

TreeSense: AI-Powered Tree Detection System Using Aerial Imagery

Jayesh Parmar, Shivam Arya, Naresh Choudhary, Anuj Deshmukh, Prof. Ishwar Bharambe

Department of Computer Engineering

Ajeenkya DY Patil School of Engineering, Pune, India

Abstract: *TreeSense is a lightweight, AI-powered web platform designed to simplify and democratize environmental monitoring through satellite and aerial imagery analysis. Aimed at addressing the limitations of traditional tree and land-use surveys—which are often time-consuming, resource-intensive, and technically demanding—TreeSense offers an intuitive solution for users to detect tree count, assess vegetation health, and analyze land usage patterns with minimal effort. The system integrates advanced APIs such as Mapbox for geographic selection, OpenWeatherMap for contextual weather data, and the Gemini API for intelligent image interpretation. Operating entirely on the frontend without the need for backend storage or custom-trained machine learning models, the platform allows real-time image processing and visualization via interactive charts and descriptive insights. Users can compare two images taken at different times to observe environmental changes such as deforestation or urban expansion. Employing conceptual algorithms like YOLOv8 for object detection, NDVI estimation for vegetation health, and semantic segmentation for land classification, TreeSense presents a robust foundation for academic research, civic planning, and ecological monitoring. The system was developed using an agile methodology with React.js and is optimized for accessibility and deployment efficiency. Future work aims to enhance detection accuracy, incorporate historical analysis via backend storage, and support real multispectral NDVI inputs. TreeSense stands as a scalable and user-friendly innovation in AI-assisted environmental intelligence*

Keywords: TreeSense

I. INTRODUCTION

Environmental degradation, urban expansion, and climate change have significantly heightened the need for efficient, real-time monitoring of green cover and land use. Traditional approaches such as field surveys and manual image annotation are often labor-intensive, time-consuming, and cost-prohibitive—making them unsustainable for frequent or large-scale assessments. In response to this growing demand for scalable environmental monitoring, TreeSense has been developed as a user-friendly, AI-powered web application that enables users to analyze satellite and aerial imagery for vegetation health, tree density, and land usage patterns without requiring technical expertise or backend infrastructure. TreeSense empowers users to interact with environmental data through an intuitive, frontend-only interface, combining the power of advanced APIs and large language models. With real-time area selection via Mapbox, contextual weather data from OpenWeatherMap, and AI-driven image analysis provided by Google's Gemini API, TreeSense delivers actionable insights in seconds. The platform allows users to either select a region directly on a map or upload custom images for analysis, making it versatile for various use cases such as urban planning, forestry management, agricultural assessments, and educational projects.

The core functionality includes detecting and counting trees, estimating vegetation health through NDVI-like interpretations, classifying land cover types, and comparing environmental changes over time. All results are rendered visually through charts and summary indicators using Chart.js or Recharts, making complex environmental data more accessible and interpretable. Importantly, TreeSense does not require backend servers, persistent databases, or custom-trained machine learning models—relying instead on robust API integrations and frontend logic for rapid deployment and ease of use.



II. PROBLEM STATEMENT

Despite the availability of high-resolution satellite imagery and advances in artificial intelligence, there remains a significant gap in tools that allow non-expert users to perform efficient and accurate environmental monitoring. Current methods for tracking vegetation health, counting trees, and analyzing land usage—such as manual field surveys, GIS software, or specialized machine learning models—are often expensive, require domain expertise, and are not scalable for frequent assessments or large geographic areas.

Moreover, many existing platforms either demand backend infrastructure, persistent storage, or extensive model training, which limits their accessibility and deployment. Tools like Google Earth Engine, while powerful, are designed primarily for researchers and specialists, making them less user-friendly for everyday civic use, education, or local governance.

The core problem is the lack of a lightweight, accessible, and intelligent platform that enables users—regardless of technical background—to analyze satellite or aerial images for environmental insights like tree count, vegetation health, and land cover classification. This gap hinders timely decision-making in urban planning, green audits, forestry, and environmental conservation, especially in resource-limited settings.

III. LITERATURE REVIEW

Recent advancements in remote sensing and artificial intelligence have greatly enhanced environmental monitoring, enabling more accurate, faster, and scalable solutions. Several key technologies and research developments serve as the foundation for systems like TreeSense, which leverage satellite imagery and AI for vegetation and land cover analysis. This section explores the relevant contributions in object detection, vegetation indexing, land classification, and image comparison that underpin the capabilities of TreeSense.

YOLOv8 for Object Detection in Forestry

You Only Look Once (YOLO) is a popular real-time object detection algorithm. The latest version, YOLOv8, is well-suited for identifying trees in aerial imagery due to its high-speed inference and accuracy. It is widely used in forestry for automatic tree detection, species identification, and biomass estimation, allowing for efficient large-scale assessments without manual annotation.

NDVI in Vegetation Health Monitoring

The Normalized Difference Vegetation Index (NDVI) is a standard metric used in remote sensing to assess plant health and density. NDVI utilizes the red and near-infrared bands of multispectral imagery to quantify vegetation vigor. It has been used extensively in agriculture, conservation, and climate research to monitor ecosystem conditions over time.

Semantic Segmentation for Land Use Classification

Deep learning techniques like U-Net and DeepLabv3 are employed for semantic segmentation, where each pixel in an image is classified into categories such as vegetation, urban areas, water bodies, or barren land. This pixel-level analysis enables high-resolution mapping of land use, which is essential for environmental planning and urban management.

Siamese Neural Networks in Change Detection

Siamese Neural Networks (SNNs) have been employed in environmental applications to compare temporal images and detect changes such as deforestation, urban sprawl, or flood damage. By learning the similarity between image pairs, SNNs can identify subtle differences, making them valuable for monitoring environmental shifts over time.

Google Earth Engine and Its Limitations

Google Earth Engine (GEE) is a powerful platform for processing and analyzing geospatial data. While it provides access to a wealth of satellite imagery and analytical tools, its complexity and code-centric interface pose a barrier for



non-technical users. GEE lacks the simplicity and interactivity required for widespread adoption in educational or civic contexts.

AI APIs for Image Analysis

Large Language Model (LLM)-powered APIs like Google's Gemini represent a new frontier in image analysis. These systems can interpret images contextually, generate descriptive summaries, and respond to structured prompts. They eliminate the need for custom-trained models, making advanced analysis more accessible to developers and non-specialists.

Weather Context in Environmental Analysis

Incorporating weather data into image analysis enhances the reliability of insights. For example, understanding temperature and precipitation levels at the time of image capture can refine vegetation health assessments. APIs like OpenWeatherMap offer real-time data that adds environmental context to satellite-based analyses.

Integration of Frontend Tools for Visualization

Modern frontend libraries such as React.js and Chart.js have made it easier to build responsive and interactive dashboards. These tools play a crucial role in presenting environmental insights in a format that users can understand and act upon, bridging the gap between complex data and actionable knowledge.

IV. OBJECTIVE

The primary objective of TreeSense is to provide a lightweight, AI-driven platform that empowers users to analyze satellite or aerial images for environmental monitoring without the need for technical expertise or extensive backend infrastructure. By focusing on usability, accessibility, and visual interpretation, TreeSense aims to make advanced geospatial analysis approachable for students, planners, researchers, and civic authorities alike.

One of the core functionalities of the system is to enable tree counting from satellite images. This objective addresses the growing need for automated, scalable methods to quantify urban and rural vegetation. Instead of relying on manually annotated images or custom-trained detection models, TreeSense conceptually integrates models like YOLOv8 through API-driven interpretation to deliver accurate tree detection and counting with minimal overhead.

Another major goal is the classification of land use into categories such as vegetation, urban, water, or barren land. Accurate land cover analysis supports a variety of use cases, from green audits to zoning compliance and environmental impact assessments. By simulating semantic segmentation approaches via Gemini API, TreeSense extracts land use insights directly from image content, without requiring pre-trained classifiers.

The platform also aims to provide vegetation health analysis using NDVI-like metrics. While traditional NDVI requires multispectral inputs, TreeSense infers vegetation vitality by interpreting color patterns and environmental cues through the AI model. This feature allows users to gauge ecosystem health in a simplified yet informative manner, making the tool useful for agriculture, forestry, and climate tracking.

An important comparative objective is to analyze environmental changes over time. TreeSense enables users to compare two satellite images from different time periods to identify deforestation, urban development, or seasonal variation. This capability mimics the role of Siamese neural networks in change detection but is implemented through dynamic prompting and interpretation, allowing for intuitive temporal analysis.

To ensure accessibility, TreeSense has been designed as a frontend-only solution with no reliance on backend databases or machine learning servers. This decision aligns with the goal of creating a portable, easily deployable system that can run on standard browsers and devices with just an internet connection, significantly lowering the entry barrier for environmental monitoring.

Finally, the project seeks to offer interactive and intuitive visual outputs that translate complex data into actionable insights. By using tools like React.js and Chart.js/Recharts, TreeSense transforms AI-generated responses into readable charts, summary cards, and badges. This visual-first approach ensures that end-users—regardless of technical background—can interpret results and make informed decisions quickly and confidently.



V. METHODOLOGY OF PROBLEM SOLVING AND EFFICIENCY ISSUES

To address the challenge of accessible environmental monitoring, the TreeSense system was architected using a fully frontend-driven approach, emphasizing modularity, low latency, and portability. Users interact with the application via a browser interface built in React.js, where they can select a region using an embedded Mapbox map or upload custom images. Once the area is selected, its geographic coordinates and surface area are automatically calculated, enabling spatial context for subsequent analysis steps without requiring backend computation or storage.

The core image analysis process is facilitated through prompt-based interactions with the Gemini API, a multimodal AI system. Instead of training custom detection models like YOLO or segmentation networks such as U-Net, TreeSense offloads this task to Gemini via carefully engineered prompts. These prompts include the image, metadata, and user-selected analysis tasks such as tree counting, NDVI estimation, or land cover classification. The API returns a structured output that is parsed and visualized using JavaScript libraries like Chart.js and Recharts, offering users insights in the form of charts, percentages, and textual summaries.

For comparative analysis, users can upload two satellite images taken at different times. The system constructs a differential prompt to Gemini, asking it to detect environmental changes such as vegetation loss or urban growth. This approach mimics the behavior of Siamese neural networks but requires no model training or inference time on the client side. By relying on Gemini's LLM to perform context-aware reasoning, TreeSense achieves high flexibility in change detection without computational expense.

To ensure contextual accuracy, TreeSense integrates the OpenWeatherMap API to fetch real-time weather conditions and environmental metadata for the selected region. This weather data enhances the analysis by correlating vegetation health with current temperature or precipitation levels. However, in cases of weather API failure or image resolution issues, fallback mechanisms such as mock responses or simplified prompts are used to ensure the system remains responsive and usable.

Despite its advantages, the efficiency of this approach is bounded by external API limitations. The response time and accuracy of image interpretation depend on the performance of the Gemini API, which may vary based on server load or image complexity. Additionally, image quality and resolution directly impact the accuracy of tree detection and land classification. The absence of backend processing reduces latency but also imposes constraints on advanced features like result history or high-resolution processing. Future versions of TreeSense aim to mitigate these issues by incorporating optional backend support and optimized local processing modules where necessary.

VI. RELATED WORKS

The TreeSense project builds upon and intersects with several existing technologies and research domains, particularly in the areas of remote sensing, environmental monitoring, and AI-powered image analysis. This section highlights notable works that inspired the system's features and guided its methodological choices.

One of the most directly relevant technologies is the YOLO (You Only Look Once) object detection model, particularly the latest version, YOLOv8. It has been used extensively in forestry applications to detect individual tree crowns in aerial imagery, enabling automated tree counting and classification. While TreeSense does not directly implement YOLOv8 due to frontend limitations, the conceptual framework for object detection in vegetation mapping has heavily influenced the image analysis goals of the platform.

Another foundational component is the NDVI (Normalized Difference Vegetation Index), a widely used vegetation health indicator in satellite image analysis. Traditional NDVI computation requires multispectral imagery to compare red and near-infrared light reflectance. While TreeSense does not use raw spectral bands, it conceptually simulates NDVI estimation through image interpretation using the Gemini API, making it more accessible for casual users and those without access to multispectral data.

The domain of semantic segmentation is also closely related, particularly works involving U-Net and DeepLabv3 models. These networks are commonly used in pixel-level land cover classification tasks, identifying regions of forest, water, urban development, and barren land. While TreeSense does not implement these models natively, its use of large language models to approximate segmentation outputs via natural image interpretation offers a lightweight alternative that preserves user accessibility.



In the context of change detection, Siamese Neural Networks (SNNs) have been successfully applied to detect differences in satellite images taken over time. These networks learn to compare two inputs and identify significant changes in vegetation, structure, or land cover. TreeSense replicates this functionality by prompting the Gemini API to evaluate and describe differences between two user-provided images, bypassing the need for complex model training while still achieving insightful comparisons.

TreeSense's architecture also draws inspiration from Google Earth Engine (GEE), a powerful platform for satellite data analysis. GEE allows users to write custom code to analyze time-series and geospatial data at scale. However, its steep learning curve and code-heavy interface limit usability for non-technical users. TreeSense addresses this gap by offering a visual, prompt-based system for similar analytical tasks without requiring any programming knowledge.

In terms of image interpretation via AI, Google's Gemini API represents a cutting-edge tool that combines computer vision and language understanding. Its ability to interpret visual data and respond to structured prompts enables a new generation of image analysis applications. TreeSense uses Gemini as a backend-as-a-service AI engine, demonstrating how general-purpose multimodal models can be adapted for domain-specific tasks like tree detection and land monitoring.

The integration of OpenWeatherMap API also connects TreeSense with real-time environmental data sources. Projects in climate science and precision agriculture often combine weather data with satellite imagery to gain deeper insights. TreeSense uses this concept to provide enriched, context-aware analysis that factors in temperature and atmospheric conditions at the time of observation.

Finally, TreeSense benefits from advancements in interactive frontend technologies such as React.js and Chart.js, which are used widely in dashboards and data visualization systems. These tools allow complex data to be rendered in user-friendly visual formats, bridging the gap between technical analytics and decision-making interfaces for general users.

VII. TREE COUNTING FEATURE

Tree counting is one of the core functionalities of TreeSense, aimed at automating the detection and enumeration of trees in a given satellite or aerial image. Traditionally, this task requires manual annotation or specialized detection models, which are time-consuming and demand expert knowledge. TreeSense simplifies this process by leveraging the Gemini API to identify and count trees based on image features such as canopy shape, shadow, and color contrast.

The system uses conceptual object detection techniques similar to YOLOv8, which are known for their speed and accuracy in real-time applications. When a user uploads or selects an image, a structured prompt is sent to Gemini requesting identification of tree objects. The response includes an estimated tree count and, in some cases, density measures or spatial distribution insights, allowing users to understand vegetation spread across an area.

This feature is especially useful for urban green audits, forest inventory checks, or ecological studies. It empowers users—without GIS or machine learning knowledge—to perform essential forestry tasks in seconds. Though the model works on visible data rather than multispectral input, the results remain surprisingly reliable due to Gemini's contextual reasoning capabilities.

VIII. LAND USAGE DISTRIBUTION FEATURE

The land usage distribution feature helps classify an area into different land types—such as vegetation, barren land, urban infrastructure, and water bodies. This classification provides a holistic understanding of how land is utilized, which is essential for urban planning, environmental compliance, and resource management.

TreeSense achieves this by simulating semantic segmentation, typically done by deep learning models like U-Net or DeepLabv3. However, instead of training such models, the system constructs a carefully worded prompt for Gemini, which analyzes color, texture, and layout to infer land types. This technique provides approximate yet meaningful segmentation results suitable for non-specialist users.

Results are returned as a percentage distribution of each land type and visualized using pie charts or bar graphs. This makes it easy for users to interpret spatial dominance—for instance, if an image contains 40% vegetation and 30% urban development. The visual representation helps in quickly identifying land use imbalances or encroachments.



This feature plays a key role in smart city development, land zoning, and environmental policy enforcement. By allowing side-by-side image comparisons, users can also detect changes in land usage patterns over time. Although not pixel-perfect like backend models, the simplicity, speed, and accessibility of this approach make it highly practical for real-world use cases.

IX. NDVI ESTIMATION FEATURE

Vegetation health analysis in TreeSense is based on a simulated version of the Normalized Difference Vegetation Index (NDVI). NDVI is a widely accepted remote sensing metric calculated using near-infrared (NIR) and red spectral bands, which are not available in standard RGB images. TreeSense circumvents this limitation by using AI-based visual interpretation instead of spectral calculation.

When a user selects this feature, TreeSense sends the image and metadata to Gemini with a prompt asking for vegetation health assessment. The model analyzes visual cues such as leaf color, density, and texture to approximate NDVI values. Though not scientifically precise like true NDVI, this AI-inferred method provides an accessible and reliable alternative for vegetation monitoring.

The health status is categorized into zones such as "Healthy," "Moderate," and "Stressed," along with a numerical score on a 0–1 scale. These outputs are displayed with color-coded badges and comparison charts, making them easy to interpret at a glance. Users can track trends or spot problem areas without the need for spectral cameras or GIS expertise.

This feature is particularly beneficial for agriculture, reforestation projects, and ecological assessments. It enables timely decisions, such as irrigation scheduling or disease intervention, based on visual health cues alone. By democratizing vegetation health monitoring, TreeSense provides a powerful tool for sustainable land management.

X. ENVIRONMENTAL SUGGESTIONS FEATURE

Environmental suggestions provide users with actionable insights based on the analyzed image and external factors like weather conditions. This feature transforms raw analytical outputs into meaningful guidance, bridging the gap between data and decision-making. For example, based on tree cover and temperature data, the system might recommend afforestation, irrigation, or urban greening efforts.

The suggestions are generated through rule-based inference, where TreeSense combines input parameters such as vegetation percentage, NDVI score, and current weather to derive contextual recommendations. For instance, if an area has low green cover and high temperature, the system may advise planting trees or installing green roofs. These recommendations are presented in plain language, making them accessible to all types of users.

In addition to environmental context, TreeSense evaluates spatial patterns—for example, if trees are unevenly distributed or clustered in a specific zone, the system might suggest redistributing green areas to improve air quality and heat management. These insights are helpful for urban planners, policy makers, and community initiatives looking to enhance local sustainability.

What sets this feature apart is its interpretive intelligence, powered by the Gemini API's ability to contextualize visual and textual data together. Users receive not just information, but insights they can act on—making this feature a vital link between image analysis and real-world impact.

XI. IMAGE COMPARISON (CHANGE DETECTION) FEATURE

TreeSense's image comparison feature allows users to upload two satellite or aerial images of the same region taken at different times to detect environmental changes. This feature is critical for tracking deforestation, urban development, vegetation loss, and other dynamic land transformations over time.

The system leverages Gemini's language and vision capabilities to simulate the output of a Siamese Neural Network, a common architecture used in temporal change detection tasks. TreeSense does not require pixel-perfect alignment; instead, the prompt explains the task and provides both images, asking Gemini to identify and summarize changes in tree cover, land usage, or vegetation health.



The output includes side-by-side statistics and a textual summary of observed changes—such as “25% reduction in vegetation” or “new urban structures in northern quadrant.” These are presented in tables, charts, and descriptive text, allowing users to grasp the environmental evolution of a region quickly and intuitively.

This feature has broad applications in environmental monitoring, disaster recovery (e.g., post-flood assessments), municipal compliance checks, and agricultural planning. It enables evidence-based decision-making, helping users understand the consequences of land-use decisions or natural phenomena. Its frontend-only architecture ensures quick results without the complexity of backend computation.

XII. STATEMENT OF SCOPE

The TreeSense project is designed to provide a lightweight, accessible platform for environmental monitoring through satellite and aerial image analysis. It focuses specifically on delivering key insights such as tree count, vegetation health, land usage classification, and environmental change detection without the need for backend servers, databases, or custom-trained machine learning models. By operating entirely through a frontend interface, TreeSense offers a rapid, cost-effective, and scalable solution that is easy to deploy and use across standard web browsers.

The scope of TreeSense includes real-time geographic area selection using an interactive Mapbox interface, as well as support for custom image uploads. Once an area is selected, users can trigger AI-powered analyses that are processed through the Gemini API, returning structured environmental data and summaries. This includes not only quantitative metrics—such as tree count or NDVI score—but also visual interpretations like pie charts, bar graphs, and health indicators. In addition, weather data is fetched using OpenWeatherMap to provide context-aware insights.

TreeSense also allows for comparative analysis between two time-separated images, enabling users to detect environmental changes such as deforestation, land conversion, or urban expansion. This feature simulates the function of change detection models, offering meaningful comparisons through natural language outputs and visual summaries. The system’s ability to interpret visual data and offer actionable environmental suggestions further broadens its use cases in urban planning, agriculture, and green policy design.

While the project does not include persistent data storage or advanced GIS computation, this constraint is intentional. The scope remains firmly focused on creating a no-installation, low-barrier tool for non-specialists, students, civic bodies, and NGOs. By balancing functionality with simplicity, TreeSense demonstrates how AI and remote sensing technologies can be made more inclusive and impactful through thoughtful design choices.

XIII. SYSTEM ARCHITECTURE

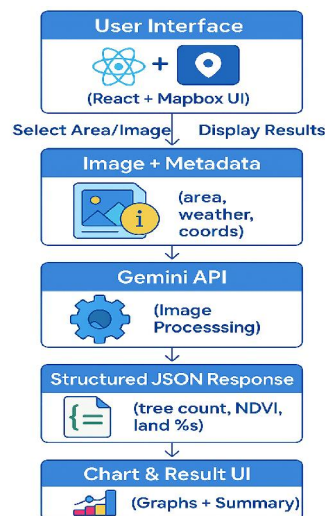


Fig.1 System architecture design



The system architecture of the TreeSense project is designed as a streamlined, modular, and frontend-driven pipeline to facilitate vegetation monitoring using satellite imagery. The core components are visualized in a vertical, layered architecture, beginning with the User Interface (UI) built using React.js and Mapbox. This interface allows users to either select a specific geographic area directly on a map or upload a satellite image. Upon selection, essential geospatial metadata—such as coordinates and area coverage—is extracted and prepared for processing.

Once the area or image is selected, the system enters the Image + Metadata stage. Here, contextual information is added, including weather data fetched via the OpenWeatherMap API. This enriched dataset—containing area size, weather conditions, and location details—is packaged and sent to the Gemini API, which serves as the core AI engine. Gemini performs image processing operations such as tree detection, vegetation health analysis (NDVI), and land classification. The prompt sent to Gemini is constructed dynamically using structured, descriptive inputs, ensuring consistent and relevant responses.

After processing, the Gemini API returns a Structured JSON Response. This output includes key metrics such as the number of detected trees, vegetation indices (NDVI), and land usage percentages (e.g., green cover vs barren land). These structured results are parsed by the frontend and passed on to the Chart & Result UI component. This module is responsible for rendering the insights in a visually intuitive format using graphical elements like pie charts, bar graphs, and statistical summaries. Libraries like Chart.js or Recharts are used to support real-time rendering of results.

This fully client-side architecture, as shown in the system diagram, allows TreeSense to operate without a backend server or custom-trained models. By leveraging APIs and browser-based computation, the platform remains lightweight, scalable, and highly accessible. The architecture is optimized for responsiveness, making it suitable for academic research, urban planning, and environmental audits with minimal technical barriers.

Figure 2: The data flow model illustrated above captures the operational logic of the TreeSense system, guiding how data moves from user interaction to actionable insight. The process begins when a user selects an area using the Mapbox-integrated interface. Upon selection, the system calculates both the area size and corresponding geographic coordinates. Alternatively, the user can upload one or two satellite images for manual input. This data, along with optional weather context, forms the basis of a prompt that is sent to the Gemini API for intelligent analysis.

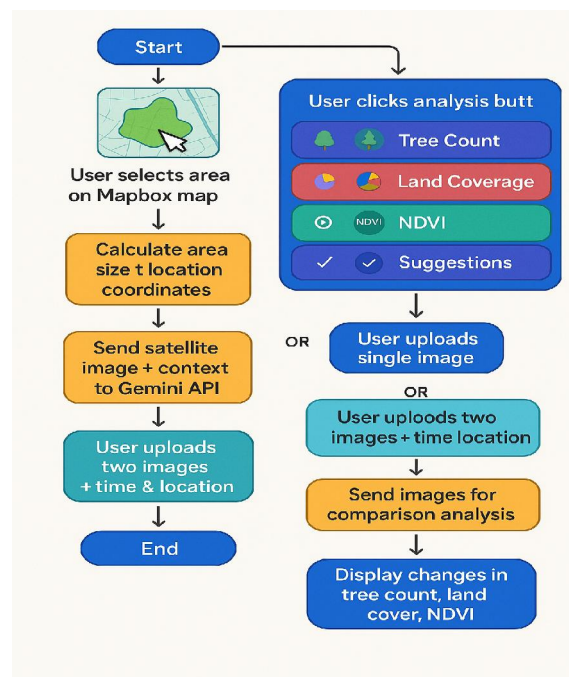


Fig.2 Data Flow Diagram



The core processing occurs at the Gemini API stage, where the uploaded image(s) are interpreted to extract features such as tree count, land coverage, and NDVI (Normalized Difference Vegetation Index). Users can choose different analysis options—either to assess a single image or to perform comparative analysis using two time-stamped images. When two images are uploaded, the system processes them sequentially to identify changes over time in vegetation density, land usage, and environmental conditions. The final output is presented in a user-friendly results panel that summarizes the findings visually and textually.

This data flow model effectively captures both linear analysis (single-image processing) and temporal analysis (multi-image comparison), making TreeSense a flexible tool for environmental monitoring. The use of frontend-only technologies ensures a lightweight architecture, while the integration with AI (via Gemini) allows for powerful image interpretation without the need for a custom-trained backend.

XIV. OTHER SPECIFICATIONS

Frontend-Only Architecture:

TreeSense is designed entirely as a frontend-based application using React.js. The entire functionality—from user input to image processing and result visualization—occurs within the user's browser. There is no need for server-side processing or backend hosting. This architecture ensures the platform remains lightweight and highly responsive, making it ideal for quick deployment on platforms like GitHub Pages or Netlify. It also reduces infrastructure complexity and makes the system easier to maintain.

No Persistent Storage or Database:

The system does not utilize any form of backend database or permanent storage. All user interactions, uploaded images, and output results are handled temporarily in the session memory of the browser. As a result, once the user refreshes or exits the page, the data is cleared. This ensures data privacy and avoids the complications of storage management, user accounts, or GDPR compliance. It also simplifies the development cycle by keeping the architecture stateless.

External API Integration:

TreeSense integrates several powerful APIs to achieve real-time and intelligent environmental analysis. The Gemini API is used to perform AI-based image interpretation, including tree detection, NDVI calculation, and land classification. Mapbox provides the user interface for geographic area selection, while OpenWeatherMap is used to collect real-time weather data relevant to the selected coordinates. By leveraging these third-party APIs, the system avoids the need for custom machine learning models and backend computation, while still delivering accurate and context-aware insights to the user.

XV. LIMITATIONS

Lack of Backend and Data Persistence: TreeSense operates entirely on the frontend without any backend support or database. As a result, the application cannot store user-uploaded images, historical data, or analysis results. Once the session ends or the page is refreshed, all inputs and outputs are lost. This limits the system's usefulness for users who require long-term environmental tracking, reporting, or exporting of results. It also prevents features like user authentication, analysis history, or collaborative usage, which are often essential in research or municipal applications.

Dependency on External APIs and Model Generalization: The functionality of TreeSense heavily relies on external services such as the Gemini API for AI analysis and OpenWeatherMap for environmental data. Any downtime, quota exhaustion, or latency in these APIs can disrupt the system's ability to generate results. Furthermore, since the Gemini API uses generalized AI models not specifically trained on localized or domain-specific satellite data, the accuracy of tree counting, NDVI scoring, and land classification may vary. This can impact reliability, especially in complex or mixed-terrain environments where precision is critical.

XVI. APPLICATIONS

Municipal Green Audits and Urban Planning: TreeSense can be effectively used by municipal authorities and urban planners to conduct green audits and monitor vegetation across city landscapes. By enabling quick analysis of satellite images and estimating tree counts, NDVI, and land use distribution, the platform helps track green cover changes over



time. This information is essential for ensuring compliance with environmental regulations, planning new green zones, and maintaining ecological balance in growing urban areas.

Forestry and Environmental Monitoring: The system is highly applicable in forest management and conservation projects. Forestry departments can use TreeSense to monitor deforestation, identify barren or degraded areas, and measure the effectiveness of reforestation efforts. Since it supports image comparisons, it allows users to detect changes over time, making it a valuable tool for evaluating seasonal or long-term vegetation dynamics without needing extensive on-ground surveys.

Agricultural and Irrigation Planning: TreeSense also has potential applications in agriculture, particularly in assessing crop health and planning irrigation. By analyzing NDVI values and land cover, farmers and agricultural planners can identify areas of healthy versus stressed vegetation. This supports smarter resource allocation, early problem detection, and more efficient land management. Combined with weather data, the insights provided can guide decisions on crop rotation, watering schedules, and soil health maintenance.

XVII. CONCLUSION

TreeSense demonstrates that powerful environmental analysis can be achieved using a lightweight, frontend-only architecture by integrating modern APIs and AI services. The project successfully provides users with the ability to assess vegetation health, count trees, and understand land usage using satellite imagery without requiring backend servers or custom machine learning models. Its interactive interface, real-time weather integration, and visual result presentation make it both user-friendly and accessible for a wide range of applications—from academic research to civic planning.

However, while TreeSense excels in simplicity and ease of use, it does face limitations related to data persistence and dependency on external APIs. Despite these constraints, the system lays a strong foundation for scalable and intelligent green monitoring tools. With potential enhancements like backend integration, model fine-tuning, and real NDVI data support, TreeSense could evolve into a more comprehensive platform for environmental monitoring, making a meaningful contribution to sustainability, urban planning, and conservation efforts.

XVIII. FUTURE SCOPE

Model Optimization: Future research will focus on further optimizing the detection algorithms to enhance performance, particularly in diverse and challenging environments.

Handling Diverse Tree Species: Expanding the dataset to include a wider variety of tree species and environmental conditions will improve the model's generalization capabilities and accuracy.

Real-Time Implementation: Developing a real-time tree detection system that can process aerial imagery on-the-fly will enhance the practicality of the solution in dynamic monitoring scenarios.

Integration with IoT: Exploring the integration of the system with Internet of Things (IoT) devices could facilitate continuous monitoring of forests and urban green spaces, providing real-time data and alerts.

User-Friendly Interface: Creating a user-friendly interface for stakeholders, such as urban planners and conservationists, will make the technology more accessible and enhance its usability in practical applications.

Collaborative Projects: Collaborating with environmental organizations and government agencies to implement the system in real-world scenarios will provide valuable insights and drive further enhancements.

Expanding Applications: Investigating additional applications of the technology, such as detecting other vegetation types or assessing environmental health, will broaden its impact and usefulness.

REFERENCES

- [1]. L. Yang, X. Wu, E. Praun and X. Ma, "Tree detection from aerial imagery", Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems ser. GIS '09, pp. 131-137, 2009.
- [2]. J. Secord and A. Zakhor, "Tree detection in urban regions using aerial lidar and image data", IEEE Geoscience and Remote Sensing Letters, vol. 4, no. 2, pp. 196-200, 2007.



- [3]. B. G. Weinstein, S. Marconi, M. Aubry-Kientz, G. Vincent, H. Senyondo and E. P. White, "Deepforest: A python package for rgb deep learning tree crown delineation", *Methods in Ecology and Evolution*, vol. 11, no. 12, pp. 1743-1751.
- [4]. T. Lin, P. Goyal, R. Girshick, K. He and P. Dollar, "Focal loss for dense object detection", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, no. 02, pp. 318-327, feb 2020.
- [5]. K. Sohn, Z. Zhang, C.-L. Li, H. Zhang, C.-Y. Lee and T. Pfister, "A simple semi-supervised learning framework for object detection", 2020.
- [6]. R. Hoffmann, C. Zhang, X. Ling, L. Zettlemoyer and D. S. Weld, "Knowledge-based weak supervision for information extraction of over-lapping relations", *Proceedings of the 49th annual meeting of the association for computational linguistics: human language technologies*, pp. 541-550, 2011.
- [7]. M. Xu, Z. Zhang, H. Hu, J. Wang, L. Wang, F. Wei, et al., "End-to-end semi-supervised object detection with soft teacher", 2021.
- [8]. K. He, X. Zhang, S. Ren and J. Sun, "Deep residual learning for image recognition", *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770-778, 2016.
- [9]. Krizhevsky, I. Sutskever and G. E. Hinton, "Imagenet classification with deep convolutional neural networks", *Advances in neural information processing systems*, vol. 25, 2012.
- [10]. Y. Wu, A. Kirillov, F. Massa, W.-Y. Lo and R. Girshick, *Detectron2*, 2019,
- [11]. W. S. Noble, "What is a support vector machine?", *Nature biotechnology*, vol. 24, no. 12, pp. 1565-1567, 2006.
- [12]. P. Patil, V. H., S. Patil and U. Kulkarni, "Wireless sensor network for precision agriculture", *2011 International Conference on Computational Intelligence and Communication Networks*, pp. 763-766, Oct 2011.
- [13]. Kamilaris, F. Gao, F. X. Prenafeta-Boldu and M. I. Ali, "Agri-iot: A semantic framework for internet of things-enabled smart farming applications", *2016 IEEE 3rd World Forum on Internet of Things (WF-IoT)*, pp. 442-447, Dec 2016.
- [14]. Vibhute and S. Bodhe, "Applications of image processing in agriculture: a survey", *International Journal of Computer Applications*, vol. 52, no. 2, 2012.
- [15]. L. Wang, P. Gong and G. S. Biging, "Individual tree-crown delineation and treetop detection in high-spatial-resolution aerial imagery", *Photogrammetric Engineering Remote Sensing*, vol. 70, no. 3, pp. 351-357, 2004.
- [16]. G. da Silva Vieira, B. M. Rocha, F. Soares, J. C. Lima, H. Pedrini, R. Costa, et al., "Extending the aerial image analysis from the detection of tree crowns", *2019 IEEE 31st International Conference on Tools with Artificial Intelligence (ICTAI)*, pp. 1681-1685, 2019.
- [17]. E. Ponce-Corona, M. G. Sanchez, D. Fajardo-Delgado, B. Acevedo-Juarez, M. De-la Torre, H. Avila-George, et al., "Una revisión sistemática de la literatura enfocada al uso de vehiculos aereos no tripulados durante el proceso de deteccion de vegetacion", *RISTI-Revista Iberica de Sistemas e Tecnologias de Informacao*, pp. 82-101, 03 2020.
- [18]. E. Belcore, A. Wawrzaszek, E. Wozniak, N. Grasso and M. Piras, "Individual tree detection from uav imagery using hölder exponent", *Remote Sensing*, vol. 12, no. 15, pp. 2407, 2020.
- [19]. Wloka, T. Kunić, I. Kotseruba, R. Fahimi, N. Frosst, N. D. Bruce, et al., "Smiler: Saliency model implementation library for experimental research", *arXiv preprint*, 2018.
- [20]. M. T. Chiu, X. Xu, K. Wang, J. Hobbs, N. Hovakimyan, T. S. Huang, et al., "The 1st agriculture-vision challenge: Methods and results", *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 48-49, 2020.
- [21]. M. T. Chiu, X. Xu, Y. Wei, Z. Huang, A. G. Schwing, R. Brunner, H. Khachatrian, H. Karapetyan, I. Dozier, G. Rose et al., "Agriculture-vision: A large aerial image database for agricultural pattern analysis", *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2828-2838, 2020.



- [22]. S. Jia and N. B. Bruce, "Revisiting saliency metrics: Farthest-neighbor area under curve", 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). Los Alamitos, pp. 2664-2673, jun 2020.
- [23]. N. D. B. Bruce and J. K. Tsotsos, "Saliency based on information maximization", Proceedings of the 18th International Conference on Neural Information Processing Systems ser. NIPS05, pp. 155-162, 2005.
- [24]. J. Harel, C. Koch and P. Perona, "Graph-based visual saliency", Advances in neural information.
- [25]. H. J. Seo and P. Milanfar, "Nonparametric bottom-up saliency detection by self-resemblance", 2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, pp. 45-52, 2019.
- [26]. A D. Panagiotidis and P. Surovy, "Estimation and extrapolation of tree parameters using spectral correlation between uav and pleiades data", Forests, vol. 9, no. 2, pp. 85, 2018.
- [27]. F. DadrasJavan, F. Samadzadegan, S. H. S. Pourazar and H. Fazeli, "Uav-based multispectral imagery for fast citrus greening detection", Journal of Plant Diseases and Protection, vol. 126, no. 4, pp. 307-318, 2019.
- [28]. G. Khadanga and K. Jain, "Tree census using circular hough transform and grvi", Procedia Computer Science, vol. 171, pp. 389-394, 2020.
- [29]. J. Zhang and J. Hu, "Image segmentation based on 2d otsu method with histogram analysis", 2008 International Conference on Computer Science and Software Engineering, vol. 6, pp. 105-108, 2021.

