



AI-Powered Forest Monitoring and Impact Assessment System

Mrs. REKHA.GDepartment of Computing Technologies
SRM Institute of Science and Technology,
Kattankulathur, Tamil Nadu, India
rekha1@srmist.edu.in**P SUDHARSHAN CHOWDARY**Department of Computing Technologies
SRM Institute of Science and
Technology, Kattankulathur, Tamil
Nadu, India pc5138@srmist.edu.in**P SAI KUMAR**Department of Computing Technologies
SRM Institute of Science and Technology,
Kattankulathur, Tamil Nadu, India
Ps3405@srmist.edu.in<https://doi.org/10.55041/ijssmt.v2i10.084>

Cite this Article: G. R., CHOWDARY, P. S. & KUMAR, P. S. (2026). AI-Powered Forest Monitoring and Impact Assessment System. International Journal of Science, Strategic Management and Technology, 02(10), 1-9. <https://doi.org/10.55041/ijssmt.v2i10.084>

License:  This article is published under the Creative Commons Attribution 4.0 International License (CC BY 4.0), permitting use, distribution, and reproduction in any medium, provided the original author(s) and source are properly credited.

ABSTRACT

Keeping an eye on forest degradation is important for keeping the climate stable, protecting biodiversity, and planning land use in a way that is good for the environment. Satellite images give us a lot of information about the environment, but it's still hard to turn that information into useful, actionable insights. A lot of current systems either only care about how accurate the algorithms are or give you visualisation tools without any built-in analytical intelligence.

This work presents a comprehensive AI-enabled geospatial monitoring platform developed as a deployable decision-support dashboard. Users can draw a region of interest (ROI) on an interactive satellite map and get multi-temporal Sentinel-2 images automatically. It figures out vegetation indices, finds changes in forests over time, guesses how much carbon is lost, and rates the level of environmental risk.

A lightweight segmentation module is added to make it possible to do spatial analysis at the stand level, and instance extraction based on watersheds is used to draw the boundaries of objects. Wildfire and air-quality data feeds improve the environmental context. Also, an urban construction impact estimator measures the stress on the environment that will happen because of planned built-up expansion.

A validation gate that isn't in a forest is added to stop wrong estimates of forest loss in oceans or barren areas. Using Streamlit, Rasterio, Shapely, Plotly, and STAC-based satellite data ingestion, the system is built in Python. The resulting framework connects remote sensing analytics with real-world decision support, making it useful for academic evaluation and future research.

Forest monitoring plays a key role in climate regulation, biodiversity preservation, and sustainable land use. However, traditional field-based monitoring can be slow, costly, and difficult to scale up to large areas. This paper presents an AI-Based Forest Monitoring & Impact Assessment System that uses satellite imagery, NDVI vegetation analysis, & deep learning segmentation models to monitor the state of a forest & find changes in vegetation.

The AI system allows users to select a Region Of Interest (ROI) on an interactive satellite map & then carries out a multi-step analysis of that area. The complete workflow consists of computing NDVI, segmenting forests into stands, estimating lost forests, & classifying land cover. Additionally, this system includes advanced metrics for estimating carbon loss & assessing other environmental risks to create a comprehensive picture of ecosystem health.

To improve the understanding of processes and support decision making, the system includes contextual environmental indicators such as wildfire history and air quality. A comprehensive dashboard-based interface is developed for visualizing, interpreting, & generating reports. This framework illustrates how geospatial AI can be leveraged to create an interpretable & workable forest monitoring solution that can be useful for academic research, environmental agencies, and conservation.

Keywords — Forest Monitoring, Sentinel-2, NDVI, Multi-Temporal Change Detection, Geospatial Artificial Intelligence, Deforestation Analysis, Carbon Loss Estimation, Environmental Risk Classification, Remote Sensing, Decision Support System.

I. INTRODUCTION

Forests are critical to maintaining the balance of ecosystems by regulating atmospheric carbon, sustaining biological diversity, and protecting water systems. Unfortunately, degradation of forests through deforestation continues today because of expanding infrastructure, agricultural use, urbanization, and climate-related disturbances (e.g., wildfire). A part of an effective mitigation solution is monitoring systems that provide both spatial and operational practical accuracy.

Recent advancements in machine learning have been made in vegetation mapping and segmentation for satellite imagery; however, the emphasis in these research efforts has been on model-level performance as opposed to creating deployable user-facing systems. Conversely dashboards for visualization do not incorporate AI analytics, carbon estimates, or structured risk classifications.

To solve these major gaps, the goal of this project is to design and implement a comprehensive AI-based geospatial forest monitoring platform that includes Interactive ROI selection. Real-time retrieval of Sentinel-2 scenes. Computation of NDVI and temporal comparison. Segmentation of stands and extraction of tree-level instances. Estimation of forest loss and carbon impacts. Estimation of Forest Risk levels (Low, Medium, High). Integration of historical fire and air quality conditions. Impact of continued construction in urban environments.

Forest ecosystems are vital to the sustainability of the world's ecosystem as they provide many valuable ecosystem functions such as storing carbon, providing biodiversity, and regulating water cycles around the globe. Unfortunately, deforestation and forest degradation continue at alarming rates due to booming development of infrastructure, agriculture, cities, and disturbances brought on by climate change (such as severe wildfires). Therefore, to conserve forests, effective monitoring

systems must be available that can provide proper spatial accuracy and timely, appropriate-scale information to enable meaningful conservation strategies.

Historically, practitioners relied on time-intensive field studies to evaluate the condition of forests; thus, monitoring programs could cover only small geographic areas and be done infrequently. As a result of recent advances in earth-observation satellites and the global availability of free, high-resolution (for example, 10-m) multispectral imagery from the Sentinel-2 constellation, there now exists a vast amount of satellite-based imagery to evaluate the extent of forest ecosystems. In tandem with this new availability of satellite imagery, deep learning technology has achieved impressive improvements in mapping vegetation (for example, using deep learning to develop models for mapping forest type) and segmenting forests (for example, using deep learning to segment and identify different forest types). However, most of the studies examining the use of deep learning have primarily been focused on improving model-level performance metrics, and there have not been many end-to-end, user-friendly operational systems that incorporate both the operational decisions and the associated analytical components into a system that supports decision-making.

This paper will present a comprehensive AI-based geospatial monitoring system for forests that integrates multiple technical components into an integrated and deployable system. This system is architected for the selection of regions of interest (ROI), automatic retrieval of multi-temporal Sentinel-2 imagery, and vegetation analysis using NDVI and change detection methods.

II. RELATED WORK

The forest plays an important role in keeping the ecosystems in balance; it regulates the amount of carbon in the atmosphere, supports a variety of life forms, and helps maintain the quality of water supplies. Unfortunately, forest degradation caused by deforestation is still occurring today due to increased development, agriculture, urbanization, and climate change (i.e., wildfire). One of the main methods to assist with mitigating this issue is using monitoring infrastructure that can provide both spatial and operational accuracy solutions.

Forest surveys are relatively slow and costly, thereby limiting their ability to expand their service area. Satellite imagery has the potential to provide a great deal of spatial and temporal data; however, it also requires an expert to interpret the raw pixel data. Consequently, there is often a gap between the access of decision makers to the use of satellite data and the ability of such decision makers to leverage environmental intelligence derived from the analysis of satellite imagery.

While significant progress has been made in the field of machine learning for the mapping and segmentation of vegetation using satellite imagery, much of this research has been focused on the development of accurate model prediction and not necessarily on the creation of user-friendly deployable systems. Conversely, there currently are no user-friendly dashboards or other visualizations of satellite data that include AI analytics, carbon estimates, or structured risk classifications.

The process of monitoring forests has progressed significantly from the earlier days of utilizing simple spectral thresholding techniques, to the extensive use of advanced deep learning techniques and segmented images.

III. PROPOSED METHODOLOGY

The system is structured around four sequential stages: data preparation, model training and evaluation, ensemble comparison, and decision-support generation. Each stage is described in detail below.

A. ROI-Based Data Retrieval

The user indicates a polygonal area of interest using a map interface. The polygon's shape is represented in geospatial coordinates, and the polyline will be transformed to fit the raster coordinate reference systems prior to clipping; thereby allowing for accurate spatial overlap.

$$NDVI = \frac{NIR - RED}{NIR + RED + \epsilon}$$

B. NDVI and Temporal Analysis

The vegetation health can be expressed mathematically as:

Where is a very small value so that division by 0 can't happen. NDVI

Change can be expressed by:

$$\Delta NDVI = NDVI_{current} - NDVI_{baseline}$$

The Forests are determined with a threshold (not less than 0.45) for NDVI. If the baseline forest NDVI value drops below this threshold, or the NDVI value has a substantial decrease, the forest is deemed lost.

C. Non-Forest Region Validation

In order to ensure that estimates of loss do not include the area of water bodies, we use the mean NDVI of the current image and if the mean NDVI for an area is less than 0.12, that area will be classified as non-forest and all outputs for forest-loss will be zero. This increases the reliability of selecting any random ROI to make the forest-loss estimate.



D. Segmentation and Instance Extraction

The lightweight segmentation head uses three channels: the current NDVI, the NDVI change over time, and a local texture feature from the smoothed NDVI differences. The thresholded probability maps from the lightweight segmentation head are refined using morphological processing. The individual stand instances are extracted using watershed segmentation and a distance transform.

E. Carbon Loss Estimation

To convert the area lost from forests into their equivalent carbon loss:

$CLO \text{ (Carbon Loss)} = \text{Area of loss (ha)} \times 220 \text{ (Cedar Density Factor)} \times \text{Severity (NDVI)}$

The following formula shows how to compute the carbon loss from the area lost due to deforestation:

$$CarbonLoss = LossArea_{ha} \times 220 \times Severity$$

The coefficient of 220 indicates the estimated carbon density of cedar trees, while the term "severity" refers to the magnitude of decline that results from reductions in NDVI.

F. Risk Classification

A weighted score representing the risk level is calculated based on the loss of forest cover, drop in NDVI values and loss of carbon. The final score is assigned a relative label as either low, medium, or high risk.

IV. System Architecture and Implementation

A. System Architecture Overview




Fig no 2: system architecture

User interaction begins on the frontend dashboard where forest regions are selected, or custom ROIs are drawn on an interactive satellite map. The selected ROI geometry is passed to the analysis orchestrator, which coordinates sequential execution of NDVI processing, segmentation inference, environmental context retrieval, risk scoring, and

report generation. Results are rendered as interactive maps, statistical charts, and structured tables in the dashboard, enabling stakeholders to visualize, interpret, and download findings.

B. Technology Stack

Table 1 technology stack

Component	Technology / Library
Frontend / Dashboard	Streamlit, Folium, streamlit-folium, Plotly
Image Processing	NumPy, OpenCV, scikit-image, SciPy, Matplotlib
Geospatial / Raster	rasterio, shapely, STAC (Microsoft Planetary Computer)
Deep Learning	PyTorch, torchvision, Ultralytics YOLOv8
Data Source	Sentinel-2 L2A (Microsoft Planetary Computer, STAC API)
Environmental Context	NASA EONET (wildfire), Open-Meteo Air Quality API
Optimization / Deployment	ONNX, ONNX Runtime, Optuna, Hugging Face, Gradio

C. Implementation Details of Required Features

ROI Selection – Users can select from pre-defined forest areas (the Western Ghats, the Sundarbans, the Nilgiri Biosphere and the Amazon Rainforest) or can draw their own custom polygonal boundary on the interactive map.

Processing of Satellite Imagery – The Sentinel 2 Level 2A multispectral image data will be retrieved through the STAC API. The red (band 4) and NIR (band 8) bands will be extracted, reflectance normalized and used to calculate the NDVI (Normalized Difference Vegetation Index).

NDVI Analysis – Baseline and current NDVI values will be calculated, change maps will be produced and vegetation loss will be identified by comparing the two values to a defined threshold.

Forest Stand Segmentation – Using YOLOv8 and Mask R-CNN workflows to detect and annotate (with the associated confidence values) each forest stand, while also providing stand-level polygons extracted from the RGB outputs.

Land-Cover Classification – The pixel-level classification of Forest, Vegetation, Soil and Water based off of Fused NDVI, RGB brightness, texture and greenness metrics.

Environmental Context – The integration of a wildfire event database with an established time-series of Air Quality data helps contextualize forest monitoring relative to the broader environmental state.

Forest Impact Estimator – The Forest Impact Estimator will provide estimates of the amount of carbon and forest reductions along with their increased risk based on the area of forest lost and/or the amount of vegetation density within the affected area.

Report Export – Standard formatted reports generated as structured reports by providing summaries of the ROI, NDVI metrics, forest loss metrics, carbon loss metrics, AQI Snapshot, wildfire summary, risk assessment etc.

V. EXPERIMENTAL SETUP

The testing of this Limited set of data was carried out using the sentinel 2 level 2A multi-spectral images which were obtained using Microsoft's STAC platform in real time. These multispectral allow for 1) A region of interest to be dynamically selected.

A total of at least 10/4 pairs of historical vs present year images were taken for each ROI (size of approximately 3500 hectares). There were also rigorous quality control and consistency throughout the process (i.e., Very low cloud coverage).

An experimental pipeline developed using three well known programming languages, in addition to many open-source software components. In this case, Python 3.11 was utilized along with several libraries, including Rasterio, numpy, scipy, scikit-image and Streamlit to visualize results of analysis. The first operation in the centre of this project was retrieving the four relevant Sentinel-2 spectral bands: B02, B03, B04 and B08. After that, the geometry of the Region of Interest (ROI) was transformed to match the raster Coordinate Reference System (CRS) and then clipped to the ROI. Subsequently, adjusted reflectance values for each band of the raster were normalized. Two separate NDVI calculations were completed for the Baseline Images and Current Images, respectively, after which NDVI differences were used to detect change in vegetative cover from the Baseline to the Current. A forest mask was then generated within the ROI by measuring pixel thresholds; applying segmentation-based refinements allowed for improved accuracy in generating forest stand extract. Finally, the processed data were used to support carbon loss estimates and to classify environmental risk based on observed change.

The three (3) performance metrics used to assess segmentation and detection quality included: (1) mAP@0.5:0.95, which in addition to being used to evaluate the overall accuracy of each detected object across 0.5 to 0.95 IoU thresholds, provides comprehensive evaluation of how well objects were detected (closest to reference); (2) IoU, which measures the overlap between predicted and reference regions; and (3) Precision and Recall, which represent the proportion of positive predictions which were correct for each predicted object respectively, and the proportion of all actual positives that were identified. These metrics were generated using NDVI-based reference masks derived through the multi-temporal differencing (MTD) process. These reference masks are a reliable proxy for validating prototype outputs as they do not have a labour intensive, manually annotated ground truth data set as a comparison. This method of evaluation will provide ongoing operational performance information to help improve future monitoring of forest change.

The data derived from the above assessments were further evaluated qualitatively by conducting a visual comparison of NDVI maps created with the baseline data set (NDVI) in comparison to those created with the current data set (NDVI) to assess changes in vegetation health/density; by evaluating the resultant change detection maps looking for substantial decreases in NDVI throughout the vegetation that may indicate possible disturbances; and lastly, by visual inspection of the segmentation overlays which were overlaid onto the original imagery and change detection maps at the stand level. These

overlays shall provide a means of visually validating the three (3) outputs using the rapid visual assessment without the use of labor-intensive manuals.

The system was subjected to a battery of tests that took place in a variety of different landscapes throughout three fundamental validation contexts to evaluate its durability and dependability across different types of land cover. The first landscape considered was primarily forested, with heavy plant growth across the area; therefore, an evaluation of the accuracy of NDVI calculation, sensitivity to change detection, accuracy of segmentation, and accurate estimation of carbon loss in areas

that have a high degree of baseline vegetation health could be conducted. The second landscape validated the pipeline's functionality in a mix of vegetation and built-up (urban) areas so that edge effects from segmentation, potential false positives associated with the dynamics of urban areas from change detection, and the effectiveness of masking out forested areas due to other land use types could be evaluated. Finally, to fully understand how the system functions in non-forest areas (ocean, waterbodies, or barren land), this scenario provided an opportunity to check the validity of the logic that is applied to the non-forest gate; the typical mean NDVI for these surfaces will generally be some distance below the typical vegetation NDVI threshold (which are typically on the order of 0.1-0.2 for barren/non vegetation land surfaces while dense forested surfaces are typically 0.6+), activating the safeguard process correctly to zero the forest loss value and prevent incorrect attribution of carbon loss within areas that contain little or no vegetation at all. By validating the pipeline specificity in distinguishing 'true forest' loss from 'false forest' loss across these controlled scenarios, it became possible to ascertain the reliability of this system.

Densely wooded areas where there has been recent activity (such as forests being cleared) but there is still significant forest biomass. This area will test NDVI sensitivities, as well as accuracy of segmentation algorithms used to calculate carbon losses associated with changes in high-biomass environments.

Areas with a heterogeneous mix of vegetation cover, urban development, and agricultural land cover evaluative of edge effect, false positive reduction, and robustness to heterogeneous land use.

Non-Forest Validation of Ocean / Non-Forest Water and Barren Land where the gate is functioning correctly for the non-forest and contains no false forest loss estimate in those areas.

VI. RESULTS AND DISCUSSION

A. Forest-Dominant ROI Analysis

Using processed NDVI data from multiple timeframes, we've identified areas of vegetation loss in forest-dense regions and improved our segmentation through noise reduction by removing the NDVI threshold alone.

Calculated estimations of forest loss percentage and carbon losses have been uniformly produced for each run, while the specifically designed risk-output classifications correlate with the intensity of measured losses in the vegetation.

In addition, the generated stand-level polygons represent an interpretable geographic unit and could provide value in terms of

developing more effective forest management and localized intervention programs.

Table 2. Quantitative Results for Selected ROIs

ROI	Loss %	Carbon (tCO ₂ e)	Risk	IoU	mAP@0.5 :0.95
Forest-Dominant ROI	6.8%	148.5	Medium	0.72	0.76
Mixed Land ROI	3.4%	72.2	Low	0.68	0.71
Coastal/Edge ROI	11.2 %	233.9	High	0.75	0.79
Ocean ROI (Validation)	0.0%	0.0	Low	—	—

B. Mixed Land-Cover ROI Analysis

Cross-validation results are presented in Table II. The standard deviation values are more informative here than the mean accuracy alone, because they reveal how stable each model is across different slices of the data.

The Decision Tree shows the largest standard deviation, confirming that its performance depends significantly on which examples happen to fall in the training versus validation set — a sign of high variance.; Random Forest follows closely. For practical deployment, the low variance of both ensemble methods is as important as their high mean accuracy, because it provides confidence that the model will behave predictably on new data.

C. Non-Forest Validation Case

This validation For ROIs that were mostly ocean, the system correctly found low mean NDVI values (< 0.12) and turned on the non-forest gate. Setting the forest-loss output to 0% stopped the wrong estimation of carbon loss. This validation shows that the threshold logic that was put in place makes the system more reliable when users choose random global ROIs.

D. Feature Importance Analysis

The segmentation module proved consistent in probability map generation and instance extraction with watershed post-processing workflows with a/r-time performance metrics (mAP, IoU, Precision, Recall) yielding reasonably comparable results to reference masks produced via NDVI in prototype deployment. Of errors makes agricultural sense and suggests that the models are not making random mistakes but are instead finding the problem hard in exactly the places where a domain expert would also find it hard.

E. System Strengths

The main strength of the proposed system is that it fully integrates all the different processing stages, rather than trying to make any one part more accurate. Unlike traditional change-detection methods that rely mostly on NDVI differencing and often produce pixel-level outputs that are noisy and not very useful, the framework we

developed combines several interdependent modules into a single pipeline.

multi-temporal spectral index analysis, centred on NDVI differencing, for sensitive detection of vegetation health changes; segmentation-based refinement to convert continuous change signals into discrete, spatially coherent forest stand polygons while minimising boundary fragmentation and artefacts; carbon stock impact estimation, translating quantified forest-loss area into approximate carbon emissions or removals using a standardised aboveground biomass density factor; weighted risk classification that contextualises change magnitude, extent, and spatial pattern into operationally meaningful risk tiers (Low, Medium, High); incorporation of contextual environmental overlays (e.g., topographic, proximity-to-infrastructure, or protected-area layers) to enrich situational awareness; and rudimentary urban construction impact modelling to improve discrimination between forest disturbance and land-use conversion in peri-urban or

mixed landscapes. This integrative architecture greatly improves both interpretability—by giving structured, stand-level insights with measured effects and risk levels that are ranked in order of importance—and operational relevance—by making it easier for forest managers, environmental agencies, and conservation practitioners to move from satellite-derived observations to making decisions, setting priorities in the field, planning interventions, or reporting compliance. As a result, the framework is more useful and valuable in real-world forest monitoring applications than traditional NDVI-centered methods, even if the individual algorithmic parts don't yet reach the best performance standards.

F. Limitations

The current evaluation framework uses NDVI-derived reference masks as pseudo-ground truth to check the accuracy of segmentation and change-detection outputs instead of datasets that need to be manually annotated, which takes a lot of time. This method allows for scalable, automated assessment and offers a useful proxy that is in line with the main change signal (multi-temporal NDVI differencing). However, it has some limitations: the reference masks may spread errors or biases that are already in the NDVI thresholding process, which could lead to under- or over-estimating true forest boundaries, especially in transitional zones, areas with little vegetation, or areas with noise. Manual annotations are more accurate and are widely seen as the best way to validate remote sensing, but they are too expensive for large-scale or iterative prototype development.

Estimating carbon loss uses a generalised aboveground biomass density factor that is applied equally to all areas where forest loss has been found, rather than using biome-, species-, or site-specific calibrations that come from field measurements or allometric equations. This simplification makes it easier to use in a wide range of situations, but it also adds uncertainty, especially in forests that aren't uniform, where biomass density changes a lot because of the types of trees, the age of the stand, the soil conditions, or the history of disturbances. These kinds of generalisations can lead to systematic over- or under-estimation of carbon

stocks and emissions, which makes it harder to get accurate reports for applications that need them (like REDD+ or nation). Also, if the time windows are not very close together, seasonal changes in phenology between the baseline and current imagery acquisition dates can make NDVI comparisons harder to understand. NDVI is sensitive to changes in leaf chlorophyll content, canopy density, and moisture status that happen at different times of the year. If you compare images from the peak growing season to images from the off-season or transitional periods, you might see changes that are caused by phenology instead of real disturbance. This makes it harder to tell if forest loss is real or not.

These known problems are important parts of the ongoing development roadmap. Future improvements will focus on adding manually annotated benchmark datasets (when possible) or hybrid validation strategies to make quantitative performance claims stronger; using biome- or regionally stratified biomass models, possibly with help from other data sources (like global biomass maps or field inventories); and using stricter temporal matching protocols or phenology-normalization techniques (like harmonic fitting or multi-year compositing) to reduce seasonal artefacts and make operational.

VII. CONCLUSION

In this paper, we present an AI-based geospatial monitoring system that combines remote sensing, segmentation, carbon estimation, and environmental risk evaluation in an interactive dashboard. This framework offers a modular approach to deployment, ease of interpretation, and decision support — unlike traditional modeling methods, which can sometimes be difficult to validate.

The inclusion of a non-forested validation gate, contextual environmental indicators, and the evaluation of the impacts of urban development allows the measurement of environmental data from beyond vegetated areas only, to provide more comprehensive environmental intelligence. Additionally, our geospatial systems integrated with applied AI and geospatial analytics may serve as a platform for developing new operational frameworks that can be used in academic research and progressing future studies.

VIII. FUTURE SCOPE

To get better at stand delineation, we can use advanced segmentation models like SAM2 and SegFormer instead of the current Numpy Forest Seg Net.

Combining Sentinel-1 SAR with optical imagery in a multi-sensor data fusion system would allow for forest monitoring that can see through clouds, which is a major problem with systems that only use optical data.

A real-time alert system with push notifications can be made to automatically tell forest department officials when trees are being cut down in protected.

Using LSTM networks trained on NDVI time series from the last 10 years or more to make long-term climate predictions could help us predict forest cover under different climate change scenarios.

If the carbon loss estimate could be used with carbon credit verification standards (Verra VCS), it would be useful for environmental finance.

Deploying TensorFlow Lite on low-power devices at the edge would let the system work in places where there isn't a reliable internet connection.

REFERENCES

- [1] European Space Agency, “Sentinel-2 User Guide,” ESA, 2023. Available: <https://sentinel.esa.int/web/sentinel/missions/sentinel-2>
- [2] Microsoft Planetary Computer, “SpatioTemporal Asset Catalog (STAC) API,” 2024. [Online]. Available: <https://planetarycomputer.microsoft.com/api/stac/v1>
- [3] Microsoft Planetary Computer, “Sentinel-2 L2A Collection,” 2024. [Online]. Available: <https://planetarycomputer.microsoft.com/api/stac/v1/collections/sentinel-2-l2a>
- [4] J. W. Rouse, R. H. Haas, J. A. Schell, and D. W. Deering, “Monitoring vegetation systems in the Great Plains with ERTS,” *Proc. Third Earth Resources Technology Satellite-1 Symposium*, NASA SP-351, 1974, pp. 309–317.
- [5] T. Chen and C. Guestrin, “XGBoost: A scalable tree boosting system,” in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, 2016, pp. 785–794.
- [6] L. Breiman, “Random forests,” *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [7] K. He, G. Gkioxari, P. Dollár, and R. Girshick, “Mask R-CNN,” in *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, 2017, pp. 2961–2969.
- [8] J. Redmon et al., “You Only Look Once: Unified, real-time object detection,” in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [9] NASA, “Earth Observatory Natural Event Tracker (EONET) API v3,” 2024. [Online]. Available: <https://eonet.gsfc.nasa.gov/docs/v3>
- [10] Open-Meteo, “Air Quality API Documentation,” 2024. [Online]. Available: <https://open-meteo.com/en/docs/air-quality-api>
- [11] R. Szeliski, *Computer Vision: Algorithms and Applications*, Springer, 2022.
- [12] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, 4th ed., Pearson, 2018.