

Through the Perspective of LiDAR: A Feature-Enriched and Uncertainty-Aware Annotation Pipeline for Terrestrial Point Cloud Segmentation

Fei Zhang^{a,*}, Rob Chancia^a, Josie Clapp^a, Amirhossein Hassanzadeh^a, Dimah Dera^a, Richard MacKenzie^b, Jan van Aardt^a

^a*Chester F. Carlson Center for Imaging Science, Rochester Institute of Technology, Rochester, NY, USA*

^b*U.S. Forest Service, USA*

Abstract

Accurate semantic segmentation of terrestrial laser scanning (TLS) point clouds is limited by costly manual annotation. We propose a semi-automated, uncertainty-aware pipeline that integrates spherical projection, feature enrichment, ensemble learning, and targeted annotation to reduce labeling effort, while sustaining high accuracy. Our approach projects 3D points to a 2D spherical grid, enriches pixels with multi-source features, and trains an ensemble of segmentation networks to produce pseudo-labels and uncertainty maps, the latter guiding annotation of ambiguous regions. The 2D outputs are back-projected to 3D, yielding densely annotated point clouds supported by a three-tier visualization suite (2D feature maps, 3D colorized point clouds, and compact virtual spheres) for rapid triage and reviewer guidance. Using this pipeline, we build *Mangrove3D*, a semantic segmentation TLS dataset for mangrove forests. We further evaluate data efficiency and feature importance to address two key questions: (1) how much annotated data are needed and (2) which features matter most. Results show that performance saturates after ~ 12 annotated scans, geometric features contribute the most, and compact nine-channel stacks capture nearly all discriminative power, with the mean Intersection over Union (mIoU) plateauing at around 0.76. Finally, we confirm the generalization of our feature-enrichment strategy through cross-dataset tests on *ForestSemantic* and *Semantic3D*. Our contributions include: (i) a robust, uncertainty-aware TLS annotation pipeline with visualization tools; (ii) the *Mangrove3D* dataset; and (iii) empirical guidance on data efficiency and feature importance, thus enabling scalable, high-quality segmentation of TLS point clouds for ecological monitoring and beyond. The dataset and processing scripts are publicly available at <https://fz-rit.github.io/through-the-lidars-eye/>.

Keywords: Terrestrial laser scanning (TLS), Semantic Segmentation, Uncertainty analysis, Spherical Projection, LiDAR Point Cloud, Forest Ecology

*Corresponding author.

Email address: fzhcis@rit.edu (Fei Zhang)

1. Introduction

LiDAR (light detection and ranging) systems are deployed across platforms ranging from satellites to unmanned aerial systems (UAS), mobile units, and terrestrial scanners. Among these, terrestrial laser scanning (TLS) provides exceptionally dense 3D point clouds (mm–cm scale), enabling detailed structural analysis in ecological environments [1, 2]. In forest ecosystems, TLS supports the retrieval of tree metrics, biomass, and habitat features [3, 4], and these applications fundamentally rely on accurate semantic segmentation to distinguish ground, stems, branches, foliage, and roots within complex scenes.

While deep learning architectures such as PointNet++ and KPConv have surpassed traditional classifiers (e.g., Random Forests, SVMs) in TLS segmentation [5, 6], their broader utility remains limited by two persistent bottlenecks. First, high-quality annotated datasets are scarce: manual labeling of full-resolution TLS scans is prohibitively labor-intensive, and ecological scenes are especially problematic due to severe occlusion, irregular geometry, and intertwined tree structures [7, 8]. Second, the few open-source datasets that do exist are biased toward urban or indoor environments, where objects are simpler and more regularly shaped [9, 10, 11, 12]. Consequently, current segmentation models often fail to generalize to ecologically complex scenes, slowing the adoption of automated analysis, despite urgent demand in forestry and environmental monitoring.

To overcome these obstacles, we introduce a human-in-the-loop annotation pipeline that reduces labeling burden, while improving segmentation fidelity in challenging forest environments. Our pipeline (i) leverages spherical projection to transform irregular 3D data into structured 2D maps, enabling efficient annotation and feature extraction in a lower-dimensional space; (ii) integrates feature-enriched segmentation to better capture radiometric, geometric, and statistical information; and (iii) incorporates uncertainty analysis to target regions of low model confidence for human correction. Together, these components accelerate annotation, enhance model robustness, and bridge the gap between small, high-effort ecological datasets and scalable, reproducible benchmarks. To this end, we release *Mangrove3D*—to our knowledge the first TLS benchmark explicitly designed for structurally complex mangrove forests—and demonstrate the pipeline’s cross-domain effectiveness through evaluations on *Mangrove3D*, *ForestSemantic*[13, 14], and *Semantic3D*[9].

In summary, this paper aims to: (a) alleviate the annotation bottleneck through a semi-automated pipeline, (b) introduce the TLS dataset tailored to mangrove forests, and (c) evaluate the effectiveness and cross-domain robustness of feature-enriched segmentation.

1.1. Related Work

1.1.1. Manual Annotation Software and Methodologies

Semantic annotation of TLS point clouds remains a major bottleneck in ecological research. Unlike urban or autonomous driving scenes, TLS data contain dense occlusions, irregular branching, and intertwined roots, making manual labeling slow and error-prone. Existing 3D annotation tools (Table 1)—such as *CloudCompare*, *3D Slicer*, *MathWorks LiDAR Labeler*, and *Segments.ai*—offer functions like region clipping and polygonal selection, with some integrating deep neural networks (DNNs) for pre-annotation. However,

Table 1: Comparison of commonly used 3D point cloud annotation software and their key features. Clip & segment: region-based clipping and segmentation; Paintbrush: brush-like manual labeling; Multi-view: annotation from multiple perspectives; RGB cross-ref.: annotation aided by associated RGB/multispectral images; DL pre-annot.: integrated deep learning pre-annotation; Typical applications: primary research or industry usage areas.

Software	Open	Clip &	Paint-	Multi-	RGB	DL Pre-	Typical
Name	Source	Segment	brush	view	Cross-ref.	annot.	Applications
CloudCompare	Yes	Yes	No	No	No	No	General research, surveying, geology
3D Slicer	Yes	Yes	Yes	Yes	Yes	No	Medical imaging, biomedical applications
LabelCloud	Yes	Yes	No	No	Yes	No	Autonomous driving, urban scene annotation
Scalabel	Yes	Yes	No	Yes	Yes	No	Autonomous driving, robotics, object detection
MathWorks LiDAR Labeler	No	Yes	No	Yes	Yes	Yes	Autonomous vehicles, robotics, industry
Segments.ai	No	Yes	Yes	Yes	Yes	Yes	Autonomous driving, urban analytics
Autodesk ReCap[15]	No	Yes	No	Yes	Yes	No	Engineering, architecture, construction
Trimble RealWorks[16]	No	Yes	No	Yes	Yes	No	Infrastructure, surveying, civil engineering
Semantic Segmentation Editor[17]	Yes	Yes	Yes	Yes	Yes	No	Autonomous driving
LiDAR360[17]	No	Yes	Yes	Yes	Yes	Yes	Forestry, surveying & mapping, general usage

most were developed for robotics or urban datasets and remain labor-intensive, requiring annotators to switch viewpoints, navigate dense 3D scenes, and often cross-reference RGB imagery that is rarely available in TLS.

To overcome these limitations, we propose a new annotation pipeline emphasizing efficiency, uncertainty-awareness, and adaptability for complex TLS datasets.

1.1.2. Spherical (Equirectangular) Projection for TLS Point Clouds

The concept of spherical projection originates from the equirectangular (plate carrée) projection, which linearly maps longitude and latitude to planar x - y coordinates. First used in cartography as early as AD 400 [18], this simple yet powerful mapping preserves angular relationships along parallels and meridians. Centuries later, the same principle finds renewed application in LiDAR point-cloud processing, where unstructured 3D measurements are transformed into structured 2D “range images.” Each point’s azimuth and elevation (or zenith) are linearly mapped to pixel indices, creating a dense angular grid commonly referred to as a spherical projection [19]. In essence, this projection provides a way of viewing the world *through the perspective of LiDAR*—capturing how the sensor perceives its surroundings across its field of view. Scalar attributes such as range, intensity, and return count are then stored as per-pixel channels, providing a compact 2D representation that preserves the geometric structure of the original 3D scene.

Early works, such as Barnea et al. [20, 21] and Mahmoudabadi et al. [22], leveraged spherical intensity and range maps in combination with classical clustering methods such as mean-shift and graph-based segmentation. These approaches require manual parameter tuning, process each feature channel independently, and are constrained by the limited capacity of traditional machine learning algorithms to integrate heterogeneous

features. More recently, DNN-based models such as RangeNet++[23], SalsaNext[24], and SqueezeSeg[25] demonstrated the effectiveness of spherical projection for semantic segmentation, particularly in autonomous driving datasets. However, these approaches were primarily designed for automotive LiDAR, which typically has a narrow vertical field of view and limited feature inputs (often a single channel such as range or intensity), making them strongly dependent on synchronized multi-sensor data.

In this study, we adapt spherical projection and DNN to TLS, which is widely used in ecological, indoor, and construction environments where camera imagery is often unavailable or unreliable. We extend the traditional single-band representation to multi-channel feature stacks that generalize pseudo-color encodings on spherical maps, serving as an intuitive aid for both human annotators and DNN-based segmentation models.

1.1.3. Active Learning and Self-Training for Efficient Annotation

To further alleviate the burden of pixel-level semantic labeling on the spherical maps, we integrate two human-efficient annotation paradigms: active learning and self-training.

Active learning (AL) iteratively queries the annotator for the most informative samples, often those with the highest uncertainty, approximated via Monte-Carlo dropout [26] or ensembles [27]. While well-studied for image classification, dense segmentation is more difficult, since annotation units (pixel, tile, or superpoint) must balance annotation cost against information gain. Recent studies [28, 29, 30] adapt AL to 2D and 3D domains, but mainly on RGB-rich indoor datasets, and still relies on slow, expertise-intensive 3D point editing.

Self-training is a semi-supervised strategy where a model trained on a seed set reuses its own predictions to expand the labeled pool. Typically, high-confidence predictions are promoted to pseudo-labels and reintroduced for training, sometimes with consistency regularization or noise injection to mitigate overfitting [31, 32]. This approach has shown success in natural images and range-image LiDAR [33], but faces well-known limitations, e.g., pseudo-label errors easily accumulate, confidence does not always imply correctness, and most pipelines omit human oversight, allowing mistakes to cascade unchecked [34, 35]. These challenges are amplified on spherical projection maps of TLS scans, where clutter, occlusion, and class imbalance produce misleading confidence estimates.

We therefore propose an efficient, semi-automatic annotation framework that integrates active learning and self-training within the spherical projection domain to address these challenges. We use an ensemble of UNet++ [36], DeepLabV3+ [37], and Segformer [38] to quantify pixel-wise epistemic uncertainty, directing annotators to the most informative regions (AL), while high-confidence pixels are automatically promoted to pseudo-labels (self-training). This integrated, uncertainty-aware pipeline reduces manual effort while still ensuring human verification, and—to our knowledge—represents the first scalable framework demonstrated on full-resolution TLS spherical projection maps.

1.1.4. Feature fusion in semantic segmentation of point clouds

A closely related approach is PointPainting and its derivatives [39, 40], which project point clouds onto segmented RGB images, attach semantic scores to points, and feed the enriched clouds into downstream networks. While effective for autonomous driving, these approaches depend on accurately registered multi-sensor data. In contrast, TLS generally operates without co-registered imagery, meaning that spectral cues must be inferred from intrinsic LiDAR signals such as intensity, range, and derived geometric attributes. Our framework is therefore designed to operate solely on LiDAR-derived features, eliminating dependence on external imagery while still capturing complementary structural and radiometric information.

Within LiDAR-only pipelines, feature fusion exploits complementary cues—geometric structure, radiometric response, and positional context—to enhance segmentation. Prior studies have shown that integrating geometric descriptors with intensity and elevation improves class separability [41, 42], and benchmarks such as Semantic3D confirm gains from combining coordinates, intensity, and shape features [9]. Accordingly, both projection-based and fully 3D networks (e.g., SqueezeSeg, PointNet++, RandLA-Net) often employ stacked multi-cue inputs [5, 25, 43]. Yet, despite broad agreement that multi-feature fusion helps, there is still no consensus on an optimal LiDAR-only feature set; redundant or correlated inputs can inflate dimensionality without consistent performance gains [44].

To address this gap, we systematically expand and evaluate feature channels on TLS spherical projection maps—starting from single feature families and progressively integrating others—to quantify their complementarity and identify combinations that yield the greatest marginal improvements. This progressive analysis results in a compact yet robust feature stack that reduces redundancy, enhances segmentation efficiency, and provides practical guidance for TLS applications, where feature selection has traditionally been ad hoc.

2. Methods and Materials

In this section we introduce a new TLS dataset curated with our pipeline, detail the pipeline design, and describe the qualitative and quantitative approaches used to assess feature enrichment and data efficiency.

2.1. A New TLS Dataset for Mangrove Forest

We introduce the *Mangrove3D* dataset. The data were collected in spring 2024 on Babeldaob Island, Palau ($7^{\circ} 31' 49''\text{N}$, $134^{\circ} 33' 53''\text{E}$), in coastal areas of *Rhizophora* mangrove characterized by dense prop-root networks and multilayered canopies.

We use the Canopy Biomass LiDAR (Version 2.0; CBL), a TLS built around a SICK LMS-151 LiDAR unit (SICK AG, Waldkirch, Germany) [45], which employs a 905 nm wavelength laser and has an effective measurement range of 0.5–50 m at 90% reflectivity. The LiDAR is attached to a rotation stage and the CBL mounted on a modified surface elevation table (SET) arm—a portable leveling device attached to a benchmark pipe, as shown in Fig. 1. For each scan, we invert the CBL from the standard tripod orientation to obtain the maximum potential ground points surrounding the SET, orienting the 90° uncovered cone




Figure 1: Map and field setup for TLS scans in Palau mangrove forests

above the scanner position, thus enabling $360^\circ \times 270^\circ$ scans while minimizing occlusion from below. Each CBL scan is acquired at an angular resolution of 0.25° and completed in 33s.

We collect eight scans at each of the benchmark sites, starting with the first scan position with the SET-arm oriented directly north of the SET and the preceding scans at successive 45° intervals rotating in a clockwise direction. Scans that suffer from obvious obstruction or operator error are eliminated to ensure data quality.

The dataset comprises 39 TLS scans from seven mangrove plots, totaling 31.3 million points. Individual scans range from 0.51M to 0.94M points (averagely 0.80M per scan), depending on site conditions and preprocessing. Each point is assigned to one of five classes—*Ground & Water*, *Stem*, *Canopy*, *Root*, and *Object*—with a sixth label, *Void*, marking empty pixels in the spherical projection image cube. For segmentation tasks, 30 scans from plots #1–5 are used for training and validation, while 9 scans from plots #6–7 are held out as a fixed test set. In our main experiment, we adopt a train/validation/test split of 27/3/9, yielding a validation-to-train ratio of approximately 0.1. The class distributions across these splits are shown in Fig. 2. A brief annotation guideline is provided in the Appendix A.

The *Mangrove3D* dataset is, to our knowledge, the first TLS benchmark for semantic segmentation




Figure 2: Class counts for the *Mangrove3D* dataset.

in mangrove forests, providing centimeter-scale 3D geometry of tidally influenced ecosystems with dense root and canopy structures. The resulting significant occlusion poses a challenging benchmark for semantic segmentation, a key prerequisite for downstream applications such as biomass estimation and blue-carbon assessments.

2.2. Semi-automated Annotation Pipeline

As illustrated in Fig. 3, we design a three-stage pipeline to facilitate efficient and consistent annotation of TLS point clouds, detailed in the following subsections.

2.2.1. Stage 1: Spherical Projection & Visualization

We unwrap the TLS point cloud into a 2D spherical domain parameterized by zenith (θ) and azimuth (ϕ), yielding multi-channel spherical projection maps where each channel encodes aforementioned various features. Fig. 4 illustrates the process of spherical projection on a CBL LiDAR scan, point cloud and feature map were colored from raw intensity values. For a 3D point (x, y, z) ,

$$\theta = \arccos\left(\frac{z}{\sqrt{x^2 + y^2 + z^2}}\right), \quad \phi = \text{mod}(\arctan 2(y, x), 2\pi),$$

which are mapped to pixel coordinates via the equirectangular projection

$$i = \left\lfloor \frac{\theta - \theta_{\min}}{\Delta\theta} \right\rfloor, \quad j = \left\lfloor \frac{\phi - \phi_{\min}}{\Delta\phi} \right\rfloor.$$

For each CBL LiDAR scan, with an angular step of 0.25° and a vertical field-of-view (FOV) = 135° and a horizontal FoV = 360° , we obtain a 540×1440 spherical grid, providing a near one-to-one mapping between beam directions and pixels, as validated in Fig. 5.

This projection provides a stable canvas for stacking per-pixel attributes and descriptors. We organize the features into three groups: 1) basic properties: radiometric intensity, range, and inverted height; 2) geometric




Figure 3: Three-stage workflow for annotating terrestrial-LiDAR scans. **Stage 1:** Spherical projection converts raw TLS points into two-dimensional feature maps and pseudo-RGB images. **Stage 2:** An iterative loop combines active learning and self-training: an ensemble segmentation model is repeatedly refined using uncertainty-guided queries and high-confidence pseudo-labels. **Stage 3:** The resulting 2-D segmentation masks are back-projected, followed by label refinement in 3D space and then reproject back to 2D, to yield a fully annotated point cloud and refined 2D segmentation mask.




Figure 4: Spherical projection workflow illustrated with a scan of the CBL LiDAR. (a) Original 3D point cloud visualized by intensity with plasma color scale. (b) Geometric illustration of the spherical projection. (c) Spherical projection map of raw intensity values.

properties: normals, curvature, anisotropy, and planarity; and 3) statistical properties: low-dimensional features obtained via PCA, with MNF and ICA considered as alternatives. Local 3D structure is quantified using eigenvalue-based descriptors derived from the covariance of neighboring points. Specifically, curvature, anisotropy, and planarity are computed from the ordered eigenvalues, providing clear visual signatures and demonstrating segmentation utility.

Optimizations. Two challenges are scale and density variation. Nearly one million points per scan can lead to excessive memory requirements, and scanner geometry leads to uneven point densities. We partition clouds




Figure 5: Point-density validation of the spherical projection map. When the map resolution matches the TLS angular resolution, most pixels contain exactly one point; a value of 0 indicates no return from that angle, 2 indicates a dual-return measurement, and values above 2 occur rarely, typically due to systematic noise.

into azimuth–elevation tiles processed in batches, padded with k_b neighbors to mitigate edge artifacts. For density variation, we use an adaptive neighborhood radius $r_i = \text{clamp}(\lambda d_{(k)}, r_{\min}, r_{\max})$ with $d_{(k)}$ is the k -th nearest neighbor distance, λ is a scaling factor (e.g., $\lambda = 1.5$), and $[r_{\min}, r_{\max}] = [0.02, 0.3]$ m. Together, these strategies enable a scalable, consistent feature computation.

Detailed preprocessing steps, additional feature map examples, correlation analyses, and alternative dimensionality-reduction methods are provided in Appendix B (Tables B.6, Figs. B.21-B.23).

2.2.2. Stage 2: Hybrid Annotation with Semi-Supervised and Active Learning





Figure 6: Examples of (a) a pseudo-label map and (b) epistemic uncertainty calculated from mutual information.

In Stage 2, we adopt a human-in-the-loop strategy that combines self-training and active learning to reduce manual labeling effort, while maintaining high annotation quality (Fig. 6).


A small subset of spherical projection image cubes—one per TLS scan—is manually annotated using Adobe Photoshop (Adobe Inc., 2024) [46]. We use these seed labels for training an ensemble of three diverse 2D semantic segmentation models: UNet++, DeepLabV3+, and Segformer. Each model adopts a unique encoder backbone—i.e., ResNet-34 [47], EfficientNet-B3 [48], and MiT-B1 [49], respectively—to promote representational diversity and diminish correlated errors. UNet++ is well-suited for boundary recovery [36], DeepLabV3+ captures multi-scale context through atrous spatial pyramid pooling [37], and SegFormer

leverages hierarchical Transformers for long-range dependencies [38].

Each image is partitioned into five vertical tiles with buffer zones (paddings) to mitigate boundary artifacts in order to augment the training set and expose the model to a broader range of local contexts. Loss and evaluation are confined to the unbuffered cores, and tile dimensions are padded to multiples of 32 for architectural compatibility. This strategy ensures tractable training while preserving spatial continuity.



(a) Multi-encoder fusion for $3 \times N$ -channel heterogeneous feature groups. Dedicated encoders preserve pretrained weights and are fused before decoding.



(b) Schematic of the ensemble-based inference pipeline. Each model is trained independently and later fused at inference using logit averaging and softmax.

Figure 7: Ensemble inference pipeline and model architecture for feature fusion.

Because the input feature maps extend beyond standard RGB, we re-architect each backbone into a multi-encoder fusion design (Fig. 7a). Each 3-channel feature group (e.g., intensity, range, Z-inv) is processed by a dedicated encoder (ResNet-34, EfficientNet-B3, or MiT-B1). These encoders are initialized with ImageNet pre-trained weights, ensuring transfer of learned low-level filters while allowing adaptation to non-RGB feature groups. Deep features are concatenated at the bottleneck and aligned by interpolation before decoding, a strategy shown to outperform naive early fusion [50, 51, 52]. This modular design accommodates any $3 \times N$ channel configuration with minimal retraining.

The ensemble combines predictions to generate segmentation masks and epistemic uncertainty maps: high-uncertainty regions are refined manually (active learning), while high-confidence predictions are retained as pseudo-labels (self-training). Fusing logits allows the networks to complement one another [53], while deep ensembles provide more reliable uncertainty estimates than a Monte-Carlo dropout approach [27, 54]. The inference-refinement cycle is repeated until all scans are fully annotated. In practice, we display the pseudo-label masks, the uncertainty maps, and the input feature maps to guide annotators. This joint view accelerates review by steering attention to the genuinely ambiguous regions, while leaving well-segmented areas largely untouched. Technical details of the Dice-CrossEntropy joint loss [55, 56, 57] and the ensemble-based

epistemic uncertainty formulation are provided in Appendix C (Eqs. C.1-C.7).

2.2.3. Stage 3: Back Projection and Refinement in 3D Space




Figure 8: Illustration of the stage 3 processes.




Figure 9: Colorized point cloud from (a) stack of preprocessed intensity-range-z value; (b) pseudo-color from normals; and (c) first three components of PCA. Animation available in Appendix E-Table E.7

Stage 3 connects the 2D annotation workflow back to the 3D domain (Fig. 8). Each 2D segmentation mask $L(i, j)$ is back-projected to its corresponding 3D point coordinates, producing an labeled point cloud. Fig. 9 illustrates examples of colorized 3D renderings from different feature groups, enabling visual inspection and cross-validation. A key challenge arises from class ambiguity along object boundaries in the 2D projection (e.g., stem-canopy transitions). After back-projection, however, each point is re-embedded within its geometric neighborhood, providing local context for disambiguation. To leverage this, we apply a two-stage refinement process in 3D space: 1) Geometric smoothing via k -nearest neighbor (kNN) majority voting to suppress small boundary errors; 2) Feature-driven repair using a Random Forest classifier trained on a reliable core set, which corrects systematic ambiguities along complex boundaries. Together, these steps yield crisp and reliable 3D semantic annotations, which we further verify through manual inspection.




Figure 10: Colorized virtual spheres from different feature stacks. Animation available in Appendix E-Table E.7

Compact Virtual Spheres. While back-projected TLS point clouds provide detailed annotations, they remain cumbersome to inspect: large file sizes slow interaction, non-uniform angular sampling introduces distortions, and simple navigation tasks such as zooming or panning can be unintuitive. We thus introduce *virtual spheres* to address these challenges: synthetically-defined spherical grids with preset and tunable angular resolution and radius, onto which 2D feature maps are re-projected. These constructs are compact, density-neutral, and resolution-controllable, while still preserving the global structure of each scan. In practice, they serve as lightweight three-dimensional “thumbnails” that enable rapid inspection of coverage, spatial proportions, and segmentation quality, without the overhead of full-resolution point clouds (Fig. 10).

Technical details of the back-projection equations, kNN voting, and Random-Forest relabeling are provided in Appendix D(Eqs. D.1–D.6), along with additional information and applications of the virtual spheres.

2.3. Performance Evaluation and Analysis

Evaluation metrics. We evaluate each feature set and its epistemic-uncertainty map with the metrics in Table 2. Segmentation quality is captured by overall accuracy (oAcc), mean class accuracy (mAcc), mean IoU (mIoU) and per-class IoU (IoU_c). Pixel-wise errors and their uncertainty are quantified with *Shannon Entropy*, while the uncertainty map’s ability to expose errors is measured by the *area under the precision–recall curve (AUPRC)* between the uncertainty map and a binary error mask.

Marginal Gains from Expanded Feature Sets. Each handcrafted feature map highlights a specific geometric or radiometric property of the point cloud. These maps clearly enhance human interpretation by providing additional visual cues. However, it is less clear whether modern DNNs gain similar benefits from such explicit descriptors, or whether they can instead learn equivalent representations directly from simpler inputs. Furthermore, adding extra feature channels unavoidably raises memory requirements and computational costs on both CPU and GPU during training and inference.

We therefore compare segmentation performance in terms of accuracy and uncertainty across progressively

Table 2: Segmentation accuracy and uncertainty evaluation metrics.¹

Group	Metric	Definition / Equation	Interpretation
Pseudo-label accuracy	Overall Accuracy (oAcc)	$\frac{\sum_c \text{TP}_c}{\sum_c (\text{TP}_c + \text{FP}_c + \text{FN}_c)}$	Fraction of correctly labelled pixels.
	Mean Accuracy (mAcc)	$\frac{1}{C} \sum_{c=1}^C \frac{\text{TP}_c}{\text{TP}_c + \text{FN}_c}$	Average per-class recall (balanced).
	Per-class IoU (IoU _c)	$\frac{\text{TP}_c}{\text{TP}_c + \text{FP}_c + \text{FN}_c}$	Overlap between prediction and truth for class <i>c</i> .
	Mean IoU (mIoU)	$\frac{1}{C} \sum_{c=1}^C \text{IoU}_c$	Mean overlap across classes.
Uncertainty maps	Shannon Entropy (<i>H</i>)	$H(\mathbf{p}) = -\sum_{k=1}^K p_k \log p_k$	Info content of error/uncertainty map; higher <i>H</i> = more dispersed.
	Area Under PR Curve (AUPRC)	$\sum_{j=1}^{J-1} (\mathcal{R}_{j+1} - \mathcal{R}_j) \mathcal{P}_{j+1}$	Ability of an uncertainty map to highlight errors (1 = perfect, 0 = random).

¹ **Notation.** TP_c, FP_c, FN_c: true-, false-positive, and false-negative pixels; *C*: number of classes. $\mathbf{p} = [p_1, \dots, p_K]$: histogram of map values binned into *K* bins. $\mathcal{P}_j, \mathcal{R}_j$: precision and recall at the *j*th threshold.

enriched input configurations in order to quantify this accuracy-versus-cost trade-off and to assess the necessity of explicit feature engineering:

- a) Basic 3-channel sets: raw {intensity, range, *z*} and their preprocessed counterparts;
- b) Additional 3-channel sets: (i) geometric descriptors - curvature, anisotropy, planarity, (ii) pseudo-RGB from surface normals, and (iii) first three components of PCA, MNF, or ICA applied to the nine-channel set (combined from preprocessed basic set, i, and ii);
- c) Six-channel combinations: the basic set concatenated with any one additional trio, or two additional trios merged;
- d) Nine- and twelve-channel stacks: the core set joined with both normals and geometric descriptors (9 ch); and the same stack further augmented by the PCA trio (12 ch).

Impact of Annotated Sample Size. To assess the data efficiency of our annotation pipeline, we perform experiments with progressively larger subsets of annotated training data. Specifically, we randomly sample training sets of size $n = \{4, 8, 12, 16, 20, 24, 28\}$ scans. For each subset size, we maintain a fixed validation ratio of 0.25, ensuring integral train/validation set splits (e.g., 3/1 for $n = 4$; 6/2 for $n = 8$, etc.). The test set remains 9 scans from plots #6–7 across all experiments.




Figure 11: **Qualitative comparison on four representative test scans.** Each column shows the input and prediction outputs of a scan. From top to bottom: (1) preprocessed Intensity–Range–Z (I.R.Z) input stack, (2) ground-truth mask, (3) predicted mask, (4) epistemic-uncertainty map (values rescaled to $[0, 1]$ and rendered with the same *hot* color map as Fig. 6b; brighter red = higher uncertainty), (5) binary error map obtained by the pixel-wise XOR between GT and prediction, and (6) precision–recall (PR) curve measuring how well the uncertainty map localizes errors (AUPRC reported in legend). **Purple** rectangles highlight regions where high uncertainty does not coincide with errors; **cyan** circles mark areas of good agreement between uncertainty and error maps.

3. Results

3.1. 2D Segmentation Result

Fig. 11 illustrates several consistent patterns across the four representative test scans. It is clear that the ensemble model captures vertical stratification well: the upper canopy (green), root zone (orange), and ground-and-water layer (purple) predictions are almost indistinguishable from the ground-truth masks, indicating that the model handles the dominant classes with high confidence.

Errors concentrate along cluttered, high-frequency boundaries, particularly in two transition zones: (i) the interface between the distant upper canopy and the fine stems that pierce it, where sparse sampling and rapid color–depth changes reduce discriminative power; and (ii) the contact between exposed roots and the muddy, debris-strewn ground, where heterogeneous textures dominate. Most of these challenging pixels are captured by the epistemic-uncertainty maps. Bright-red patches in the uncertainty layer generally

coincide with the white error pixels beneath, yielding AUPRC scores of 0.30–0.40. The initial segment of each precision–recall curve starts near unity, showing that the top 3–5 % most-uncertain pixels are predominantly true errors—valuable for targeted manual inspection in an active-learning workflow. Despite this alignment, prediction reliability gaps remain. As highlighted by purple rectangles, discrepancies are most pronounced in scans with dense root tangles (e.g., site07), where AUPRC drops from 0.403 to 0.302.

3.2. Effect of feature enrichment




Figure 12: Effect of feature enrichment on segmentation performance. (a) Evolution of global accuracy metrics. (b) Class-specific IoU across the same feature sets, with dashed lines separating single-channel, 3-channel, 6-channel, and 9/12-channel groups.

Fig. 12 summarizes segmentation performance across 18 combinations of gradually expanded feature groups. Overall accuracy ranges between 0.80 - 0.89, mean accuracy between 0.78 - 0.87, and mean IoU between 0.68 - 0.76.

Across single-channel inputs—raw intensity, range, Z, their preprocessed versions, and per-point geometry (curvature, anisotropy, planarity)—mean IoU ranges between 0.702 - 0.731. Most three-channel combinations perform similarly, but the contrast-enhanced *I.R.Z* stack (intensity, range, inverse Z) lifts mIoU to 0.745, confirming the value of the histogram stretch in preprocessing. Merging feature groups into a six-channel input pushes performance further (0.754–0.761), yet adding still more channels yields no real gain: results level off at ≈ 0.76 . The best configuration, *IRZ_n3_CAP*, tops out at 0.768—three points better than raw intensity alone and six points better than raw range. In practice, therefore, a compact six-channel feature set captures nearly all available discriminative power on the *Mangrove3D* dataset, whereas larger stacks add computation complexity with little return.

Examining performance per class, *Ground & Water* and *Void* classes consistently achieve high IoU scores (≥ 0.80) across nearly all configurations, demonstrating their ease of separation irrespective of feature complexity. Conversely, the *Stem* class remains consistently challenging (IoU ranging between 0.47 - 0.53), highlighting persistent segmentation difficulties, due to fewer sampling points as distance increases and its thin, partially occluded structures. The *Canopy*, *Roots*, and *Objects* classes notably benefit from enriched

Table 3: PointNet++ segmentation performance on the *Mangrove3D* test set for ten incremental feature-channel configurations. Within each metric column, the top three scores are shaded (dark = best, medium = second, light = third). The XYZ baseline is shown in light gray. “Extra feat.” = number of channels added on top of XYZ.

Extra feat.	Feature group	Global metrics			IoU per class				
		oAccu.	mAccu.	mIoU	Ground	Stem	Canopy	Roots	Objects
0	baseline-xyz	0.848	0.741	0.634	0.851	0.401	0.663	0.736	0.516
1	xyz_i0	0.852	0.744	0.636	0.867	0.411	0.656	0.739	0.508
3	xyz_IRZ	0.850	0.760	0.640	0.873	0.384	0.655	0.733	0.553
3	xyz_CAP	0.859	0.806	0.689	0.867	0.424	0.676	0.748	0.731
3	xyz_N3	0.855	0.834	0.704	0.861	0.430	0.655	0.737	0.837
3	xyz_PCA	0.848	0.770	0.669	0.852	0.383	0.644	0.736	0.728
6	xyz_IRZ_N3	0.861	0.823	0.705	0.872	0.424	0.671	0.748	0.811
6	xyz_N3_CAP	0.866	0.832	0.712	0.870	0.456	0.685	0.752	0.796
9	xyz_IRZ_N3_CAP	0.861	0.808	0.699	0.865	0.446	0.686	0.742	0.758
12	xyz_IRZ_N3_CAP_PCA	0.855	0.785	0.678	0.856	0.397	0.670	0.745	0.720

feature sets, with improvements of 0.07–0.10 IoU points observed when moving from single- to six-channel inputs. The best overall balance is achieved with the six-channel *IRZ_PCA*, providing the highest IoU for four of the six semantic classes and achieving the top mIoU (0.766).

Benchmarking with PointNet++. We benchmark PointNet++ on the *Mangrove3D* dataset using 30 scans for training and nine scans for testing, following the split protocol in Section 2.1. Models are trained for 40 epochs with the Adam optimizer (initial learning rate 0.01, cosine-annealing schedule), block size 2 m, batch size 32, and 4096 points per block; the checkpoint with the highest validation mIoU is retained.

As shown in Table 3, the XYZ-only baseline achieves 0.634 mIoU, reflecting common *Stem–Canopy* confusion. Adding geometric descriptors—surface normals and *C.A.P* (six additional channels)—boosts performance to 0.712 mIoU, the best configuration observed. Class-level trends mirror this pattern: *Objects* show the largest gain (0.516 to 0.837), while *Ground* is already strong and improves only marginally. Beyond this six-extra-channel setup, additional features do not yield further benefits, indicating feature saturation once orientation and local-shape cues are included.

3.3. Data efficiency

Fig. 13 shows how training sample size influences ensemble model performance under four representative feature configurations. Accuracy improves steadily with additional samples, with performance largely saturating around 12 scans. Fig. 14 depicts the corresponding changes in entropy and AUPRC. As training




Figure 13: Prediction performance of the ensemble model trained with increasing numbers of train/val scans (4–28) across four feature configurations.




Figure 14: Uncertainty quantification trends for ensemble models trained with different feature configurations and varying numbers of train/val scans (4–28). Top row (a–d): Entropy of the total uncertainty map and mutual information map (blue), and prediction error map (red). Bottom row (e–h): Area under the precision-recall curve (AUPRC) measuring how well the mutual information map (orange) and total uncertainty map (blue) correlate with actual prediction errors.

size increases, all feature configurations exhibit consistent reductions in the entropies of total uncertainty, epistemic uncertainty (as measured by mutual information), and error maps—indicating greater model confidence and stability. Higher-dimensional inputs (6-Ch and 12-Ch) maintain lower entropy and stronger alignment between uncertainty estimates and true errors, particularly once trained with 12 or more scans.

With fewer than 12 samples, models display elevated entropy and more variable AUPRC values. Beyond this threshold, AUPRC values converge across configurations, suggesting that all feature groups achieve

comparably stable alignment between uncertainty and prediction error.

3.4. Generalization to Open-Source TLS Datasets

We apply our feature-enriched, uncertainty-aware pipeline to two public datasets to assess transferability beyond mangroves: *ForestSemantic* (boreal forest scenes) and *Semantic3D* (urban scenes).

3.4.1. *ForestSemantic*

The *ForestSemantic* dataset comprises ~ 720 million points sampled from six nominal $32\text{ m} \times 32\text{ m}$ plots in Evo, Finland (61.19° N , 25.11° E). Each plot is labeled into six classes: Ground, Trunk, First-order branch, Higher-order branch, Foliage, and Miscellany. At the time of writing (Sept 2025) only three plots—#1, #3, and #5—were publicly available on Zenodo [14]; our experiments therefore focused on these.

Visual inspection of the ground surface reveals five circular voids per plot, indicating that each plot is a registration of five individual scans. To prepare the dataset for our pipeline, we first detect the centroids of these voids and treat them as *pseudo TLS positions*. We then randomly subsample the point cloud at 5% density (e.g., from 20 million per plot to 1 million points per plot) for computational efficiency, then recenter each point by subtracting the nearest pseudo-scanner center in (x, y) and the global mean z . This yields five re-centered pseudo scans per plot and fifteen in total.

Qualitative Transfer. Figure 15 shows different feature channels and trio stacks from the spherical projection of the preprocessed *ForestSemantic* dataset. Each composite highlights complementary structural cues—from ground texture and mid-story complexity to fine-scale canopy patterns. The accompanying correlation matrix reflects moderately high positive correlations, driven largely by the substantial zero-valued (void) regions in the maps.

Figure 16 renders the 3D point clouds colored by the same feature groups. Compared with raw intensity alone, the *I.R.Z* stack reveals clearer stratification between classes; normals and geometric descriptors (*C.A.P*) accentuate stem boundaries and canopy surfaces; PCA produces smoother yet contextually meaningful color gradients, especially the colors of the stem and foliage, looking very similar to the manually assigned color of the ground truth labels. Figure 17 visualizes virtual spheres back-projected from the colorized feature maps.

Together, these visualizations qualitatively confirm that the multi-group feature framework in our pipeline transfers effectively to new TLS forest environments.

Quantitative trend. As shown in Table 4, the comparison of feature groups on the boreal-forest *ForestSemantic* plots reproduces the additive pattern. Replacing the raw intensity–range–height triplet with its preprocessed counterpart (*I.R.Z*) yields only a marginal lift over the baseline (mIoU 0.455 vs 0.453), indicating that simple radiometric normalization alone is insufficient for these cluttered, canopy-rich plots. Pure geometry (*C.A.P*) is far more valuable, adding +2.6 pts to mIoU (0.479) and +3.3 pts to mean class accuracy. Fusing normals with radiometry (*N3*) helps trunk discrimination, but does not surpass the geometric stack globally. The turning point again comes from feature fusion: adding curvature to radiometry (*IRZ_CAP*) pushes mIoU to

Table 4: Comparison of feature–input combinations for semantic segmentation on *ForestSemantic*. Within each metric column, the top three scores are highlighted with progressively darker blue shading (dark = best, medium = second, light = third); the baseline *irz_raw* row is shown in light gray for reference.

Feat.#	Feature	Global Metrics			IoU per class							mIoU Void excl.
		oAcc.	mAcc.	mIoU	Void	Ground	Trunk	1 st -order branch	Higher-order branch	Foliage	Misc.	
3	irz_raw	0.870	0.537	0.453	0.952	0.748	0.483	0.090	0.010	0.633	0.254	0.393
3	I.R.Z	0.876	0.540	0.455	0.958	0.771	0.490	0.082	0.002	0.642	0.238	0.397
3	C.A.P	0.881	0.570	0.479	0.971	0.758	0.548	0.131	0.013	0.655	0.280	0.421
3	N3	0.878	0.547	0.461	0.957	0.767	0.561	0.039	0.009	0.640	0.251	0.403
3	PCA	0.885	0.538	0.460	0.971	0.766	0.541	0.053	0.002	0.670	0.213	0.406
3	MNF	0.789	0.466	0.363	0.826	0.610	0.420	0.011	0.000	0.506	0.170	0.309
3	ICA	0.845	0.509	0.411	0.899	0.701	0.489	0.012	0.000	0.598	0.166	0.362
6	IRZ_CAP	0.895	0.583	0.496	0.986	0.809	0.579	0.130	0.010	0.679	0.280	0.441
6	IRZ_N3	0.890	0.567	0.484	0.969	0.810	0.582	0.098	0.001	0.669	0.260	0.432
6	N3_CAP	0.890	0.581	0.495	0.977	0.786	0.600	0.120	0.015	0.674	0.293	0.439
9	IRZ_N3_CAP	0.896	0.582	0.501	0.982	0.815	0.606	0.132	0.013	0.683	0.275	0.450
12	IRZ_N3_CAP_PCA	0.902	0.595	0.511	0.990	0.833	0.640	0.133	0.010	0.693	0.280	0.462

Feature codes: *IRZ_raw* = intensity, range, Z (raw); *I.R.Z* = pre-processed intensity, range, Z_{inv} ; *C.A.P* = curvature, anisotropy, planarity; *N3* = pseudo-RGB from normals; *PCA* = first three principal components of *PCA*; *MNF* = first three components of *MNF*; *ICA* = first three components of *ICA*; *IRZ_CAP* = *I.R.Z* + *C.A.P*; *IRZ_N3* = *I.R.Z* + *N3*; *N3_CAP* = *N3* + *C.A.P*; *IRZ_N3_CAP* = *I.R.Z* + *N3* + *C.A.P*; *IRZ_N3_CAP_PCA* = *IRZ_N3_CAP* + *PCA* of that set.

0.496, while the nine-channel union (*IRZ_N3_CAP*) breaks the 0.50 barrier (mIoU = 0.501). Appending a three-component PCA projection (*IRZ_N3_CAP_PCA*) delivers the overall best scores (oAcc = 0.902, mAcc = 0.595, mIoU = 0.511), confirming that modest dimensionality reduction can still help when the scene is dominated by fine-scale, self-occluding vegetation.

Multi-channel stacks excel at woody material: *IRZ_CAP* attains the best *Foliage* IoU (0.679) and, together with *N3_CAP*, boosts *Higher-order branch* discrimination beyond 0.015—triple the value achieved by geometry or radiometry alone. Normal information remains critical for *Trunk* (0.640 with *IRZ_N3_CAP_PCA*) and *Ground* (0.833). Interestingly, the PCA-extended model improves *Foliage* again to 0.693 and edges out competitors in seven of the eight valid classes, while maintaining the strongest overall accuracy. Taken together, the results show that in dense forest plots the geometric features (*C.A.P* and normals) are indispensable, and that a compact PCA head can reconcile these diverse cues without diluting their class-specific strengths.

3.4.2. Semantic3D

The *Semantic3D* dataset contains nearly 1.96 billion points of eight classes - 1: man-made terrain, 2: natural terrain, 3: high vegetation, 4: low vegetation, 5: buildings, 6: hardscape, 7: scanning artifacts, 8:

cars. In our experiments, only the training set was used, as the ground-truth labels for the original test set are not publicly available. We split the 15 training scenes into training, validation, and test subsets in an 11:1:3 ratio.

Since the *Semantic3D* has an apparent long-tail class imbalance, we applied a two-stage hybrid subsampling approach. This approach first applies stratified sampling that preferentially samples rare classes ($\leq 1\%$ frequency) with a 5% allocation from the 12M target points, preventing loss of minority classes like scanning artifacts and cars that are critical for urban scene understanding. Subsequent voxel grid downsampling (0.01 m) ensures spatial uniformity while maintaining the enhanced class distribution, enabling scalable processing of large urban point clouds for semantic segmentation applications.

Qualitative transfer. Figure 18 presents the 2D spherical-projection maps derived from a preprocessed *Semantic3D* scan. Stacking the basic LiDAR channels—intensity, range, and inverse height—already discriminates facades, ground, and hardscape. The surface-normal pseudo-color stack offers the most intuitive visual cue for human inspection, clearly distinguishing buildings and street furniture by orientation. Geometric descriptors (anisotropy, curvature, planarity) further sharpen structural edges, while the statistical transforms (PCA, MNF, ICA) distill the dominant variance into three-band composites that reveal subtle material transitions not evident in the single channels. Pearson correlation coefficients $|\rho| \leq 0.4$ demonstrate that the different feature maps convey largely complementary information, supporting their subsequent fusion.

Figure 19 renders the 3D point clouds colored by the original RGB values and the proposed feature groups. Similar to the observations on 2D feature maps, the proposed feature maps provide additional insights into object characteristics compared to original RGB or Intensity alone, such as ground vs. non-ground objects (*I.R.Z* stack), edges (*C.A.P* stack), orientations (Normals stack), and combinations of these enhancements (PCA stack).

Figure 20 visualizes virtual 3D “rings” back-projected from the colored feature maps - with the top and bottom cropped ($Zenith \in [45^\circ, 136^\circ]$) from what used to be “spheres”.

Quantitative trend. As summarized in Table 5, the comparative study on the urban *Semantic3D* benchmark echoes—and strengthens—the trends we observed in forest scenes.

Under-matched training conditions, replacing the RGB baseline (mIoU = 0.414) with the preprocessed stack of *I.R.Z* - which is essentially depth-aware radiometry, already lifts mean IoU by +3.7 pts and mean class accuracy by +3.2 pts, showing that well-normalized LiDAR channels can serve as a practical substitute when cameras are absent or their colors are unreliable. Pure geometry features - *C.A.P* - deliver a slightly lower accuracy (mIoU = 0.406). Surface-normal colorization is most impactful: the three-channel *N3* stack jumps to mIoU = 0.562—an absolute +14.8 pts over RGB—while posting the second-best overall accuracy (0.893). Fusing all three sources (*IRZ_N3_CAP*, nine channels) yields column-best scores in every global metric (mIoU = 0.563, oAcc = 0.898, mAcc = 0.660). The statistical feature groups - p3, m3, and c3 - showed various improvements compared with baseline, +8.5pts, +1 pt, and +2.6 pts, respectively. Adding a PCA

Table 5: Comparison of feature–input combinations for semantic segmentation on *Semantic3D*. Within each metric column, the top three scores are highlighted with progressively darker blue shading (dark = best, medium = second, light = third).

Feat. #	Feature	Global Metrics			IoU per class									mIoU Void excl.
		oAcc.	mAcc.	mIoU	Void	Man-made terr.	Natural terr.	High veg.	Low veg.	Buildings	Remain. hardscape	Scanning artifacts	Cars & trucks	
3	baseline-RGB	0.790	0.519	0.414	0.911	0.616	0.282	0.769	0.091	0.653	0.116	0.000	0.286	0.352
3	irz_raw	0.784	0.495	0.388	0.918	0.594	0.322	0.650	0.054	0.644	0.173	0.003	0.091	0.317
3	I.R.Z	0.818	0.551	0.451	0.914	0.651	0.455	0.740	0.019	0.672	0.131	0.044	0.432	0.393
3	C.A.P	0.814	0.521	0.406	0.918	0.680	0.441	0.786	0.066	0.614	0.128	0.016	0.342	0.348
3	N3	0.893	0.646	0.562	0.927	0.823	0.767	0.837	0.181	0.675	0.183	0.047	0.612	0.516
3	PCA	0.848	0.622	0.499	0.917	0.739	0.557	0.828	0.090	0.666	0.145	0.033	0.515	0.447
3	MNF	0.810	0.547	0.424	0.891	0.672	0.486	0.758	0.138	0.571	0.179	0.041	0.084	0.369
3	ICA	0.822	0.551	0.440	0.872	0.717	0.502	0.772	0.146	0.629	0.113	0.017	0.170	0.386
6	IRZ_CAP	0.870	0.584	0.493	0.938	0.780	0.680	0.793	0.077	0.626	0.104	0.061	0.382	0.438
6	IRZ_N3	0.888	0.611	0.501	0.932	0.823	0.760	0.825	0.208	0.703	0.187	0.043	0.032	0.448
6	N3_CAP	0.877	0.615	0.507	0.934	0.798	0.698	0.814	0.207	0.669	0.100	0.048	0.292	0.453
9	IRZ_N3_CAP	0.898	0.660	0.563	0.948	0.826	0.759	0.830	0.197	0.692	0.262	0.065	0.487	0.515
12	IRZ_N3_CAP_PCA	0.877	0.614	0.512	0.939	0.787	0.701	0.822	0.178	0.635	0.166	0.043	0.333	0.458

Feature codes: *baseline-RGB* = original RGB values; *irz_raw* = intensity, range, *Z* (raw); *I.R.Z* = pre-processed intensity, range, Z_{inv} ; *C.A.P* = curvature, anisotropy, planarity; *N3* = pseudo-RGB from surface normals; *PCA* = first three principal components of *PCA*; *MNF* = first three components of *MNF*; *ICA* = first three components of *ICA*; *IRZ_CAP* = *I.R.Z* + *C.A.P*; *IRZ_N3* = *I.R.Z* + *N3*; *N3_CAP* = *N3* + *C.A.P*; *IRZ_N3_CAP* = *I.R.Z* + *N3* + *C.A.P*; *IRZ_N3_CAP_PCA* = *I.R.Z_N3_cap* + *PCA* of that set.

projection (*IRZ_N3_CAP_PCA*) offers no further benefit, indicating that the nine-channel fusion already captures the complementary information.

Normals excel on organic shapes: *N3* alone tops *High vegetation* (IoU = 0.837) and *Cars & trucks* (0.612). Geometric stacks aid fine, man-made structures: *IRZ_N3* pushes *Buildings* to IoU = 0.703. The nine-channel fusion wins or ties in seven of nine classes and more than triples IoU for *Scanning artifacts* class (0.065 vs. 0.017–0.048 elsewhere). Together, these results demonstrate that our feature-enrichment strategy scales beyond mangrove scenes, providing consistent, cross-domain improvements, while preserving the pipeline’s low-annotation ethos.




Figure 15: Spherical-projection views of a *ForestSemantic* pseudo-scan, organized by feature type. (a–c) Basic features: preprocessed intensity, range, and Z-Inv; (d–f) Geometric features: anisotropy, curvature, and planarity; (g–i) Combined geometric and appearance features: normals, I.R.Z., and C.A.P; (j–l) Statistical features: first three components of PCA, MNF, and ICA; (m) Ground truth semantic segmentation mask projected into 2D from the 3D point cloud labels; (n) The corresponding correlation matrix of the 9 channels mentioned in (a)–(g).




Figure 16: Feature-driven colorization of the *ForestSemantic* 3D point cloud. Each inset shows a full overview of the entire plot to provide spatial context for the zoomed-in region. Animation available in Appendix E-Table E.7




Figure 17: Colorized virtual spheres of an example scan in the *ForestSemantic* dataset with three different feature groups. Animation available in Appendix E-Table E.7




Figure 18: Spherical-projection maps of a *Semantic3D* scan, organized by feature type.




Figure 19: Feature-driven colorization of a *Semantic3D* point cloud. Each inset shows a full overview of the entire plot to provide spatial context for the zoomed-in region.




Figure 20: Virtual spheres of an example *Semantic3D* scan colored with three different feature groups. Animation available in Appendix E-Table E.7

4. Discussion

High-quality 2D and 3D visualizations. A key outcome of this study is the demonstration that multi-channel visualizations—spanning 2D and 3D representations—provide an interpretable medium for TLS scene recognition. Our outputs form a three-tier suite: (i) *2D spherical feature maps* that stack basic features, geometric features, and statistical descriptors to reveal vertical stratification, edges, and orientation patterns; (ii) *3D colorized point clouds* that enable detailed inspection of boundaries, occlusions, and scan artifacts; and (iii) *3D virtual spheres* that act as compact, visually striking “thumbnails” of entire scans. These products serve complementary purposes: 2D maps enable segmentation in 2D spherical space, 3D point clouds allow fine-grained error verification, and virtual spheres provide global summaries at a fraction of the data size. Together, this visualization strategy accelerates dataset triage, guides annotators, and supports cross-scene comparison, enhancing clarity in complex environments.

Data efficiency. Our experiments on the *Mangrove3D* dataset reveal that ensemble segmentation performance saturates after approximately 12 annotated scans, regardless of the specific feature configuration. This finding underscores the data efficiency of the proposed pipeline, particularly when combined with uncertainty-guided sample selection. By focusing annotation effort on the most uncertain regions—as identified by ensemble-derived mutual information maps—it is possible to approach peak performance with substantially fewer labeled samples. In practical terms, this reduces both the cost and time required to develop a high-performing TLS segmentation model, making the approach attractive for large-scale or time-sensitive monitoring campaigns. The observed saturation point also provides guidance for field data annotation strategies: beyond a certain threshold (e.g., 12 scans for *Mangrove3D*), manually annotate additional scans may yield diminishing returns, unless the new samples introduce substantially novel scene geometries or conditions.

Feature importance. Across all three benchmarked datasets—*Mangrove3D*, *ForestSemantic*, and *Semantic3D*—feature enrichment consistently improves segmentation accuracy, but the gains are not uniform across feature types. Surface normals ($N3$) emerge as the most consistently valuable, delivering substantial improvements in class-level IoUs for geometrically complex structures such as tree stems, roots, and small urban objects. Basic features ($I.R.Z$) provide more modest but still meaningful gains, particularly for classes with strong vertical separation or height-dependent geometric–radiometric patterns. Importantly, the optimal trade-off between accuracy and computational cost is achieved with compact, nine-channel configurations (e.g., IRZ_N3_CAP), which capture nearly all available discriminative power, while avoiding the redundancy and increased training time associated with larger feature stacks. This suggests that thoughtful feature selection, rather than maximal channel inclusion, is critical for efficient TLS segmentation pipelines.

Generalization across TLS datasets. The proposed pipeline demonstrates consistent cross-domain robustness when transferred from the mangrove-dominated *Mangrove3D* dataset to the structurally distinct *ForestSemantic* (boreal forest) and *Semantic3D* (urban) benchmarks. In all cases, feature enrichment improves

segmentation accuracy, and the relative ranking of feature groups is broadly preserved despite differences in vegetation type, terrain complexity, and object composition.

Several design choices render the proposed framework both scalable and adaptable. First, all features are derived exclusively from LiDAR data, ensuring applicability to TLS scans regardless of the availability or reliability of RGB imagery. Second, feature fusion is performed at the bottleneck stage, enabling straightforward substitution of different encoder–decoder backbones (e.g., CNNs or Transformers) without altering the overall pipeline. Third, the patch size is determined adaptively, allowing the model to accommodate varying input resolutions and scene extents while maintaining computational efficiency. Fourth, the framework accommodates a flexible number of input channels under a triadic grouping scheme, thereby enabling seamless integration of diverse radiometric, geometric, or statistical descriptors without architectural modification. Finally, the *pseudo-scanner center* concept introduced for *ForestSemantic* can be extended to mobile or UAS-based LiDAR platforms that cover larger areas at lower point densities, though further optimization may be required to determine ideal center placements, spherical map resolutions, and truncation distances.

Together, the cross-dataset results and modular design show that the pipeline can be deployed in diverse TLS environments with minimal adjustment—an essential feature for large-scale monitoring programs where site conditions and sensor setups vary widely.

Limitations, future directions, and broader implications. Persistent errors in fine-scale, self-occluding structures (e.g., mid-story stems in dense canopies) and performance drops in visually heterogeneous regions indicate that enriched geometric cues and current feature stacks remain insufficient to capture the full variability of complex natural scenes. While multi-model ensembles improve accuracy and uncertainty estimation, their added computational cost may limit scalability without substantial computational resources. Future work should investigate more robust and generic preprocessing strategies, along with lighter backbone architectures that reduce redundancy, while preserving segmentation performance. Efforts also are needed to narrow the reliability gap between uncertainty maps and actual error distributions by adopting improved uncertainty estimation methods.

5. Conclusion

We present a feature-enriched, uncertainty-aware pipeline for semantic segmentation of TLS point clouds, supported by a modular design that integrates diverse feature groups, flexible encoder backbones, and ensemble-based uncertainty estimation. Across mangrove, boreal forest, and urban benchmarks, the framework consistently improves segmentation accuracy, with compact nine-channel configurations capturing nearly all available discriminative power. The inclusion of complementary visualization products—2D spherical feature maps, 3D colorized point clouds, and compact virtual spheres—proves valuable for both annotator guidance and rapid scene assessment, enabling efficient, targeted annotation and dataset triage. The pipeline’s reliance solely on LiDAR-derived features, coupled with its demonstrated cross-domain robustness, positions it as

a scalable and portable solution for large-scale monitoring in forestry, ecology, and urban environments. Future developments in preprocessing, backbone efficiency, uncertainty evaluation, and mobile deployment are expected to further expand its applicability and impact.

While sensor fusion is becoming increasingly popular, our feature-enrichment results highlight the untapped potential of LiDAR point clouds alone. Careful exploitation of LiDAR-derived features can, in some cases, surpass early (or “raw”) LiDAR–camera fusion in the spherical projection space, revealing a distinct and informative view *through the perspective of LiDAR*. We therefore conclude that for downstream tasks such as semantic segmentation, rigorous feature design, and preprocessing of LiDAR data are no less important than integrating multi-source sensing.

Acknowledgments

The authors are grateful to Dr. Nidhal Carla Bouaynaya and Dr. Bartosz Krawczyk for their valuable insights on uncertainty evaluation and semantic segmentation methods. We thank Mr. Brett Matzke for his assistance in setting up the computing resources. We also acknowledge RIT Research Computing [58] for providing access to NVIDIA A100 computing resources.

CRedit authorship contribution statement

Fei Zhang: Conceptualization, Methodology, Software, Validation, Formal Analysis, Investigation, Data Curation, Visualization, Writing – original draft, Project administration, Supervision. **Fei Zhang** and **Rob Chancia:** Conceptualization, Experimental design. **Fei Zhang:** Experiment execution. **Fei Zhang** and **Josie Clapp:** Data Annotation (*Mangrove3D* dataset). **Rob Chancia**, **Richard MacKenzie**, and **Jan van Aardt:** Data Collection (*Mangrove3D* dataset). **Rob Chancia** and **Amirhossein Hassanzadeh:** Validation, Writing – review & editing, Software. **Jan van Aardt:** Writing – review & editing, Supervision. **Dimah Dera:** Writing – review & editing.

Funding

This work was supported by the USDA Forest Service [project numbers 20-JV-11272136-016].

Conflicts of Interest

The authors declare that there is no conflict of interest regarding the publication of this article.

Data and Code Availability

The *Mangrove3D* dataset has been deposited on Zenodo (doi:10.5281/zenodo.16933584) and is currently under embargo; it will be made publicly accessible upon publication. The accompanying preprocessing,

feature engineering, training, and evaluation code (including environment specification and reproducibility instructions) is available at <https://fz-rit.github.io/through-the-lidars-eye/>.

Supplementary Materials

Appendix A. Supplementary Material for Mangrove3D Dataset

Annotation guideline can be downloaded at https://drive.google.com/file/d/1S8VgY1TX00Rs9HeCIrMU0g3UhkRZFU1P/view?usp=drive_link.

Appendix B. Supplementary Material for Stage 1

Preprocessing details. Table S1 lists preprocessing steps for all feature maps, including intensity, range, Z-inv, anisotropy, curvature, planarity, and statistical stacks.

Table B.6: Preprocessing steps for the feature maps shown in Figures B.21 and B.23.

Feature Group	Fig.	Map	Meaning	Preprocessing
Basic	Fig. B.21(a)	Intensity	Radiometric intensity	Global histogram stretch; normalize to [0.01, 1.0] $H = Z - Z_{\min}$; negate H ; normalize to [0.01, 1.0]
	Fig. B.21(b)	Range	Distance from scanner	
	Fig. B.21(c)	Z-Inv	Inverted height above ground	
Geometric	Fig. B.21(d)	Anisotropy	Directional structure	None
	Fig. B.21(e)	Curvature	Local surface dev.	
	Fig. B.21(f)	Planarity	Planar alignment	
Direct Stack	Fig. B.23(a)	Normals	X/Y/Z of surface normals	Convert to azimuth/elevation; map to HSV (azimuth $\rightarrow$ hue, elevation $\rightarrow$ value, saturation = 0.6); convert to RGB
	Fig. B.23(b)	I.R.Z stack	Intensity, Range, Z-Inv $\rightarrow$ R/G/B	Channel order: I $\rightarrow$ R, R $\rightarrow$ G, Z $\rightarrow$ B
	Fig. B.23(c)	C.A.P stack	Curvature, Anisotropy, Planarity $\rightarrow$ R/G/B	Channel order: C $\rightarrow$ R, A $\rightarrow$ G, P $\rightarrow$ B
Statistical Stack	Fig. B.23(d)	PCA	First 3 PCs	Normalize each to [0,1]; stack to RGB
	Fig. B.23(e)	MNF	First 3 MNFs	
	Fig. B.23(f)	ICA	First 3 ICs	

Additional feature maps. Fig. B.21 illustrates single-channel maps from scan #site1_01 of the Mangrove3D dataset (intensity, range, Z-inv, anisotropy, curvature, planarity).




Figure B.21: Single channel feature maps generated from the 2D spherical projection of the terrestrial LiDAR point cloud. Each map represents a different basic or geometric attribute. For visualization consistency, all maps were normalized to the $[0,1]$ range and visualized using the plasma colormap, where deep blue indicates low values and bright yellow denotes high values.

Eigenvalue-based Geometric Descriptors. For each point i , the local covariance matrix was computed as

$$\mathbf{C}_i = \frac{1}{N-1} \sum_{j=1}^N (\mathbf{x}_j - \bar{\mathbf{x}})(\mathbf{x}_j - \bar{\mathbf{x}})^\top, \quad (\text{B.1})$$

where $\mathbf{x}_j \in \mathbb{R}^3$ denotes the coordinates of the j -th neighbor, $\bar{\mathbf{x}}$ is the neighborhood centroid, and N is the number of points in the neighborhood.

Eigen-decomposition of $\mathbf{C}_i$ yields ordered eigenvalues $\lambda_1 \leq \lambda_2 \leq \lambda_3$, which describe the variance along orthogonal axes. From these, we define three commonly used geometric descriptors:

$$\kappa = \frac{\lambda_1}{\lambda_1 + \lambda_2 + \lambda_3}, \quad A = \frac{\lambda_3 - \lambda_2}{\lambda_3}, \quad P = \frac{\lambda_2 - \lambda_1}{\lambda_3}. \quad (\text{B.2})$$

Here, curvature (κ) reflects how “linear” or “planar” a neighborhood is, approaching zero for flat surfaces and increasing when points are distributed along a sharp edge or corner. Anisotropy (A) quantifies the dominance of one principal axis over the others, with higher values indicating elongated, line-like structures (e.g., stems or branches). Planarity (P) measures the degree to which points lie within a 2D plane, with higher values indicating flat surfaces (e.g., ground or large leaves). These features provide compact measures of local surface geometry and have been widely used for point-cloud classification and segmentation [59, 60, 61].

Feature correlations. Fig. B.22 presents the Pearson correlation matrix among feature maps, validating inter-feature relationships.




Figure B.22: Correlation matrix of the feature maps.

Stacked channels. Fig. B.23 shows pseudo-color images from stacked feature channels (I.R.Z, C.A.P, and normal-based encodings).

Additional Dimensionality-Reduction Methods. In addition to PCA, we explored two other dimensionality-reduction techniques commonly applied in remote sensing.

- **MNF** [62]: a two-step transform that first whitens data using an estimated noise covariance matrix, then applies PCA on the noise-whitened data. The resulting components are ranked by signal-to-noise ratio.
- **ICA** [63]: a statistical technique that seeks latent components which are mutually independent and non-Gaussian, thus capturing higher-order relationships beyond variance and correlation.




Figure B.23: Three-channel feature maps generated from the 2D spherical projection of the terrestrial LiDAR point cloud.

Appendix C. Supplementary Material for Stage 2

Appendix C.1. Annotation guideline

A concise annotation protocol is provided here [details].

Appendix C.2. Loss functions

We trained models using a weighted combination of Dice loss and Cross-Entropy loss:

$$\mathcal{L} = 0.5 \mathcal{L}_{\text{Dice}} + 0.5 \mathcal{L}_{\text{CrossEntropy}}. \quad (\text{C.1})$$

Dice loss [55, 56]:

$$\mathcal{L}_{\text{Dice}} = 1 - \frac{2 \sum_i p_i g_i}{\sum_i p_i + \sum_i g_i}, \quad (\text{C.2})$$

where p_i and g_i are the predicted probability and ground-truth label for pixel i .

Cross-Entropy loss [57]:

$$\mathcal{L}_{\text{CrossEntropy}} = - \sum_i \sum_c g_{i,c} \log(p_{i,c}), \quad (\text{C.3})$$

where $p_{i,c}$ is the predicted probability of class c at pixel i , and $g_{i,c}$ is the corresponding ground-truth indicator.

Appendix C.3. Ensemble inference and uncertainty estimation

For logits $\mathbf{Z}^{(m)} \in \mathbb{R}^{C \times H \times W}$ from the m -th model, the ensemble-averaged probability map is:

$$\bar{\mathbf{Z}} = \frac{1}{M} \sum_{m=1}^M \mathbf{Z}^{(m)}, \quad \bar{\mathbf{P}} = \text{softmax}(\bar{\mathbf{Z}}). \quad (\text{C.4})$$

The predictive entropy of the averaged probability distribution is:

$$H[\bar{\mathbf{P}}] = - \sum_{c=1}^C \bar{P}_{c,i,j} \log \bar{P}_{c,i,j}. \quad (\text{C.5})$$

The expected entropy across the M models is:

$$\mathbb{E}[H[\mathbf{P}]] = \frac{1}{M} \sum_{m=1}^M \left(- \sum_{c=1}^C P_{c,i,j}^{(m)} \log P_{c,i,j}^{(m)} \right). \quad (\text{C.6})$$

Finally, epistemic uncertainty is estimated as the mutual information [27, 54]:

$$U_{\text{ep}}(i, j) = H[\bar{\mathbf{P}}] - \mathbb{E}[H[\mathbf{P}]]. \quad (\text{C.7})$$

These equations quantify model disagreement at each pixel, providing a principled measure of epistemic uncertainty widely adopted in semantic segmentation.

Appendix D. Supplementary Material for Stage 3

Appendix D.1. Back projection

Each 3D point (x, y, z) is assigned the label of its corresponding pixel (i, j) from the 2D spherical map:

$$\ell(x, y, z) = L(i, j). \quad (\text{D.1})$$

In compact form:

$$\ell(x, y, z) = L\left(\left\lfloor \frac{\theta(x, y, z) - \theta_{\min}}{\Delta\theta} \right\rfloor, \left\lfloor \frac{\phi(x, y, z) - \phi_{\min}}{\Delta\phi} \right\rfloor\right), \quad (\text{D.2})$$

where θ, ϕ are spherical angles and $(\Delta\theta, \Delta\phi)$ are grid steps.

Appendix D.2. Geometric smoothing

For a point i at $\mathbf{p}_i$, the k nearest neighbors are

$$\mathcal{N}_k(i) = \arg \text{top-k}_{j \neq i} \|\mathbf{p}_i - \mathbf{p}_j\|_2. \quad (\text{D.3})$$

The label is updated by majority vote:

$$\hat{y}_i = \arg \max_{c \in \{1, \dots, C\}} \sum_{j \in \mathcal{N}_k(i)} \mathbf{1}[y_j = c]. \quad (\text{D.4})$$

Appendix D.3. Random-Forest relabeling

A reliable core set is defined as

$$\mathcal{R} = \{ i \mid y_i = \hat{y}_i \wedge y_i \neq 0 \}. \quad (\text{D.5})$$

A balanced Random Forest f_θ is trained on multiscale features (XYZ, normals, etc.). For suspect points, final labels are:

$$y_i^{\text{final}} = \begin{cases} \arg \max_c p_i(c), & \text{if } \max_c p_i(c) \geq \tau, \\ \hat{y}_i, & \text{otherwise,} \end{cases} \quad (\text{D.6})$$

with confidence threshold $\tau = 0.8$.

Appendix D.4. Virtual spheres

Virtual spheres re-project 2D feature maps onto synthetic, uniformly sampled spherical grids rather than irregular TLS point distributions. This produces a compact, standardized abstraction of each scan that preserves global structure while dramatically reducing data volume.

At 1° angular resolution, for example, an average 800K-point CBL TLS scan compresses to $\sim 48K$ points ($\approx 1/16$ of the memory) with minimal structural loss. Because the representation is density-neutral, it mitigates distortions from uneven TLS sampling, and resolution can be flexibly tuned to balance fidelity against storage. Importantly, the number of points in a virtual sphere is determined solely by angular resolution and field of view, not by the density of the original point cloud. Thus, while a Semantic3D scan may contain up to 200M points, its virtual sphere at 0.2° angular resolution yields only $\sim 0.8M$ points—cutting memory demands by a factor of 250 while preserving global structure.

These properties make virtual spheres broadly useful. They enable rapid quality control of scan coverage, occlusion, and segmentation outputs; serve as lightweight metadata summaries for large TLS archives; and support visual benchmarking across sites, campaigns, or ecological conditions. Their clarity and compactness also extend to education and outreach, where visually intuitive overviews of complex forest structures can engage non-expert audiences.

Appendix E. Animations of the Colorized Point Clouds and Virtual Spheres

Table E.7: Figures and corresponding animation links.

Figure	Animation Link
Figure 9: Mangrove3D point clouds	View Animations
Figure 10: Mangrove3D virtual spheres	View Animations
Figure 16: ForestSemantic point clouds	View Animations
Figure 17: ForestSemantic virtual spheres	View Animations
Figure 20: Semantic3D virtual spheres	View Animations

References

- [1] K. Calders, G. Newnham, A. Burt, M. Disney, P. Raunonen, M. Herold, et al., Terrestrial laser scanning in forest ecology: Expanding the horizon, *Remote Sensing of Environment* 251 (2020) 112102.
- [2] M. I. Disney, M. Boni Vicari, A. Burt, K. Calders, S. L. Lewis, P. Raunonen, P. Wilkes, Weighing trees with lasers: advances, challenges and opportunities, *Interface Focus* 8 (2) (2018) 20170048.
- [3] X. Liang, V. Kankare, J. Hyypä, Y. Wang, A. Kukko, I. Häggström, M. Vastaranta, Terrestrial laser scanning in forest inventories, *ISPRS Journal of Photogrammetry and Remote Sensing* 115 (2016) 63–77.
- [4] S. Penman, P. Lentini, B. Law, A. York, An instructional workflow for using terrestrial laser scanning (tls) to quantify vegetation structure for wildlife studies, *Forest Ecology and Management* 548 (2023) 121405.
- [5] C. R. Qi, L. Yi, H. Su, L. J. Guibas, Pointnet++: Deep hierarchical feature learning on point sets in a metric space, *Advances in neural information processing systems* 30 (2017).
- [6] H. Thomas, C. R. Qi, J.-E. Deschaud, B. Marcotegui, F. Goulette, L. J. Guibas, Kpconv: Flexible and deformable convolution for point clouds, in: *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 6411–6420.
- [7] M. Kulicki, C. Cabo, T. Trzciński, J. Będkowski, K. Stereńczak, Artificial intelligence and terrestrial point clouds for forest monitoring, *Current Forestry Reports* 11 (1) (2024) 5.
- [8] A. Bornand, M. Abegg, F. Morsdorf, N. Rehus, Completing 3d point clouds of individual trees using deep learning, *Methods in Ecology and Evolution* 15 (11) (2024) 2010–2023.
- [9] T. Hackel, N. Savinov, L. Ladicky, J. D. Wegner, K. Schindler, M. Pollefeys, Semantic3d. net: A new large-scale point cloud classification benchmark, *arXiv preprint arXiv:1704.03847* (2017).

- [10] J. Walczak, Inlut3d: Challenging real indoor dataset for point cloud analysis, arXiv preprint arXiv:2408.03338 (2024).
- [11] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, M. Nießner, Scannet: Richly-annotated 3d reconstructions of indoor scenes, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2017, pp. 5828–5839.
- [12] M. Mézizette, N. Audebert, P. Kervella, J. Verdun, 3dses: an indoor lidar point cloud segmentation dataset with real and pseudo-labels from a 3d model, arXiv preprint arXiv:2501.17534 (2025).
- [13] X. Liang, H. Qi, X. Deng, J. Chen, S. Cai, Q. Zhang, Y. Wang, A. Kukko, J. Hyypä, Forestsemantic: a dataset for semantic learning of forest from close-range sensing, *Geo-Spatial Information Science* 28 (1) (2025) 185–211.
- [14] M. Lab, Forestsemantic: A dataset for semantic learning of forest from close-range sensing (Apr. 2025). doi:10.5281/zenodo.15193973.
URL <https://doi.org/10.5281/zenodo.15193973>
- [15] I. Autodesk, Autodesk recap, <https://www.autodesk.com/products/recap/overview>, accessed: 2025-06-05 (2025).
- [16] T. Inc., Trimble realworks, <https://geospatial.trimble.com/products-and-solutions/trimble-realworks>, accessed: 2025-06-05 (2025).
- [17] S. S. E. Contributors, Semantic segmentation editor, <https://github.com/Hitachi-Automotive-And-Industry-Lab/semantic-segmentation-editor>, accessed: 2025-06-05 (2025).
- [18] J. P. Snyder, Flattening the earth: two thousand years of map projections, University of Chicago Press, 1997.
- [19] J. Mao, S. Shi, X. Wang, H. Li, 3d object detection for autonomous driving: A comprehensive survey, *International Journal of Computer Vision* 131 (8) (2023) 1909–1963.
- [20] S. Barnea, S. Filin, Segmentation of terrestrial laser scanning data by integrating range and image content, in: The Proceedings of XXIth ISPRS Congress, 2008.
- [21] S. Barnea, S. Filin, Segmentation of terrestrial laser scanning data using geometry and image information, *ISPRS Journal of Photogrammetry and Remote Sensing* 76 (2013) 33–48, terrestrial 3D modelling. doi:<https://doi.org/10.1016/j.isprsjprs.2012.05.001>.
URL <https://www.sciencedirect.com/science/article/pii/S0924271612000809>

- [22] H. Mahmoudabadi, M. J. Olsen, S. Todorovic, Efficient terrestrial laser scan segmentation exploiting data structure, *ISPRS Journal of Photogrammetry and Remote Sensing* 119 (2016) 135–150. doi:<https://doi.org/10.1016/j.isprsjprs.2016.05.015>.
URL <https://www.sciencedirect.com/science/article/pii/S0924271616301083>
- [23] A. Milioto, I. Vizzo, J. Behley, C. Stachniss, Rangenet++: Fast and accurate lidar semantic segmentation, in: 2019 IEEE/RSJ international conference on intelligent robots and systems (IROS), IEEE, 2019, pp. 4213–4220.
- [24] T. Cortinhal, G. Tzelepis, E. Erdal Aksoy, Salsanext: Fast, uncertainty-aware semantic segmentation of lidar point clouds, in: *International Symposium on Visual Computing*, Springer, 2020, pp. 207–222.
- [25] B. Wu, A. Wan, X. Yue, K. Keutzer, Squeezeseg: Convolutional neural nets with recurrent crf for real-time road-object segmentation from 3d lidar point cloud, in: 2018 IEEE international conference on robotics and automation (ICRA), IEEE, 2018, pp. 1887–1893.
- [26] Y. Gal, Z. Ghahramani, Dropout as a bayesian approximation: Representing model uncertainty in deep learning, in: *international conference on machine learning*, PMLR, 2016, pp. 1050–1059.
- [27] B. Lakshminarayanan, A. Pritzel, C. Blundell, Simple and scalable predictive uncertainty estimation using deep ensembles, in: *Advances in Neural Information Processing Systems*, 2017, pp. 6402–6413.
- [28] O. Sener, S. Savarese, Active learning for convolutional neural networks: A core-set approach, in: *International Conference on Learning Representations (ICLR)*, 2018.
- [29] F. Shao, Y. Luo, P. Liu, J. Chen, Y. Yang, Y. Lu, J. Xiao, Active learning for point cloud semantic segmentation via spatial-structural diversity reasoning, in: *Proceedings of the 30th ACM International Conference on Multimedia*, 2022, pp. 2575–2585.
- [30] Z. Xu, B. Yuan, S. Zhao, Q. Zhang, X. Gao, Hierarchical point-based active learning for semi-supervised point cloud semantic segmentation, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 18098–18108.
- [31] Q. Xie, M.-T. Luong, E. Hovy, Q. V. Le, Self-training with noisy student improves imagenet classification, in: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 10687–10698.
- [32] Y. Zou, Z. Zhang, H. Zhang, C.-L. Li, X. Bian, J.-B. Huang, T. Pfister, Pseudoseg: Designing pseudo labels for semantic segmentation, *arXiv preprint arXiv:2010.09713* (2020).
- [33] B. Wu, X. Zhou, S. Zhao, X. Yue, K. Keutzer, Squeezesegv2: Improved model structure and unsupervised domain adaptation for road-object segmentation from a lidar point cloud, in: 2019 international conference on robotics and automation (ICRA), IEEE, 2019, pp. 4376–4382.

- [34] Y. Zou, Z. Yu, X. Liu, B. Kumar, J. Wang, Confidence regularized self-training, in: Proceedings of the IEEE/CVF international conference on computer vision, 2019, pp. 5982–5991.
- [35] L. El Mendili, S. Daniel, T. Badard, Distribution-aware contrastive learning for domain adaptation in 3d lidar segmentation, *Computer Vision and Image Understanding* (2025) 104438.
- [36] Z. Zhou, M. M. Rahman Siddiquee, N. Tajbakhsh, J. Liang, Unet++: A nested u-net architecture for medical image segmentation, in: Deep learning in medical image analysis and multimodal learning for clinical decision support: 4th international workshop, DLMIA 2018, and 8th international workshop, ML-CDS 2018, held in conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, proceedings 4, Springer, 2018, pp. 3–11.
- [37] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, H. Adam, Encoder-decoder with atrous separable convolution for semantic image segmentation, in: Proceedings of the European conference on computer vision (ECCV), 2018, pp. 801–818.
- [38] E. Xie, W. Wang, Z. Yu, A. Anandkumar, J. M. Alvarez, P. Luo, Segformer: Simple and efficient design for semantic segmentation with transformers, *Advances in neural information processing systems* 34 (2021) 12077–12090.
- [39] S. Vora, A. H. Lang, B. Helou, O. Beijbom, Pointpainting: Sequential fusion for 3d object detection, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2020, pp. 4604–4612.
- [40] Z. Dong, H. Ji, X. Huang, W. Zhang, X. Zhan, J. Chen, Pep: A point enhanced painting method for unified point cloud tasks, *arXiv preprint arXiv:2310.07591* (2023).
- [41] N. Chehata, L. Guo, C. Mallet, Airborne lidar feature selection for urban classification using random forests, in: *Laserscanning*, 2009.
- [42] J. Zhang, X. Lin, X. Ning, Svm-based classification of segmented airborne lidar point clouds in urban areas, *Remote Sensing* 5 (8) (2013) 3749–3775. doi:10.3390/rs5083749.
URL <https://www.mdpi.com/2072-4292/5/8/3749>
- [43] Q. Hu, B. Yang, L. Xie, S. Rosa, Y. Guo, Z. Wang, N. Trigoni, A. Markham, Randla-net: Efficient semantic segmentation of large-scale point clouds, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2020, pp. 11108–11117.
- [44] M. Weinmann, B. Jutzi, S. Hinz, C. Mallet, Semantic point cloud interpretation based on optimal neighborhoods, relevant features and efficient classifiers, *ISPRS Journal of Photogrammetry and Remote Sensing* 105 (2015) 286–304.

- [45] SICK AG, LMS151 Laser Measurement Sensor, SICK AG, datasheet, accessed August 30, 2025 (2017).
URL <https://www.sick.com/ae/en/products/lidar-and-radar-sensors/lidar-sensors/lms1xx/lms151-10100/p/p141840>
- [46] Adobe Inc., Adobe photoshop, <https://www.adobe.com/products/photoshop.html>, accessed: 2025-06-28 (2024).
- [47] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.
- [48] M. Tan, Q. V. Le, Efficientnet: Rethinking model scaling for convolutional neural networks, International Conference on Machine Learning (ICML) (2019).
- [49] X. Yu, J. Wang, Y. Zhao, Y. Gao, Mix-vit: Mixing attentive vision transformer for ultra-fine-grained visual categorization, Pattern Recognition 135 (2023) 109131.
- [50] S. Saidi, S. Idbraim, Y. Karmoude, A. Masse, M. Arbelo, Deep-learning for change detection using multi-modal fusion of remote sensing images: A review, Remote Sensing 16 (20) (2024). doi:10.3390/rs16203852.
URL <https://www.mdpi.com/2072-4292/16/20/3852>
- [51] Y. Lei, D. Peng, P. Zhang, Q. Ke, H. Li, Hierarchical paired channel fusion network for street scene change detection, IEEE Transactions on Image Processing 30 (2021) 55–67. doi:10.1109/TIP.2020.3031173.
- [52] C. Zhang, P. Yue, D. Tapete, L. Jiang, B. Shangguan, L. Huang, G. Liu, A deeply supervised image fusion network for change detection in high resolution bi-temporal remote sensing images, ISPRS Journal of Photogrammetry and Remote Sensing 166 (2020) 183–200. doi:<https://doi.org/10.1016/j.isprsjprs.2020.06.003>.
URL <https://www.sciencedirect.com/science/article/pii/S0924271620301532>
- [53] T. G. Dietterich, Ensemble methods in machine learning, in: International workshop on multiple classifier systems, Springer, 2000, pp. 1–15.
- [54] A. Kendall, Y. Gal, What uncertainties do we need in bayesian deep learning for computer vision?, Advances in neural information processing systems 30 (2017).
- [55] F. Milletari, N. Navab, S.-A. Ahmadi, V-net: Fully convolutional neural networks for volumetric medical image segmentation, in: 2016 fourth international conference on 3D vision (3DV), Ieee, 2016, pp. 565–571.
- [56] C. H. Sudre, W. Li, T. Vercauteren, S. Ourselin, M. Jorge Cardoso, Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations, in: Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: Third International Workshop, DLMIA 2017,

and 7th International Workshop, ML-CDS 2017, Held in Conjunction with MICCAI 2017, Québec City, QC, Canada, September 14, Proceedings 3, Springer, 2017, pp. 240–248.

- [57] A. Mao, M. Mohri, Y. Zhong, Cross-entropy loss functions: Theoretical analysis and applications, in: International conference on Machine learning, PMLR, 2023, pp. 23803–23828.
- [58] R. I. of Technology, Research computing services (2019). doi:10.34788/OS3G-QD15.
URL <https://www.rit.edu/researchcomputing/>
- [59] H. Harshit, S. Kushwaha, K. Jain, Geometric features interpretation of photogrammetric point cloud from unmanned aerial vehicle, ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences 10 (2022) 83–88.
- [60] Z. Duran, K. Ozcan, M. E. Atik, Classification of photogrammetric and airborne lidar point clouds using machine learning algorithms, Drones 5 (4) (2021). doi:10.3390/drones5040104.
URL <https://www.mdpi.com/2504-446X/5/4/104>
- [61] M. E. Atik, Z. Duran, D. Z. Seker, Machine learning-based supervised classification of point clouds using multiscale geometric features, ISPRS International Journal of Geo-Information 10 (3) (2021). doi:10.3390/ijgi10030187.
URL <https://www.mdpi.com/2220-9964/10/3/187>
- [62] A. A. Green, M. Berman, P. Switzer, M. D. Craig, A transformation for ordering multispectral data in terms of image quality with implications for noise removal, IEEE Transactions on geoscience and remote sensing 26 (1) (1988) 65–74.
- [63] A. Hyvärinen, E. Oja, Independent component analysis: algorithms and applications, Neural networks 13 (4-5) (2000) 411–430.