remote sensing

MDPI

Article

Fallen Tree Detection and Mapping from Airborne Laser Scanning Point Clouds Using Physics-Guided Supervoxel Learning

Xueru Liang [1], Yueqian Shen [1,*], Jinhu Wang [2], Jinguo Wang [1], Yanming Chen [1] and Dong Chen [3]

[1] School of Earth Sciences and Engineering, Hohai University, Nanjing 210098, China; xrliang@hhu.edu.cn (X.L.)
[2] Institute for Biodiversity and Ecosystem Dynamics (IBED), University of Amsterdam, 1012 WX Amsterdam, The Netherlands
[3] College of Civil Engineering, Nanjing Forestry University, Nanjing 210037, China
* Correspondence: y.shen_lidar@hhu.edu.cn

Highlights

What are the main findings?

- A physics-guided supervoxel framework was developed for fallen tree detection from operational ALS point clouds.
- Under point-level evaluation across seven independent ALS test plots, the full pipeline achieved a mean fallen tree F1-score of 85.39%.

What are the implications of the main findings?

- Supervoxel-level learning can produce interpretable fallen tree maps under medium-to-low ALS point densities of 10–15 pts/m^2.
- The results support ALS-based near-ground forest structure analysis, while dense overlap and understory occlusion remain key limits.

Abstract

Accurate fallen tree detection from airborne laser scanning (ALS) point clouds is important for forest structural assessment and disturbance monitoring. However, fallen trees remain difficult to extract from operational ALS data because near-ground stems are often fragmented by understory occlusion, uneven sampling, and geometric confusion with shrubs, terrain surfaces, and other low vegetation. This study proposes a physics-guided supervoxel framework for fallen tree detection from ALS point clouds. The framework represents ALS observations as spatially coherent supervoxels and describes each supervoxel using local geometric features, cylindrical morphology, and graph-based neighborhood context. Cylindrical descriptors encode the elongated stem structure of fallen trees, while graph-topological features represent spatial continuity among adjacent supervoxels. A stacking ensemble classifier is trained with SHAP-based feature selection, followed by graph-constrained spatial refinement to improve local label consistency. The framework was trained using 30 plots and evaluated on seven held-out forest plots from the Dutch Actueel Hoogtebestand Nederland 4 (AHN4) dataset, with point densities of 10–15 pts/m^2. Using point-level reference labels, the stacking classifier achieved a mean fallen tree F1-score of 83.86% across the test plots, which increased to 85.39% after spatial refinement. Ablation results show that cylindrical descriptors provided the largest gain, improving the fallen tree F1 score by 11.77%. These results indicate that physics-guided supervoxel learning can produce interpretable fallen tree maps from operational ALS data and support subsequent analyses of near-ground forest structure and disturbance effects.

Check for updates

Academic Editors: Ningning Zhu, Zhenyang Hui, Wenxia Dai and Linyuan Li

Received: 15 September 2026
Revised: 11 October 2026
Accepted: 15 October 2026
Published: 16 October 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

Keywords: airborne laser scanning; ALS point cloud; ensemble learning; fallen tree; feature extraction; supervoxel

1. Introduction

Fallen trees are an ecologically important component of coarse woody debris (CWD). They provide habitat for a wide range of organisms, contribute to nutrient cycling, and act as a long-term carbon pool in forest ecosystems [1]. The amount and spatial distribution of CWD also reflect disturbance history, stand development, and forest recovery processes [2,3]. Consequently, spatially explicit fallen tree mapping can support forest structural assessment, disturbance monitoring, and CWD-related carbon accounting [4]. However, fallen trees are difficult to map consistently over large areas because they occur close to the ground and are frequently obscured by understory vegetation, canopy remnants, and terrain variation.

Field inventories can provide detailed measurements of fallen stems, but their spatial coverage is limited and repeated surveys are labor-intensive. Airborne laser scanning (ALS) provides three-dimensional observations of forest structure over extensive areas and has therefore become an important data source for forest inventory. ALS has been used to estimate forest height and related stand attributes [5], retrieve leaf area index [6], characterize three-dimensional vegetation structure [7], detect standing dead trees, and estimate tree-level stem volume [8] and biomass-related attributes [9]. In principle, ALS can also capture near-ground woody structures when laser pulses penetrate the canopy and understory. For fallen tree detection, useful cues include height above ground, local linearity, horizontal orientation, cylindrical morphology, and continuity along the stem direction.

Previous studies have explored several approaches for detecting fallen trees or CWD from airborne remote-sensing data. Raster- and object-based methods have used ALS-derived surface products to delineate downed logs in relatively open or disturbed forests [10,11]. Other studies have exploited the elongated geometry of fallen stems through echo-width filtering and morphological image analysis [12] or line-template matching [13]. These approaches include rule-based procedures that identify candidate stems using predefined filtering thresholds, shape criteria, or geometric templates. Such procedures can be effective where stems retain recognizable linear structures, but their performance depends on how well the selected criteria represent the observed objects. Irregular or overlapping stems may depart from the assumed geometry, while non-target linear structures can also satisfy the detection criteria. For example, Lindberg et al. [13] reported that ditches and roads produced detected lines, illustrating the difficulty of distinguishing fallen stems using linearity alone. At plot scale, area-based ALS metrics have also been used to predict CWD volume [14]; however, such approaches do not directly identify the spatial locations of individual fallen tree structures.

Learning-based methods provide an alternative by learning the relationship between descriptive features and reference labels from annotated samples, rather than relying exclusively on predefined detection rules. For ALS point clouds, Polewski et al. [15] combined point-level and primitive-level shape descriptors to classify candidate fallen-stem segments. Their subsequent work used simulated training samples to learn the similarity function and stopping criterion for a Normalized Cut procedure that merged short segments into whole stems [16]. Polewski et al. [17] further incorporated contextual constraints through a conditional random field, accounting for spatial interactions between adjacent segment candidates. Alongside these developments, Heinaro et al. [18] combined iterative Hough line detection with region growing to map individual fallen trees from

ALS point clouds, demonstrating the influence of stem size, surrounding vegetation, and decay state on detectability. Heinaro et al. [19] subsequently examined how point-cloud characteristics and filtering decisions affect fallen tree detection performance. Beyond ALS point-cloud methods, Lopes Queiroz et al. [20] applied random forest classification within an object-based framework to map CWD from centimetric aerial imagery, combining spectral and spatial features with LiDAR-derived height variables. Related applications have also extended to stream environments, where Dakin Kuiper et al. [21] developed an automated ALS-based approach for detecting instream wood in small coastal streams. Together, these studies illustrate the diversity of methods and application settings for woody debris mapping, while highlighting the need to consider data characteristics and environmental conditions when interpreting detection performance.

Machine learning has been widely investigated for 3D point-cloud classification and semantic interpretation, with particular attention to feature selection and neighborhood representation [22]. Deep learning has also advanced semantic segmentation by learning feature representations directly from data [23]. Nevertheless, its application to fallen tree detection from operational ALS data remains challenging. Near-ground fallen stems are often sparsely sampled, fragmented by occlusion, and underrepresented in labelled training data. In addition, their local geometry can be confused with low vegetation, terrain fragments, roots, and crown debris. Under these conditions, an interpretable representation that explicitly incorporates stem morphology and neighborhood structure remains valuable.

Three methodological considerations are particularly relevant to operational ALS-based fallen tree detection. First, the use of point-level and segment-level shape representations in previous studies highlights the importance of describing stem morphology beyond isolated geometric measurements [15,16]. This motivates examining whether explicit cylindrical descriptors provide additional information beyond generic local shape features. Second, adjacent stem fragments should be considered in their spatial context. The constrained conditional random field developed by Polewski et al. [17] demonstrates the relevance of interactions between neighboring candidates, motivating the joint representation of local morphology and structural continuity. Third, detection performance must be interpreted in relation to acquisition conditions and reference definitions. Previous studies have reported sensitivity to vegetation conditions, laser penetration, point-cloud characteristics, and filtering decisions [12]. The effects of feature representations and spatial constraints should therefore be evaluated under clearly defined data and reference-label settings. These considerations motivate the combination of stem-shape priors and neighborhood context adopted in this study.

This study proposes a physics-guided supervoxel framework for fallen tree detection from ALS point clouds. The framework partitions the point cloud into spatially coherent supervoxels and represents each supervoxel using local geometric descriptors, random sample consensus (RANSAC)-based cylindrical descriptors, and graph-topological neighborhood features. A stacking ensemble classifier performs supervoxel-level semantic classification using point-level reference labels aggregated to the supervoxel representation. Shapley Additive exPlanations (SHAP)-based feature selection identifies informative descriptors, and graph-constrained spatial refinement improves local label consistency. The framework is trained using 30 forest plots and independently evaluated using seven held-out plots from the Actueel Hoogtebestand Nederland 4 (AHN4) dataset, with ALS point densities of 10–15 pts/m^2.

The contributions of this study are threefold. First, we formulate fallen tree detection as a supervoxel-level semantic mapping task supported by point-level reference labels. Second, we combine local geometric features, cylindrical stem descriptors, and graph-topological context to improve the separation of fallen trees from near-ground background

objects. Third, we evaluate the framework across independent forest plots and quantify the effects of feature groups, spatial refinement, supervoxel parameters, and complex forest conditions on fallen tree detection performance.

2. Materials and Methods

2.1. Study Area and ALS Data

This study used airborne laser scanning (ALS) point clouds from the Actueel Hoogtebestand Nederland 4 (AHN4), the fourth-generation national elevation dataset of the Netherlands [24]. The ALS data used in this study had operational point densities of approximately 10–15 pts/m^2. All plots had horizontal extents greater than 100 m, allowing the analysis to capture within-plot variation in near-ground forest structure and fallen tree configuration.

The experimental dataset comprised 30 development plots and seven independent test plots. The development plots were used for supervoxel-parameter selection, feature screening, SHAP-based feature selection, and stacking ensemble development. The seven test plots were spatially separate from the development plots and were excluded from all model-development procedures. They were used exclusively for final independent performance evaluation.

The seven test plots are distributed across the Netherlands (Figure 1), and representative ALS point-cloud visualizations are shown in Figure 2. The plots are denoted as Plot 1 to Plot 7 throughout the manuscript. Their locations, stand characteristics, dominant species, and site conditions relevant to fallen tree detection are summarized in Table 1.

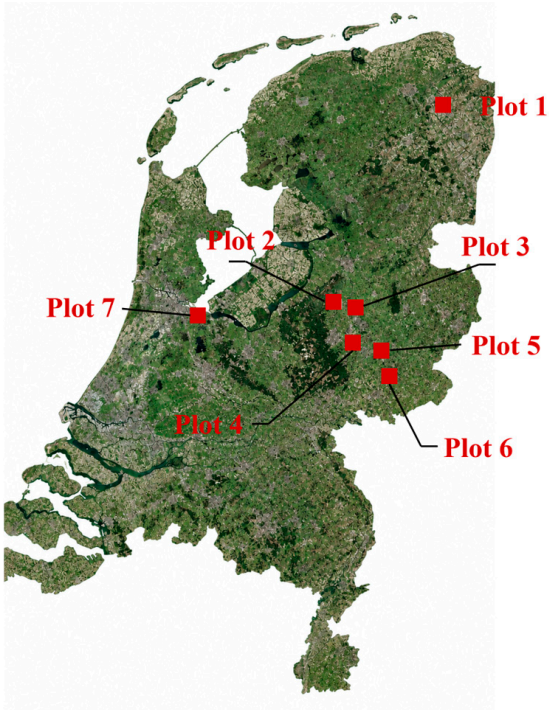
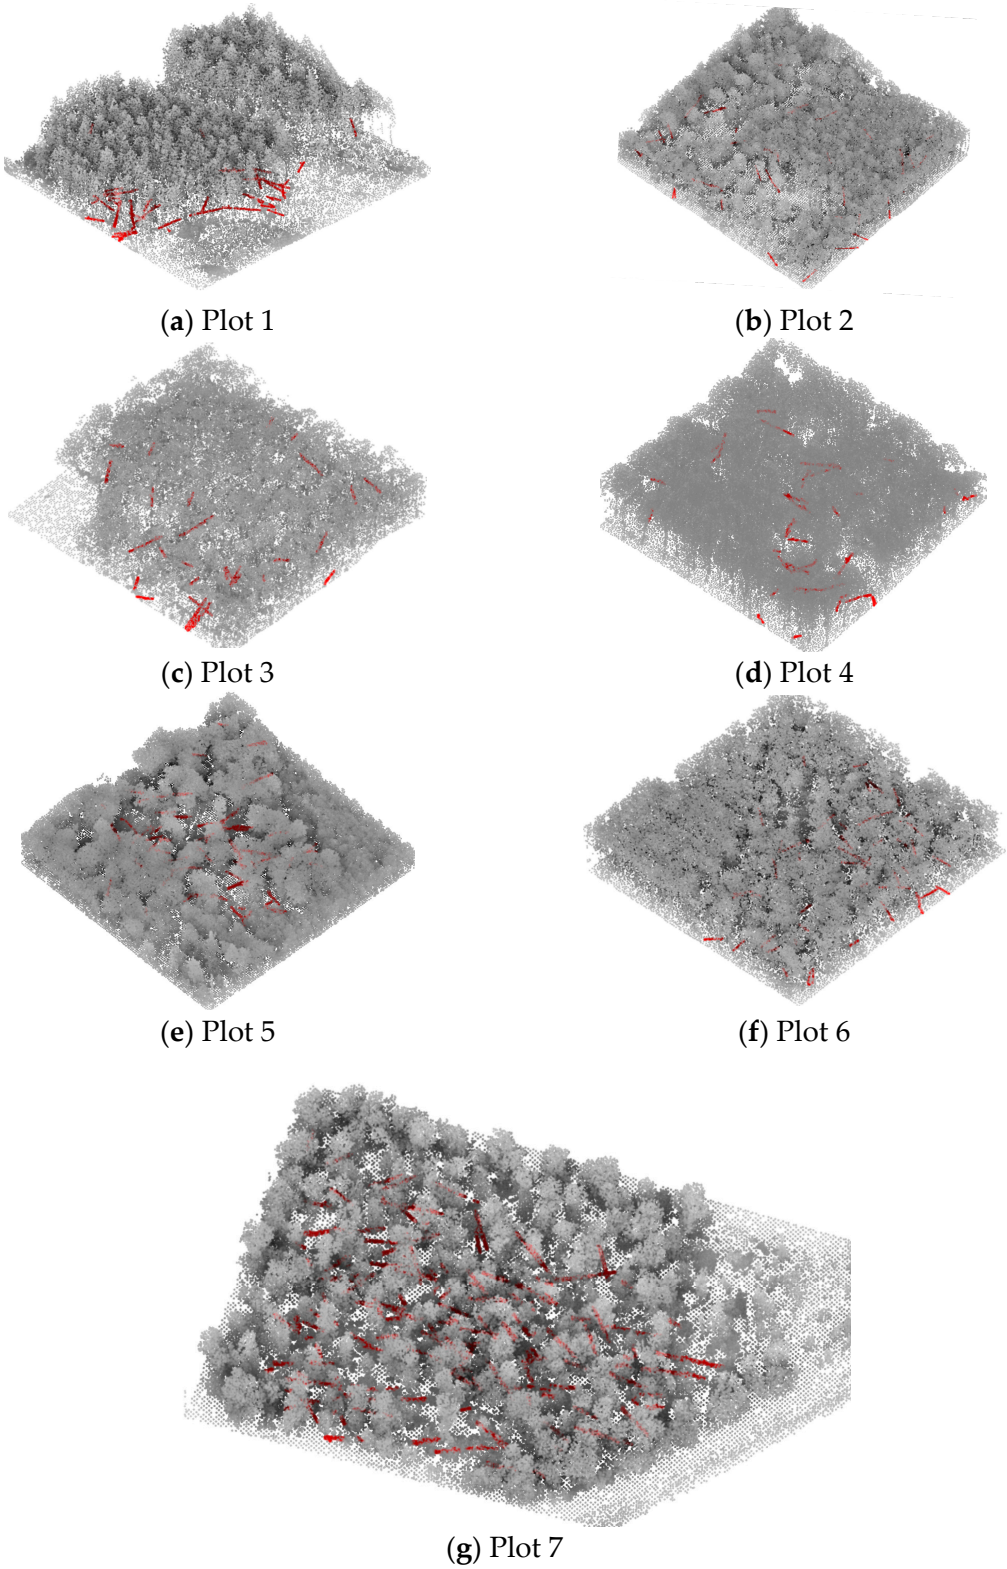


Figure 1. Locations of the seven spatially independent ALS test plots across the Netherlands.

The test plots span coniferous plantations, mixed forests, old-growth broadleaf stands, a forest-heathland ecotone, and a polder plantation. They represent a range of near-ground structural conditions, including dense understory vegetation, variable decay stages, windthrow disturbance, irregular stem morphology, upturned root plates, pit-and-mound microtopography, low-hanging live branches, and densely overlapping fallen stems. This variation provides a heterogeneous evaluation setting for examining fallen tree detection performance under contrasting forest structures and sources of near-ground geometric ambiguity.



(**a**) Plot 1 (**b**) Plot 2

(**c**) Plot 3 (**d**) Plot 4

(**e**) Plot 5 (**f**) Plot 6

(**g**) Plot 7

Figure 2. Visualization of the ALS point cloud datasets for the seven experimental plots. Red points represent fallen trees, while gray points represent other objects.

https://doi.org/10.3390/rs18203558

Table 1. Characteristics of the seven independent ALS test plots, including site conditions relevant to fallen tree detection.

Plot ID	Location/Name	Forest Type	Dominant Species	Site Conditions Relevant to Detection
Plot 1	Drents-Friese Wold (Drenthe/Friesland)	Coniferous plantation	Mixed conifers	Dense shrub understory and severe sub-canopy occlusion
Plot 2	De Hoge Veluwe	Mixed coniferous-broadleaf	Conifers and broadleaf	Heterogeneous decay stages and windthrow disturbance
Plot 3	Leuvenumse Bos	Old-growth broadleaf	Sessile oak, European beech	Irregular non-cylindrical stem morphologies
Plot 4	Garderen	Coniferous plantation	Douglas fir, European larch	Pathological stem breakage with large upturned root plates
Plot 5	Kroondomein Het Loo	Mixed uneven-aged	Highly mixed species	High geometric heterogeneity and wildlife-disturbed ground
Plot 6	Hoog Soeren	Ecotone (forest–heathland transition)	Mixed broadleaf and conifer	Low-hanging live branches and absence of humus layer
Plot 7	Pampushout, Almere	Polder plantation	Poplar (*Populus* spp.)	Severely overlapping and stacked fallen tree configurations

2.2. Point-Level Fallen Tree Reference Labels

Reference labels were manually delineated at the point level in the ALS point clouds. The fallen tree class included points belonging to visually identifiable fallen stem structures. Root plates, crown fragments, shrubs, low vegetation, and terrain-related objects were assigned to the other class unless they formed part of a discernible fallen stem. This definition was applied consistently to the training and test plots.

Point-level labels were used to support both model development and independent accuracy assessment. For supervoxel-level classification, the point-level labels of all constituent points were aggregated by majority voting. A supervoxel was assigned the fallen tree label when more than half of its constituent points were labelled as fallen tree; otherwise, it was assigned to the other class. During evaluation, the predicted label of each supervoxel was transferred to its constituent points and compared with the point-level reference labels. This design preserves point-level reference information while allowing the model to operate on spatially coherent supervoxel units.

2.3. Overview of the Proposed Framework

This study proposes a physics-guided supervoxel framework for fallen tree semantic mapping from ALS point clouds. The framework treats fallen tree detection as a supervoxel-level semantic classification task supported by point-level reference labels. It incorporates structural priors associated with fallen stems, including elongated geometry, approximate cylindrical morphology, and spatial continuity among neighboring point-cloud segments.

As illustrated in Figure 3, the workflow consists of six stages. First, ALS point clouds are preprocessed through ground filtering and height normalization to obtain near-ground candidate points. Second, the candidate points are partitioned into spatially coherent supervoxels, which are further organized into an adjacency graph and locally convex connected patches (LCCPs). Third, local geometric descriptors, RANSAC-based cylindrical descriptors, and spatial and graph-topological descriptors are extracted for each supervoxel. Fourth, SHAP-based feature selection and a stacking ensemble classifier are used to

generate initial supervoxel labels. Fifth, graph-constrained spatial refinement improves local label consistency using region-level and axis-consistency constraints. Finally, the refined predictions are evaluated against point-level reference labels using quantitative accuracy metrics and visual inspection.

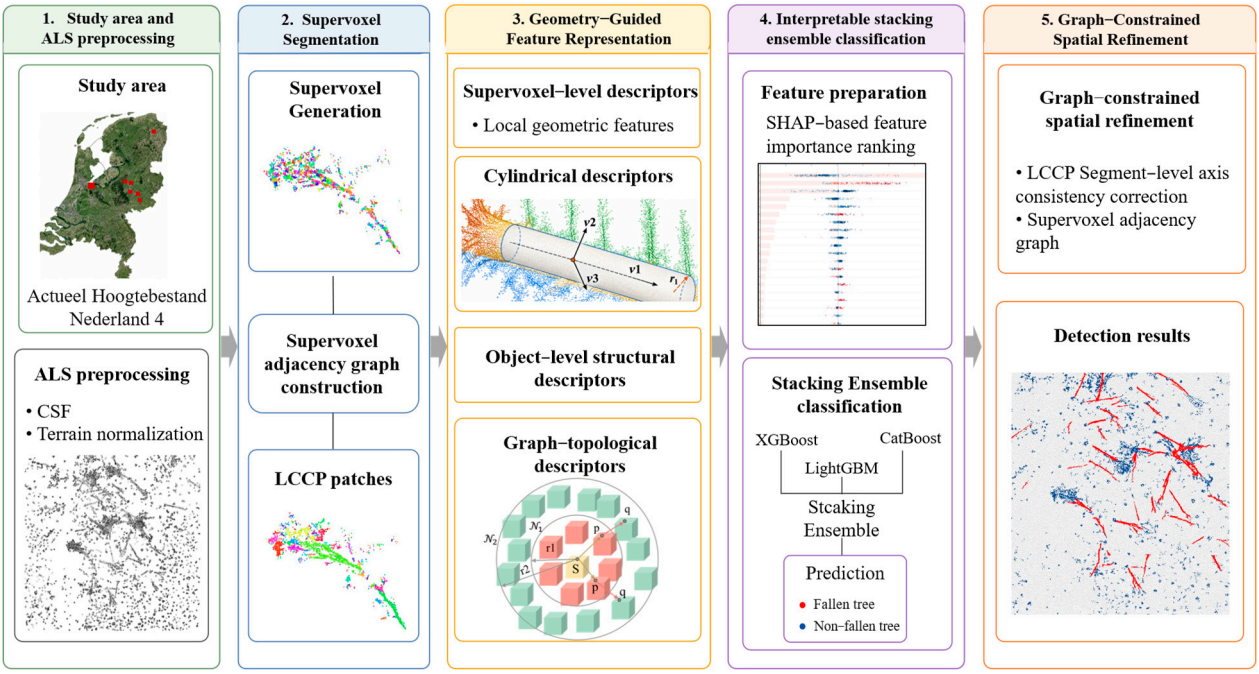


Figure 3. Overview of the proposed method.

The subsequent sections describe ALS preprocessing and supervoxel construction (Section 2.4), physics-guided feature representation (Section 2.5), stacking ensemble classification and feature selection (Section 2.6), graph-constrained spatial refinement (Section 2.7), and accuracy assessment (Section 2.8).

2.4. ALS Preprocessing and Supervoxel Construction

2.4.1. ALS Preprocessing and Terrain Normalization

All ALS point clouds were processed using an identical preprocessing workflow to generate near-ground candidate points. Ground and non-ground points were first separated using the Cloth Simulation Filtering (CSF) algorithm [25]. A digital terrain model (DTM) was then reconstructed from the ground points and used to normalize point elevations. Non-ground points were subsequently filtered according to height above ground, and points between 0 and 1.2 m were retained as the near-ground candidate set. This height threshold was used only to restrict the search space and did not directly determine the fallen tree class. The resulting candidate point clouds were used for supervoxel construction, feature extraction, and subsequent classification.

2.4.2. VCCS-Based Supervoxel Generation

Direct point-level interpretation is difficult in operational ALS data because fallen stems are often represented by sparse and discontinuous returns. Canopy and understory occlusion, uneven sampling, terrain variation, and geometric similarity to low vegetation can reduce the reliability of isolated-point or fixed-neighborhood representations. Therefore, the near-ground candidate point clouds were represented using spatially coherent supervoxels.

A supervoxel aggregates spatially adjacent points with similar local geometric properties and provides an intermediate representation between individual ALS points and complete fallen tree instances. This representation retains local three-dimensional structure while reducing sensitivity to irregular point sampling [26]. In this study, supervoxels served as the basic units for feature extraction, graph construction, and semantic classification. They were not assumed to correspond to individual fallen tree segments. Instead, each supervoxel was classified according to whether its structural characteristics were consistent with the fallen tree class or the other class.

The preprocessed near-ground candidate point cloud was partitioned into supervoxels using voxel cloud connectivity segmentation (VCCS) [27]. As illustrated in Figure 4, VCCS first organizes the point cloud through octree-based voxelization and initializes a set of seed points. The seeds are subsequently expanded through region growing according to a weighted distance measure that combines spatial proximity and surface-normal similarity. This process groups neighboring points with similar local geometry into compact and spatially coherent clusters.

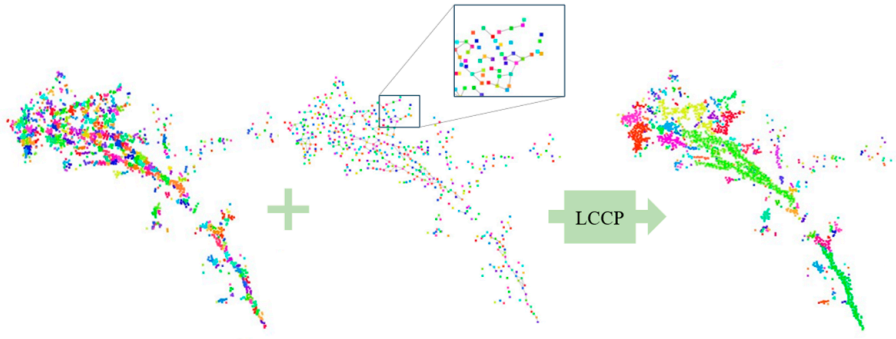


Figure 4. Schematic illustration of VCCS-based supervoxel generation and spatial adjacency graph construction for near-ground ALS point clouds.

VCCS was used as an over-segmentation procedure rather than as a direct fallen tree delineation method. The resulting supervoxels preserve local structural boundaries and provide a stable spatial representation for subsequent descriptor calculation and classification. The effects of the supervoxel parameters on fallen tree detection performance are examined in Section 3.3.

2.4.3. Supervoxel Adjacency Graph Construction

Following the supervoxel generation, a spatial adjacency graph was constructed to describe neighborhood relationships among supervoxels. Specifically, in this graph, each node represents an individual supervoxel, and each edge represents spatial adjacency between two neighboring supervoxels [28,29]. Two supervoxels are considered adjacent when they are spatially contiguous or share a local boundary in three-dimensional space. Based on this strict definition, an undirected graph model $G = (V,E)$ is formulated, where nodes V represent individual supervoxels and edges E encode their spatial adjacency relationships. Because fallen trees typically form continuous linear structures extending along a specific principal axis, the undirected graph model efficiently captures such long-range spatial dependency. This representation supports subsequent graph-based analyses, including neighborhood statistical analysis and geometry-constrained label propagation.

2.4.4. LCCP-Based Structural Patch Clustering

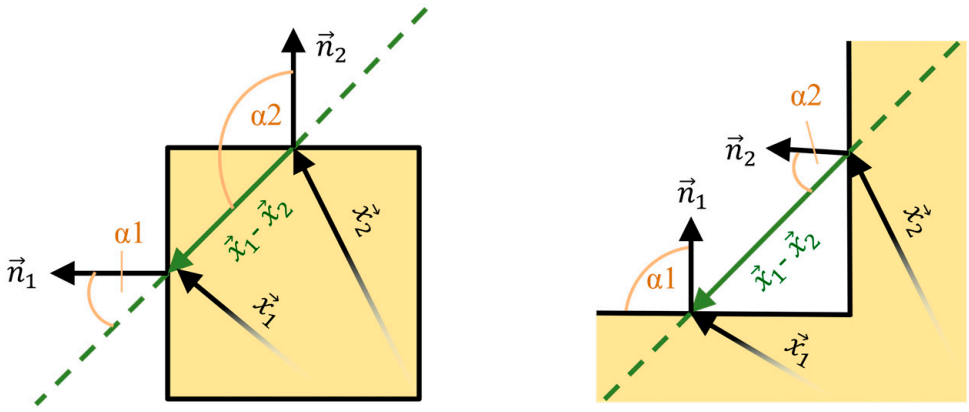
Based on the supervoxel adjacency graph, locally convex connected patches (LCCPs) are further introduced to cluster adjacent supervoxels into structurally coherent regions,

as shown in Figure 4. LCCP evaluates the convexity or concavity of boundaries between neighboring supervoxels by comparing normal vector variation and relative spatial displacement [30]. The physical discriminative mechanism of LCCPs is illustrated in Figure 5. Specifically, Figure 5a represents a convex boundary condition where surface normals diverge, prompting LCCPs to merge the adjacent supervoxels into a coherent patch. Conversely, Figure 5b illustrates a concave boundary condition characterized by converging normals, which signifies a structural discontinuity and effectively inhibits the merging process. The specific mathematical criteria for determining convex or concave boundaries are defined as:

$$\hat{d} = \frac{\vec{x}_1 - \vec{x}_2}{\left\| \vec{x}_1 - \vec{x}_2 \right\|} \tag{1}$$

$$\Delta\alpha = \vec{n_1} \cdot \hat{d} - \vec{n_2} \cdot \hat{d} \tag{2}$$

where $\vec{x}_1$ and $\vec{x}_2$ denote the centroid positions of two adjacent supervoxels, and $\vec{n_1}$ and $\vec{n_2}$ represent their corresponding estimated normal vectors. The variable $\Delta\alpha$ characterizes the boundary type between the supervoxels: a positive value ($\Delta\alpha > 0$) indicates a convex relationship, signifying that the angle between the normal vector of the first supervoxel and the directional vector connecting their centroids is acute, whereas a negative value ($\Delta\alpha < 0$) indicates a concave boundary, reflecting a structural discontinuity between the adjacent supervoxels.



(**a**) convex geometric relationship (**b**) concave geometric relationship

Figure 5. Schematic diagram of the convexity–concavity evaluation criterion.

Through this process, adjacent supervoxels were clustered into segmentation units with improved structural coherence [31]. In forest scenes, these structurally coherent patches help reduce fragmentation caused by local over-segmentation and provide a basis for subsequent region-level feature extraction and spatial refinement.

2.5. Physics-Guided Feature Representation

Building on the supervoxel representation, a multiscale morphology-aware feature framework is developed to characterize fallen woody structures from ALS point clouds. The feature design captures three complementary properties of fallen stems: local geometric organization, approximately cylindrical morphology, and larger-scale spatial structure. Descriptors are extracted at the supervoxel, LCCP structural-patch, and local raster-context scales. Principal component analysis (PCA)-based features characterize local shape, RANSAC-based descriptors quantify cylindrical stem morphology, and height-, orientation-

, and spatial-context features describe the position and extent of candidate structures. The complete feature set and its spatial scales are summarized in Table 2.

Table 2. Computation of morphology-aware structural descriptors at different spatial scales. Scale ID 1–3 correspond to the supervoxel, LCCP, and local raster-context scales, respectively.

Feature Group	Descriptor	Feature Formula	Scale ID
PCA-based shape descriptors	Linearity	$(\lambda_1 - \lambda_2)/\lambda_1$	1, 2
	Planarity	$(\lambda_2 - \lambda_3)/\lambda_1$	1, 2
	Scattering	λ_3/λ_1	1, 2
	Curvature	$\lambda_3/(\lambda_1 + \lambda_2 + \lambda_3)$	1, 2
	Anisotropy	$(\lambda_1 - \lambda_3)/\lambda_1$	1, 2
	Omnivariance	$(\lambda_1 \times \lambda_2 \times \lambda_3)^{1/3}$	1, 2
	Eigenentropy	$-\sum_{i=1}^{3} \lambda_i \ln(\lambda_i)$	1
	Verticality	$1 - \mathrm{nz}$	1, 2
RANSAC cylindrical descriptors	Cylinder radius	r	1
	Cylinder fitting residual	$\sqrt{\frac{1}{N}\sum_{i=1}^{N} d_i^2}$	1
	Cylinder-axis direction	$\mathbf{a} = (a_x, a_y, a_z)$	1
	Cylinder fitting confidence	$N_{\mathrm{inlier}}/N_{\mathrm{total}}$	1
Height, orientation, and spatial distribution descriptors	Local point density	$N_{\mathrm{points}}/V_{\mathrm{local}}$	1
	Raster density	$N_{\mathrm{points}}/A_{\mathrm{cell}}$	3
	Internal height range	$z_{\max} - z_{\min}$	1
	Relative height	$z_{\mathrm{centroid}} - z_{\mathrm{ref}}$	1
	Normal-vector angle	$\arccos(\mathbf{n} \cdot \mathbf{e}_z)$	1
	Raster height range	$z_{\max^r} - z_{\min^r}$	3
	Raster elevation standard deviation	$\sigma(z^r)$	3
	2D eigenvalue ratio	$\lambda_{2,2D}/\lambda_{1,2D}$	1
	Major-axis length	$\max(p_i) - \min(p_i)$	2
	Aspect ratio	λ_1/λ_3	2
	Axis horizontality	$\sqrt{a_x^2 + a_y^2}/\lvert\mathbf{a}\rvert$	1

2.5.1. PCA-Based Local Geometric Descriptors

To describe the local structure of each supervoxel, basic spatial statistics and PCA-based descriptors were extracted. Basic descriptors included local point density, raster density, internal height range, relative height, normal vector angle, raster height difference, and raster elevation standard deviation. For each supervoxel, 3D PCA was applied to the covariance matrix of its internal points. The resulting eigenvalues were used to calculate shape descriptors, including linearity, planarity, scattering, curvature, anisotropy, omnivariance, eigenentropy, and verticality, as detailed in Table 2. These descriptors characterize whether a supervoxel is more linear, planar, or volumetric in local geometry. Furthermore, horizontal 2D eigenvalue ratios and regional raster elevation variations are incorporated to suppress extreme terrain undulations.

2.5.2. RANSAC-Based Cylindrical Descriptors

Fallen trees often approximate finite-length cylindrical structures. This morphological characteristic provides a physically meaningful prior for distinguishing fallen stems from

shrubs, terrain surfaces, and irregular background objects. To encode this prior, RANSAC-based [32] cylinder fitting is applied to the internal point cloud of each supervoxel, as illustrated in Figure 6.

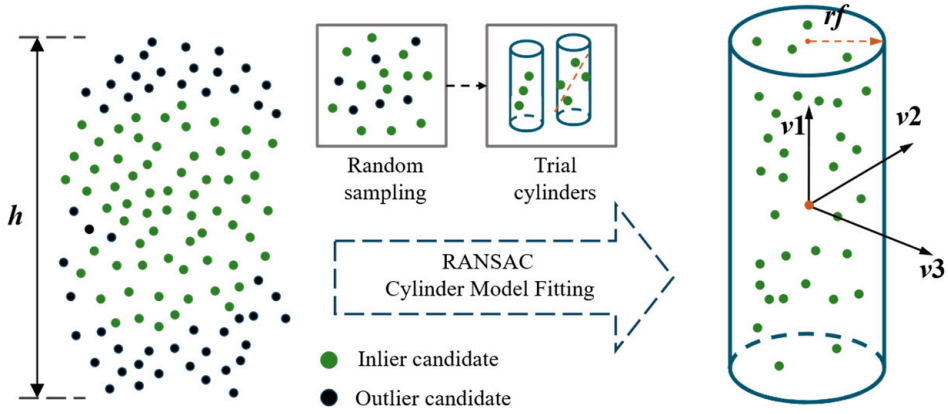


Figure 6. Schematic illustration of the RANSAC-based cylindrical descriptors used to characterize fallen tree morphology.

Four cylindrical descriptors are derived from the fitted model: cylinder radius, fitting residual, axis direction, and fitting score. The fitting radius represents the estimated radius of the optimal cylinder. The fitting residual was calculated as the root-mean-square error between the internal points and the fitted cylindrical surface. The axis direction describes the three-dimensional orientation of the fitted cylinder. The fitting score was calculated as the proportion of inlier points and was used to evaluate the reliability of the cylindrical structure.

2.5.3. Structural and Spatial-Context Descriptors

Additional descriptors were calculated to represent the position, spatial extent, and local context of each candidate structure. At the supervoxel scale, these features included local point density, internal height range, relative height, normal-vector angle, and the two-dimensional eigenvalue ratio. Local raster-context descriptors included raster point density, raster height range, and raster elevation standard deviation. These variables describe near-ground position, horizontal elongation, local sampling density, and elevation variability in the surrounding area.

To represent structure beyond individual supervoxels, LCCPs were further characterized using patch-level point count, PCA-based shape descriptors, major-axis length, and aspect ratio. Axis horizontality was also calculated to describe the orientation of the dominant structural direction. These descriptors complement local geometry and cylindrical fitting by representing the spatial extent and environmental context of candidate fallen tree structures. The class-wise distributions of representative descriptors are examined in Section 3.1.

2.5.4. Graph-Topological Descriptors

Spatial continuity is an important structural characteristic of fallen tree structures. Therefore, relying only on isolated supervoxel features may be insufficient for distinguishing fallen stems from shrubs, branches, and terrain fragments. To incorporate neighborhood context, we propose the graph-based topological descriptors that are derived from the supervoxel adjacency graph.

For a given supervoxel SV_i, its first-order neighborhood $\mathcal{N}_1$ is defined as $\mathcal{N}_1(i) = \{SV_j | (i,j) \in E\}$ and the second-order neighborhood $\mathcal{N}_2$ is defined as $\mathcal{N}_2(i) = \bigcup_{j \in \mathcal{N}_1(i)} \mathcal{N}_1(j) \setminus \{i\}$.

These descriptors complement the morphology-aware features by introducing contextual information from neighboring structural units, as illustrated in Figure 7. Specifically, statistical aggregations are computed over these neighborhood sets for key geometric properties, including linearity, planarity, scattering, curvature, anisotropy, the sum of eigenvalues, eigenentropy, and verticality. Additionally, the maximum linearity, mean normal angle, and mean elevation difference within each neighborhood level are incorporated to capture geometric consistency.

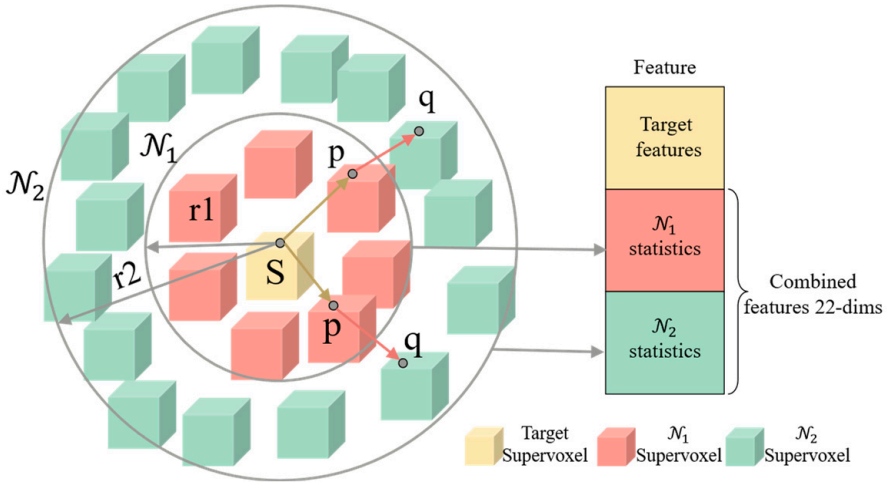


Figure 7. Schematic illustration of first-order and second-order neighborhoods in the supervoxel adjacency graph.

Beyond simple statistical aggregation, graph-based topological continuity is explicitly measured. High-linearity connectivity count is defined as the total number of reachable supervoxels exhibiting a linearity value exceeding 0.5. This is determined by traversing the adjacency graph via a Breadth-First Search (BFS) initiated from the target supervoxel. This descriptor captures the longitudinal continuity of fallen stems, which is difficult to represent using isolated geometric features. This topological feature critically characterizes the spatially continuous linear chains formed by fallen trees. Whereas an isolated linear supervoxel may correspond to a branch tip or noise, a supervoxel associated with a high chain-length L_{chain} value strongly indicates membership in an intact fallen tree stem.

2.6. Stacking Ensemble Classification and SHAP-Based Feature Selection

2.6.1. Base Learners and Class-Imbalance Handling

The descriptors described in Section 2.5 were used as inputs for supervoxel-level fallen tree classification. A stacking ensemble was constructed from three gradient-boosting base learners: XGBoost [33], LightGBM [34], and CatBoost [35]. These models were selected because they provide complementary tree-based learning strategies for structured feature data.

Fallen tree supervoxels represented a minority of the training samples. To address this imbalance, class weighting was applied during base-learner training. For XGBoost and LightGBM, the background-class weight was fixed at 1, and the positive-class weight was specified using the "scale_pos_weight" parameter and calculated as $w_{\text{fallen}} = 1 + 0.5\left(\frac{N_{\text{background}}}{N_{\text{fallen}}} - 1\right)$. Here, $N_{\text{background}}$ and N_{fallen} denote the numbers of background and fallen tree supervoxels in the training set, respectively. The coefficient 0.5 provided partial compensation for class imbalance. For CatBoost, automatic class weighting was enabled by setting the "auto_class_weights" parameter to "Balanced". The resulting base-learner probability estimates were subsequently combined by the meta-classifier.

Random Forest (RF) [36] was included as an independent baseline to complement the gradient-boosting classifiers. RF combines decision trees trained on bootstrap samples with random feature selection, providing a bagging-based comparison for evaluating the stacking framework. It was used only as a standalone baseline and was not included among the three base learners of the stacking architecture.

2.6.2. Stacking Ensemble Training

Stacking (Stacked Generalization) is a hierarchical ensemble strategy in which a meta-learner is trained using predictions generated by multiple learners [37]. In this study, the first layer independently trained XGBoost, LightGBM, and CatBoost on the supervoxel feature set and generated class-probability outputs. These classifiers were selected because they represent complementary tree-based ensemble learning strategies widely used for structured tabular data. The second layer used these probability outputs as input features for a logistic regression meta-learner, denoted by M, which estimated the fallen tree probability of each supervoxel. The resulting probability was then thresholded to obtain the initial class label.

The model-development dataset comprises 30 training plots, whereas the seven evaluation plots are retained as an independent test set and are not involved in model training or feature selection. Let the training set be defined as $\{(x_i, y_i)\}_{i=1}^{N}$, where x_i denotes the supervoxel feature vector and y_i the ground truth label. To prevent data leakage, a k-fold out-of-fold strategy is employed during the base-layer training. The output prediction of the k-th base learner in the first layer can be expressed as:

$$h_{\mathrm{j}}(x_i), \mathrm{j} = 1, 2, 3 \tag{3}$$

The concatenated input vector to the second-layer meta-learner is then constructed as:

$$\mathbf{z}_i = [h_1(x_i), h_2(x_i), h_3(x_i)] \tag{4}$$

The logistic regression meta-learner combines these inputs to estimate the fallen tree probability as follows:

$$\hat{y}_i = M(\mathbf{z}_i) = \frac{1}{1 + \exp\left[-\left(\beta_0 + \boldsymbol{\beta}^{\mathrm{T}}\mathbf{z}_i\right)\right]} \tag{5}$$

where M denotes the logistic regression meta-learner, β_0 is the intercept, and $\boldsymbol{\beta}$ is the vector of coefficients associated with the three base-learner probability outputs. The quantity $\hat{y}_i$ represents the predicted probability that supervoxel i belongs to the fallen tree class. A threshold of 0.45 was applied to obtain the initial class label, as illustrated in Figure 8.

This out-of-fold training strategy ensured that the inputs to the meta-learner were generated from samples not used to train the corresponding base learners. In addition, the spatial separation between the training data and the seven test plots provided an independent basis for evaluating cross-plot generalization. The detailed implementation workflow is illustrated in Figure 8.

2.6.3. SHAP-Based Feature Selection

The morphological descriptors derived from PCA and cylindrical characterization were integrated to construct a comprehensive representation of fallen woody structures. Because different structural descriptors contribute differently to classification performance, feature importance analysis was performed to identify the dominant variables associated with fallen tree recognition.

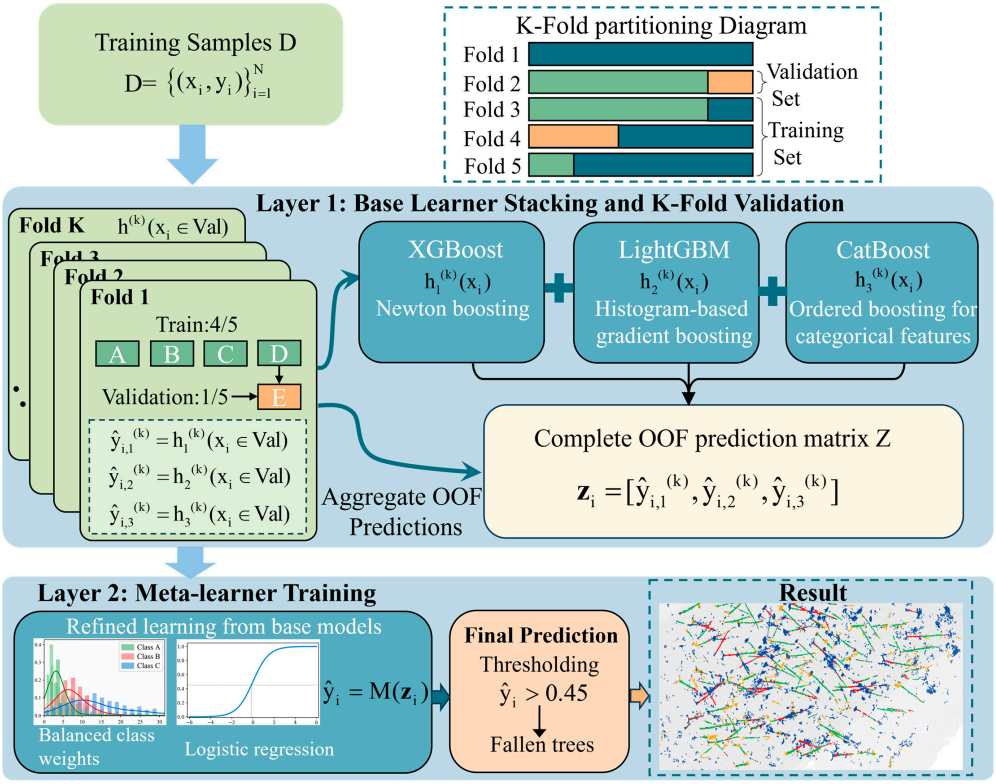


Figure 8. Schematic diagram of the stacking ensemble framework.

To quantify the contribution of each feature to supervoxel classification and improve model interpretability, SHAP was used to interpret the contribution of individual structural descriptors to model predictions and reveal the structural characteristics associated with fallen tree occurrence. Grounded in cooperative game theory, the SHAP method calculates the marginal contribution of each feature across all possible permutations, thereby providing a unified, consistent model-interpretation framework for a refined feature subset. For a model utilizing a complete feature set F, the SHAP value of the *i-th* feature is defined as:

$$\phi_i = \sum_{S \subseteq \mathcal{F} \setminus \{i\}} \frac{|S|!(n-|S|-1)!}{n!} [f(S \cup i) - f(S)] \quad (6)$$

where S denotes any feature subset of F that excludes feature, n = $|F|$ is the total number of features, $|S|$ denotes the cardinality of subset S, $f(S)$ denotes the model output using feature subset S, and $f(S \cup i)$ denotes the model output when feature i is additionally included. The weighting term $\frac{|S|!(n-|S|-1)!}{n!}$ accounts for all possible orderings in which feature i can be introduced into the subset.

SHAP values were calculated using the 30 model-development plots only. Global importance was quantified as the mean absolute SHAP value across all training supervoxels. The descriptors were ranked according to this measure, and the 20 highest-ranked descriptors were retained as the final feature subset. The stacking pipeline was then refitted using the selected features before independent evaluation on the seven held-out test plots.

2.7. Graph-Constrained Spatial Refinement

Although the stacking ensemble learning model effectively integrates multidimensional features for initial supervoxel classification, localized misclassifications inevitably occur in complex forest scenes due to highly heterogeneous land covers, overlapping local features, and noise interference. These errors are common near transitional boundaries among fallen trees, ground surfaces, and low-lying vegetation, where feature-driven clas-

sification struggles to guarantee spatial consistency. To address this, a graph-constrained post-processing is introduced to refine the supervoxel labels and improve spatial label consistency. Unlike conventional spatial smoothing, the proposed refinement considers structural connectivity among neighboring supervoxels rather than only spatial proximity.

By incorporating the supervoxel adjacency graph and explicit geometric constraints, this post-processing stage refines the initial predictions, thereby advancing from local supervoxel-level discrimination towards global structural consistency.

During the supervoxel segmentation phase, an undirected graph structure, denoted as $G = (V, E)$, is established. Here, nodes V represent individual supervoxels and edges E denote their spatial adjacency relationships. Based on this topological foundation, spatially adjacent supervoxels belonging to the same physical object are assumed to exhibit high categorical consistency.

First, a localized neighborhood voting mechanism is implemented to iteratively update the classification labels. Let the initial predicted label of a target supervoxel be y_i, and its corresponding spatial neighborhood set be denoted by $\mathcal{N}_i$. The updated class label is formulated as follows:

$$y_i = \underset{c \in \mathcal{C}}{\operatorname{argmax}} \sum_{j \in \mathcal{N}_i} \mathbb{I}(y_j = c) \tag{7}$$

where $\mathcal{C}$ = {fallen tree, other} is the set of candidate classes, c denotes a candidate class, and $\mathcal{N}_i$ is the set of supervoxels adjacent to supervoxel i. The label of neighboring supervoxel j is denoted by y_i. The indicator function $\mathbb{I}(y_j = c)$ equals 1 when $y_j = c$ and 0 otherwise. The summation therefore counts the neighbors assigned to each class, and y_i is the class receiving the largest number of votes. This local voting operation promotes label consistency among neighboring supervoxels.

Second, region-level consistency correction is applied using LCCP. Contiguous supervoxels with convex structural continuity are aggregated into unified region clusters S_k. At this regional scale, categorical consistency is strictly enforced via majority voting:

$$\hat{y}_k = \underset{c \in \mathcal{C}}{\operatorname{argmax}} \sum_{SV_i \in S_k} \mathbb{I}(y_i = c) \tag{8}$$

The resulting majority label was assigned to the supervoxels belonging to the same LCCP.

Finally, graph-based label propagation was applied to restore continuity among fragmented fallen tree supervoxels. High-confidence fallen tree supervoxels were initialized as seed nodes. Breadth-first search was then performed on the adjacency graph to identify neighboring candidates that satisfied the predefined geometric-consistency constraints. These constraints included thresholds on local linearity and the angular deviation between the principal axes of adjacent supervoxels. Labels were propagated only along graph paths satisfying these constraints, thereby limiting propagation into adjacent shrubs, terrain fragments, and other non-target structures.

The refinement procedure produced the final supervoxel labels used for fallen tree semantic mapping. Its effect on classification accuracy is evaluated in Section 3.5.

2.8. Accuracy Assessment

To evaluate the performance of the proposed method for the fallen tree detection task, four standard evaluation metrics are used: Accuracy, Precision, Recall, and F1-score [23]. These metrics are calculated as follows:

$$\begin{cases} Accuracy = \frac{TP+TN}{TP+TN+FP+FN} \\ Precision = \frac{TP}{TP+FP} \\ Recall = \frac{TP}{TP+FN} \\ F_1 score = 2 \times \frac{Precision \times Recall}{Precision+Recall} = \frac{2TP}{2TP+FP+FN} \end{cases} \quad (9)$$

where TP, FP, TN, and FN denote the number of true positives, false positives, true negatives, and false negatives, respectively.

Class-specific metrics were calculated for both the fallen tree and other classes. The fallen tree F1 score was treated as the primary metric because fallen tree points represented the minority class. Quantitative evaluation was complemented by visual inspection of representative test plots with contrasting forest structures and error patterns.

3. Results

3.1. Feature Separability and Feature Importance

To examine whether the physics-guided descriptors capture distinguishable structural characteristics of fallen trees, representative supervoxels corresponding to fallen trees, shrubs, and ground surfaces are first compared in the feature space. As shown in Figure 9a,b, fallen tree supervoxels are predominantly distributed toward high linearity and high axis horizontality, whereas shrub supervoxels exhibit a broader distribution across both dimensions. This contrast is consistent with the elongated and predominantly horizontal morphology of fallen stems, while the irregular branching structure of shrubs produces greater geometric variability.

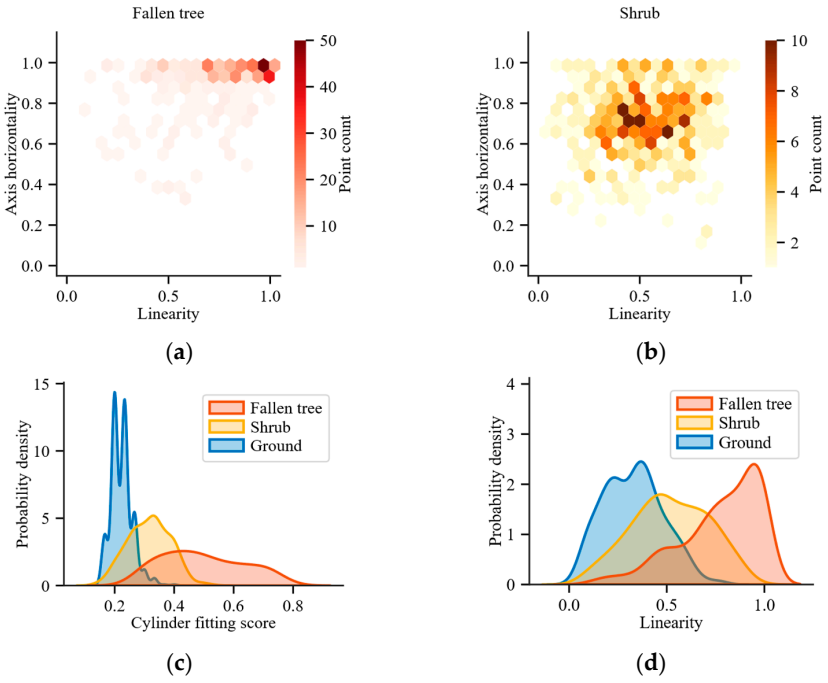


Figure 9. Distributional characteristics of supervoxel features for different object types proximate to the ground surface. (**a**) Hexbin density distribution of fallen trees in the joint feature space of linearity and axis horizontality. (**b**) Hexbin density distribution of shrubs in the same joint feature space. (**c**) Probability density distribution of the cylinder fitting score for the three land cover categories. (**d**) Probability density distribution of linearity for the three categories.

The univariate distributions further illustrate differences among the three object classes. As shown in Figure 9c, fallen tree supervoxels generally exhibit higher cylinder fitting scores than ground supervoxels, whereas shrubs occupy an intermediate and partially

overlapping range. Similarly, Figure 9d shows that the linearity distribution of fallen trees is concentrated toward high values, while shrubs and ground surfaces exhibit broader distributions at lower values. Nevertheless, overlap remains among the classes in parts of the feature space, indicating that individual geometric descriptors alone are insufficient for complete separation. These results therefore support the combined use of complementary morphological and contextual descriptors for fallen tree classification.

Prior to SHAP-based ranking, univariate AUC screening removed nine low-information variables, leaving 51 candidate descriptors. SHAP analysis of the final stacking pipeline ranked the descriptors according to their mean absolute contribution to the predicted probability of the fallen tree class (Figure 10). LCCP major-axis length, patch linearity, and patch point count were the three most influential descriptors, accounting for 27.8%, 18.3%, and 13.7% of the total mean absolute SHAP importance, respectively. The SHAP distributions further showed how low and high feature values contributed differently to individual supervoxel predictions. These results indicate that patch-scale structural extent, local linearity, and point-cloud support were dominant components of the final classification decisions.

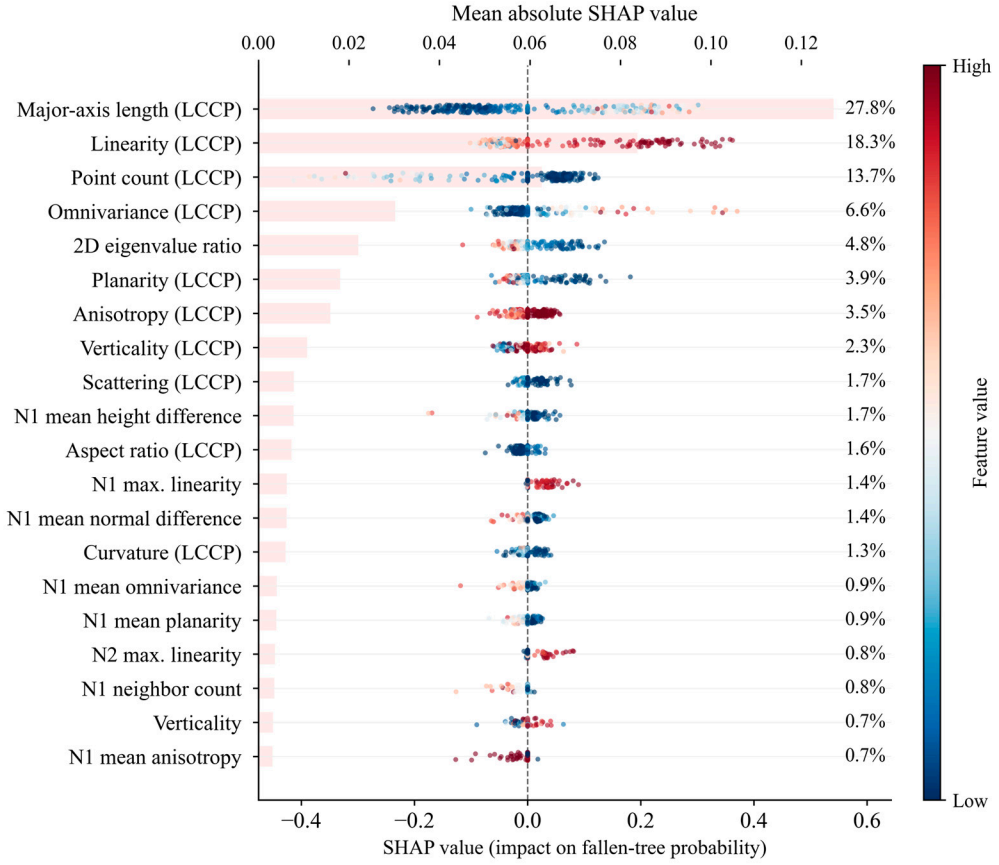


Figure 10. Importance ranking of the top 20 features based on SHAP values.

Feature subsets containing 5, 10, 20, 30, 40, and 51 descriptors were evaluated using out-of-fold validation on the model-development data. The OOF fallen tree F1 score increased as the number of retained descriptors increased from 5 to 20 and changed only marginally when additional descriptors were included (Figure 11). The top 20 descriptors were therefore retained as the final feature subset. Although Plot 1 and Plot 7 are shown for descriptive comparison, feature dimensionality was determined exclusively from OOF validation on the 30 training plots; the independent test plots were not used for feature selection.

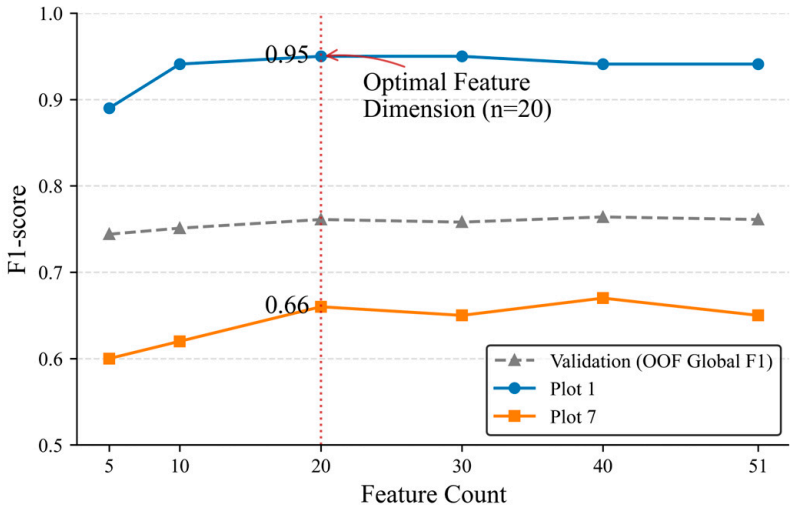


Figure 11. Variation in F1-score across different numbers of retained features. The OOF validation curve is used to determine the feature dimensionality, while Plot 1 and Plot 7 are shown for descriptive comparison.

Figure 12 summarizes pairwise interactions among the 20 retained descriptors, complementing the global importance ranking in Figure 10. Nodes represent features, with their size and color indicating global SHAP importance, while edges represent pairwise interactions, with their thickness and color indicating interaction strength. The network provides a feature-level overview of the relationships among patch-scale geometry and neighborhood descriptors, rather than showing their contributions to individual predictions.

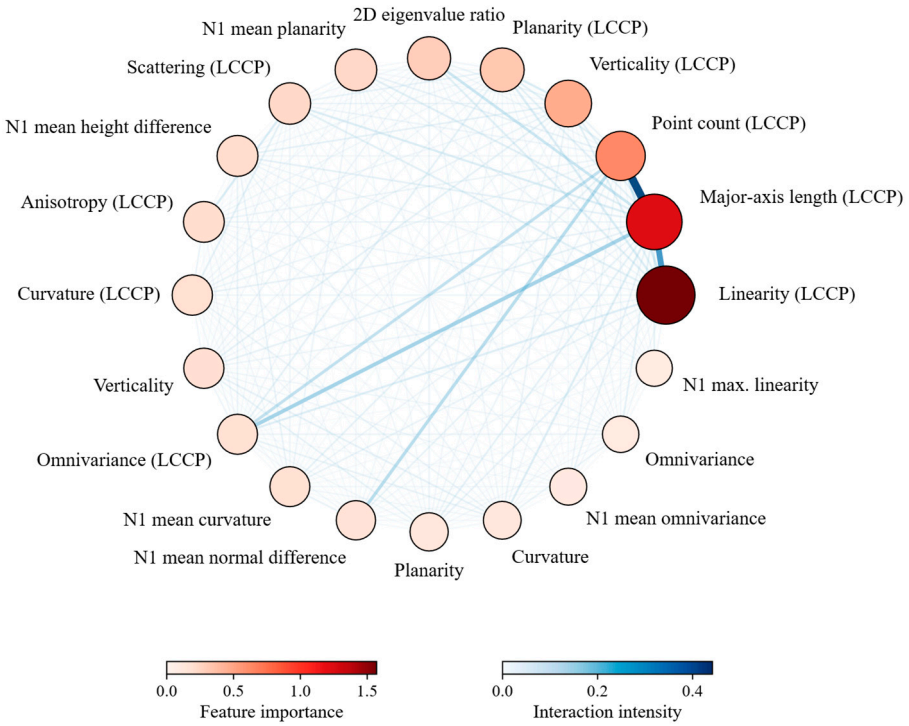


Figure 12. Feature-level SHAP interaction network of the 20 retained descriptors. Node size and color indicate global importance; edge thickness and color indicate pairwise interaction strength.

Figure 13 presents the distribution of SHAP values across all supervoxels for the final stacking pipeline. Supervoxels are ordered along the horizontal axis according to the model output $f(x)$, from high values on the left to low values on the right. Features are arranged

vertically by global importance. Each cell represents the SHAP contribution of one feature to the prediction for one supervoxel: red indicates a positive contribution to the model output, whereas blue indicates a negative contribution. The marginal bars on the right summarize the global importance of the corresponding features.

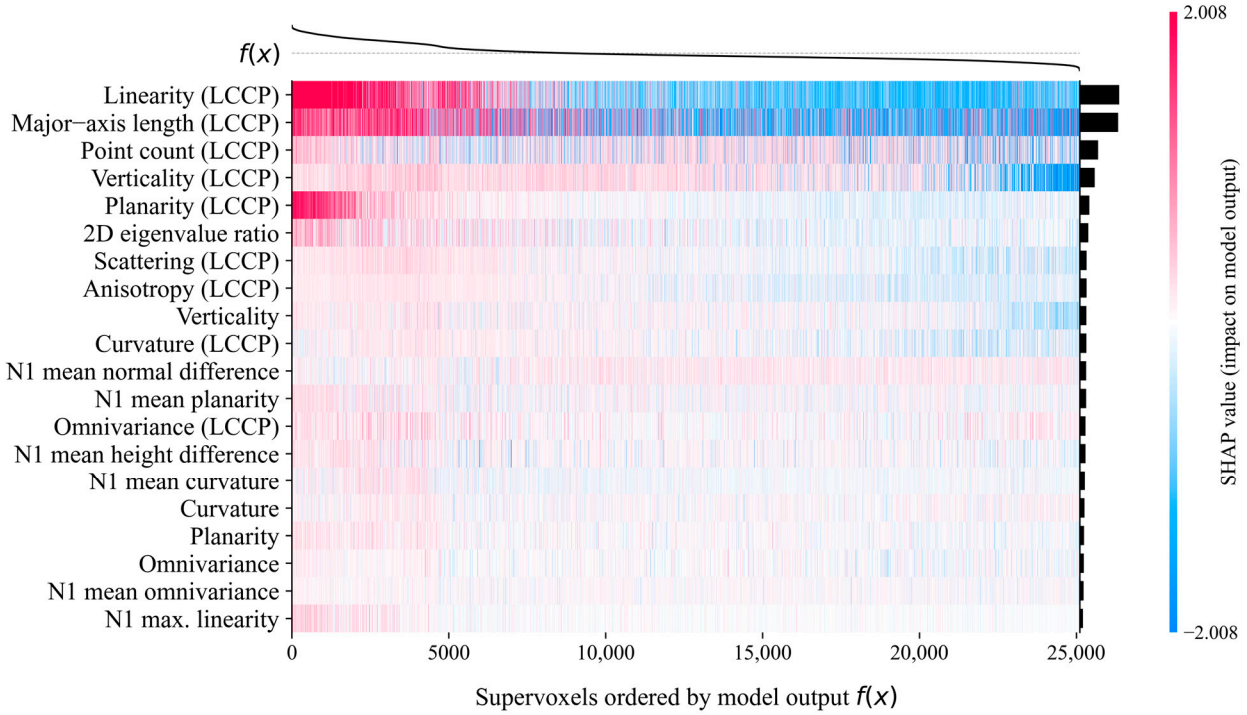


Figure 13. Sample-level SHAP contribution heatmap for the retained descriptors. Columns represent supervoxels ordered by model output, and rows represent features. Red and blue indicate positive and negative contributions, respectively.

The heatmap shows that the highest-ranked LCCP-derived descriptors, particularly Linearity (LCCP), Major-axis length (LCCP), and Point count (LCCP), contributed most strongly to the separation between supervoxels with high and low fallen tree model outputs. Supervoxels at the high-output end were characterized mainly by positive contributions from these descriptors, whereas their contributions became predominantly negative toward the low-output end. This pattern is consistent with the importance of elongated and spatially extended structural patches for identifying fallen trees. Other LCCP-derived descriptors, including Verticality (LCCP), Planarity (LCCP), the 2D eigenvalue ratio, Scattering (LCCP), and Anisotropy (LCCP), showed more moderate but spatially coherent contributions across the ordered supervoxels. Neighborhood descriptors, such as N1 mean normal difference, N1 mean planarity, N1 mean height difference, and N1 max linearity, were ranked lower overall but still contributed to individual predictions, indicating that local topological context complemented patch-level geometric information.

Figures 12 and 13 provide complementary interpretations of the retained feature set. Figure 12 summarizes pairwise interactions among descriptors at the feature level, whereas Figure 13 shows the direction and variation in their contributions across individual supervoxels. Together, these analyses indicate that classification decisions were driven primarily by LCCP-derived structural descriptors, with neighborhood features providing additional contextual support. They complement the global importance ranking in Figure 10 and help interpret the 20-feature subset selected through the OOF validation presented in Figure 11.

3.2. Fallen Tree Detection Accuracy Across Test Plots

To validate the effectiveness of the proposed stacking ensemble, its predictive performance was systematically compared against four baseline classifiers, namely CatBoost, LightGBM, XGBoost, and RF, using the optimal 20-dimensional feature subset. The quantitative results across the seven test plots are summarized in Table 3.

Table 3. Performance comparison of different classifiers across the experimental plots.

Dataset		Stacking		CatBoost		LightGBM		XGBoost		RF	
		Other	Fallen Tree	Other	Fallen Tree	Other	Fallen Tree	Other	Fallen Tree	Other	Fallen Tree
Plot1	Pre	0.96	0.88	0.92	0.83	0.89	0.84	0.9	0.85	0.9	0.77
	Rec	0.95	0.9	0.93	0.81	0.94	0.75	0.94	0.76	0.89	0.77
	F1	0.95	0.89	0.92	0.82	0.91	0.79	0.92	0.8	0.9	0.77
	Acc	0.93		0.89		0.88		0.89		0.86	
Plot2	Pre	0.95	0.77	0.92	0.75	0.9	0.78	0.91	0.79	0.91	0.76
	Rec	0.86	0.91	0.88	0.85	0.88	0.81	0.89	0.82	0.87	0.83
	F1	0.9	0.83	0.89	0.8	0.89	0.79	0.9	0.81	0.89	0.79
	Acc	0.88		0.86		0.86		0.87		0.86	
Plot3	Pre	0.94	0.9	0.93	0.93	0.9	0.94	0.91	0.93	0.92	0.87
	Rec	0.93	0.92	0.95	0.89	0.96	0.84	0.96	0.85	0.92	0.87
	F1	0.94	0.9	0.94	0.91	0.93	0.89	0.93	0.89	0.92	0.87
	Acc	0.92		0.93		0.92		0.92		0.9	
Plot4	Pre	0.95	0.85	0.95	0.85	0.91	0.87	0.91	0.88	0.95	0.86
	Rec	0.92	0.93	0.9	0.93	0.92	0.86	0.93	0.85	0.91	0.92
	F1	0.93	0.89	0.93	0.89	0.92	0.87	0.92	0.87	0.93	0.89
	Acc	0.91		0.91		0.9		0.9		0.91	
Plot5	Pre	0.99	0.72	0.9	0.71	0.89	0.75	0.9	0.77	0.91	0.7
	Rec	0.88	0.98	0.91	0.69	0.93	0.64	0.94	0.66	0.91	0.71
	F1	0.93	0.83	0.91	0.7	0.91	0.69	0.92	0.71	0.91	0.71
	Acc	0.9		0.86		0.86		0.87		0.86	
Plot6	Pre	0.96	0.82	0.93	0.81	0.91	0.83	0.91	0.83	0.94	0.83
	Rec	0.9	0.92	0.9	0.87	0.92	0.81	0.92	0.82	0.91	0.88
	F1	0.93	0.87	0.91	0.84	0.91	0.82	0.92	0.83	0.92	0.85
	Acc	0.91		0.89		0.88		0.89		0.9	
Plot7	Pre	0.86	0.59	0.83	0.58	0.8	0.63	0.81	0.63	0.81	0.55
	Rec	0.69	0.76	0.73	0.72	0.82	0.61	0.82	0.62	0.73	0.66
	F1	0.77	0.66	0.78	0.64	0.81	0.62	0.81	0.63	0.76	0.6
	Acc	0.72		0.73		0.75		0.75		0.7	

Overall, the stacking framework demonstrates superior and stable generalization across heterogeneous forest plots. It achieves the highest fallen tree F1-score in five plots (Plots 1, 2, 5, 6, and 7) and joint-highest in Plot 4, peaking at 0.89 in Plots 1 and 4. Furthermore, it attains the highest or joint-highest overall Accuracy in five plots (Plots 1, 2,

4, 5, and 6), reaching 0.93 in Plot 1. These results show that combining multiple gradient boosting learners improved the balance between Precision and Recall compared with individual classifiers.

The individual classifiers show greater variation in performance across the seven test plots. CatBoost achieves the highest fallen tree F1 score in Plot 3 (0.91), whereas RF shows the lowest performance in Plot 7, with a fallen tree F1 score of 0.60. LightGBM and XGBoost showed a clear trade-off between overall accuracy and fallen tree detection performance. In the highly challenging Plot 7, LightGBM and XGBoost achieved the highest overall accuracy of 0.75, but their fallen tree F1 scores of 0.62 and 0.63 were lower than the stacking ensemble value of 0.66. Their fallen tree recalls were also lower, at 0.61 and 0.62 compared with 0.76 for stacking. This pattern indicates that individual learners may favor the dominant background class under strong class imbalance, leading to lower recall for fallen trees. Similarly, in Plot 2, XGBoost achieved the highest precision (0.79), but its Recall (0.82) was notably lower than that of the stacking ensemble (0.91), indicating a more conservative prediction pattern with more missed fallen tree samples.

Figure 14 presents representative classification results under low-, medium-, and high-complexity conditions. To facilitate qualitative assessment, supervoxels are color-coded according to their confusion matrix outcomes. Green points denote correctly detected fallen tree supervoxels (TP), blue points represent correctly identified background supervoxels (TN), orange points indicate false positives (FPs), and red points indicate false negatives (FNs).

In the low-complexity plot, as illustrated in Figure 14a, fallen trees are sparsely distributed and spatially separated. Under these conditions, the model captured the linear, elongated morphology of fallen tree stems, producing continuous true-positive segments with only limited false positives and false negatives. However, some isolated fallen tree segments with extremely low point density or spatial discontinuity remained fragmented, indicating that sparse local sampling can still weaken supervoxel-level representation. In the moderately complex plot, as shown in Figure 14b, fallen trees exhibit a denser distribution with localized overlapping. False positives were mainly concentrated near boundaries between fallen tree tips and low-lying vegetation, whereas false negatives occurred more frequently at mid-stem sections with reduced point density. Despite these localized errors, the overall classification results maintain relatively continuous spatial patterns for most fallen tree structures.

In the high-complexity plot, as shown in Figure 14c, misclassifications are more concentrated and correspond closely to the extreme structural heterogeneity of the scene. It is particularly noteworthy that the fallen tree F1-scores of all evaluated methods decline substantially in Plot 7, concentrating within the range of 0.60 to 0.66 and exhibiting a pronounced performance gap relative to the other plots. This severe degradation can be attributed to two main sources of complexity. First, dense and overlapping fallen tree accumulations caused individual supervoxels to contain mixed geometric signatures from multiple adjacent stems. This reduced the separability of local geometric descriptors. Second, dense understory shrubs introduced strong occlusion, making fallen tree point clouds sparse and discontinuous. Under these conditions, key descriptors such as linearity and axis length became less reliable, and some fallen tree supervoxels showed geometric responses similar to shrubs or ground cover. These results indicated that dense stem overlap and severe understory occlusion remain the main limitations of the current framework.

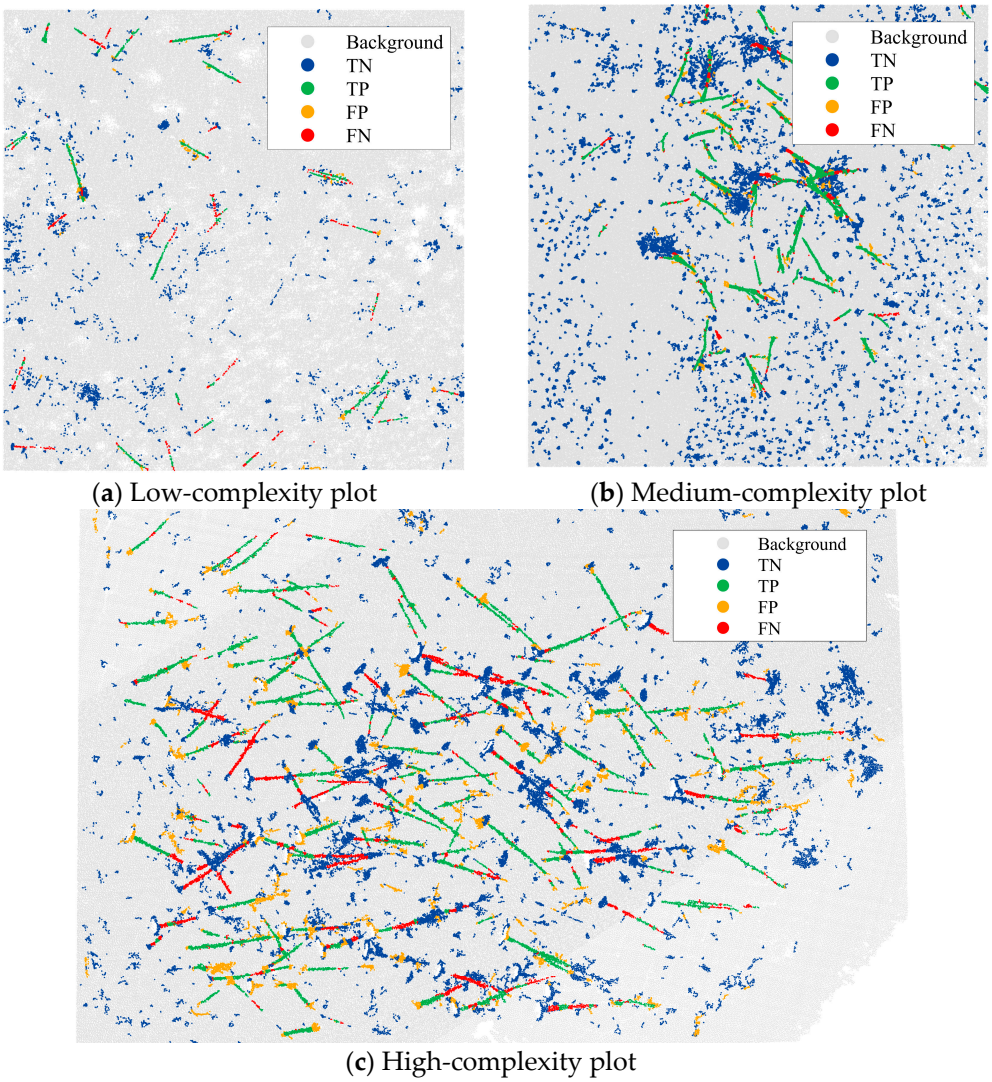


Figure 14. Point cloud classification results for three representative plots of increasing scene complexity.

3.3. Sensitivity to Supervoxel Parameters

In the proposed framework, supervoxel segmentation parameters affect the structural consistency and stability of the baseline analytical units. Key parameters include voxel resolution, seed resolution, color weight, spatial weight, and normal weight. Their functional descriptions are provided in Table 4. Because the ALS point clouds utilized in this study lack valid RGB information, the color weight is fixed at 0. The spatial and normal weights are set to 1.0 and 2.0, respectively, to emphasize geometric continuity and surface-normal consistency in forest point cloud segmentation.

Among these parameters, voxel resolution and seed resolution have the strongest influence on segmentation quality and are therefore selected for sensitivity analysis. Fallen trees in the study plots showed substantial morphological variation, with stem lengths typically ranging from 6 to 17 m and diameters of approximately 0.4 to 0.8 m. Inappropriate parameterization can affect segmentation quality. Specifically, excessively small parameter values may cause over-segmentation, fragmenting continuous fallen stems into many small units and weakening the representation of high-linearity connectivity. Conversely, excessively large values may cause under-segmentation, erroneously merging fallen trees with adjacent ground or understory shrub and critically weakening inter-class separability. All supervoxel-parameter sensitivity analyses and parameter selections were performed using the 30 development plots only. The seven independent test plots were not used to

select voxel resolution, seed resolution, feature subsets, or classifier settings. Tables 5 and 6 summarize the parameter sensitivity analysis, and Figure 15 provides representative segmentation visualizations.

Table 4. Descriptions of supervoxel segmentation parameters.

Parameter	Description
Voxel resolution	Defines the size of the voxels in space. A smaller value captures finer geometric details but increases computational cost, whereas a larger value reduces complexity but may lose structural information.
Seed resolution	Controls the spacing between supervoxel seeds. A smaller value produces a denser seed distribution and finer supervoxels, whereas a larger value produces sparser seeds and coarser structural units.
Color importance	Determines the weight of color information. A higher value places more emphasis on color variation during segmentation.
Spatial importance	Adjusts the influence of spatial location. Greater weights enhance the continuity of spatial structure.
Normal importance	Modulates the influence of surface normal vectors. Higher values promote smoother geometric segmentation.

Table 5. Impact of seed point resolution on segmentation performance, evaluated with a fixed supervoxel resolution of 0.35 m.

R_{seed} (m)	F1 Fallen Wood	F1 Other
1.2	0.74	0.87
1.4	0.83	0.91
1.5	0.87	0.94
1.6	0.85	0.93
1.8	0.82	0.92
2.0	0.84	0.93

Table 6. Impact of supervoxel resolution on segmentation performance, evaluated with a fixed seed point resolution of 1.5 m.

R_{voxel} (m)	F1 Fallen Wood	F1 Others
0.20	0.75	0.88
0.25	0.81	0.91
0.30	0.84	0.93
0.35	0.87	0.94
0.40	0.85	0.94

As quantitatively detailed in Tables 5 and 6 and visually corroborated in Figure 15, the classification performance followed an inverted-U-shaped response to both parameters. Regarding voxel resolution, the fallen tree F1-score reached its highest value of 0.87 when R_{voxel} = 0.35 m. At finer resolutions such as 0.20 m, severe over-segmentation and structural fragmentation occur due to increased sensitivity to local noise, whereas at coarser resolutions such as 0.40 m, supervoxel boundaries span across fallen trees and ground surfaces, resulting in under-segmentation. Regarding seed resolution, the highest F1-score of 0.87 is achieved at R_{seed} = 1.5 m. At denser seed distributions such as 1.2 m, supervoxels become too small to adequately capture the linear prior of fallen trees, while at sparser distributions such as 2.0 m, intersecting fallen trees or adjacent shrubs are erroneously merged into single units. Consequently, the optimal parameters R_{voxel} = 0.35 m and R_{seed} = 1.5 m are selected as the best trade-off for subsequent experiments.

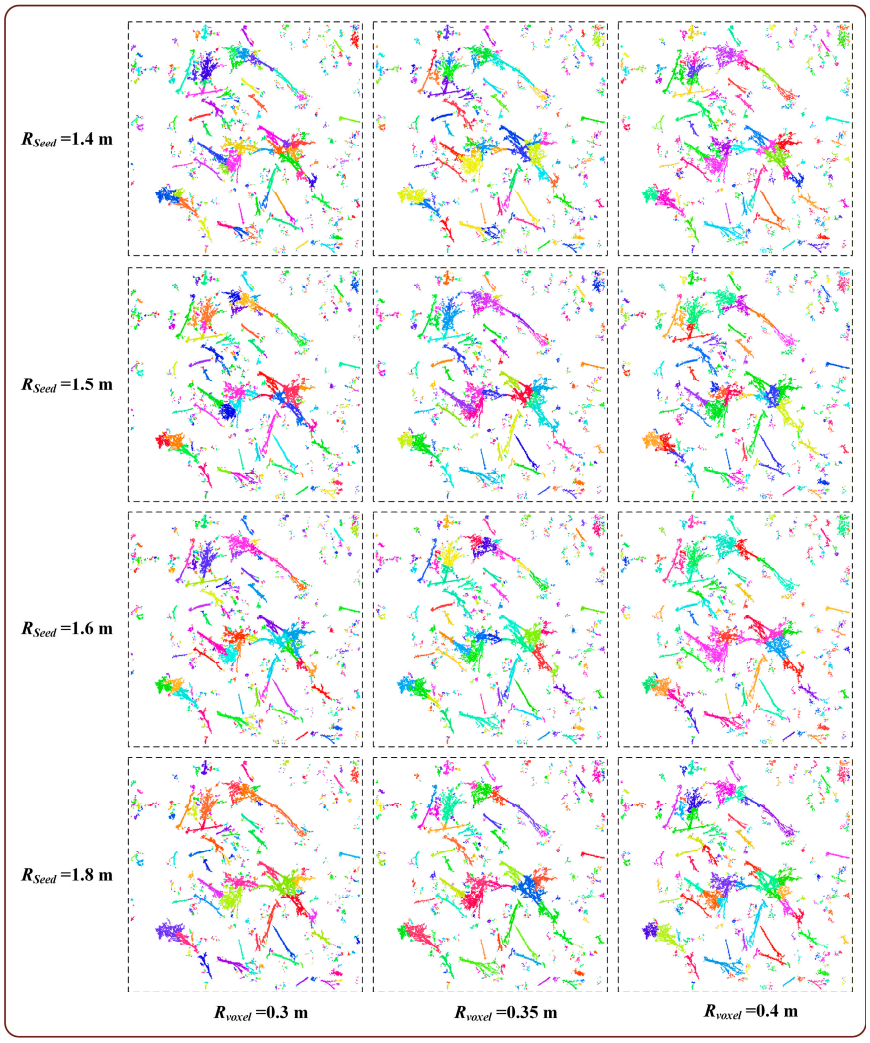


Figure 15. Visualization of supervoxel segmentation results under varying combinations of voxel resolution and seed resolution.

3.4. Ablation Study of Feature Modules

To evaluate the contribution of the major feature modules, a progressive ablation experiment is conducted using the same classifier architecture. Group A uses PCA-based local geometric features as the baseline. Group B additionally incorporates adjacency-graph topological descriptors and the high-linearity connectivity count. Group C further introduces the RANSAC-based cylindrical descriptors. The quantitative results are summarized in Table 7.

Table 7. Ablation study of the contributions of graph-topological and cylindrical feature modules to fallen tree classification performance.

Metric	A	B	C
F1 fallen tree	69.73%	71.12%	82.89%
F1 other	90.11%	91.00%	93.01%
Precision	86.01%	86.71%	90.17%

The fallen tree F1 score increases from 69.73% in Group A to 71.12% in Group B, corresponding to an improvement of 1.39% after adding graph-topological information. When the cylindrical descriptors are further introduced in Group C, the fallen tree F1 score increases to 82.89%, representing an additional improvement of 11.77% relative to Group

B and a cumulative increase of 13.16% relative to Group A. Among the evaluated feature modules, the cylindrical descriptors therefore provide the largest incremental improvement in this ablation setting.

Figure 16 provides a qualitative comparison among the three configurations. The addition of graph-topological and cylindrical descriptors reduces some isolated false detections and produces more continuous fallen tree predictions. However, errors remain in structurally complex regions where fallen stems overlap or are partially obscured by near-ground vegetation. Overall, the ablation experiment shows that cylindrical descriptors contribute strongly to classification performance, while graph-topological descriptors provide complementary contextual information.

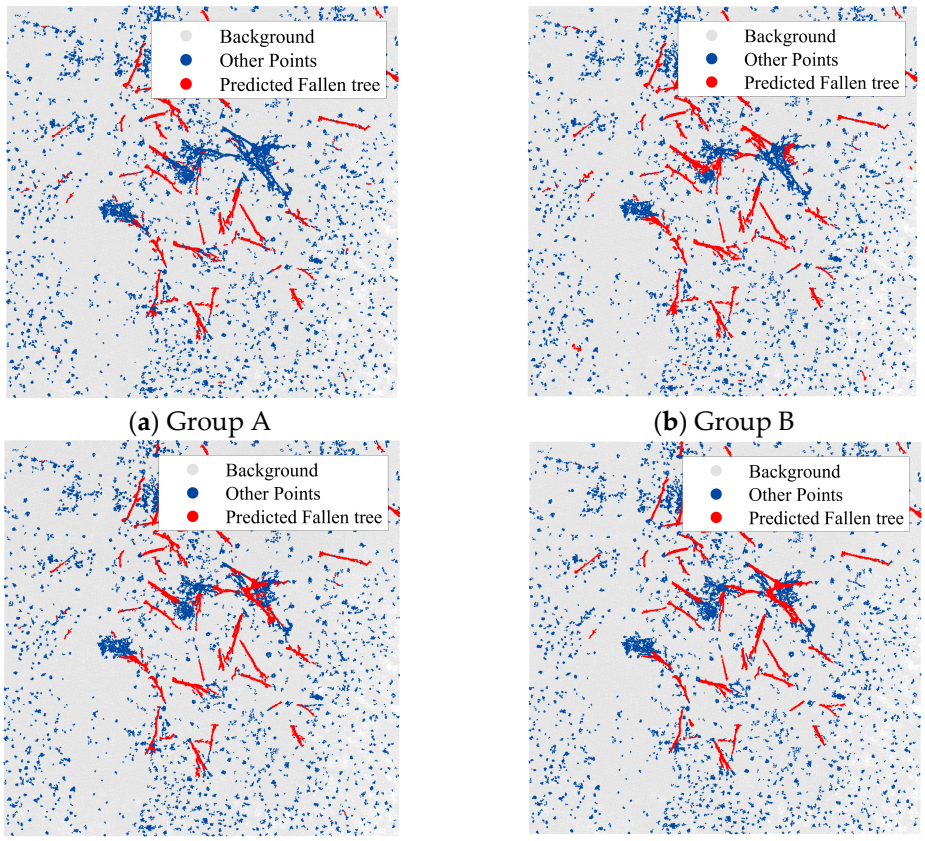


Figure 16. Visualization of classification results with and without neighborhood topological and cylindrical features.

3.5. Effect of Graph-Constrained Spatial Refinement

To assess the contribution of graph-constrained spatial refinement, three post-processing configurations were evaluated using the same initial predictions from the stacking ensemble, as described in Section 2.7. Performance was summarized using the mean classification metrics across the seven independent test plots.

The Baseline configuration used the raw stacking predictions binarized with a probability threshold of 0.45. The Neighborhood Voting configuration additionally applied local label smoothing over first-order graph neighborhoods. The Full Pipeline further incorporated LCCP-based region-level majority voting and geometry-constrained label propagation.

The results are summarized in Table 8. Relative to the baseline, neighborhood voting increased Recall from 90.29% to 92.30% but reduced Precision from 79.00% to 77.56%, indicating that unconstrained local smoothing increased the detection of fallen tree points at the expense of additional false positives. The fallen tree F1-score changed only marginally,

from 83.86% to 83.93%. In contrast, the Full Pipeline achieved a fallen tree F1-score of 85.39%, corresponding to an improvement of 1.53 % over the baseline. Overall Accuracy also increased from 88.14% to 89.85%. These results indicate that region-level correction and geometry-constrained label propagation were more effective than neighborhood voting alone for improving spatial label consistency.

Table 8. Quantitative comparison of mean classification metrics (%) across progressive post-processing stages for the seven test plots.

Stage	Precision (Fallen)	Recall (Fallen)	F1 (Fallen)	F1 (Other)	Accuracy
Baseline	79.00%	90.29%	83.86%	90.71%	88.14%
Neighborhood Voting	77.56%	92.30%	83.93%	90.41%	88.11%
Full Pipeline	81.84%	89.49%	85.39%	92.15%	89.85%

3.6. Error Patterns Under Complex Forest Conditions

To further investigate the source of misclassifications, Figure 17 shows enlarged views of representative error regions. In areas where multiple fallen trees intersect and overlap spatially, as shown in the right inset of Figure 17, supervoxels may contain points from several adjacent stems. These mixed supervoxels do not fully match the linear geometry of a single cylindrical stem or the dispersed structure of typical background classes. As a result, false positives and false negatives were concentrated near these intersection regions. This indicates that dense fallen tree accumulations remain difficult to resolve using local geometrical features alone.

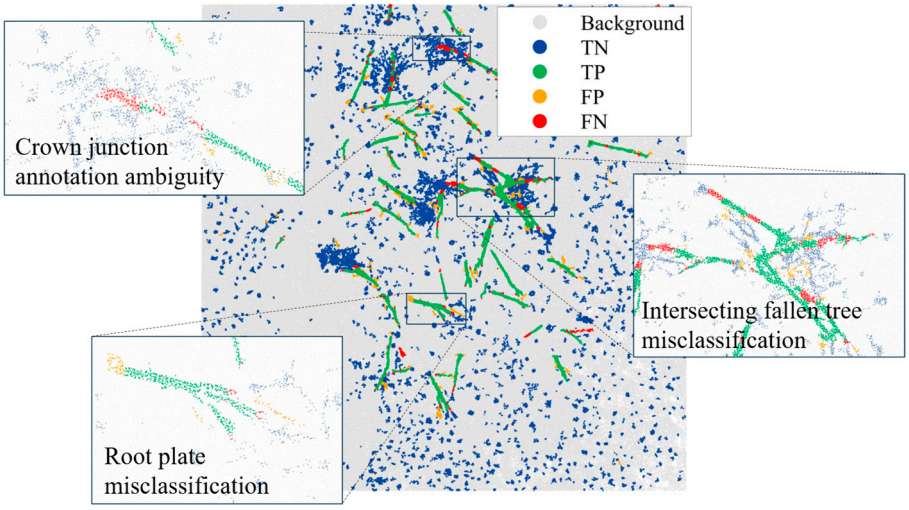


Figure 17. Representative misclassification regions in the high-complexity plot.

Another source of error was boundary ambiguity at the structural extremities of fallen trees. Localized false negatives occurred where the main trunk was intertwined with disorganized crown point clouds. Although the reference manual annotations delineate a straight, continuous stem, the model struggles to isolate the linear trunk extension from the cluttered crown structure, resulting in structural fragmentation at the crown junction. False positives were commonly observed at root extremities, where small upturned root plates exhibit geometric continuity and morphological similarity to the main trunk. These structures were sometimes classified as fallen tree points, although they were excluded from the reference annotation. These errors reflect the difficulty of aligning discrete annotation rules with the complex physical structure of forest point clouds.

4. Discussion

The discussion first examines the contributions of cylindrical and topological priors, then compares the proposed framework with existing fallen tree detection studies. It subsequently considers performance variation across forest conditions and outlines the main limitations and directions for future work.

4.1. Contribution of Cylindrical and Topological Priors

The ablation and refinement results indicate that the main contribution of the proposed framework arises from combining stem-shape priors with spatial context. Among the evaluated feature modules, the RANSAC-based cylindrical descriptors produced the largest performance gain. Adding these descriptors increased the fallen tree F1 score by 11.77% relative to the configuration containing local geometric and graph-topological descriptors. This result suggests that local cylindrical consistency provides information that is not fully represented by generic PCA-based shape descriptors.

The importance of cylindrical descriptors is consistent with the approximately cylindrical and elongated structure of fallen stems. Earlier work by Polewski et al. [38] combined point-level and segment-level shape descriptors to identify fallen-stem segments, demonstrating the usefulness of three-dimensional shape information for fallen tree recognition. The present framework explicitly represents local stem morphology through RANSAC-based cylinder radius, fitting residual, axis direction, and fitting confidence. The 11.77-percentage-point gain observed in our ablation study provides evidence that these descriptors contribute information beyond the local geometric and graph-topological features already included in the framework. The SHAP analysis further highlights the importance of structural information, with patch-scale major-axis length, linearity, and point count among the most influential descriptors. These findings support the use of complementary morphological representations, although they do not establish the superiority of cylindrical fitting over the shape descriptors used in previous studies. Here, physics-guided refers to the incorporation of explicit structural priors associated with fallen stems, rather than to a physical simulation of stem formation or forest disturbance processes.

Graph-topological descriptors provided a smaller but complementary contribution, increasing the fallen tree F1 score by 1.39% in the feature-ablation experiment. These descriptors characterize whether a supervoxel is embedded within a locally continuous structure, helping distinguish fallen-stem fragments from isolated branches, shrub elements, or terrain artifacts. The relevance of spatial context is consistent with the constrained conditional random field approach of Polewski et al. [17], which incorporated pairwise spatial interactions between adjacent stem-segment candidates rather than considering each candidate independently. Our framework represents neighborhood structure through supervoxel adjacency descriptors and subsequently applies geometry-constrained spatial refinement. Although the representations and inference procedures differ, both approaches incorporate spatial relationships to complement local appearance information. The smaller gain from graph-topological descriptors in our experiment nevertheless indicates that neighborhood context alone provides limited additional discrimination when local stem morphology is weak or ambiguous.

The refined experiment further showed that spatial context must be constrained by geometry. Neighborhood voting alone increased Recall but reduced fallen tree Precision, while the F1-score remained nearly unchanged, increasing slightly from 83.86% to 83.93%. In contrast, the Full Pipeline increased the fallen tree F1-score from 83.86% to 85.39%. This result suggests that region-level LCCP voting and geometry-constrained label propagation were more effective when applied together than simple neighborhood smoothing alone.

Taken together, these results support a complementary role for the two priors. Cylindrical descriptors provide the primary local evidence for fallen-stem morphology, whereas graph-based descriptors and constrained refinement improve label consistency among fragmented but structurally compatible supervoxels. This combination remains limited in areas where stem overlap, crown junctions, root plates, or understory occlusion obscure both local cylindrical structure and reliable spatial continuity.

4.2. Comparison with Existing Fallen Tree Detection Studies

Table 9 places the present study in the context of representative ALS-based fallen tree detection approaches. Direct numerical comparisons should be interpreted cautiously because the studies differ in ALS acquisition characteristics, point density, forest structure, target definition, prediction unit, and evaluation protocol. The table is therefore intended to compare methodological settings and reported performance ranges rather than to establish a direct ranking among methods.

Table 9. Comparison between the proposed framework and representative fallen tree detection methods.

Method	Point Density	Evaluation Protocol	Test Scenario	Key Metrics
Raster filtering with echo-width features	29.4 pts/m^2 full-waveform	Object-level	Single deciduous forest plot	Precision: 90% Recall: 76%
Normalized Cut graph-based clustering	30 pts/m^2	Object-level	5 mixed mountain forest plots	Precision: 47–97% Recall: 34–71%
Iterative Hough transform-based line detection	15 pts/m^2 (ALS)	Object-level	Boreal coniferous forest	Best F1: 0.39 (ALS)
Proposed method	10–15 pts/m^2 (AHN4)	Supervoxel-level	7 plots across diverse ecosystem types	Mean F1: 0.84

Note: The proposed framework generates supervoxel-level predictions but is evaluated against point-level reference labels. Direct numerical comparisons should be interpreted cautiously because the studies differ in target definitions, prediction units, ALS characteristics, forest conditions, and evaluation protocols.

The studies summarized in Table 9 differ in how they identify and organize candidate stem structures. Mücke et al. [12] used echo width, echo type, and normalized height to filter full-waveform ALS observations and construct a digital height model in which downed stems appeared as line-like features. Image-processing operations were then used to delineate and refine the candidate regions. Polewski et al. [16,17] instead represented fallen stems as collections of short segments and applied Normalized Cut clustering to merge them. The segment-similarity function and clustering stopping criterion were learned using simulated fallen tree configurations. Their subsequent constrained conditional random field approach incorporated interactions between adjacent segment candidates to improve the suitability of the detected fragments for subsequent merging.

Heinaro et al. [18,19] evaluated a line-detection-based workflow combining near-ground filtering, iterative Hough-transform line detection, and candidate-segment delineation and classification. This approach used linear structures to generate stem candidates and additional filtering to reduce false detections. In comparison, the present framework performs semantic classification on supervoxels using local geometry, explicit cylindrical descriptors, and graph-topological context, followed by geometry-constrained refinement. It therefore uses stem morphology and spatial continuity as inputs to classification rather than requiring complete stems to be delineated before semantic labeling. These methodological differences provide context for the performance values in Table 9, but do not

remove the differences in target definitions and evaluation protocols that prevent direct numerical ranking.

The present study was evaluated using AHN4 data with an operational point density of approximately 10–15 pts/m^2, which is lower than that used in several previous studies summarized in Table 9. Rather than relying on individual points as independent observations, the proposed method aggregates locally coherent points into supervoxels and integrates geometric, cylindrical, and graph-topological descriptors. Across the seven spatially independent test plots, the framework achieved a mean fallen tree F1-score of 0.84. Importantly, although classification was performed at the supervoxel level, performance was assessed against point-level reference labels. This evaluation protocol provides a spatially granular assessment and is sensitive to partial stem omission, boundary displacement, and localized false detections; it is therefore not directly equivalent to object-level detection metrics reported in some previous studies.

Forest structure, particularly coniferous vegetation and near-ground clutter, has also been identified as an important source of performance variation. Heinaro et al. [19] reported a best F1-score of 0.39 for an iterative Hough-transform method using 15 pts/m^2 ALS data and found that additional vegetation returns could increase false detections. Similarly, Polewski et al. [16] reported reduced performance in coniferous plots with canopy closure above 50%. In the present study, the two coniferous test plots (Plots 1 and 4) each achieved a fallen tree F1-score of 0.89 despite dense understory interference in Plot 1 and substantial stem breakage and root-plate structures in Plot 4. These results suggest that the combination of supervoxel aggregation, cylindrical descriptors, and topological context can provide stable discrimination under the coniferous conditions represented by the present test sites.

Nevertheless, the results should not be interpreted as evidence that the method is universally robust across all forest types or ALS acquisitions. Performance declined in the plot containing severely overlapping fallen trees, and errors remained common where stems were obscured by understory vegetation, intertwined with crown material, or connected to complex root structures. The framework should therefore be viewed as an interpretable approach for point-level fallen tree mapping under the tested operational ALS conditions, rather than as a directly comparable replacement for methods evaluated using different target definitions or object-level protocols.

4.3. Performance Variation Across Forest Conditions

To examine performance variation across forest conditions, the seven independent test plots were grouped according to stand type and landscape setting: coniferous plantation plots (Plots 1 and 4), mixed, broadleaf, and transitional forest plots (Plots 2, 3, 5, and 6), and a polder plantation plot (Plot 7). The plots differed in terrain, stand structure, understory complexity, and fallen tree configuration (Table 1), providing a heterogeneous basis for assessing the transferability of the proposed framework.

The coniferous plantation plots yielded consistently high fallen tree F1-scores. Plots 1 and 4 both achieved an F1-score of 0.89, with overall accuracies of 0.93 and 0.91, respectively (Table 3). These results are notable because the two plots presented distinct challenges: Plot 1 contained dense shrub understory and substantial sub-canopy occlusion, whereas Plot 4 included broken stems and large upturned root plates. Thus, high performance in these plots cannot be attributed simply to unobstructed near-ground observations. Instead, the results suggest that the combined geometric, cylindrical, and topological representation remained informative where recognizable stem segments were present, even when surrounding near-ground structures were complex.

The mixed, broadleaf, and transitional forest plots showed more heterogeneous performance, although F1-scores remained within a relatively narrow range of 0.83–0.90.

Plot 3 achieved the highest F1-score among the seven test plots (0.90), despite its old-growth broadleaf structure and irregular fallen-stem morphology. This result indicates that broadleaf stand structure and stem-shape variability do not necessarily preclude reliable fallen tree mapping when local structural evidence is sufficiently preserved. By contrast, Plot 2 achieved an F1-score of 0.83, which is consistent with the greater geometric ambiguity introduced by windthrow disturbance, stem breakage, and variable decay conditions. Plots 5 and 6 also achieved F1-scores of 0.83 and 0.87, respectively, indicating that the framework remained applicable across mixed and transitional forest settings with heterogeneous ground and understory structures.

Plot 7 was the most challenging case, with a fallen tree F1-score of 0.66. This plot contained densely overlapping and stacked fallen trees in a poplar-dominated polder plantation. Under these conditions, a supervoxel may include returns from multiple stems, vegetation, and ground surfaces, weakening the local cylindrical and linear signatures on which the classifier relies. The method nevertheless achieved a recall of 0.76, indicating that a substantial proportion of fallen tree points were identified. However, precision declined to 0.59, showing that false-positive suppression remains difficult when fallen stems overlap extensively or are obscured by dense near-ground vegetation.

Overall, the cross-plot results suggest that the framework is applicable across the coniferous, mixed, broadleaf, and transitional forest conditions represented in this study, while remaining sensitive to complex near-ground structures. Heinaro et al. [19] similarly demonstrated that fallen tree detection performance depends on point-cloud characteristics and the treatment of non-target structures. Their line-detection approach was sensitive to noise, and filtering reduced false detections at the cost of removing some true detections. This precision–recall trade-off provides relevant context for the reduced precision observed in Plot 7. In our framework, densely overlapping stems and nearby vegetation can produce mixed supervoxel geometry, weakening the cylindrical and linear signatures used for classification. The studies therefore highlight a related difficulty in separating fallen stems from surrounding structures, despite differences in their detection methods. However, our point-level evaluation is not directly comparable to the individual-tree detection metrics reported by Heinaro et al. [19]. Furthermore, the present cross-plot comparison does not isolate the independent effects of species composition, canopy closure, terrain, understory density, and fallen tree arrangement. The observed performance differences should consequently be interpreted as site-specific patterns rather than independent estimates of the effects of individual forest conditions.

4.4. Limitations and Future Work

The framework was developed and evaluated using AHN4 data from the Netherlands, with operational point densities of approximately 10–15 pts/m^2. Although the test plots covered contrasting forest settings, all data originated from the same national ALS programme. The transferability of the method to other acquisition systems, point densities, scan geometries, seasons, and forest regions therefore remains to be tested. In particular, the effectiveness of the cylindrical and topological descriptors may vary when stem surfaces are represented by substantially fewer returns or when near-ground visibility is altered by acquisition conditions.

Second, the reference data were manually delineated at the point level and defined fallen stems as the target class, while excluding root plates, crown fragments, shrubs, and other near-ground objects unless they formed part of an identifiable stem. This rule provides a clear basis for supervised classification, but transitions between stems, crowns, and root structures are not always physically discrete in ALS point clouds. Consequently, some apparent classification errors may reflect ambiguity at class boundaries rather than

fully separable geometric categories. Future work could combine point-level labels with field-verified individual-stem references to evaluate both semantic mapping accuracy and object-level completeness.

Third, the current framework depends on near-ground filtering, supervoxel segmentation, and graph-based refinement. These operations improve structural coherence but may also propagate local segmentation errors. In particular, when multiple fallen stems overlap within a small spatial region, one supervoxel may contain mixed observations from several objects. The resulting feature vector no longer represents a single stem segment, which can reduce precision, as observed in Plot 7. An instance-aware approach that explicitly separates intersecting or stacked stems, potentially through multi-scale segmentation and local stem-axis hypotheses, would be a useful extension.

Finally, the present evaluation establishes cross-plot generalization within the tested dataset, but it does not quantify the separate contributions of point density, canopy condition, species composition, decay stage, or understory cover. A future multi-site experiment with deliberately varied acquisition settings and independently field-validated reference data could determine the operating conditions under which the framework remains reliable and identify when alternative representations are required.

5. Conclusions

This study presented a physics-guided supervoxel framework for fallen tree detection from ALS point clouds. The method represents ALS data at the supervoxel level and integrates local geometric descriptors, RANSAC-based cylindrical descriptors, graph-topological context, stacking ensemble classification, SHAP-based feature selection, and graph-constrained spatial refinement. By combining stem-shape priors with neighborhood context, the framework improves the separation of fallen trees from shrubs, terrain fragments, and other near-ground objects in operational ALS data.

Experiments on seven held-out AHN4 forest plots showed that the proposed framework achieved a mean fallen tree F1-score of 83.86% under point densities of 10–15 pts/m^2, which increased to 85.39% after graph-constrained refinement. The ablation analysis further showed that cylindrical descriptors were the dominant feature group, increasing the fallen tree F1 score by 11.77%. Graph-topological descriptors and spatial refinement provided additional gains in local consistency, although their effects were smaller than those of the cylindrical descriptors.

The results suggest that physics-guided supervoxel learning can support interpretable fallen tree semantic mapping from ALS point clouds. The current framework is evaluated with point-level reference labels and does not yet resolve individual fallen tree instances or estimate log-level geometry. Dense fallen tree overlap, crown junctions, root plates, and understory occlusion remain the main sources of error. Future work should extend the framework toward instance-level delineation, object-level matching, and uncertainty-aware aggregation from detected fallen tree points to plot-level forest structure indicators.

Author Contributions: Conceptualization, X.L.; methodology, X.L.; software, J.W. (Jinguo Wang); validation, Y.S.; formal analysis, J.W. (Jinhu Wang); investigation, Y.C.; resources, D.C.; data curation, X.L.; writing—original draft preparation, X.L.; writing—review and editing, Y.S.; visualization, X.L.; supervision, Y.S.; project administration, J.W. (Jinguo Wang); funding acquisition, Y.S. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the National Natural Science Foundation of China (Grant No. 42571513).

Data Availability Statement: The ALS point cloud datasets (AHN4) used in this study are publicly available and can be accessed via the official Dutch national elevation model repository https://www.pdok.nl/ (accessed on 14 September 2025). The ground-truth annotations generated for this study and the implementation code for the proposed physics-guided supervoxel framework are available from the author upon reasonable request.

Acknowledgments: The authors are grateful to the High-Performance Computing platform at Hohai University for processing point cloud data in this paper.

Conflicts of Interest: The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

1. Harmon, M.E.; Franklin, J.F.; Swanson, F.J.; Sollins, P.; Gregory, S.V.; Lattin, J.D.; Anderson, N.H.; Cline, S.P.; Aumen, N.G.; Sedell, J.R.; et al. Ecology of Coarse Woody Debris in Temperate Ecosystems. In *Advances in Ecological Research*; MacFadyen, A., Ford, E.D., Eds.; Academic Press: Cambridge, MA, USA, 1986; Volume 15, pp. 133–302. [CrossRef]
2. Ståhl, G.; Ringvall, A.; Fridman, J. Assessment of Coarse Woody Debris: A Methodological Overview. *Ecol. Bull.* **2001**, *49*, 57–70.
3. Woodall, C.W. Biodiversity in Dead Wood. *For. Sci.* **2013**, *59*, 598. [CrossRef]
4. Pan, Y.; Birdsey, R.A.; Fang, J.; Houghton, R.; Kauppi, P.E.; Kurz, W.A.; Phillips, O.L.; Shvidenko, A.; Lewis, S.L.; Canadell, J.G.; et al. A Large and Persistent Carbon Sink in the World's Forests. *Science* **2011**, *333*, 988–993. [CrossRef] [PubMed]
5. Næsset, E. Determination of mean tree height of forest stands using airborne laser scanner data. *ISPRS J. Photogramm. Remote Sens.* **1997**, *52*, 49–56. [CrossRef]
6. Morsdorf, F.; Kötz, B.; Meier, E.; Itten, K.I.; Allgöwer, B. Estimation of LAI and fractional cover from small footprint airborne laser scanning data based on gap fraction. *Remote Sens. Environ.* **2006**, *104*, 50–61. [CrossRef]
7. Hyyppä, J.; Holopainen, M.; Olsson, H. Laser Scanning in Forests. *Remote Sens.* **2012**, *4*, 2919–2922. [CrossRef]
8. Yao, W.; Krzystek, P.; Heurich, M. Tree species classification and estimation of stem volume and DBH based on single tree extraction by exploiting airborne full-waveform LiDAR data. *Remote Sens. Environ.* **2012**, *123*, 368–380. [CrossRef]
9. Kankare, V.; Räty, M.; Yu, X.; Holopainen, M.; Vastaranta, M.; Kantola, T.; Hyyppä, J.; Hyyppä, H.; Alho, P.; Viitala, R. Single tree biomass modelling using airborne laser scanning. *ISPRS J. Photogramm. Remote Sens.* **2013**, *85*, 66–73. [CrossRef]
10. Antonarakis, A.S.; Richards, K.S.; Brasington, J. Object-based land cover classification using airborne LiDAR. *Remote Sens. Environ.* **2008**, *112*, 2988–2998. [CrossRef]
11. Blanchard, S.D.; Jakubowski, M.K.; Kelly, M. Object-Based Image Analysis of Downed Logs in Disturbed Forested Landscapes Using Lidar. *Remote Sens.* **2011**, *3*, 2420–2439. [CrossRef]
12. Mücke, W.; Deák, B.; Schroiff, A.; Hollaus, M.; Pfeifer, N. Detection of fallen trees in forested areas using small footprint airborne laser scanning data. *Can. J. Remote Sens.* **2013**, *39*, S32–S40. [CrossRef]
13. Lindberg, E.; Hollaus, M.; Mücke, W.; Fransson, J.E.S.; Pfeifer, N. Detection of lying tree stems from airborne laser scanning data using a line template matching algorithm. *ISPRS Ann. Photogramm. Remote Sens. Spat. Inf. Sci.* **2013**, *II-5/W2*, 169–174. [CrossRef]
14. Pesonen, A.; Maltamo, M.; Eerikäinen, K.; Packalèn, P. Airborne laser scanning-based prediction of coarse woody debris volumes in a conservation area. *For. Ecol. Manag.* **2008**, *255*, 3288–3296. [CrossRef]
15. Polewski, P.; Yao, W.; Heurich, M.; Krzystek, P.; Stilla, U. Detection of fallen trees in ALS point clouds of a temperate forest by combining point/primitive-level shape descriptors. In *Proceedings of the Gemeinsame Tagung 2014 der DGfK, der DGPF, der GfGI und des GiN, Hamburg, Germany, 26–28 March 2014*; DGPF Tagungsband 23; DGPF: München, Germany; pp. 1–12.
16. Polewski, P.; Yao, W.; Heurich, M.; Krzystek, P.; Stilla, U. Detection of fallen trees in ALS point clouds using a Normalized Cut approach trained by simulation. *ISPRS J. Photogramm. Remote Sens.* **2015**, *105*, 252–271. [CrossRef]
17. Polewski, P.; Yao, W.; Heurich, M.; Krzystek, P.; Stilla, U. Learning a constrained conditional random field for enhanced segmentation of fallen trees in ALS point clouds. *ISPRS J. Photogramm. Remote Sens.* **2018**, *140*, 33–44. [CrossRef]
18. Heinaro, E.; Tanhuanpää, T.; Yrttimaa, T.; Holopainen, M.; Vastaranta, M. Airborne laser scanning reveals large tree trunks on forest floor. *For. Ecol. Manag.* **2021**, *491*, 119225. [CrossRef]
19. Heinaro, E.; Tanhuanpää, T.; Vastaranta, M.; Yrttimaa, T.; Kukko, A.; Hakala, T.; Mattsson, T.; Holopainen, M. Evaluating Factors Impacting Fallen Tree Detection from Airborne Laser Scanning Point Clouds. *Remote Sens.* **2023**, *15*, 382. [CrossRef]
20. Lopes Queiroz, G.; McDermid, G.J.; Castilla, G.; Linke, J.; Rahman, M.M. Mapping Coarse Woody Debris with Random Forest Classification of Centimetric Aerial Imagery. *Forests* **2019**, *10*, 471. [CrossRef]
21. Dakin Kuiper, S.; Coops, N.C.; Jarron, L.R.; Tompalski, P.; White, J.C. An automated approach to detecting instream wood using airborne laser scanning in small coastal streams. *Int. J. Appl. Earth Obs. Geoinf.* **2023**, *118*, 103272. [CrossRef]

22. Weinmann, M.; Jutzi, B.; Hinz, S.; Mallet, C. Semantic point cloud interpretation based on optimal neighborhoods, relevant features and efficient classifiers. *ISPRS J. Photogramm. Remote Sens.* **2015**, *105*, 286–304. [CrossRef]
23. Garcia-Garcia, A.; Orts-Escolano, S.; Oprea, S.; Villena-Martinez, V.; Garcia-Rodriguez, J. A review on deep learning techniques applied to semantic segmentation. *Appl. Soft Comput.* **2018**, *70*, 41–63. [CrossRef]
24. PDOK. Actueel Hoogtebestand Nederland 4 (AHN4). Available online: https://www.pdok.nl/introductie/-/article/actueel-hoogtebestand-nederland-ahn (accessed on 14 September 2025).
25. Zhang, W.; Qi, J.; Wan, P.; Wang, H.; Xie, D.; Wang, X.; Yan, G. An Easy-to-Use Airborne LiDAR Data Filtering Method Based on Cloth Simulation. *Remote Sens.* **2016**, *8*, 501. [CrossRef]
26. Lin, Y.; Wang, C.; Zhai, D.; Li, W.; Li, J. Toward better boundary preserved supervoxel segmentation for 3D point clouds. *ISPRS J. Photogramm. Remote Sens.* **2018**, *143*, 39–47. [CrossRef]
27. Papon, J.; Abramov, A.; Schoeler, M.; Wörgötter, F. Voxel Cloud Connectivity Segmentation–Supervoxels for Point Clouds. In *Proceedings of the 2013 IEEE Conference on Computer Vision and Pattern Recognition, Portland, OR, USA, 23–28 June 2013*; IEEE: Piscataway, NJ, USA, 2013; pp. 2027–2034. [CrossRef]
28. Weinmann, M.; Urban, S.; Hinz, S.; Jutzi, B.; Mallet, C. Distinctive 2D and 3D features for automated large-scale scene analysis in urban areas. *Comput. Graph.* **2015**, *49*, 47–57. [CrossRef]
29. Zhou, Y.; Yu, Y.; Lu, G.; Du, S. Super-Segments Based Classification of 3D Urban Street Scenes. *Int. J. Adv. Robot. Syst.* **2012**, *9*, 248. [CrossRef] [PubMed]
30. Stein, S.C.; Schoeler, M.; Papon, J.; Worgotter, F. Object Partitioning Using Local Convexity. In *Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition*; IEEE: Piscataway, NJ, USA, 2014; pp. 304–311. [CrossRef]
31. Wang, W.; Wu, Y.; Jia, S.; Xie, L.; Guo, R.; Tang, S. Leaf–Wood Separation in Tree Point Clouds via Unlearning-Driven Supervoxel Random Forest. *IEEE Trans. Geosci. Remote Sens.* **2025**, *63*, 1–19. [CrossRef]
32. Fischler, M.A.; Bolles, R.C. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM* **1981**, *24*, 381–395. [CrossRef]
33. Chen, T.; Guestrin, C. XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, 13–17 August 2016*; Association for Computing Machinery: New York, NY, USA, 2016; pp. 785–794. [CrossRef]
34. Ke, G.; Meng, Q.; Finley, T.; Wang, T.; Chen, W.; Ma, W.; Ye, Q.; Liu, T.-Y. LightGBM: A Highly Efficient Gradient Boosting Decision Tree. In *Proceedings of the Advances in Neural Information Processing Systems 30*; NeurIPS: San Diego CA, USA, 2017; pp. 3146–3154.
35. Prokhorenkova, L.; Gusev, G.; Vorobev, A.; Dorogush, A.V.; Gulin, A. CatBoost: Unbiased boosting with categorical features. In *Proceedings of the 32nd Conference on Neural Information Processing Systems (NeurIPS 2018)*; Curran Associates Inc.: Red Hook, NY, USA, 2018; pp. 6638–6648.
36. Breiman, L. Random forests. *Mach. Learn.* **2001**, *45*, 5–32. [CrossRef]
37. Wolpert, D.H. Stacked generalization. *Neural Netw.* **1992**, *5*, 241–259. [CrossRef]
38. Polewski, P.; Yao, W.; Heurich, M.; Krzystek, P.; Stilla, U. Detection of fallen trees in ALS point clouds by learning the Normalized Cut similarity function from simulated samples. *ISPRS Ann. Photogramm. Remote Sens. Spat. Inf. Sci.* **2014**, *II-3*, 111–118. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.