© 2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. This is the author's accepted version of the article accepted for publication in IEEE Transactions on Visualization and Computer Graphics (IEEE VIS 2026). The final version of this record will be available on IEEE Xplore.

ForestVis: Visualization of 3D Forestry Point Clouds

Benjamin Powley, Petra Horváth, Basam Dahy, Nivan Ferreira, Johan E. S. Fransson, Andreas Kerren, Claudio D. G. Linhares, and Amilcar Soares

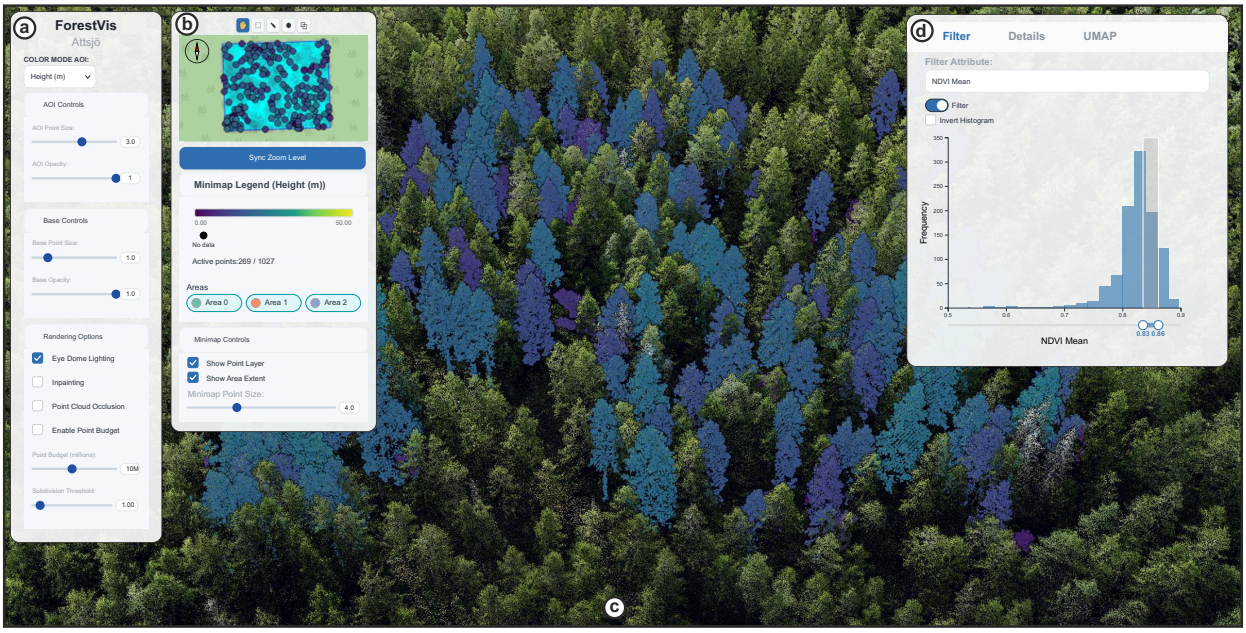


Figure 1: ForestVis is a web-based visualization system for tree-level forestry analysis. Panel (a) provides rendering controls, including color scale and point sizes for the background and AOI. Panel (b) shows the minimap and legend. The point cloud view (c) compares tree height and NDVI using a filtered height layer, while panel (d) filters data by mean NDVI.

Abstract—Forestry presents complex, irregular environments that are underrepresented in the visualization literature, despite the increasing availability of large-scale LiDAR and multispectral datasets. Extracting meaningful information from this data requires analyzing subtle structural and physiological variations among individual trees, a task that remains challenging due to the irregular, multi-layered nature of forests. In this paper, we present a visualization-driven approach for the interactive exploration and analysis of 3D forestry point clouds at the tree level, enriched with Normalized Difference Vegetation Index (NDVI) metrics. Our web-based system, entitled ForestVis, enables users to navigate dense forest environments, query individual tree attributes, and analyze vegetation metrics within their spatial context. The system design is informed by a semi-structured interview with forestry domain experts, from which we derive a set of analytical tasks emphasizing tree-centric abstractions, interactive exploration, and scalability to large datasets. Through illustrative use cases, we show how the system supports applications such as precision forestry, biodiversity monitoring, and detection of tree stress. More broadly, this work highlights the potential of interactive visualization providing novel tools and frameworks for ecological research and sustainable forest management. We evaluate ForestVis through two expert studies: a qualitative study of usability and analytical reasoning, and a quantitative study reporting objective task measures together with benchmarks of rendering, streaming, and interaction performance across two hardware configurations. Our evaluation suggested that our tool is effective in helping domain experts explore specific forest tasks, in large forest point cloud datasets, while maintaining computational performance. Our findings suggest that different point cloud rendering styles are suited to different tasks, whether analyzing individual trees or groups of trees, and interactive filtering helped experts relate vegetation indices to the point-cloud instance segmentation and to tree height.

Index Terms—Forestry Visualization, 3D Point Clouds, LiDAR, NDVI, Vegetation Monitoring.

1 Introduction

- *B. Powley, P. Horváth, C. D. G. Linhares, and A. Soares are with the Dep. of Computer Science and Media Technology, Linnaeus University, Sweden. E-mail: {benjamin.powley, amilcar.soares, claudio.linhares}@lnu.se*
- *B. Dahy, and J. Fransson are with the Dep. of Forestry and Wood Technology, Linnaeus University, Sweden. E-mail: {bassam.saeed, johan.fransson}@lnu.se*
- *N. Ferreira is with the Federal University of Pernambuco, Brazil. E-mail: nivan@cin.ufpe.br*
- *A. Kerren is with the Dep. of Science and Technology, Linköping University and with Linnaeus University, Sweden. E-mail: andreas.kerren@{liu/lnu}.se*

Forests play a critical role in sustaining global ecosystems [68], supporting biodiversity, regulating climate [7], and providing essential resources, yet they are increasingly threatened by environmental change, disease, and human activity [68]. Advances in remote sensing technologies, such as Light Detection and Ranging (LiDAR) and multispectral imaging, have enabled the large-scale acquisition of detailed three-dimensional (3D) representations of forest environments in the form of point clouds [61]. These datasets capture geometric properties from the LiDAR and spectral properties from multispectral imagery at the level of individual trees, creating new opportunities for fine-grained forest analysis. However, their scale and complexity require effective visualization and interactive exploration techniques. The increasing

availability of 3D point cloud data has made its visualization a central topic in visualization research [13, 43]. Existing techniques address challenges such as scalable rendering, abstraction, and interaction for exploring dense spatial data [44, 60]. In applied contexts, however, most work has focused on urban environments, where structured objects such as buildings and roads support clearly defined analytical tasks, including navigation, semantic exploration, and city-scale modeling [11, 13, 25, 26, 49, 57, 67].

Forests are composed of irregular, biologically driven structures with high variability, and meaningful analysis often requires reasoning about individual trees rather than aggregated or standardized primitives [36, 42]. While drone-acquired 3D data (including LiDAR and optical imagery) provides unprecedented coverage and detail, extracting actionable insights from these large datasets remains challenging. Existing systems have largely focused on data storage, processing, or rendering performance, with limited support for task-driven, exploratory analysis [14]. Forestry analysis presents several key challenges for interactive visualization. Extracting meaningful information from large-scale 3D forest data requires identifying subtle structural and physiological variations among individual trees [28, 73]. Detecting stressed or dying trees in dense stands combines geometric, spectral, and contextual evidence [48], while deriving structural metrics and comparing vegetation characteristics across forest regions requires reasoning over highly irregular and heterogeneous vegetation distributions [5, 24, 34, 50]. These challenges are inherently visual and call for new abstractions and interaction techniques that support tree-level querying, comparison, and exploration within vast 3D datasets.

In this paper, we present ForestVis (Fig. 1), a visualization-driven approach to support tree-level analysis of large-scale forest point clouds designed in collaboration with domain experts. Specifically, we:

- Identify a set of analytical tasks for tree-level forest analysis from semi-structured interviews with forestry domain experts, focusing on tree-level querying, comparison, and multi-scale exploration.
- Design and implement a web-based visualization system for interactive analysis of 3D forest point clouds enriched with vegetation indices.
- Design tree-centric visual representations and coordinated interactions that combine established rendering, filtering, and projection techniques to support the analysis of structural and physiological attributes in spatial context.
- Perform a qualitative expert evaluation of the usability, and a quantitative assessment of computational performance (expert task measures and benchmarking) of our system.

2 Related Work

2.1 Visualization of Forestry Data

Forestry visualization systems represent forest environments as three-dimensional models to support the analysis of structural properties derived from remote sensing data. Depending on the analytical objective, forests can be studied at multiple levels of detail [33], including the canopy level (i.e., the upper layer of vegetation formed by tree crowns), forest stands (groups of trees within a defined area), or at the level of individual trees as distinct analytical units. These levels reflect common forestry practices, where both aggregated and tree-level analyses are required to understand forest structure, health, and dynamics.

Forest data can be represented using different modeling approaches, including surface-based meshes and data-driven point clouds. Meshes approximate forest structure by inferring surfaces, for example through convex or concave hulls around tree crowns [62], while point clouds directly represent individual measurements captured by LiDAR sensors. This distinction is critical, as point clouds preserve fine-grained geometric detail and enable the inspection of acquisition quality, including gaps or noise in the data. Point clouds have been widely used to visualize Terrestrial Laser Scanning (TLS) LiDAR data and support tasks such as tree segmentation and the extraction of structural attributes, including Diameter at Breast Height (DBH), height, and crown volume [62]. Alternative approaches use parameterized or procedurally generated models to represent trees based on attributes such as species, size, and location [23, 37, 58]. While these representations can support communication and visual realism, they abstract away the exact geometry and spectral properties of individual trees. Realistic or efficient rendering of dense vegetation is itself a well-established problem, addressed for example by level-of-detail and aggregate-detail simplification techniques [9]; this goal is distinct from the interactive, attribute-driven analysis of measured tree-level data that we target.

Visualization techniques have been applied to analyze forest attributes across spatial and attribute dimensions. Kernel Density Estimation (KDE) has been used to reveal patterns in vegetation indices from remote sensing data [17], while dimensionality reduction techniques such as UMAP [35] have been applied to high-dimensional remote sensing data, e.g., to relate topography to vegetation dynamics in a nature reserve [74] as well as to reduce and visualize PolInSAR features [52]. Similarly, interaction techniques such as attribute-based filtering and scented widgets support exploratory analysis of data distributions [11, 71]. However, these approaches are typically applied in isolation and are not integrated into systems that support interactive, tree-level analysis of geospatial point cloud data.

In this work, we address this gap by combining point cloud visualization with vegetation indices, segmentation, and interactive filtering, enabling analysis at the level of individual trees. To the best of our knowledge, no prior system supports interactive filtering of vegetation indices aggregated per tree within point cloud visualizations.

2.2 Forestry and Environmental Point Clouds

Outside visualization venues, a large body of work has investigated the use of 3D point clouds for forestry and environmental analysis [15, 19, 32–34, 41, 42, 69]. These studies focus on extracting structural and biophysical information from LiDAR and photogrammetric data, enabling tasks such as tree segmentation, canopy modeling, and biomass estimation. This body of work demonstrates that point clouds support detailed, tree-level analysis and provide valuable insights into forest structure and vegetation dynamics. However, visualization in these approaches is typically limited to presenting results, rather than supporting interactive analytical reasoning. A complementary line of work addresses the technical challenges of handling large-scale environmental point clouds, including web-based delivery, data streaming, and scalable rendering [10, 11, 13, 31, 51, 54]. While these systems enable efficient access to massive datasets, their primary focus lies on performance, infrastructure, and data management, with limited emphasis on user tasks, interaction design, or visual analysis.

Other studies make visualization a more central component of domain-specific forestry applications, such as vegetation monitoring, immersive exploration, and decision support [20, 30, 59, 62, 64–66]. Although these works demonstrate the practical relevance of 3D visualization in forestry, they typically rely on predefined views, static metrics, or tightly coupled processing pipelines, rather than supporting open-ended, exploratory analysis. Recent surveys highlight the growing importance of 3D data in forestry and environmental sciences, particularly for applications such as precision forestry, ecosystem monitoring, and education [36, 42, 73]. At the same time, they emphasize that visualization remains underdeveloped as a research topic, with limited attention to user tasks, interaction techniques, and design principles tailored to forest data. Consequently there remains a gap between forestry-oriented point cloud analysis and visualization research that explicitly supports interactive, tree-level exploration.

2.3 Point Cloud Visualization in Visualization Research

Within visualization venues, research on point clouds has primarily focused on methodological and domain-agnostic challenges [10, 12, 44, 60]. Much of this work addresses foundational problems such as registration, alignment, and feature extraction, emphasizing algorithmic accuracy and robustness [60]. While these contributions are essential for preparing and processing 3D data, visualization is typically positioned as a supporting tool for inspection or validation rather than as a medium for interactive analysis.

A complementary line of work explores abstract representations of point clouds to support general analytical tasks. Techniques based on density estimation, topological summaries, or dimensionality reduction aim to reveal global structures, clusters, or anomalies in high-dimensional data [44]. However, these approaches deliberately abstract away spatial context, object identity, and domain semantics. As a result, they provide limited support for analysis tasks that depend on interpretable entities and real-world meaning.

Overall, visualization research on point clouds has prioritized general-purpose methods and abstract analytical techniques over application-driven design. Such application-driven work has mostly targeted urban environments [38], where structured objects such as buildings provide stable semantic units; natural environments like forests, by contrast, are characterized by irregular, biologically defined structures and ambiguous boundaries. These characteristics are well studied in remote sensing and vegetation rendering [9] but have received comparatively little attention as interactive, task-driven visualization problems in visualization venues. This gap highlights the need for domain-specific, task-driven visualization approaches that support interactive, tree-level analysis within complex ecological datasets.

3 Domain Requirements and Tasks

While forests are widely studied through remote sensing and forestry pipelines, they are underexplored in visualization research in terms of well-defined analytical tasks [4, 38–40]. While prior work has established task frameworks for domains such as urban analytics [38, 72], these do not directly transfer to forest environments. As a result, there is a lack of task-oriented guidance for the design of interactive visualization systems for forest data. To address this gap, we conducted a semi-structured interview with a forestry domain expert, grounded in real-world workflows such as forest inventory, vegetation health assessment, and ecological monitoring. The discussion focused on how forest data is collected, analyzed, and interpreted in practice, as well as the limitations of existing tools. From this process, we derived a set of core analytical tasks **(T1–T3)** that capture recurring needs in forestry analysis and inform the design of our visualization system.

T1 – Interactive multi-scale data exploration from individual trees to areas. Forestry analysis requires interactive exploration across multiple spatial scales, ranging from individual trees to small areas and larger forest regions. This requirement was derived from the domain expert's interest in fitting models to evaluate forest health. Similar to multi-scale exploration tasks described in urban planning and site analysis [38], forestry workflows involve navigating between levels of detail to identify patterns and support reasoning across scales. In practice, forest attributes are measured or estimated at the level of individual trees and aggregated to represent larger areas. Tree-level metrics, such as NDVI derived from reflectance-based multispectral imagery or canopy height, can be combined into higher-level spatial summaries while retaining the ability to drill down to individual trees. Likewise, LiDAR-based sampling and extrapolation strategies for large-area forest characterization [72] rely on comparing localized sample areas with broader regions. Supporting this task is challenging due to the size and density of point cloud data, as well as the lack of explicit hierarchical structures in forest environments. Visualization must therefore enable seamless navigation across scales while preserving spatial context, for example through coordinated views, scale transitions, and mechanisms that link local observations to larger spatial patterns.

T2 – Tree-level querying and comparison of attributes. Forestry analysis requires the ability to compare structural and physiological attributes across individual trees or small groups. These attributes are derived from multiple data sources, including remote sensing imagery, LiDAR point clouds, and in-field measurements. After identifying trees of interest, experts analyze properties such as height and vegetation indices (e.g., NDVI) to assess variation in tree condition and health. This task is conceptually related to unit-level comparison tasks in urban analytics [38], but is more challenging in forest environments due to the lack of regular geometry and the high natural variability between trees. Differences in tree size, shape, and spatial arrangement make direct comparison less straightforward and require careful contextual interpretation. Supporting this task requires visualization techniques that enable interactive querying and comparison of tree attributes while preserving spatial context. Analysts must be able to select individual trees, compare their attributes, and relate these differences to their surrounding environment or expected conditions.

T3 – Area-level comparison of forest attributes. Forestry analysis also requires comparing structural and physiological attributes across larger spatial regions, such as forest stands or selected areas. These comparisons are used to identify patterns in vegetation condition, growth, or environmental variation, based on aggregated attributes derived from tree-level data, including metrics such as canopy height and vegetation indices (e.g., NDVI). This task is conceptually related to area-based comparison tasks in urban analytics [38], but is more challenging in forest environments due to the absence of clearly defined boundaries and the heterogeneous distribution of vegetation. Forest regions are often irregular and continuous, making it difficult to delineate meaningful areas and compare them consistently. Supporting this task requires visualization techniques that enable users to define, select, and compare regions while preserving spatial context. Analysts must be able to aggregate tree-level attributes, explore differences between areas, and relate regional patterns to underlying environmental or structural factors.

4 ForestVis

The ForestVis system (Fig. 2) comprises a data processing pipeline that transforms data from point clouds and multispectral imagery (i.e., data collection) into formats suitable for visualization. The visualization interface loads formats exported from the data pre-processing as data sources and enables interactive web-based analysis.

4.1 Data Collection

The study was conducted at the Attsjö super test site in Southern Sweden, a 2,300-hectare boreal forest research infrastructure intended for multi-scale forest mapping and monitoring. The site integrates Unmanned Aerial Vehicle (UAV)- based remote sensing with extensive field measurements, including hundreds of sample plots and detailed tree-level attributes, providing a robust reference framework for analysis. Data acquisition was carried out in August 2024 using UAV-based multispectral and LiDAR sensors. Multispectral imagery was collected using a DJI Mavic 3 Multispectral platform, capturing green, red, red-edge, and near-infrared bands at a ground sampling distance of approximately 5.5 cm. In parallel, high-density LiDAR data were acquired using a DJI Zenmuse L2 LiDAR sensor mounted on a DJI Matrice 350 RTK platform, producing point clouds with average densities exceeding 200 points/m². The combined acquisition setup enabled detailed capture of both spectral responses and 3D forest structure across 37 subareas covering the study site. The point cloud visualized spans the three forest areas used in the use cases and comprises over 22 billion points in total (≈200 GB of pre-processed COPC/LAZ data), streamed on demand via COPC.

All datasets were collected under consistent flight parameters to ensure uniform quality and comparability. The LiDAR data capture preserved fine-grained canopy geometry and understory complexity, while multispectral imagery enabled the derivation of vegetation indices, including NDVI and related metrics. These data sources provide a rich representation of forest conditions, combining structural and physiological information at the individual-tree level.

4.2 Data Pre-processing

This section describes the data processing pipeline and the methods for formatting the point clouds for data extraction, extracting areas of interest, running segmentation, extracting raster elevation, extracting raster segmentation, calculating canopy height and mean NDVI and formatting the point cloud and multispectral data for the interface.

The first stage of pre-processing is to format the point cloud as Cloud Optimized Point Cloud (COPC), and extract the area of interest for segmentation. The forest point cloud data was provided as folders of LiDAR Aerial Survey (i.e., laz files); each area of interest was provided as a vector polygon boundary. Each folder of the forest

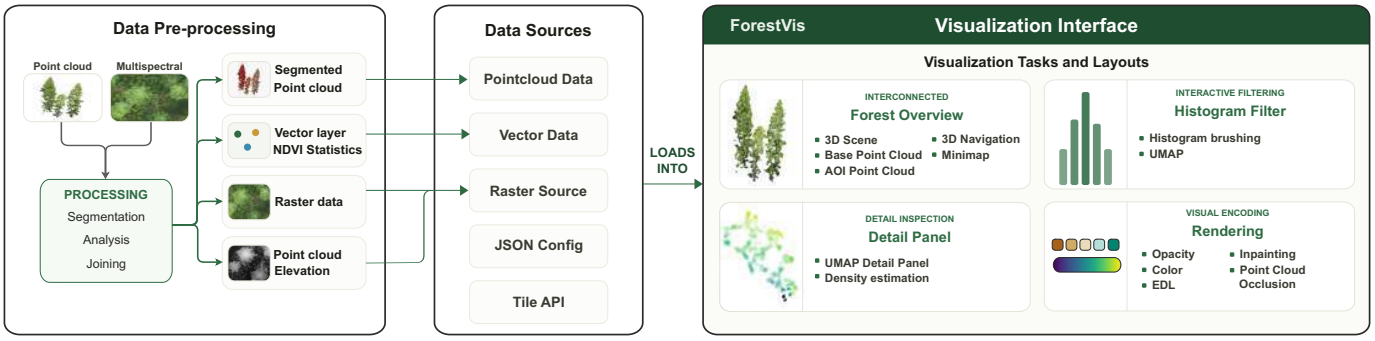


Figure 2: ForestVis pipeline describing the pre-processing back end, data sources, visualization tasks, and visualization layout.

laz data corresponded to an area of the forest, and each folder was converted to COPC format using untwine [22] to make data processing more efficient prior to running the pipeline. In order to extract regions from the COPC point clouds, a virtual point cloud layer was created from the COPC files. The vector polygon outlines were used to extract point clouds by filtering the VPC layer using PDAL/PDAL Wrench [46]. The point clouds for the areas of interest were then input into the data processing pipeline. The multispectral data were provided in GeoTIFF format. Data preprocessing may introduce minor uncertainties related to classification and spatial alignment, which are inherent to LiDAR data processing, and should be considered when interpreting the resulting data. Vegetation indices, NDVI, Leaf Chlorophyll Index (LCI), Green NDVI (GNDVI), Optimized Soil Adjusted Vegetation Index (OSAVI), Normalized Difference Red Edge Index (NDRE), were input into the processing pipeline as raster GeoTIFF ortho-images calculated from multispectral bands (cf. Fig. 2).

Segmentation, Analysis, Joining. The point cloud areas of interest were segmented using the Segment-Any-Tree algorithm [70]. Segmentation is performed as an offline pre-processing step prior to loading the data into ForestVis. Consequently, the computational overhead associated with tree instance segmentation is incurred only once during data preparation and does not affect the interactive performance of the visualization system. ForestVis currently assumes that segmented point clouds are available as input and does not perform segmentation dynamically during analysis sessions. Segmented point clouds contain a custom dimension identifying each tree instance.

After segmenting the point clouds, they were colored using the aerial imagery by projecting the multispectral bands into the point cloud with a colorization filter. The canopy height was calculated from the point cloud and rendered into a GeoTIFF raster image, and an ortho-image of the segmentation of the forest data was rendered into a GeoTIFF raster using the canopy height to filter and remove data points occluded by the canopy. The segmentation raster and the vegetation indices ortho-imagery were used to calculate the summary statistics of the indices for each tree by sampling the ortho-images for each pixel on the segmentation raster. The summary statistics (mean, median, Q25, Q75, standard deviation, maximum, minimum) were calculated. Mean NDVI were spatially sampled from orthomosaics and assigned to corresponding points in the point cloud to assist with visualization. A nominal *nodata* value was chosen for the trees not visible from the ortho-imagery, this nominal value was written into the point cloud, and a VLR (Variable Length Record) record was inserted into the point cloud metadata to document the *nodata* value in JSON format. Centroids were calculated for each tree and stored on a GeoJSON vector layer. The summary statistics from the vegetation indices were written into the vector layer. The centroid stores a bounding box around the tree (X, Y, and Z dimensions), and the tree height.

Running the pre-processing pipeline results in a folder of segmented point cloud, vector, and raster data sources, which are provided to the visualization interface, with JSON configuration for configuring the file paths, and user interface preferences such as color scales. The minimap uses OSM data [45] from a tile API.

The system is motivated by Shneiderman's information-seeking mantra [56] and supports multi-scale drill-down exploratory visual analysis. The user interface features a prominent overview (Fig. 1(c)) with an interconnected minimap for navigation (Fig. 1(b)), and visualization of individual tree statistics.

The **Interconnected Forest Overview** comprises a 3D point cloud model of the forest, with an area-of-interest point cloud composited into the scene. The minimap provides an orthographic overview of the 3D scene. The minimap shows a base layer from a raster tile API (rendering OSM Tiles [45]), vector outlines of the areas of interest, and areas can be toggled. There is a point for each tree in the dataset; the points are colored according to attributes stored in the vector point layer. Points on the minimap can be clicked, and this highlights the corresponding 3D points in the point cloud to perform a single tree selection. Navigation is synchronized between the minimap and the 3D scene: dragging rotates and pans the point cloud, the mouse wheel zooms both views, and a toggle links their zoom levels. The point cloud is hierarchically tiled, so new tiles stream from the server as the view zooms in.

Interactive Filtering. The right panel (Fig. 1 (d)) has three different tabs: filter, details, and UMAP. The filter tab provides an interconnected histogram filter with a bidirectional range selector; a range can be selected, and applied to the minimap and the 3D point cloud area of interest to filter trees by an attribute such as mean NDVI. The UMAP tab shows a UMAP projection of the tree attributes. Bar charts summarizing attributes for each area are visualized above the UMAP projection (Fig. 4 (b)). The UMAP allows users to make box and free-hand polygon selections, pan and zoom on the projected visualization. Use of the UMAP will be described in Section 5.

Detail Inspection. The minimap provides details on demand on hover through the use of tooltips which show a summary of the individual tree statistics on each data point. When the minimap is clicked, further details about each tree are shown on a details tab. The details tab shows density distributions of the NDVI, Red, Green, and Blue channels for a selected tree on line charts. Use of the density distributions will be described in Section 5.

Visual Encoding. The left design panel provides features for manipulating the visual encoding of the data such as point size, opacity, and color scales (Fig. 1 (a)). The area of interest point cloud color scale can be configured to show different dimensions of the point cloud. The size and opacity for the area of interest, and the base layer could be adjusted with sliders. The rendering options for the point cloud include Eye Dome Lighting (EDL) [6], inpainting [47], point cloud occlusion [47], a maximum point budget for rendering, and a subdivision threshold controlling the distance where new point cloud tiles are loaded. By default, EDL uses strength 0.7 and radius 1.5 px (8 directions), inpainting a 7×7 kernel for 2 steps with depth weight 0.5, occlusion a threshold of 0.9, the subdivision (screen-space-error) threshold is 1.0, and the point budget slider is set 10M but disabled by default. The depth weight prevents inpainting from bleeding across depth discontinuities, and the subdivision threshold sets when finer tiles load; the complete parameters and rationale are given in the Supplemental Material (Section 2 – Implementation Details). A selection of different rendering styles is

shown in Fig. 3, these will be discussed in Section 6.1.

4.3 Design Decisions

The following section discusses the design decisions made when selecting visualization and interaction methods for the user interface. We decided to represent the 3D forest as a point cloud. The point cloud allows expert users to assess data quality by inspecting dense and sparse regions. Alternative designs included representing the forest as a mesh surface [65]; however, we chose point clouds over mesh surfaces because they capture the recorded LiDAR return positions and can be post-processed to optionally render the point cloud as a surface. We considered progressive, stochastic sampling techniques [55] for level-of-detail management. Although these techniques improve rendering performance by reducing the number of rendered points, they are orthogonal to mechanisms for out-of-core storage and streaming of large point clouds [54]. Because our system relies on the COPC tile-streaming format for web-based visualization, we adopted a rendering approach that is directly compatible with COPC.

We considered alternative designs for navigating the 3D forest point clouds. Previous 3D forest visualizations implemented navigation through walking/flying [23, 66]. These navigation methods are motivated by gamification [66] and creating stakeholder engagement through providing an interactive experience [23]. A first-person view was considered for navigating the forest point cloud [51], and gamified flying could provide an interface to zoom in to different scales from a first-person perspective [23]. However, in our application, we focus on data analysts as our user group, rather than stakeholder engagement, and our objective was to enable multi-scale analysis. The orthographic perspective would be familiar to analysts who work with ortho-imagery in tools such as QGIS, so we decided to implement a multi-connected-view navigation technique to enable efficient transitions between scales and assist with orthographic comparison, with the 3D View synchronized to the minimap. Both the 3D view and the minimap could be interacted with to navigate.

Dimensionality reduction diagrams were chosen as a visualization technique for multidimensional data for investigating patterns and relationships within the remote sensing and point cloud data. UMAP is widely used for dimensionality reduction and allows users to search for clusters, or analyze data within localized neighborhoods [16], e.g., trees clumped close together. UMAP was applied to NDVI and height data for a nature reserve [74] to relate elevation and tree vitality; it was even applied to radiometry data to classify land use types, including trees [52]. We considered parallel coordinates, but the width of the PCP limits the number of variables that can be displayed, and analyzing non-adjacent attributes may require additional visualization techniques [18]. Thus, we finally chose dimensionality reduction due to the large number of features being input and the difficulty of comparing variables far apart on a PCP.

The histogram filter (Fig. 1 (d)) had a bidirectional range selector integrated, similar to scented widgets [71]. Users could select ranges by clicking and dragging a selection over the histogram or using the range selector handles. Existing point-cloud filtering methods operate on individual points, using either semantic/land-cover categories [11] or LiDAR attribute sliders [54]. Because our analysis is performed at the single-tree level, we filter entire tree instances instead of individual points. We chose a histogram range selector because, unlike plain sliders, it summarizes the distribution of values, helping users pick informed ranges.

The details tab presents per-tree NDVI/RGB distributions as density curves for individual-tree health assessment from ortho-image pixels [17]. We use density curves rather than histograms, which clutter when overlaying multiple bands (RGB), and compute them with FastKDE [21], which gives near-exact estimates at a fraction of the cost of exact KDE, enabling interactive recomputation.

Developed web-tool. The system architecture comprises a Python-based back-end and a statically hosted web-based front-end for interactive visualization. The front-end architecture comprises data sources and visualization components. Data sources include point cloud, vector, and raster data, as well as a map tile layer. The visualization components in the presentation layer, including the sliders, histogram filters, distribution curves, and UMAP visualization, interact with the point cloud visualization, the Three.js scene, and the minimap. Front-end software libraries include: Giro3D for point cloud visualization and layers; OpenLayers for the minimap; Three.js for scene representation; D3.js for 2D visualization; Bootstrap for UX components; GeoTIFF.js for loading raster data; FastKDE for density distribution curves; Chroma.js for color scales. Large point clouds are streamed by Giro3D using the COPC format: the file's clustered octree is fetched on demand through HTTP range requests, with a view-dependent level of detail governed by the point budget and subdivision threshold, LAZ chunks decoded in web workers, and decoded nodes held in an in-memory cache. These and the other standard building blocks of the tool, point rendering (EDL, inpainting), dimensionality reduction (UMAP), and density estimation (FastKDE), are described in full, with parameters and references, in the Supplemental Material (Section 2 – Implementation Details).

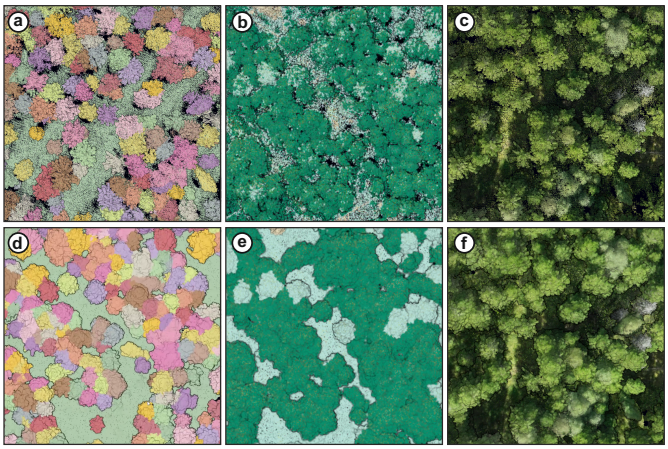


Figure 3: Comparison of 3D forest point cloud rendering styles: (a,b,c) Eye Dome Lighting (EDL) with instance, NDVI, and RGB coloring, respectively; (d,e,f) EDL with inpainting and point cloud occlusion using instance, mean NDVI, and RGB coloring, respectively.

5 Use Cases

We present four use cases in this manuscript. The first one compares NDVI with tree height. The analyst inspects a single area of interest containing some taller and some shorter trees (see Fig. 1). Using the ForestVis interface, they zoom and pan the 3D view to focus on the area of interest. They want to inspect groupings in tree height with NDVI range, so they click the filter panel and use the drop-down menu to select NDVI. The analyst draws a box selection on the histogram to perform the filter, highlighting trees with high NDVI (Fig. 1(d)). The design panel (Fig. 1(a)) is used to configure the point cloud to color by tree height, allowing for comparison of the clumping of the selected trees. They drag the box selection to analyze different NDVI ranges.

The second use case compares the NDVI value and the species of individual trees. The analyst starts with the area of interest centered, then configures the rendering style on the design panel to set a suitable point size, enabling eye dome lighting, disabling inpainting, and point cloud occlusion to show the points as dots without distortion in shape from inpainting (Fig. 3(c)). The analyst zooms into the area of interest to view the trees in great detail, then filters by NDVI to highlight those with high NDVI. They compare the shapes of the tree crowns from an orthographic perspective. Distinguishing the star shapes of the spruce trees in the point cloud.

The third use case explores high-dimensional attribute patterns for forest areas. The analyst configures the UMAP with AvgNDVI (mean), minimum NDVI, maximum NDVI, z-minimum, z-maximum, and height. Taller trees are grouped together at the bottom left of one cluster, so they select trees within the box (Fig. 4(c)), turn off Area 2 on the mini map, and zoom in, confirming the group in Area 0. Inspecting

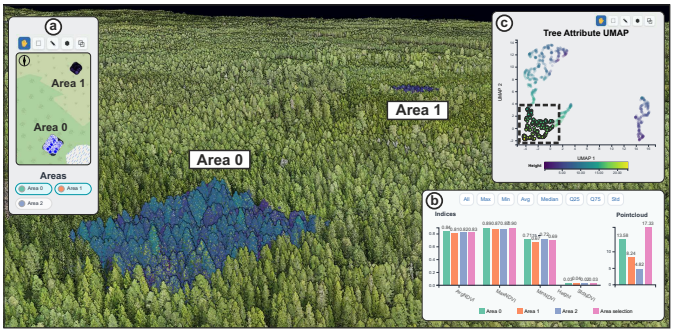


Figure 4: Investigating the location of the tall trees inside a cluster on the UMAP visualization: (a) the minimap, (b) bar charts, and (c) UMAP selection of tall trees with the mean height of the trees.

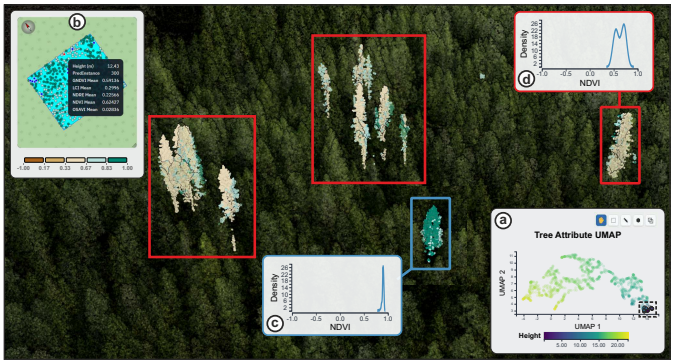


Figure 5: Exploration of unusual NDVI distributions: (a) box selection of an outlier cluster on the height-colored UMAP; (b) minimap with per-tree attribute details on hover; and a comparison of the NDVI distribution of a (c) healthy tree (sharp peak) and a (d) unhealthy tree (more dispersed).

the bar chart shows a higher average height for the selection than the area averages. Furthermore, the analysts inspects the bar chart and notices the attribute distribution for the selected (taller) trees.

The final use case inspects outliers with unusual NDVI distributions. In the UMAP panel, the analyst configures the embedding to use the mean, minimum, and maximum of the NDVI attribute and recomputes it; with the points colored by tree height, a small outlier cluster of low trees stands out at the periphery of the projection, and a box selection highlights it and its local neighborhood (Fig. 5(a)). Hovering a tree on the minimap reveals its per-attribute statistics (Fig. 5(b)), and the linked density curves let the analyst contrast a healthy tree, a narrow NDVI peak near 0.9 (Fig. 5(c)), with an unhealthy one whose NDVI is lower and more dispersed (Fig. 5(d)). Coloring the point cloud by NDVI reveals each tree's structure, helping spot segmentation errors.

6 Qualitative Evaluation

An expert evaluation was performed with four domain experts, recruited through existing contacts with co-authors. The domain experts were recruited for their knowledge of forest data and their prior experience with remote sensing imagery/vegetation indices and point cloud data. The evaluation included a pre-study briefing during which participants were introduced to the study and presented with information and consent forms. The pre-study briefing was followed by a tutorial given by the researcher in which participants were introduced to the tool and required to complete tasks to learn the interface functionality. During the evaluation the participants were encouraged to talk to the researcher, think aloud, and permitted to ask questions about the system. The tutorial was followed by an open ended data exploration/analysis task where participants could explore all features of ForestVis. The task provided involved using the features of the tool to explore relationships in location, point cloud and the vegetation indices. There were three different forest areas inside a forest point cloud to explore. After completing the task participants were presented with post study questions about their experiences with the interface.

6.1 Expert Feedback

In this section, we discuss the feedback provided by experts while using the web-based visualization system. Expert feedback is compared to the visualization tasks.

Expert demographic. All four participants were either studying for, or had a PhD in a relevant field. Experts had performed tasks such as drone operation (2/4), pre-processing ortho-imagery (4/4), Generating a true ortho-photo from photographs (1/4), developing machine learning models based on forest data (3/4), annotation (1/4), terrain generation/photogrammetry (2/4). Experience ranges from 1-5 years with forest data (e.g., ortho-imagery or point clouds), one participant reported 8 years with geospatial data outside forestry. Reported use of software relevant to geospatial data processing included CloudCompare (3/4), TerraDJI (3/4), PDAL (2/4), GDAL (3/4). QGIS (3/4). One participant reported their experience with satellite data.

Interactive multi scale data exploration (T1). Multi-scale data exploration requires navigating to different areas, with features such as zooming and rotation. Navigation was useful for transitioning between detail levels, transitioning between areas, and choosing point cloud viewing angles. When analyzing an overview of the dataset, one participant wanted to locate the ground and height of trees, and find out if trees were hidden by taller trees. One strategy was to rotate the point cloud to inspect the underside.

"I just want to detect the individual trees. So I'm looking for their height and their ground. I want to see small trees because small trees can be under [tall] trees."

Interface navigation features were useful when applying this strategy, since rotation allowed them to inspect looking down at the canopy, and from below the ground. Rotation was used for quality checking by looking for gaps in the point cloud from the LiDAR being obscured by other objects. Participants could zoom out to visualize all three areas at once, zoom in to a single area, or zoom to individual trees. Two participants liked navigating the point cloud, with one finding the interface smooth and responsive when zooming. *"It's fairly smooth in the transition from overview perspective to small, that's pretty nice."* The movement of the point cloud was useful for visualizing the forest point cloud clearly as the perspective changed. *"This movement allows me to visualize the forest but in a very specific way, in a very clear way."* Incorporating all areas within the same point cloud visualization was useful for navigation, since users could easily transition between the areas of interest. One participant liked the ability to move to multiple areas of the point cloud within the same visualization. *"It was good, you can easily see and you can easily go to another area."*

Tree level querying and comparison (T2). Tree-level querying and comparison were supported by methods such as selecting individual trees from the minimap, highlighting trees by attributes, using rendering styles to reveal details in the shape of the trees, and using density curves to analyze per-pixel distributions for NDVI and RGB data. Three participants were observed visualizing the shape of trees in the point cloud. One participant with expertise in tree segmentation was particularly interested in visualizing different tree species in the point cloud and inspecting data about tree health to make a comparison between species and health. This participant turned on RGB coloring to inspect the canopy, and found that when inspecting the shape from an off-nadir angle, it was not particularly easy to distinguish tree species with inpainting, eye dome lighting, and point cloud occlusion enabled. However, with eye dome lighting enabled, inpainting disabled, and point cloud occlusion disabled, they found that it was easy to distinguish spruce from pine trees by looking top down. A different participant with expertise in tree segmentation was observed inspecting the segmentation with eye dome lighting, inpainting, and point cloud occlusion enabled, they were able to detect errors in the segmentation of the point cloud where small trees were categorized as ground (Fig. 3 (d)). A third participant visualized NDVI with eye dome lighting style

in (Fig. 3 (b)), this contrasts to the style for visualizing mean NDVI with EDL, inpainting and point cloud occlusion when inspecting trends across clumps of trees inside an area (Fig. 3 (e)). The application of rendering styles by different participants demonstrates that different expressive rendering styles could be useful at area level or tree level depending on task.

One participant with drone data collection experience suggested that the 3D point cloud interface could assist people familiar with a forest because they would be able to compare their on-site experiences with the model. *"You could probably connect better your experience of the place with what is going on here. So it's less abstracted."*

When reading the distribution curves to inspect NDVI for individual trees, they noted differences in the width of the distributions. *"In the tree itself. It's very narrow. While in the ones that have poor NDVI, probably parts of it has still very good NDVI"*.The same participant then hypothesized visually stressed trees could potentially have bark borer. They suggested that someone in forest management may be interested in knowing this. *"So you're gonna see some trees in this area that probably have some bark borer ... someone in forestry may find this relevant."* The analysis of the tree appearance to hypothesize about health suggests that the visualization could be useful for precision forest management, biodiversity analysis.

One participant with experience developing segmentation models was inspecting the quality of tree segmentation. They looked at the area with the most saplings and noticed that the segmentation model had not segmented all of the trees correctly. They used NDVI Filters, tree instance coloration, and navigation to analyze, in 3D, how the shapes of individual trees were represented. They were able to detect errors in segmentation and related the errors to the data capture method and segmentation model *"It's all about how the forest data is being captured, so I'm still wondering [if] the segmentation is being really aligned with the data, which is interesting as well."* They reflected on the suitability of the segmentation model for forest analysis, indicating that the interactive visualization system could be an important tool to visually assess segmentation models.

Area level comparison of forest attributes (T3). Area level comparison was supported by the interconnected minimap, per-area summary statistics presented a bar chart attached to the UMAP panel (Fig. 4), box selections on the UMAP, and selections on the histogram filter to compare attributes across areas. Layers such as the averages for the vegetation indices and the height layer were useful for analyzing patterns across areas.

One participant liked to compare the minimap NDVI means with the NDVI point cloud to associate the 3D view with their prior knowledge of the forest while inspecting NDVI. *"If you look at the mean, you see more"*. The NDVI averages were reported as easier to work with, as they visually simplified the point cloud for identifying healthy or unhealthy areas, compared with the complexity of LiDAR colored by NDVI. Filtering the visualization was found useful by two participants. One participant used the filter to find relationships between tree height and health. One participant configured the point cloud to color by height, and filtered trees by NDVI values. They discovered that trees with a high NDVI value range were short. Filtering the point cloud for individual trees received positive responses. *"It's [the system is] really powerful to filter these individual trees"*.

Point cloud filtering based on individual trees helped participants understand multispectral data (NDVI and vegetation indices) and relate them to forest structure. *"I like using optical data and also the LiDAR, you merge them together. So when I filter them ... I can see their relationship distributions,"*

One participant with drone data collection experience inspected the forest using the UMAP to filter; they configured UMAP with mean, maximum, and minimum values for the NDVI, then recalculated the UMAP. Comparing the UMAP colors enabled them to identify groupings of trees based on NDVI distributions. Selecting groups of trees that had unusual maximum, minimum, and mean, with the box selection tool, uncovered mis-segmented trees. *"It's difficult to know the data quality of the NDVI from this multispectral because it's still optical data ... that's something that could be explored there."*

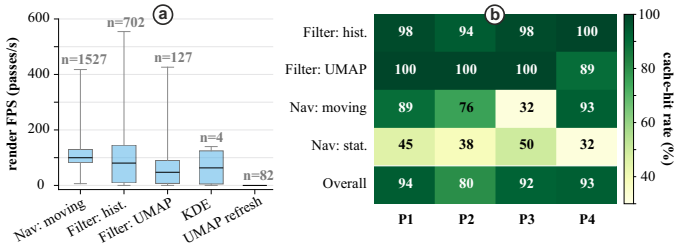


Figure 6: Expert evaluation results for: (a) FPS during each interaction, and (b) network traffic by file resources.

Visualizing tree attributes with the bar charts was useful, Two participants liked to use bar charts to compare area attributes. *"It's more useful for me [UMAP feature bar charts] because I want to see tree attributes, So if I use these things, it will be easy rather than coloring, but filter and height and the other, indexes are more useful."*

"I will create a machine learning model and just feed these indices of this data, so I can pick the relevant features from looking at this data"

One participant suggested features to show the attribute distributions of the tree areas, similar distributions for individual trees. They also suggested filtering to remove the ground when looking at NDVI distributions. Both the filter feature and the UMAP feature could be used as a selection tool. The UMAP was also used to select small localized areas inside the data. One participant found the interface user-friendly for inspecting relationships among attributes. *"It's so user-friendly. You can easily understand what's going on."*

7 Quantitative Evaluation

The quantitative evaluation was conducted with a separate group of four domain experts who had not participated in the qualitative study and were therefore unfamiliar with ForestVis, avoiding learning or carryover effects on the timed, task-based measures. Following the same protocol (briefing, consent, tutorial, think-aloud; see Sec. 6), the session logged network traffic, frame rate, interaction start/end events, and rendered point counts while participants completed timed questionnaire tasks.

7.1 Expert Evaluation

Navigation while moving resulted in the highest median frame rate out of each interaction (Fig. 6(a)); the next highest was histogram filtering, KDE generation, UMAP filtering, and UMAP refresh had the lowest frame rate. The UMAP refresh involved computing UMAP on a server, since the user does not move the interface, and the point cloud does not change, there are few redraw events. The render FPS drops to zero if the user completely stops camera movement and interactions. While the user is idle, the point cloud continues to progressively load and update tiles as they become available, resulting in a network activity spike while stationary and a low cache hit rate (Fig. 6(b)).

Network activity by interaction shows that navigation dominates on-the-wire traffic for every participant (75% for P1, 87–91% for P2–P4; Fig. 7(a)). The resource breakdown mirrors this: the base point cloud makes up 80–92% of bytes, as free camera movement continuously streams new octree nodes on demand (Fig. 7(b)). Much of this arrives while the camera is *stationary* (8.1 of P4's 14.9,GB), as the renderer refines the view with higher-resolution nodes after the camera settles, so P4's lead reflects level-of-detail refinement, not more interaction. Filtering, by contrast, is cheap on the wire: the histogram filter redecodes nodes already resident in the front-end cache (94–99.9% hit rate). So, although it is the main filtering interaction for P1–P3, it adds only a small share of their totals (2.3, 0.8, 0.5,GB). P4 instead relied on UMAP-selection filtering, which for P1–P3 narrowed to already-loaded trees and was served almost entirely from cache (≈100%, ≈0 wire bytes); P4's selections hit the cache less often (89.5%), fetching new nodes and becoming its largest filtering cost (1.4,GB). Overall, on-the-wire cost is driven by point-cloud streaming during navigation, while filtering and KDE computation remain largely client-side.

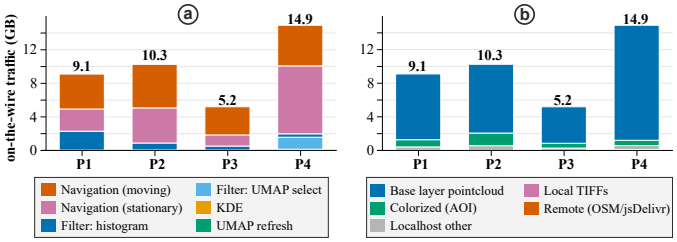


Figure 7: Expert evaluation results for: (a) proportion of logged point-cloud load events serviced without sending network traffic (cache-hit) for each interaction, and (b) network traffic by activity.

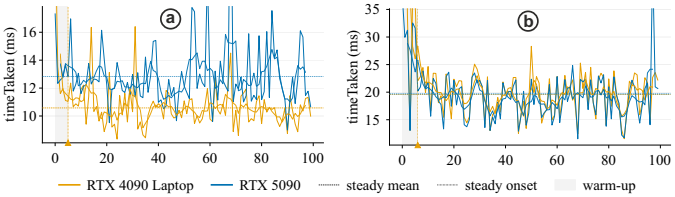


Figure 8: Benchmark results for (a) Histogram filter with point-cloud update latency, classification: Laptop ='warm-up', RTX 5090 = 'flat', (b) UMAP filter with point cloud update latency, classification: Laptop = 'warm-up', RTX 5090 = 'flat'.

Quantitative questions were provided to the participants. When asked to read the number of trees with Mean NDVI ≤ 0.65, all participants were within 1 tree from the correct answer, and two of them gave the correct number. The question required participants to move the point cloud because a small canopy from a large tree was in the area. When asked to report the number of trees ≥ 21, all participants were within 1 tree from the correct answer, and two participants gave the correct result. When asked to read the area with the shortest mean average height from the UMAP histogram, all participants answered correctly. And when asked to find the area with the highest mean average NDVI, all participants also answered correctly. Participants were asked to select a cluster in the UMAP and report the location of the trees. One of the trees was an outlier in a different area from the others. All participants identified the predominant area, and two did not find the outlier. Additional results are available in Supplemental Material.

7.2 Benchmark

Benchmarks were performed using Playwright to automatically run timed trials for User Interface (UI) interactions. UI Features and the point cloud tile loader were instrumented with logging to measure filtering latency until the point cloud update. The PELT algorithm [1] was applied to detect change-points using the ruptures library [63] and a BIC/SIC penalty [53, 75] was applied with a robust noise estimate of the variance (Mean Absolute Deviation), then steady-state detection was performed [1].

Filtering with the histogram performed faster than the UMAP filter; there are more visual elements to update in the UMAP compared to the histogram. This result is consistent with the higher frame-rate of the histogram filter in the expert evaluation (Fig. 7). Generally, in benchmarks, the laptop reached a steady state sooner than the desktop (NVIDIA RTX 5090). For the FastKDE benchmark, the median latency (all iterations) on the laptop was 29.234ms (IQR 26.619 - 31.46), and on the desktop it was 55.824ms (IQR 15.254 - 62.173). The desktop condition was classified as 'no-steady state'. The laptop condition was classified as 'warm-up,' which could account for the performance gap and suggests differing JavaScript runtime behavior in the two trials. Both performance scores indicate low latency. The Mersenne Twister generator (PRNG) was seeded and used to sample selections. Identical seeds were used to minimize variation between hardware from differences in sampled selections. Client memory also scales predictably: the JavaScript heap grows correlated with the displayed point count on both machines. During field testing, the median number of displayed points on the base layer was 10.2 million, indicating that the benchmark reflects realistic conditions. Additional statistics and computer specifications are provided in Supplemental Material.

8 Discussion and Limitations

This section discusses the results of the expert study in comparison to the tasks and related work. The point cloud visualization provided an overview of all of the areas of interest with navigation features for interactive multi-scale data exploration from individual trees to areas (**T1**). The navigation features were reported effective for comparing different areas and inspecting the point cloud data. The ability to zoom out allowed for the comparison of different areas, identifying different areas of interest, and inspecting individual trees. The navigation features enabled users to rotate when analyzing the structure of the point cloud and segmentation, and comparing the segmentation to NDVI (**T2**). Flipping the point cloud was a useful strategy for inspecting holes and small trees. Comparing the tree crowns from an orthographic perspective was useful for distinguishing species, and comparing with the NDVI values. Other strategies for navigation considered included movement by walking/flying [23, 54].

Our study found that simplified rendering of data points with inpainting, point occlusion, and EDL was useful when visualizing groupings of trees (**T3**). In previous research, meshes were applied to visualize forest canopy as a simplification of data [65]. The combination of inpainting/point cloud occlusion was applied to cities and cultural heritage visualization [47]. Inpainting/point cloud occlusion was found beneficial for the analysis of city datasets because the shapes of buildings were simplified when data points behind the building were removed to avoid confusion in the analysis of depth information. Our research found that application of the inpainting/point cloud occlusion may be task-dependent, since, similar to meshes, the inpainting/point cloud occlusion removes gaps in point clouds. Gap removal was beneficial when visualizing individual tree attributes and comparing the color of trees over areas by simplifying the background. However, distorted tree shapes increased the difficulty of distinguishing star shaped crowns when the point cloud was being inpainted after occluded point removal (**T2**). Alternative designs that we considered included the representation of geometry as meshes, the participant responses to navigation where participants were interested in looking at the gaps in the point cloud and trees that were obscured by inspecting the underneath indicate that a point cloud representation assisted with their understanding of data quality, compared to meshes which fill gaps. These observations on rendering-style suitability come from qualitative expert feedback rather than a controlled comparison. A systematic, task-based evaluation that quantifies the trade-off between gap removal and crown-shape fidelity, and compares inpainting/occlusion against alternatives such as meshes, remains an interesting direction for future work.

Our visualization allowed for the density of NDVI/RGB pixels in remote sensing imagery to be visualized using FastKDE. Density estimates calculated for remote sensing pixel data was previously applied to compare vegetated and un-vegetated areas [27], and NDVI pixel density was calculated to compare individual tree crowns for seasonal variations [17]. The inclusion of information outside the tree crown was identified as a potential source of error in the estimation of mortality with vegetation assessments [3]. In our research, we visualized NDVI/RGB pixel density for tree health comparison. Our visualization system enabled users to compare NDVI/RGB distributions with the structure/segmentation of individual trees to assess accuracy of the segmentation (**T2**). We found FastKDE was beneficial for visualizing NDVI/RGB distributions when assessing point cloud quality and comparing individual trees. Participant experience for two participants indicated that hyperspectral and satellite bands were compared with other data, one specified that they analyze areas on the point cloud in comparison to what they could inspect on satellite imagery, so future work could expand ForestVis to other data sources.

Participants could interactively filter the point cloud to find trees of interest using two methods, a histogram with a bidirectional range

selector, and a UMAP projection. UMAP has previously been applied to remote sensing data to investigate trends in elevation and NDVI for a nature reserve [74]; we did not find any research applying UMAP to remote sensing data at the individual tree scale. Previous point cloud visualization systems allow for the filtering of point cloud data by classification [54], or by attributes specific to LiDAR such as the number of returns using sliders or check-boxes. Our system allows users to filter with a histogram bidirectional range selector based on vegetation indices and tree height; or by selection from a UMAP projection. The tree filtering based on the histogram bidirectional range selector assisted with the analysis of tree height over areas (T3). Combining the vegetation indices with the point cloud allowed users to extract insights about the point cloud segmentation by detecting errors in the segmentation when selecting trees in a localized neighborhood of the UMAP. When reading the UMAP one participant questioned how the algorithm functioned, they were unfamiliar with UMAP and did not fully understand it. Participants stated that they did not know how UMAP worked, so future research may be required to make dimensionality reduction more understandable, such as done with t-viSNE [8].

Our system incorporates UMAP to support the exploration of high-dimensional tree attributes through a two-dimensional embedding. While UMAP is widely used for exploratory analysis and clustering of high-dimensional datasets [2, 35], its integration as an interactive selection interface linked to spatial 3D visualizations remains relatively unexplored. However, our research has shown that UMAP can enhance analysts experiences as a selection tool. A limitation of this approach is that dimensionality reduction techniques may distort relationships in the original feature space, particularly at the global scale, meaning that distances in the embedding should be interpreted as approximate rather than exact similarity measures. Furthermore, the structure of the embedding depends strongly on the set of input features used to generate it, which may influence the clusters and patterns users observe. Future work could investigate alternative dimensionality reduction techniques, progressive embedding methods, or tighter coupling between feature-space exploration and spatial context.

Unlike previous visualization systems, ForestVis supports exploratory tree level analytical reasoning, as participants could relate their experience with segmentation, or drone data collection to assess model quality from exploring the data. One participant reflected on the quality of the results being presented by ForestVis, rather than static presentation of results. Another example of interactive analytical reasoning was provided by one participant with experience at drone data collection. They could use their experience at drone collection combined with the data visualization and interactivity to reflect on possible causes for why trees with low NDVI estimates also looked unhealthy to provide a bark borer assessment.

Participants generally had experience with tools such as CloudCompare, or QGIS which was applied to the visualization of point clouds. Our system allows for the streaming of large point cloud data over the web and supports users to apply interactive visualizations techniques to filter, compare trees, and inspect details. User responses indicated that the experience of using interactive navigation was smooth, and that the application of the interface to data analysis tasks was user friendly.

One limitation of our study was the small sample size; since we required domain experts, it was challenging to recruit participants with the necessary level of experience, so we opted for a highly qualitative study design. We acknowledge that this limits the statistical generalizability of our findings. However, in visualization research, especially in early-stage system evaluation, small-N expert studies are a well-established and widely accepted methodology for obtaining in-depth, high-quality feedback from domain specialists [29]. By using open-ended, informal tasks, we were able to identify evidence of highly analytical reasoning and reflection grounded in participants' prior professional experience. The subsequent quantitative evaluation was conducted with a second, independent group of four domain experts (i.e., no overlap with the qualitative study), broadening the evidence base to eight experts in total and reducing familiarity bias, even though the per-study sample remains small.

9 Conclusions and Future Work

We developed a set of forest-analysis tasks (**T1–T3**) in consultation with forestry domain experts, and designed and evaluated ForestVis, a web-based system for interactive tree-level analysis of 3D forest point clouds enriched with vegetation indices, height, and elevation. The expert evaluation showed that the interconnected minimap and 3D view supported multi-scale navigation from individual trees to areas (**T1**); that selectable rendering styles, histogram and UMAP filtering, and per-tree density curves supported tree-level querying and comparison (**T2**); and that per-area bar charts and UMAP/histogram selection supported area-level comparison (**T3**), which experts found suited to different tasks. Combining vegetation indices with the segmented point cloud made segmentation errors visible and let analysts ground their reasoning in prior field experience. To our knowledge, ForestVis is the first system to integrate vegetation indices, density curves, and a UMAP projection for interactive exploratory analysis of remote-sensing and point-cloud data at the individual-tree scale, supporting an open-ended analytical workflow relevant to precision forestry, biodiversity monitoring, and ecological research.

For future work, we plan to enrich our analysis by considering additional attributes of the point cloud, e.g., DBH, basal area, and crown/stem metrics. Furthermore, although ForestVis currently relies on pre-computed tree segmentation generated during an offline preprocessing stage, future work could investigate integrating dynamic or incremental segmentation approaches to support continuously expanding datasets and enable interactive updates as new data become available. Considering usability, some analysts reported difficulties interpreting the UMAP projection. Therefore, future work should also investigate approaches to overcome this issues, for example, by integrating explainable AI methods or by using other projection methods and/or providing better visual tools to interact with these plots (e.g., [8]). In addition, the ability of ForestVis to reveal segmentation errors suggests opportunities for developing a more formal visual model debugging framework, integrating interactive visualization with quantitative quality metrics to support validation, diagnosis, and refinement of automated forestry analysis pipelines.

Acknowledgments

We evaluated the developed visualization software in user studies with adult volunteers who performed usability and analytical tasks while thinking aloud and/or completing questionnaires. We applied standard usability-testing ethics: responses cannot be linked back to participants, although such studies require no university ethics board approval in Sweden; all volunteers were informed in advance of the research's nature, purpose, and design, and of our anonymity protection, were free to ask questions and to withdraw at any time, and signed a consent form. Before starting, the published results do not reveal any personal information, and the anonymity and privacy of all people and organizations involved are preserved.

The authors acknowledge the use of AI tools (Anthropic Claude, OpenAI GPT models, Minimax-m3, Google Gemma and Gemini) for proofreading, editorial suggestions, partial software development, and OpenAI Whisper large-v2 for audio transcription (all transcripts were manually verified). AI tools were not used to generate scientific content, results, data analyses, or conclusions. This work was partially supported through the ELLIIT environment for strategic research in Sweden, InfraVis, the Knowledge Foundation (grant 20200153), the Future Forest Management in Southern Sweden (FRAS II) program, The Bridge collaboration between Södra, IKEA, and Linnaeus University, the OptiForValue project funded by the Circular Bio-based Europe Joint Undertaking (Grant Agreement No. 101157658, funded by the European Union), the Brazilian National Council for Scientific and Technological Development (CNPq, grant 311425/2023-2), and the Swedish Research Council (Vetenskapsrådet, Reg. No. 2025-06537); views and opinions expressed are those of the authors and do not necessarily reflect those of the European Union or CBE JU.

References

[1] E. Barrett, C. F. Bolz-Tereick, R. Killick, S. Mount, and L. Tratt. Virtual machine warmup blows hot and cold. *Proc. ACM Program. Lang.*, 1(OOPSLA), art. no. 52, 27 pp., Oct. 2017. doi: 10.1145/3133876 8

[2] E. Becht, L. McInnes, J. Healy, C.-A. Dutertre, I. W. H. Kwok, L. G. Ng et al. Dimensionality reduction for visualizing single-cell data using umap. *Nature Biotechnology*, 37(1):38–44, Jan 2019. doi: 10.1038/nbt.4314 9

[3] K. O. Bergmüller and M. C. Vanderwel. Predicting tree mortality using spectral indices derived from multispectral uav imagery. *Remote Sensing*, 14(9), 2022. doi: 10.3390/rs14092195 8

[4] F. Biljecki, H. Ledoux, J. Stoter, and J. Zhao. Formalisation of the level of detail in 3d city modelling. *Computers, Environment and Urban Systems*, 48:1–15, 2014. doi: 10.1016/j.compenvurbsys.2014.05.004 3

[5] T. Blaschke, D. Tiede, and M. Heurich. 3d landscape metrics to modelling forest structure and diversity based on laser scanning data. *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, XXXVI:129–132, 01 2004. https://www.isprs.org/proceedings/xxxvi/8-w2/BLASCHKE.pdf. 2

[6] C. Boucheny. *Visualisation scientifique de grands volumes de données : Pour une approche perceptive*. Theses, Université Joseph-Fourier - Grenoble I, Feb. 2009. https://theses.hal.science/tel-00438464. 4

[7] E. Brockerhoff, L. Barbaro, B. Castagneyrol, D. Forrester, B. Gardiner, J. R. González-Olabarria et al. Forest biodiversity, ecosystem functioning and the provision of ecosystem services. *Biodiversity and Conservation*, 26:3005–3035, 2017. doi: 10.1007/s10531-017-1453-2 1

[8] A. Chatzimparmpas, R. M. Martins, and A. Kerren. t-viSNE: Interactive assessment and interpretation of t-SNE projections. *IEEE Transactions on Visualization and Computer Graphics*, 26(8):2696–2714, 2020. doi: 10.1109/TVCG.2020.2986996 9

[9] R. L. Cook, J. Halstead, M. Planck, and D. Ryu. Stochastic simplification of aggregate detail. In *ACM SIGGRAPH 2007 Papers*, SIGGRAPH '07, pp. 79–es. Association for Computing Machinery, New York, NY, USA, 2007. doi: 10.1145/1275808.1276476 2, 3

[10] D. Deibe, M. Amor, and R. Doallo. Big data storage technologies: a case study for web-based lidar visualization. In *2018 IEEE International Conference on Big Data (Big Data)*, pp. 3831–3840, 2018. doi: 10.1109/BigData.2018.8622589 2

[11] S. Discher, R. Richter, and J. Döllner. A scalable webgl-based approach for visualizing massive 3d point clouds using semantics-dependent rendering techniques. In *Proceedings of the 23rd International ACM Conference on 3D Web Technology*, Web3D '18, art. no. 19. Association for Computing Machinery, New York, NY, USA, 2018. doi: 10.1145/3208806.3208816 2, 5

[12] S. Discher, R. Richter, and J. Döllner. A scalable webgl-based approach for visualizing massive 3d point clouds using semantics-dependent rendering techniques. In *Proceedings of the 23rd International ACM Conference on 3D Web Technology*, Web3D '18, art. no. 19. Association for Computing Machinery, New York, NY, USA, 2018. doi: 10.1145/3208806.3208816 2

[13] S. Discher, R. Richter, and J. Döllner. Concepts and techniques for web-based visualization and processing of massive 3d point clouds with semantics. *Graphical Models*, 104:101036, 2019. doi: 10.1016/j.gmod.2019.101036 2

[14] E. Dumic and L. A. da Silva Cruz. Three-dimensional point cloud applications, datasets, and compression methodologies for remote sensing: A meta-survey. *Sensors*, 25(6), 2025. doi: 10.3390/s25061660 2

[15] R. A. Díaz-Varela and E. González-Ferreiro. 3d point clouds in forest remote sensing. *Remote Sensing*, 13(15), 2021. doi: 10.3390/rs13152999 2

[16] M. Espadoto, R. M. Martins, A. Kerren, N. S. T. Hirata, and A. C. Telea. Toward a quantitative survey of dimension reduction techniques. *IEEE Transactions on Visualization and Computer Graphics*, 27(3):2153–2173, 2021. doi: 10.1109/TVCG.2019.2944182 5

[17] D. Fawcett, J. Bennie, and K. Anderson. Monitoring spring phenology of individual tree crowns using drone-acquired ndvi data. *Remote Sensing in Ecology and Conservation*, 7, 11 2020. doi: 10.1002/rse2.184 2, 5, 8

[18] A. Haara, J. Pykäläinen, A. Tolvanen, and M. Kurttila. Use of interactive data visualization in multi-objective forest planning. *Journal of Environmental Management*, 210:71–86, 2018. doi: 10.1016/j.jenvman.2018.01.002 5

[19] H. Hamraz, M. A. Contreras, and J. Zhang. Vertical stratification of forest canopy for segmentation of understory trees within small-footprint airborne lidar point clouds. *ISPRS Journal of Photogrammetry and Remote Sensing*, 130:385–392, 2017. doi: 10.1016/j.isprsjprs.2017.07.001 2

[20] T. Hanousek, J. Novotný, B. Navrátilová, M. Švik, J. Krejza, and R. Janoutová. Complete workflow for detailed 3d forest reconstruction: from terrestrial laser scanning to complex 3d radiative transfer modelling. *in silico Plants*, 7(2), 09 2025. doi: 10.1093/insilicoplants/diaf019 2

[21] J. Heer. Fast & accurate gaussian kernel density estimation. In *Proc. IEEE VIS Short Papers*, 2021. doi: 10.1109/VIS49827.2021.9623323 5

[22] I. Hobu. *Untwine*, 2026. https://github.com/hobuinc/untwine, Accessed: 31-03-2026. 4

[23] S. Holm and J. Schweier. Virtual forests for decision support and stakeholder communication. *Environmental Modelling & Software*, 180:106159, 2024. doi: 10.1016/j.envsoft.2024.106159 2, 5, 8

[24] N. Homainejad, S. Zlatanova, and N. Pfeifer. A voxel-based method for the three-dimensional modelling of heathland from lidar point clouds: First results. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, V-3-2022:697–704, 2022. doi: 10.5194/isprs-annals-V-3-2022-697-2022 2

[25] Q. Hu, B. Yang, S. Khalid, W. Xiao, N. Trigoni, and A. Markham. Towards semantic segmentation of urban-scale 3d point clouds: A dataset, benchmarks and challenges. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4975–4985, 2021. doi: 10.1109/CVPR46437.2021.00494 2

[26] Q. Hu, B. Yang, S. Khalid, W. Xiao, N. Trigoni, and A. Markham. Sensaturban: Learning semantics from urban-scale photogrammetric point clouds. *Int. J. Comput. Vision*, 130(2):316–343, Feb. 2022. doi: 10.1007/s11263-021-01554-9 2

[27] Q. Hu, J. Yang, B. Xu, J. Huang, D.-M. S. Memon, G. Yin et al. Evaluation of global decametric-resolution lai, fapar and fvc estimates derived from sentinel-2 imagery. *Remote Sensing*, 12:912, 03 2020. doi: 10.3390/rs12060912 8

[28] S. Illarionova, D. Shadrin, P. Tregubova, V. Ignatiev, A. Efimov, I. Oseledets et al. A survey of computer vision techniques for forest characterization and carbon monitoring tasks. *Remote Sensing*, 14(22), 2022. doi: 10.3390/rs14225861 2

[29] T. Isenberg, P. Isenberg, J. Chen, M. Sedlmair, and T. Möller. A systematic review on the practice of evaluating visualization. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2818–2827, 2013. doi: 10.1109/TVCG.2013.126 9

[30] M. Kalacska, J. P. Arroyo-Mora, and O. Lucanus. Comparing uas lidar and structure-from-motion photogrammetry for peatland mapping and virtual reality (vr) visualization. *Drones*, 5(2), 2021. doi: 10.3390/drones5020036 2

[31] A. Krooks, J. Kahkonen, L. Lehto, P. Latvala, M. Karjalainen, and E. Honkavaara. Webgl visualisation of 3d environmental models based on finnish open geospatial data sets. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XL-3:163–169, 2014. doi: 10.5194/isprsarchives-XL-3-163-2014 2

[32] M. Kulicki, C. Cabo, T. Trzciński, J. Będkowski, and K. Stereńczak. Artificial intelligence and terrestrial point clouds for forest monitoring. *Current Forestry Reports*, 11(1):5, Dec 2024. doi: 10.1007/s40725-024-00234-4 2

[33] S. K. P. Kushwaha, H. Hristova, C. Fol, T. Hands, A. Murtiyoso, A. Coltekin et al. Trends and developments in the realm of virtual forest visualization. *Current Forestry Reports*, 12(1):3, 2026. doi: 10.1007/s40725-025-00265-5 2

[34] D. Laino, C. Cabo, C. Ordóñez, R. Bolanos, R. Janvier, F. Giulioni et al. Segmentedforests: a labelled dataset of terrestrial lidar point clouds for semantic segmentation of forests. *Forestry: An International Journal of Forest Research*, p. cpaf062, 10 2025. doi: 10.1093/forestry/cpaf062 2

[35] L. McInnes, J. Healy, and J. Melville. UMAP: Uniform Manifold Approximation and Projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018. doi: 10.48550/arXiv.1802.03426 2, 9

[36] M. Meitner, R. Gandy, and S. Sheppard. Reviewing the role of visualization in communicating and understanding forest complexity. In *Ninth International Conference on Information Visualisation (IV'05)*, pp. 121–128, 2005. doi: 10.1109/IV.2005.110 2

[37] M. J. Meitner, S. R. Sheppard, D. Cavens, R. Gandy, P. Picard, H. Harshaw et al. The multiple roles of environmental data visualization in evaluating alternative forest management strategies. *Computers and Electronics in Agriculture*, 49(1):192–205, 2005. Decision Support Systems for Forest Management. doi: 10.1016/j.compag.2005.03.002 2

[38] F. Miranda, T. Ortner, G. Moreira, M. Hosseini, M. Vuckovic, F. Biljecki et al. The state of the art in visual analytics for 3d urban data. *Computer*

Graphics Forum, 43(3):e15112, 2024. doi: 10.1111/cgf.15112 3

[39] G. Moreira, M. Hosseini, M. N. Alam Nipu, M. Lage, N. Ferreira, and F. Miranda. The urban toolkit: A grammar-based framework for urban visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):1402–1412, 2024. doi: 10.1109/TVCG.2023.3326598 3

[40] R. Mota, N. Ferreira, J. D. Silva, M. Horga, M. Lage, L. Ceferino et al. A comparison of spatiotemporal visualizations for 3d urban analytics. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):1277–1287, 2023. doi: 10.1109/TVCG.2022.3209474 3

[41] A. Murtiyoso, C. Cabo, A. Singh, D. P. Obaya, W. Cherlet, J. Stoddart et al. A review of software solutions to process ground-based point clouds in forest applications. *Current Forestry Reports*, 10(6):401–419, Dec 2024. doi: 10.1007/s40725-024-00228-2 2

[42] A. Murtiyoso, S. Holm, H. Riihimäki, A. Krucher, H. Griess, V. C. Griess et al. Virtual forests: a review on emerging questions in the use and application of 3d data in forestry. *International Journal of Forest Engineering*, 35(1):29–42, 2024. doi: 10.1080/14942119.2023.2217065 2

[43] M. Münzinger, N. Prechtel, and M. Behnisch. Mapping the urban forest in detail: From lidar point clouds to 3d tree models. *Urban Forestry & Urban Greening*, 74:127637, 2022. doi: 10.1016/j.ufug.2022.127637 2

[44] P. Oesterling, C. Heine, H. Janicke, G. Scheuermann, and G. Heyer. Visualization of high-dimensional point clouds using their density distribution's topology. *IEEE Transactions on Visualization and Computer Graphics*, 17(11):1547–1559, 2011. doi: 10.1109/TVCG.2011.27 2, 3

[45] OpenStreetMap. Open street map. https://www.openstreetmap.org/, 2026. Accessed: 20-03-2026. 4

[46] PDAL contributors. *PDAL: The Point Data Abstraction Library*, 2026. doi: 10.5281/zenodo.10884408 4

[47] R. Pintus, E. Gobbetti, and M. Agus. Real-time rendering of massive unstructured raw point clouds using screen-space operators. In *Proceedings of the 12th International Conference on Virtual Reality, Archaeology and Cultural Heritage*, VAST'11, pp. 105–112. Eurographics Association, Goslar, DEU, 2011. doi: 10.2312/VAST/VAST11/105-112 4, 8

[48] A. U. Rahman, E. Heinaro, M. Ahishali, and S. Junttila. Dual-task learning for dead tree detection and segmentation with hybrid self-attention u-nets in aerial imagery. *International Journal of Applied Earth Observation and Geoinformation*, 144:104851, 2025. doi: 10.1016/j.jag.2025.104851 2

[49] R. Richter, M. Behrens, and J. Döllner. Object class segmentation of massive 3d point clouds of urban areas using point cloud topology. *International Journal of Remote Sensing*, 34(23):8408–8424, 2013. doi: 10.1080/01431161.2013.838710 2

[50] M. S. R. Saimun and M. M. Rahman. A comprehensive review of tree cover mapping using satellite sensor data. *Discover Geoscience*, 3(1):90, Aug 2025. doi: 10.1007/s44288-025-00201-x 2

[51] A. Salcedo-Navarro, G. Montalban-Faet, J. Segura-Garcia, and M. Garcia-Pineda. A scalable web system for multi-index 3d point cloud visualization and real-time sensor monitoring in precision agriculture. *SoftwareX*, 32:102443, 2025. doi: 10.1016/j.softx.2025.102443 2, 5

[52] S. Schmitz, U. Weidner, H. Hammer, and A. Thiele. Evaluating uniform manifold approximation and projection for dimension reduction and visualization of polinsar features. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, V-1-2021:39–46, 2021. doi: 10.5194/isprs-annals-V-1-2021-39-2021 2, 5

[53] G. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464, 1978. https://www.jstor.org/stable/2958889. 8

[54] M. Schütz. *Potree: Rendering Large Point Clouds in Web Browsers*. PhD thesis, 09 2016. doi: 10.34726/hss.2015.27710 2, 5, 8, 9

[55] M. Schütz, G. Mandlburger, J. Otepka, and M. Wimmer. Progressive real-time rendering of one billion points without hierarchical acceleration structures. *Computer Graphics Forum*, 39(2):51–64, 2020. doi: 10.1111/cgf.13911 5

[56] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings of the 1996 IEEE Symposium on Visual Languages*, VL '96, p. 336. IEEE Computer Society, USA, 1996. doi: 10.1109/VL.1996.545307 4

[57] U. Stilla and Y. Xu. Change detection of urban objects using 3d point clouds: A review. *ISPRS Journal of Photogrammetry and Remote Sensing*, 197:228–255, 2023. doi: 10.1016/j.isprsjprs.2023.01.010 2

[58] SVS. Stand visualization system. http://forsys.cfr.washington.edu/winsvslatest.html, 2026. Accessed: 24-03-2026. 2

[59] N. Takhtkeshha, L. Bocaux, L. Ruoppa, F. Remondino, G. Mandlburger, A. Kukko et al. 3d forest semantic segmentation using multispectral lidar and 3d deep learning. *PFG – Journal of Photogrammetry, Remote Sensing and Geoinformation Science*, Dec 2025. doi: 10.1007/s41064-025-00369-4 2

[60] W. Tao, R. Wang, X. Hua, J. Liu, X. Chen, Y. Zang et al. Handcrafted local feature descriptor-based point cloud registration and its applications: A review. *IEEE Transactions on Visualization and Computer Graphics*, 31(10):9424–9443, 2025. doi: 10.1109/TVCG.2025.3569894 2

[61] N. Thapa, L. L. Narine, and A. E. Wilson. Forest aboveground biomass estimation using airborne lidar: A systematic review and meta-analysis. *Journal of Forestry*, 123(4):389–412, 2025. doi: 10.1007/s44392-025-00029-w 1

[62] J. Trochta, M. Krůček, T. Vrška, and K. Král. 3d forest: An application for descriptions of three-dimensional forest structures using terrestrial lidar. *PLOS ONE*, 12(5):1–17, 05 2017. doi: 10.1371/journal.pone.0176871 2

[63] C. Truong, L. Oudre, and N. Vayatis. Selective review of offline change point detection methods. *Signal Processing*, 167:107299, 2020. doi: 10.1016/j.sigpro.2019.107299 8

[64] N. Tsoulias, K. K. Saha, and M. Zude-Sasse. In-situ fruit analysis by means of lidar 3d point cloud of normalized difference vegetation index (ndvi). *Computers and Electronics in Agriculture*, 205:107611, 2023. doi: 10.1016/j.compag.2022.107611 2

[65] J. Vauhkonen and R. Ruotsalainen. Reconstructing forest canopy from the 3d triangulations of airborne laser scanning point data for the visualization and planning of forested landscapes. *Annals of Forest Science*, 74(1):9, Feb 2017. doi: 10.1007/s13595-016-0598-6 2, 5, 8

[66] H. Wang, X. Chen, N. Polys, and P. Sforza. A web3d forest geovisualization and user interface evaluation. In *Proceedings of the 22nd International Conference on 3D Web Technology*, Web3D '17, art. no. 9, 9 pp. Association for Computing Machinery, New York, NY, USA, 2017. doi: 10.1145/3055624.3075956 2, 5

[67] R. Wang, J. Peethambaran, and D. Chen. Lidar point clouds to 3-d urban models: a review. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 11(2):606–627, 2018. doi: 10.1109/JSTARS.2017.2781132 2

[68] J. Watson, T. Evans, O. Venter, B. Williams, A. Tulloch, C. L. Stewart et al. The exceptional value of intact forest ecosystems. *Nature Ecology & Evolution*, 2:599 – 610, 2018. doi: 10.1038/s41559-018-0490-x 1

[69] H. Weiser, J. Schäfer, L. Winiwarter, N. Krašovec, F. E. Fassnacht, and B. Höfle. Individual tree point clouds and tree measurements from multi-platform laser scanning in german forests. *Earth System Science Data*, 14(7):2989–3012, 2022. doi: 10.5194/essd-14-2989-2022 2

[70] M. Wielgosz, S. Puliti, B. Xiang, K. Schindler, and R. Astrup. Segmentanytree: A sensor and platform agnostic deep learning model for tree segmentation using laser scanning data. *Remote Sensing of Environment*, 313:114367, 2024. doi: 10.1016/j.rse.2024.114367 4

[71] W. Willett, J. Heer, and M. Agrawala. Scented widgets: Improving navigation cues with embedded visualizations. *IEEE Trans. Visualization & Comp. Graphics (Proc. InfoVis)*, 13:1129–1136, 2007. doi: 10.1109/TVCG.2007.70589 2, 5

[72] M. A. Wulder, J. C. White, R. F. Nelson, E. Næsset, H. O. Ørka, N. C. Coops et al. Lidar sampling for large-area forest characterization: A review. *Remote Sensing of Environment*, 121:196–209, 2012. doi: 10.1016/j.rse.2012.02.001 3

[73] B. Xiang, M. Wielgosz, T. Kontogianni, T. Peters, S. Puliti, R. Astrup et al. Automated forest inventory: Analysis of high-density airborne lidar point clouds with 3d deep learning. *Remote Sensing of Environment*, 305:114078, 2024. doi: 10.1016/j.rse.2024.114078 2

[74] B. Xu, J. Li, X. Pei, L. Bian, T. Zhang, G. Yi et al. Dominance of topography on vegetation dynamics in the mt. qomolangma national nature reserve: A umap and pls-sem analysis. *Forests*, 14(7), 2023. doi: 10.3390/f14071415 2, 5, 9

[75] Y.-C. Yao. Estimating the number of change-points via schwarz' criterion. *Statistics & Probability Letters*, 6(3):181–189, 1988. doi: 10.1016/0167-7152(88)90118-6 8