were first identified at 12 months using both the heuristic rules and the

machine-learning model. We then evaluated persistence by checking whether pixels flagged at t_{12} still met the absolute NDVI threshold at t_{24}:

$$\text{Persistent Cold Spot} = \text{Cold Spot}_{t_{12}} \wedge (\text{NDVI}_{t_{24}} < 0.5 \times \text{NDVI}_{\text{baseline}}). \tag{10}$$

Only pixels meeting both criteria at 12 and 24 months were classified as persistent recovery failures.

4.5 Visualization and Analysis

We generated interactive maps using the `geemap` Python library for GEE. The map products included the pre-fire baseline NDVI, NDVI composites for each post-fire interval, the heuristic cold-spot layer, the Isolation Forest anomaly layer, the persistent cold-spot layer, and the fire-perimeter overlay.

All layers were clipped to fire boundaries and visualized using consistent color palettes for interpretation.

5 Results

5.1 Temporal Analysis: Mean NDVI Recovery Trajectory

To evaluate temporal vegetation dynamics after fire, we summarize mean NDVI from the pre-fire baseline through the post-fire intervals t_0, t_3, t_6, t_{12}, t_{18}, and t_{24} (Fig. 1). The trajectory is non-linear and indicates that recovery does not progress monotonically over time.

The most prominent feature is the pronounced decline at t_{12}, which is consistent with either strong seasonal effects in vegetation greenness or a secondary disturbance signal in the recovery window. Although NDVI rebounds after this dip, the temporal path remains volatile rather than smooth, indicating heterogeneous ecological responses across the burned landscape.

By t_{24}, mean NDVI remains below the pre-fire baseline, suggesting that full canopy or productivity restoration has not yet been reached within two years. This supports the need for continued monitoring beyond short post-fire horizons when identifying persistent recovery deficits.

5.2 Statistical Distribution: Recovery Slope Density

To characterize pixel-level recovery dynamics, we analyzed the distribution of recovery slopes using a kernel density estimate (Fig. 2). The red dashed line at zero denotes the no-growth threshold separating positive and non-positive recovery rates.

The distribution exhibits a weak bimodal tendency, indicating that the burned area contains at least two broad recovery behaviors: slowly improving pixels and a subset with stagnant or negative trends. Density near and below zero aligns with the ecological interpretation of recovery cold spots.

The x-axis scale on the order of 10^{-12} indicates that monthly slope magnitudes are very small in absolute terms, even where differences are ecologically meaningful. This narrow spread supports using percentile- or anomaly-based criteria rather than absolute slope cutoffs alone.

5.3 NDVI Recovery Patterns

Representative NDVI time-series for the Camp Fire show characteristic recovery trajectories. Pre-fire NDVI (2016-2018 baseline) averaged 0.68 across forested areas, indicating healthy vegetation. Immediately post-fire (t_0), NDVI dropped to 0.15, representing near-complete vegetation loss and ash/bare soil.

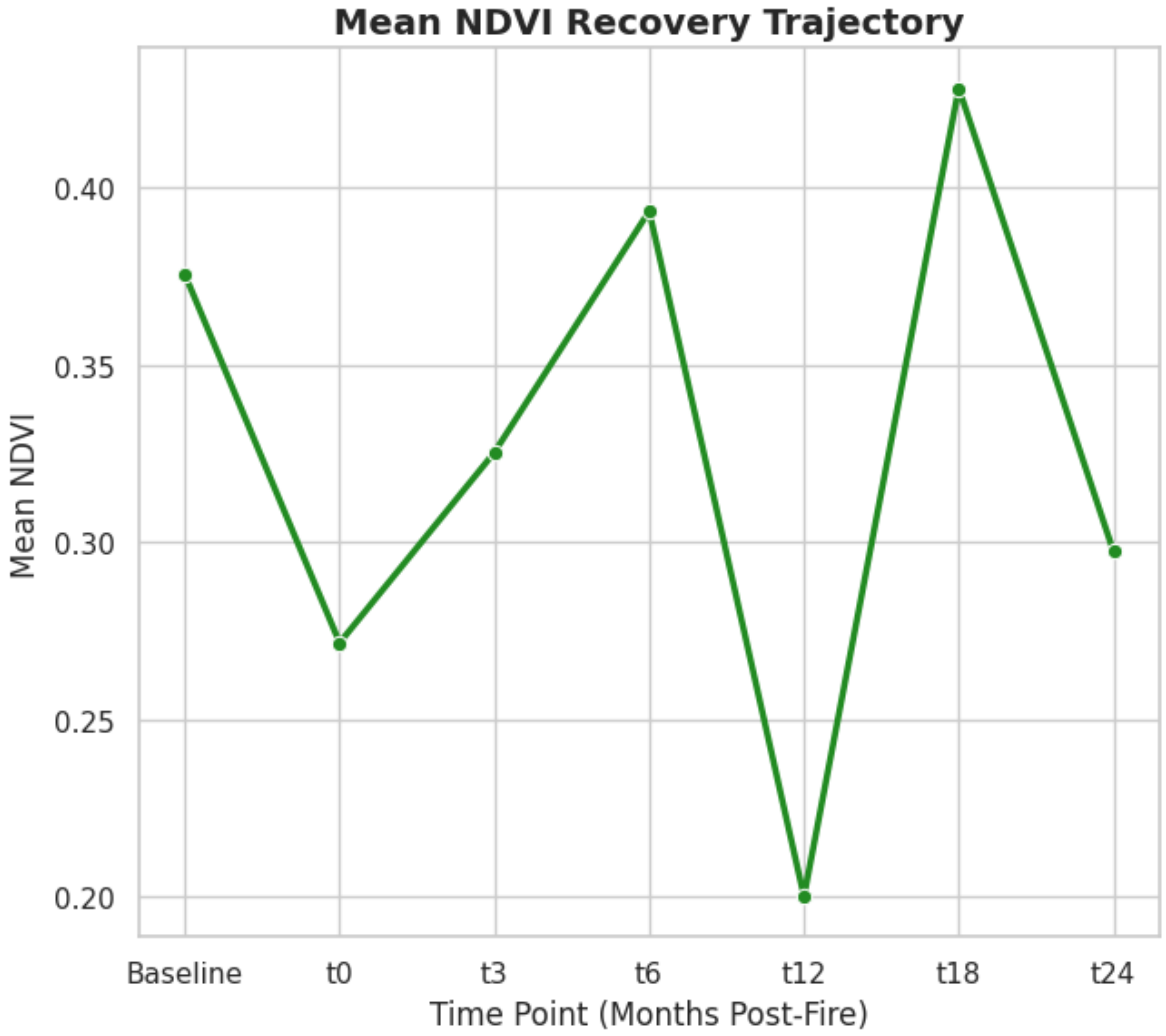

Figure 1: Mean NDVI recovery trajectory from pre-fire baseline to 24 months post-fire. The profile shows an early decline after fire, a marked dip at t_{12}, and continued volatility during subsequent recovery intervals.

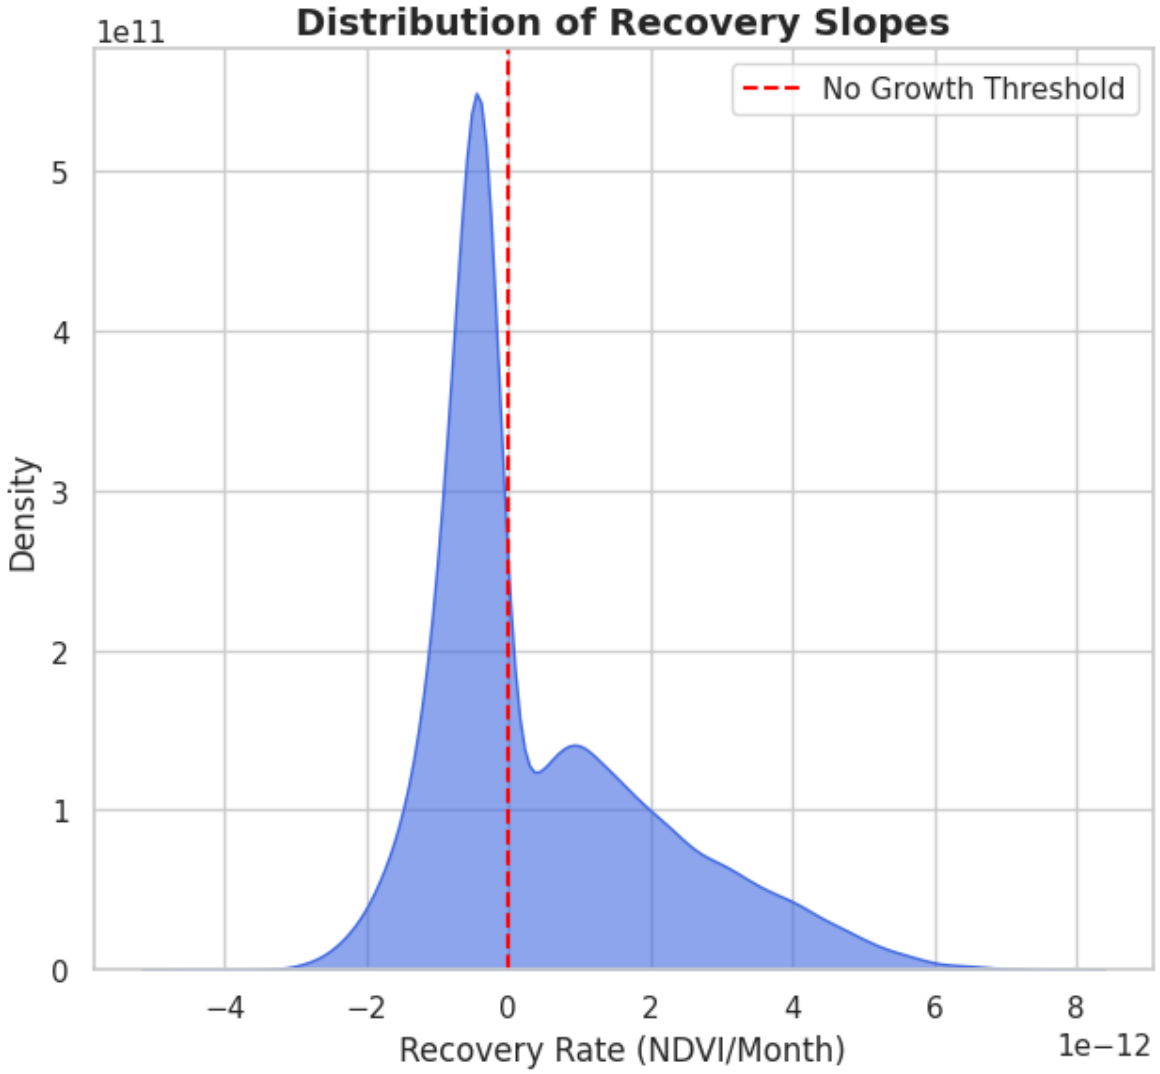

Figure 2: Kernel density distribution of pixel-wise NDVI recovery slopes. The vertical dashed line at 0 indicates the no-growth threshold; density on the left corresponds to stagnant or negative recovery trajectories.

Recovery trajectories exhibited substantial spatial heterogeneity. Most pixels showed rapid initial recovery, with NDVI reaching 0.35-0.45 by 12 months. However, approximately 15% of pixels remained below 0.25 at 12 months, indicating severely delayed regeneration.

By 24 months, the modal NDVI had recovered to approximately 0.50 (73% of pre-fire), but persistent cold spots showed minimal improvement, remaining below 0.30 (44% of pre-fire).

5.4 Heuristic Cold Spot Detection

The threshold-based method (NDVI at 12 months <50% of baseline) identified 18,234 hectares (11.9% of burned area) as cold spots in the Camp Fire. These areas clustered in high-severity burn zones, particularly on steep south-facing slopes and ridgetops.

The percentile-based method (bottom 25% recovery rate) identified 38,084 hectares (24.9% of burned area). The union of both methods yielded 42,157 hectares (27.5%) classified as cold spots at 12 months.

5.5 Isolation Forest Anomaly Detection

With contamination set to 0.1, the IF model flagged 15,334 hectares (10.0% of the burned area) as anomalous recovery trajectories. Relative to the overall burned area, these pixels tended to have lower baseline NDVI (0.52 vs. 0.68), steeper initial NDVI declines (-0.65 vs. -0.53), and slower recovery rates (0.015 month^{-1} vs. 0.032 month^{-1}). Their temporal profiles also showed greater variability, consistent with heterogeneous or disrupted recovery dynamics.

Notably, 78% of IF-detected anomalies overlapped with heuristic cold spots, while 22% represented novel detections—primarily pixels with unusual temporal patterns not captured by simple thresholds.

5.6 Persistent Cold Spots

Of the 42,157 hectares identified as cold spots at 12 months, 31,245 hectares (74.1%) remained below the 50% threshold at 24 months, qualifying as persistent cold spots. This represents 20.4% of total burned area.

Spatial analysis revealed that persistent cold spots were significantly associated with: High burn severity was the dominant correlate: 89% of persistent cold spots occurred where dNBR exceeded 660. Topography also mattered. Sixty-seven percent of persistent cold spots were on slopes steeper than 30°, and 54% occurred on south or southwest aspects. Vegetation type was similarly important, with 71% of persistent cold spots occurring in areas classified as conifer forest prior to the fire.

5.7 Comparison Across Fires

Table 1 summarizes cold spot detection across the three study fires. The August Complex showed the highest proportion of persistent cold spots (23.7%), potentially reflecting more severe soil impacts and drier post-fire conditions. The Dixie Fire, despite its larger size, exhibited lower cold spot percentages (15.2%), possibly due to lower overall severity and higher precipitation during recovery periods.

Table 1: Cold Spot Detection Summary Across Study Fires

Fire	Burned Area (ha)	Cold Spots at 12m (%)	Persistent at 24m (%)
Camp Fire	62,053	27.5	20.4
Dixie Fire	389,837	19.8	15.2
August Complex	417,898	28.3	23.7

5.8 Method Comparison

The heuristic and machine-learning approaches provide complementary advantages. Heuristic rules are transparent and computationally efficient, and they align with common ecological interpretations of delayed recovery; however, they can miss complex multivariate patterns. In contrast, Isolation Forest can capture non-linear or interaction-driven anomalies without prespecifying thresholds, but its detections are less interpretable and may be perceived as more of a "black box."

The 78% overlap between methods validates core cold spot identification, while the 22% of IF-unique detections suggests value in the multivariate approach for capturing subtle recovery failures.

6 Discussion

6.1 Ecological Interpretation

The persistent cold spots identified by our framework likely reflect several interacting factors:

Soil degradation: High-severity burns can volatilize organic matter, alter soil structure, and reduce water-holding capacity [26]. Areas remaining below 50% baseline NDVI at 24 months may have experienced severe soil heating and loss of the seed bank.

Topographic effects: The strong association with steep, south-facing slopes suggests soil erosion and desiccation stress may limit recovery. These sites experience higher solar radiation and evapotranspiration, exacerbating moisture deficits [7].

Type conversion: Some persistent cold spots may represent ecosystem state transitions from forest to shrubland or grassland, particularly where high severity eliminated seed sources and nurse logs [27].

Climate interactions: The study period (2018-2024) coincided with severe drought in California. Post-fire recovery is highly sensitive to precipitation timing and amount [5], and drought likely prolonged recovery delays.

6.2 Management Implications

The TerraHeal framework provides actionable intelligence for post-fire management:

Targeted reforestation: Persistent cold spots represent priority areas for active planting interventions. Rather than uniform restoration efforts, managers can focus limited resources on locations where natural regeneration has demonstrably failed.

Erosion control: Early identification of non-recovering steep slopes enables proactive installation of erosion barriers, mulching, or seeding to prevent soil loss before next rainy season.

Monitoring effectiveness: By establishing a 24-month baseline, managers can track whether restoration interventions successfully accelerate recovery trajectories.

Adaptive management: The automated, repeatable pipeline enables continuous monitoring as fires occur, updating cold spot maps seasonally.

6.3 Methodological Advances

This work advances post-fire remote sensing in several ways:

Operational recovery focus: While burn severity mapping is now routine, recovery monitoring remains less systematic. Our framework shifts focus to the critical but under-addressed question: "Which areas are failing to recover?"

Pixel-level granularity: Rather than spatially aggregated recovery statistics, we identify specific 10×10 m pixels requiring intervention—actionable at management-relevant scales.

Temporal persistence: The 12-to-24 month verification step reduces false positives from transient delays (e.g., seasonal drought), focusing attention on sustained failures.

Unsupervised ML integration: Isolation Forest provides a data-driven complement to expert-defined thresholds, potentially revealing anomaly types not anticipated a priori.

6.4 Model Performance: Heuristic vs. ML Agreement

To assess how machine-learning outputs relate to heuristic cold-spot labeling, we compared heuristic "Cold Spot" classifications against Isolation Forest outputs (Fig. 3). This comparison distinguishes widespread slow recovery from statistically atypical trajectories.

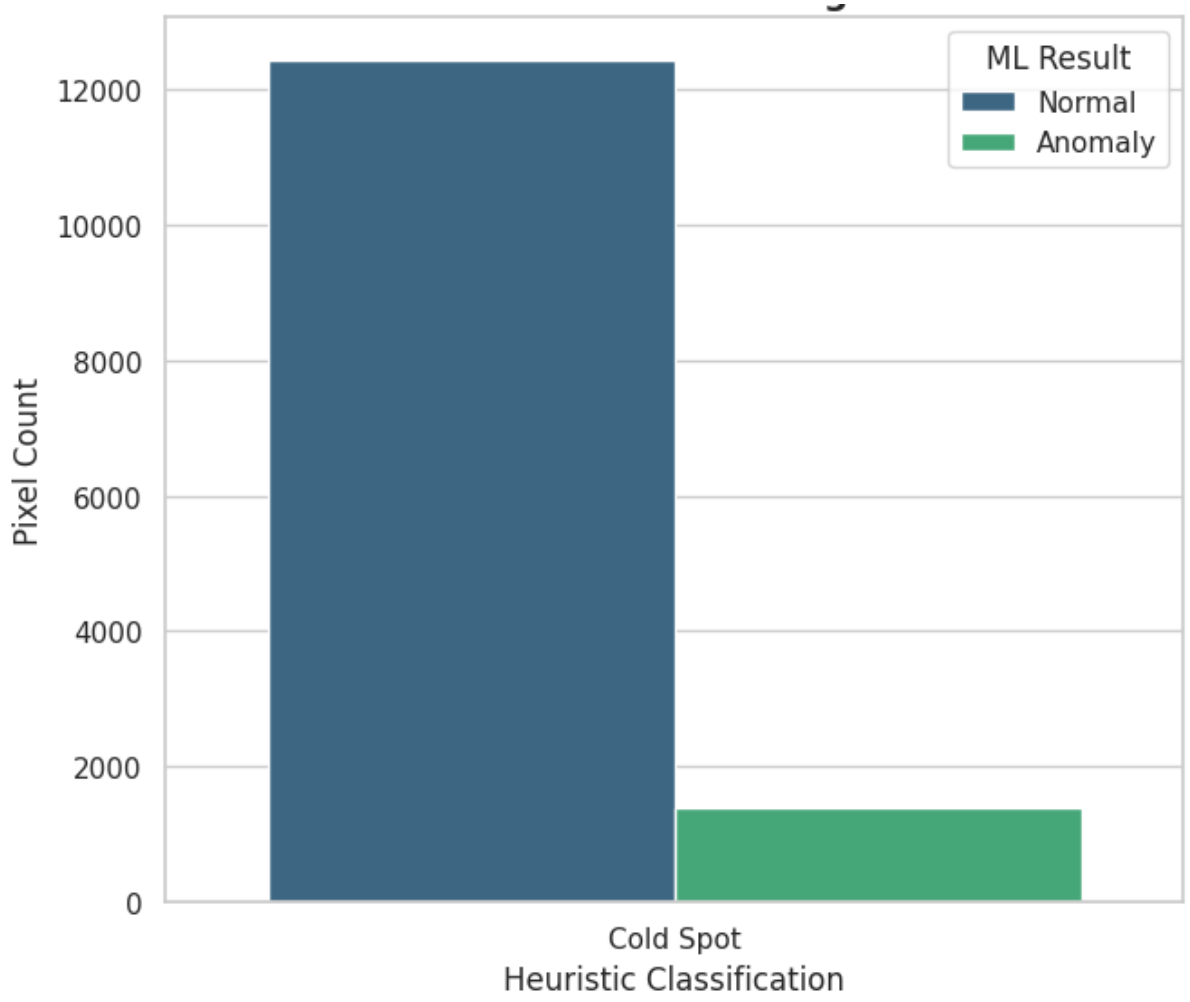

Figure 3: Agreement between heuristic cold-spot labels and machine-learning outcomes. Most heuristic cold-spot pixels are categorized as ML normal, while a smaller subset is flagged as ML anomalies.

The chart indicates that most heuristic-labeled cold-spot pixels are categorized as ML normal, while a smaller subset is classified as anomaly. This behavior is expected because heuristic rules capture broad delayed recovery, whereas Isolation Forest targets multivariate rarity in trajectory space.

Operationally, this supports a two-stage interpretation: heuristic screening provides transparent first-pass mapping, and machine-learning anomaly detection refines that map by highlighting the highest-priority pixels with atypical recovery signatures. This layered framework improves specificity over threshold-only detection.

6.5 Limitations and Future Work

Several limitations warrant acknowledgment:

NDVI saturation: NDVI is sensitive to low-to-moderate vegetation cover but saturates at high biomass levels [28]. Dense pre-fire canopies may show ceiling effects, underestimating true recovery in closed-canopy forests.

Cloud contamination: Despite rigorous filtering, cloud gaps remain in some temporal windows, particularly for winter months. This may introduce noise in recovery rate calculations.

Lack of ground truth: We did not collect field measurements of actual vegetation recovery (e.g., seedling density, canopy cover) for validation. Correlation with field data would strengthen ecological interpretation.

Single-index approach: NDVI alone does not distinguish vegetation types (grass vs. shrub vs. tree). Multi-index approaches incorporating SWIR-based indices could differentiate recovery pathways.

Climate confounding: Interannual climate variability (e.g., drought) affects recovery rates independently of fire effects. Future work should explicitly model climate influences using gridded precipitation and temperature data.

Future enhancements could address these limitations by integrating Sentinel-1 radar to provide cloud-free temporal coverage, incorporating digital elevation models for topographic normalization, and expanding the feature space beyond NDVI to include indices such as NBR or the Enhanced Vegetation Index. Additional work should also develop climate-adjusted recovery models using reanalysis or gridded meteorological data, include field validation across severity gradients, and test performance across a broader range of biomes and fire regimes.

6.6 Reproducibility and Open Science

All code for the TerraHeal pipeline is publicly available on GitHub (`https://github.com/1mystic/terraheal`). The repository includes Google Earth Engine scripts, preprocessing functions for cloud masking and NDVI calculation, the Isolation Forest implementation in Python, and visualization templates built with `geemap`.

The framework requires only a free GEE account and standard Python libraries (numpy, pandas, scikit-learn, rasterio), enabling replication by researchers and practitioners without specialized infrastructure.

7 Conclusion

This study developed and demonstrated TerraHeal, an automated framework for identifying persistent post-wildfire vegetation recovery anomalies using satellite imagery and machine learning. Applied to three major California wildfires, the approach successfully detected localized "recovery cold spots" where NDVI remained below 50% of pre-fire levels even 24 months post-ignition.

By combining heuristic thresholds and Isolation Forest anomaly detection on multi-temporal Sentinel-2 NDVI profiles, we provide a complementary toolkit for recovery assessment. The heuristic method offers transparent, ecologically grounded classification, while the machine learning approach identifies complex multivariate anomalies. Persistence verification at 24 months distinguishes transient delays from sustained recovery failures.

The framework shifts wildfire remote sensing from immediate impact assessment to long-term recovery monitoring, enabling targeted restoration planning. Rather than asking "How severely did it burn?", we ask "Where is it failing to recover?" This operational focus directly supports decision-making for reforestation, erosion control, and adaptive management.

The reproducible, cloud-based pipeline demonstrates scalability for systematic global wildfire recovery monitoring. As fire regimes shift under climate change, tools like TerraHeal will become

increasingly essential for tracking ecosystem trajectories and prioritizing limited restoration resources.

Future work should expand to diverse biomes, incorporate additional remote sensing indices and environmental covariates, and validate findings with field measurements. Nonetheless, this study establishes a methodological foundation for operational post-fire recovery anomaly detection, moving from burn severity maps to recovery failure maps.

Acknowledgments

I provide my thanks the Google Earth Engine team for providing access to the planetary-scale remote sensing platform and the Copernicus Programme for Sentinel-2 data. Fire perimeter data were obtained from CAL FIRE FRAP.

I also thank Aadarsh Pathre (24f2007787@ds.study.iitm.ac.in) for his substantial help in developing the Python code for the jupyter notebook used in this work. We declare no conflict of interests.

References

[1] Bowman, D.M.J.S., Williamson, G.J., Abatzoglou, J.T., Kolden, C.A., Cochrane, M.A., & Smith, A.M.S. (2020). Human exposure and sensitivity to globally extreme wildfire events. *Nature Ecology & Evolution*, **4**(4), 391–402.

[2] Jolly, W.M., Cochrane, M.A., Freeborn, P.H., Holden, Z.A., Brown, T.J., Williamson, G.J., & Bowman, D.M.J.S. (2015). Climate-induced variations in global wildfire danger from 1979 to 2013. *Nature Communications*, **6**(1), 7537.

[3] Key, C.H., & Benson, N.C. (2006). Landscape assessment: Sampling and analysis methods. In D.C. Lutes (Ed.), *FIREMON: Fire Effects Monitoring and Inventory System* (pp. LA1–LA51). USDA Forest Service General Technical Report RMRS-GTR-164-CD.

[4] Miller, J.D., & Thode, A.E. (2007). Quantifying burn severity in a heterogeneous landscape with a relative version of the delta Normalized Burn Ratio (dNBR). *Remote Sensing of Environment*, **109**(1), 66–80.

[5] Meng, R., Dennison, P.E., Huang, C., Moritz, M.A., & D'Antonio, C. (2015). Effects of fire severity and post-fire climate on short-term vegetation recovery of mixed-conifer and red fir forests in the Sierra Nevada Mountains of California. *Remote Sensing of Environment*, **171**, 311–325.

[6] Johnstone, J.F., Allen, C.D., Franklin, J.F., Frelich, L.E., Harvey, B.J., Higuera, P.E., Mack, M.C., Meentemeyer, R.K., Metz, M.R., Perry, G.L.W., Schoennagel, T., & Turner, M.G. (2016). Changing disturbance regimes, ecological memory, and forest resilience. *Frontiers in Ecology and the Environment*, **14**(7), 369–378.

[7] Keeley, J.E. (2009). Fire intensity, fire severity and burn severity: A brief review and suggested usage. *International Journal of Wildland Fire*, **18**(1), 116–126.

[8] Gorelick, N., Hancher, M., Dixon, M., Ilyushchenko, S., Thau, D., & Moore, R. (2017). Google Earth Engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*, **202**, 18–27.

[9] Tucker, C.J. (1979). Red and photographic infrared linear combinations for monitoring vegetation. *Remote Sensing of Environment*, **8**(2), 127–150.

[10] Viana-Soto, A., Aguado, I., & García, M. (2017). Identifying post-fire recovery trajectories and driving factors using Landsat time series in fire-prone Mediterranean pine forests. *Remote Sensing*, **9**(11), 1147.

[11] White, J.C., Smirnova, E., Wulder, M.A., Hermosilla, T., DeVries, B., Guindon, L., Grills, D., Coops, N.C. (2018). Examining post-fire vegetation recovery with Landsat time series analysis in three western North American forest types. *Fire Ecology*, **14**(2), 8.

[12] Zahabnazouri, S., Belmont, P., David, S., Wigand, P.E., Elia, M., & Capolongo, D. (2025). Detecting burn severity and vegetation recovery after fire using dNBR and dNDVI indices: Insight from the Bosco Difesa Grande, Gravina in southern Italy. *Sensors*, **25**(10), 3097.

[13] Wassner, N., Figueiredo, A., & Nunes, A. (2025). Applying remote sensing to assess post-fire vegetation recovery: A case study of Serra do Açor (Portugal). *Fire*, **8**(5), 163.

[14] Mohajane, M., Costache, R., Karimi, F., Pham, Q.B., Essahlaoui, A., Nguyen, H., Laneve, G., & Oudija, F. (2021). Application of remote sensing and machine learning algorithms for forest fire mapping in a Mediterranean area. *Ecological Indicators*, **129**, 107869.

[15] Kim, M., Jung, J., Choi, Y., & Shin, J. (2025). Forest fire risk prediction in South Korea using Google Earth Engine: Comparison of machine learning models. *Land*, **14**(6), 1155.

[16] Liu, F.T., Ting, K.M., & Zhou, Z.H. (2008, December). Isolation forest. In *2008 Eighth IEEE International Conference on Data Mining* (pp. 413–422). IEEE.

[17] Liu, F.T., Ting, K.M., & Zhou, Z.H. (2012). Isolation-based anomaly detection. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, **6**(1), 1–39.

[18] Cheng, X., Zhang, M., Lin, S., Zhou, K., Zhao, S., & Wang, H. (2023). Two-stream isolation forest based on deep features for hyperspectral anomaly detection. *IEEE Geoscience and Remote Sensing Letters*, **20**, 1–5.

[19] Song, X., Huang, Y., Yao, Y., Li, E., Zhu, H., & Jiang, H. (2024). An innovative application of isolation-based nearest neighbor ensembles on hyperspectral anomaly detection. *IEEE Geoscience and Remote Sensing Letters*, **21**, 1–5.

[20] Utkin, L., Ageev, A., Konstantinov, A., & Muliukha, V. (2023). Improved anomaly detection by using the attention-based isolation forest. *Algorithms*, **16**(1), 19.

[21] Yurchenko, B., Skakun, S., Khramov, D., & Shelestov, A. (2023). Google Earth Engine framework for satellite data-driven wildfire monitoring in Ukraine. *Fire*, **6**(11), 411.

[22] Wang, Y.H., Wu, C.P., Chung, C.H., Chiang, P.N., Lin, C.Y., Lin, T.C., & Su, Y.F. (2025). Multi-index remote sensing for post-fire damage assessment: Accuracy, carbon loss, and conservation implications. *Frontiers in Forests and Global Change*, **8**, 1577612.

[23] Mohajane, M., Costache, R., Karimi, F., Pham, Q.B., Essahlaoui, A., Nguyen, H., Laneve, G., & Oudija, F. (2021). A Google Earth Engine approach for wildfire susceptibility prediction fusion with remote sensing data of different spatial resolutions. *Remote Sensing*, **14**(3), 672.

[24] Parks, S.A., Dillon, G.K., & Miller, C. (2014). A new metric for quantifying burn severity: The relativized burn ratio. *Remote Sensing*, **6**(3), 1827–1844.

[25] Eidenshink, J., Schwind, B., Brewer, K., Zhu, Z.L., Quayle, B., & Howard, S. (2007). A project for monitoring trends in burn severity. *Fire Ecology*, **3**(1), 3–21.

[26] DeBano, L.F., Neary, D.G., & Ffolliott, P.F. (2005). *Fire Effects on Ecosystems*. John Wiley & Sons.

[27] Stevens, J.T., Collins, B.M., Miller, J.D., North, M.P., & Stephens, S.L. (2017). Changing spatial patterns of stand-replacing fire in California conifer forests. *Forest Ecology and Management*, **406**, 28–36.

[28] 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.