

A Generative Upsampling Framework for Reconstructing High-Density Tree Structures from Low-Density Airborne Lidar

Erfan Hasanpour Zaryabi¹, Liam Bennett¹, Josh Qixuan Sun², Mark Crowley², Laura Chasmer³, Christopher Hopkinson³, Jeff Boisvert¹

¹ Department of Civil and Environmental Engineering, Faculty of Engineering, University of Alberta, Edmonton, AB, Canada – (ehasanpo, lgbenet, jbb)@ualberta.ca

² Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON, Canada – (josh.q.sun, mcrowley)@uwaterloo.ca

³ Department of Geography and Environment, Western University, Ontario, Canada – (lchasmer, chris.hopkinson)@uwo.ca

Keywords: Point cloud Upsampling, Tree Reconstruction, Airborne Laser Scanning, UAV, Lidar, Deep Generative Networks.

Abstract:

Airborne Laser Scanning (ALS) enables efficient large-scale mapping of three-dimensional forest structure due to its extensive spatial coverage, but its low point density often limits the accurate representation of forest structural attributes required for applications such as wildfire management and forest health monitoring. This study proposes a deep generative upsampling framework to enhance low-density ALS tree point clouds and generate dense representations for improved structural analysis. To improve tree reconstruction, the framework incorporates a synthetic pretraining strategy that exposes the model to complete tree geometries before fine-tuning on real data. Experimental results demonstrate that the proposed framework generates high-density tree point clouds with strong spatial similarity to ground-truth point clouds. Synthetic pretraining further improves reconstruction performance, achieving lower Chamfer Distance and higher F1-scores across multiple upsampling rates. In addition, the reconstructed point clouds increase estimation accuracy of key structural attributes, including tree height, crown area, and tree volume. Qualitative experiments on real ALS data confirms the ability of the proposed framework to recover realistic tree structures from sparse point clouds. Overall, the proposed method provides a practical approach for bridging the wide spatial coverage of ALS with the detailed structural information typically obtained from high-density lidar data.

1. Introduction

Light Detection and Ranging (lidar) has become an essential tool to quantify forest structure in three dimensions, allowing the extraction of tree-level structural metrics such as height, tree volume, diameter at breast height (DBH), and biomass across various landscapes (Balestra et al., 2024). These metrics are applicable to various domains, including wildfire management, biodiversity assessment, forest health monitoring, and timber management (Moran et al., 2020; Toivonen et al., 2023). Terrestrial Laser Scanning (TLS) and Uncrewed Aerial Vehicle (UAV) lidar are widely used tools for acquiring dense point clouds (thousands of points per square meter) of forest structure (Terry et al., 2022). This data allows for the characterization of individual tree structure (Abdollahi and Yebra, 2025; Xiang et al., 2024). However, both TLS and UAV lidar have limited spatial coverage which makes them neither time- nor cost-effective for mapping extensive forested regions (Abdollahi and Yebra, 2025). Airborne Laser Scanning (ALS) is more practical for large-scale uses as it collects data from much higher altitudes compared to UAV lidar, enabling broad area coverage. However, this increased coverage comes at the cost of lower point density, typically ranging from 1 to 100 points/m² (Zhang et al., 2024), which is insufficient for characterizing forest structure at the individual tree level (Hamraz et al., 2017). A practical solution that addresses the limitations of both TLS/UAV lidar (restricted spatial coverage) and ALS (low sampling density) is needed for detailed forest structure quantification across large spatial extents. This study proposes an upsampling framework based on a deep generative network to densify low-density ALS tree point clouds and improve structural representation.

Point cloud upsampling refers to the task of generating a dense, high-resolution point cloud from a sparse, low-resolution input by creating new points that follow the underlying surface

geometry while preserving local features and patterns (Yu et al., 2018). The first learning-based upsampling framework is introduced in PU-Net (Yu et al., 2018). PU-Net upsamples point clouds in feature space by extracting multi-level features from 3D coordinates. However, PU-Net shows limitations when dealing with complex surface geometries and arbitrary upsampling rates. These limitations are later addressed by subsequent methods, including PU-GAN and PU-GCN, which improve feature extraction and expansion mechanisms (Li et al., 2019; Qian et al., 2021).

Most existing upsampling methods are developed and evaluated using synthetic point clouds generated through uniform sampling from the surfaces of simple CAD models to mimic real low-resolution inputs. In real-world scenarios, however, low-density point clouds often result from occlusions, limited sensor sampling density, or a combination of both. Occlusions arise from structural complexity in the environment; for example, in forest ecosystems, dense canopies naturally block the view of understory vegetation from aerial sensors (Kükenbrink et al., 2017; Schneider et al., 2019). Sensor sampling density refers to the number of points acquired per unit area and is influenced by factors such as pulse rate, scan pattern, flight altitude, and flight speed (Petrás et al., 2023). For instance, ALS data collected at higher flight altitudes often exhibit lower sampling densities compared to UAV lidar data acquired at lower altitudes, and this limitation becomes more pronounced in forest environments where occlusions further reduce point availability. In such cases, point cloud upsampling can serve as a valuable tool for forestry applications by leveraging the wide coverage of ALS data while improving the effective resolution of forest structural information.

Recent studies demonstrate the potential of upsampling and completion techniques in forestry applications. For example, Liu

et al. (2024) and Liu, Zhong, et al. (2025b, 2025a) apply point cloud upsampling to improve the quality of UAV lidar point clouds and report improved accuracy in individual tree species classification. Liu, Yang, et al. (2025) develop a completion network to reconstruct the crowns of *Populus tomentosa* trees from UAV lidar point clouds and show that incorporating species-specific structural traits significantly improves completion performance. Similarly, Su et al. (2025) present a transformer-based completion framework and demonstrate that upsampling improves diameter at breast height (DBH) estimation in forestry applications.

These studies collectively demonstrate that point cloud upsampling and completion improves the reconstruction of tree structure and the accuracy of downstream forestry applications. However, these approaches are primarily developed and evaluated using high-density point clouds acquired from UAV lidar or TLS sensors rather than low-density ALS data, which remains the most practical technology for large-scale forest monitoring. To address this gap, this study introduces a deep generative upsampling framework that enhances low-density ALS-like point clouds (containing 100 points per tree) and increases their density to a level comparable to UAV lidar (8,164 points per tree) while preserving tree structural characteristics.

2. Materials and Methods

2.1 Dataset and Preparation

The study area is a mature Jack pine (*Pinus banksiana*) stand (91–97 years old) located at the Boreal Ecosystems Research and Monitoring Sites (BERMS), north of Prince Albert, Saskatchewan, where coincidental UAV lidar and ALS data were collected within two days of each other, on July 29 (ALS) and July 31 (UAV lidar), 2022. The UAV lidar data has an average point density of 2,000 points m^{-2} , while the ALS data has a density of 6 points m^{-2} . The study site, covering approximately 1 ha (100 m $\times$ 100 m), contains 1,350 individual trees. We apply a spatial hold-out strategy to partition the site into two data sets: one section includes 1,200 selected trees for training, and the other includes 150 trees for validation. From UAV lidar data, paired dense and sparse point clouds are generated for training, where the sparse point clouds are produced by resampling the dense UAV lidar data to simulate ALS-like distributions. For testing, ALS tree point clouds and their corresponding UAV lidar counterparts are used. Individual trees are first extracted from high-density UAV lidar point clouds. Although each segmented UAV lidar tree point cloud initially contains approximately 20,000 points, we reduce each to 8,192 points to obtain a representative yet computationally manageable resolution for training. These individual trees are then resampled to generate low-resolution point clouds that realistically replicate ALS data in both point density and spatial distribution: for each high-density individual tree, a low-density tree is randomly selected from ALS data, and resampling is performed based on the spatial distribution profile of the ALS point cloud. This process produces paired point clouds consisting of 8,192 points (ground truth) and 128 points (sparse input) for training and evaluation.

In addition to real-world data, 15,000 synthetic tree point clouds are generated (14,000 training, 1,000 test) to allow the network to observe fully intact tree structures. This complements UAV lidar data, which often lacks lower crown details due to occlusion and sensor limitations. These synthetic trees point clouds are created by first modeling trees using the Arbaro tree simulation software and then virtually scanning them with Helios++ to produce realistic lidar point clouds (Weber and Penn, 1995;

Winiwarter et al., 2022). Individual trees are scanned using the specifications of a TLS system (Riegl VZ-400 model, 0.3 mrad beam divergence, 0.05° angular resolution, and 1.5 m elevation above ground) to capture the complete structure of the trees. The same resampling procedure is applied to the synthetic point clouds to generate paired sparse and dense point clouds for training and validation.

2.2 Upsampling Network

ODGNet (Cai et al., 2024), originally developed for point cloud completion, is adapted in this study for the upsampling task. Specifically, the network is trained to perform 64 $\times$ upsampling of low-density ALS tree point clouds, enabling the reconstruction of detailed tree structures across the entire tree rather than only filling localized missing regions. ODGNet begins with a seed generation stage, where a U-Net-based network produces an initial coarse representation of the tree by generating seed points that capture the main tree skeleton, including the stem and primary branches. The second stage refines and densifies these seeds to recover fine-scale canopy and branch details.

In the first stage, feature representation of tree point clouds are generated through extracting global, mid-level, and local features in the encoder stage and then passed through concatenation to the decoder. The decoder then generates the initial seed representation of the tree point cloud by translating features back to geometric space. To mitigate information loss, Dictionary Guidance Modules that incorporate learned shape priors are used alongside concatenation operations. Specifically, each Dictionary module measures the similarity between input features and prior geometric primitives and selects the highest score to guide the refinement process. Finally, ODGNet utilizes Upsample Transformers (Zhou et al., 2022) to generate dense, high-resolution point clouds.

ODGNet uses Chamfer Distance (CD) as the primary loss function for both seed generation and upsampling stages. CD measures the similarity between predicted and ground-truth point clouds by averaging nearest-neighbor distances in both directions. Two variants, CD-L1 and CD-L2, compute this discrepancy using L1 and L2 norms, respectively. In this study, CD-L2 is used as the loss function. It is defined as:

$$CD(U, G) = \frac{1}{|U|} \sum_{u \in U} \min |u - g| + \frac{1}{|G|} \sum_{g \in G} \min |g - u| \quad (1)$$

The final loss function jointly accounts for both the seed generation and upsampling stages and is defined as:

$$L = CD_{TreeSeed} + \lambda CD_{Upsample} + \beta \sum L_{Orth} \quad (2)$$

Here, $CD_{TreeSeed}$ denotes the loss associated with the Tree Seed Generation stage, while $CD_{Upsample}$ corresponds to the loss for the upsampling stage. The term $\sum L_{Orth}$ represents an orthogonality constraint that encourages feature diversity within the learnable dictionary. The coefficients λ and β are hyperparameters that balance the contributions of the respective loss terms and are both set to 1 in this study.

2.3 Evaluation Metrics

Two categories of metrics are considered to evaluate the performance of the proposed framework: geometric metrics and structural metrics. Geometric metrics measure the spatial similarity between the upsampled and ground-truth tree point clouds, while structural metrics assess the accuracy of tree

structural attributes derived from the reconstructed tree point clouds. The geometric quality of the reconstructed point clouds is evaluated using three commonly adopted metrics in point cloud upsampling: CD-L1, CD-L2, and the F1-score. The F1-score evaluates reconstruction accuracy as the harmonic mean of precision and recall based on the proportion of points within a predefined distance threshold. In this study, the threshold is 0.01 (in sphere-normalized units).

To assess structural characteristics, three metrics are computed for each individual tree: tree height, crown area, and tree volume. Tree height is defined as the maximum height above ground among all tree points, where height above ground represents the vertical distance of each point from the local ground surface. Ground points are removed using the Cloth Simulation Filter (CSF) method to isolate tree points (Zhang et al., 2016). Crown area is computed by projecting all tree points onto the horizontal XY plane and fitting a two-dimensional convex hull to the resulting point set, with the enclosed area of this hull taken as the crown area. Tree volume is estimated by constructing a three-dimensional convex hull over all tree points and computing the volume of the resulting polytope. This provides a geometric approximation of tree volume derived from the point cloud. Both convex hull operations are performed using the Qhull algorithm as implemented in SciPy's ConvexHull routine (Virtanen et al., 2020).

3. Experiments

The performance of ODGNet is evaluated under three training strategies: (1) trained solely on synthetic data, (2) trained solely on real data (ODGNet), and (3) pre-trained on synthetic data and subsequently fine-tuned on real data (ODGNet*). Training strategy (1) is employed solely to initialize the model for subsequent experiments and is not intended for standalone evaluation or comparison on real data. All models are trained for 250 epochs, with the best weights selected based on validation performance. Training uses the AdamW optimizer (Loshchilov & Hutter, 2019) with an initial learning rate of 0.0004, halved every 50 epochs, and a batch size of 16. Table 1 presents the quantitative results of training on synthetic data across three input resolutions. The networks trained on synthetic data are then fine-tuned on real data to assess the role of synthetic pre-training as a means of encoding prior knowledge of tree geometry, which is subsequently leveraged for point cloud upsampling on real-world data. The geometric evaluation of ODGNet under training strategies (2) and (3) is reported in Table 2.

Number of Points	Up. Rate	CD-L1	CD-L2	F1 score
128	64×	11.2	0.491	0.498
256	32×	8.59	0.278	0.639
512	16×	7.18	0.180	0.729

Table 1. Quantitative upsampling results for ODGNet on synthetic dataset at 16×, 32×, and 64× upsampling. CD-L1 and CD-L2 are reported as scaled normalized distances ($\times 10^{-3}$), while the F1-score is unitless.

Input Points	Up. Rate	Method	CD-L1	CD-L2	F1 score
128	64×	ODGNet	11.7	0.679	0.575
		ODGNet*	11.0	0.607	0.603
256	32×	ODGNet	8.44	0.328	0.712
		ODGNet*	7.73	0.296	0.743
512	16×	ODGNet	6.78	0.188	0.805
		ODGNet*	6.41	0.183	0.814

Table 2. Quantitative comparison of ODGNet performance under training strategies (2) and (3) at 16×, 32×, and 64× upsampling: real-data training (ODGNet) and synthetic pretraining followed by fine-tuning (ODGNet*). CD-L1 and CD-L2 are reported as scaled normalized distances ($\times 10^{-3}$), while the F1-score is unitless.




Figure 1. Top and side views of ODGNet and ODGNet* results on UAV lidar tree point clouds, showing upsampled trees from sparse inputs (128 points: 100 tree, 28 ground) compared with ground-truth (GT) point clouds (8,192 points: 8,164 tree, 28 ground).




Figure 2. Top and side views of ODGNet and ODGNet* applied to ALS tree point clouds. Inputs retain their original point counts (including 28 ground points) rather than being resampled to 128 points. Upsampled trees are compared with UAV lidar ground-truth point clouds (8,192 points: 8,164 tree, 28 ground) and aligned using Iterative Closest Point (ICP) to correct positional shifts (Besl and McKay, 1992).

Input Points	Up. Rate	Method	Tree Height (m)			Crown Area (m ²)			Tree Volume (m ³)		
			R ²	RMSE	RMSE%	R ²	RMSE	RMSE%	R ²	RMSE	RMSE%
128	-	Sparse Input	0.926	0.364	2.3	0.114	3.52	39.6	-0.602	44.1	65.7
	64×	ODGNet	0.971	0.227	1.4	0.863	1.38	15.6	0.796	15.7	23.5
	64×	ODGNet*	0.976	0.209	1.3	0.873	1.33	15.0	0.781	16.3	24.3
256	-	Sparse Input	0.968	0.239	1.5	0.538	2.54	28.6	-0.023	35.2	52.5
	32×	ODGNet	0.982	0.180	1.1	0.938	0.934	10.5	0.878	12.1	18.1
	32×	ODGNet*	0.979	0.196	1.2	0.942	0.898	10.1	0.843	13.8	20.6
512	-	Sparse Input	0.979	0.194	1.2	0.780	1.57	19.7	0.291	29.3	43.7
	16×	ODGNet	0.991	0.129	0.8	0.979	0.543	6.1	0.909	10.5	15.6
	16×	ODGNet*	0.991	0.126	0.8	0.971	0.642	7.2	0.889	11.6	17.3

Table 3. Evaluation of structural metrics derived from reconstructed trees at 16×, 32×, and 64× upsampling rates. Results compare sparse inputs, ODGNet and ODGNet*. Metrics (R² and RMSE) are computed by comparing the estimated tree attributes with the corresponding ground-truth values for tree height, crown area, and tree volume.




Figure 3. Scatter plots comparing estimates of tree structural attributes (tree height, crown area, and tree volume) derived from sparse inputs and reconstructed point clouds (ODGNet and ODGNet*) against corresponding estimates derived from ground-truth point clouds. Tree height is shown in (a–c), crown area in (d–f), and tree volume in (g–i). Panels (a, d, g), (b, e, h), and (c, f, i) correspond to 128, 256, and 512 input points (64×, 32×, and 16× upsampling rates), respectively. The dashed line indicates the 1:1 agreement with ground-truth estimates.

3.1 Geometric Evaluation of Experiments

Table 1 reports geometric metrics for ODGNet trained on synthetic data. The learned weights are subsequently used as pretraining prior to fine-tuning on real data; therefore, the synthetic results are not intended as the final benchmark for comparison. Instead, they provide an initial evaluation of the model's reconstruction capability. Results obtained from 1,000 synthetic tree point clouds show stable performance across all resolutions, with improved geometric accuracy as the input resolution increases.

Results on real-world data, as shown in Table 2, indicate that pretraining on synthetic data followed by fine-tuning on real data (ODGNet*) consistently improves performance compared to training solely on real data (ODGNet) across all upsampling rates. ODGNet* achieves lower CD-L1 and CD-L2 values and higher F1-scores than ODGNet at the same upsampling rates. For both training strategies, the network performance improves as the input resolution increases. The performance gain from synthetic pretraining is noticeable at higher upsampling rates.

Figure 1 presents qualitative upsampling results for UAV lidar data (sparse and dense) under training strategies (2) and (3) at a 64× upsampling rate. Visual inspection shows that both strategies generate upsampled trees that are structurally much closer to the ground truth compared to the sparse inputs, which is consistent across all three samples. The main differences between the two strategies are observed in the lower parts of the trees. Trees generated by ODGNet* exhibit structural details that more closely resemble the ground truth, whereas the results from ODGNet show greater dispersion in the lower crown regions. This difference likely arises because ODGNet* is pretrained on a large number of synthetic trees with complete structures, allowing the model to learn full tree geometry. In contrast, ODGNet is trained only on real-world UAV lidar data, which often contain occlusions in the lower parts of trees.

To further demonstrate the practical applicability of the proposed upsampling network, the trained models are applied to real ALS scans of trees co-located with UAV lidar data (Figure 2). The models are trained using UAV lidar data. Since ALS and UAV lidar data are acquired by different sensors at different times, spatial shifts and structural differences exist between corresponding trees in the two datasets, making quantitative comparison challenging. Therefore, the results are evaluated qualitatively. As shown in Figure 2, both ODGNet and ODGNet* successfully reconstruct trees from sparse ALS inputs with strong overall structural similarity to the dense UAV lidar references. For Tree 1 and Tree 2, ODGNet* produces more detailed branching patterns and a better representation of the crown and lower canopy. In contrast, a large gap in the middle of the canopy and several smaller gaps in the upper crown can be observed in the reconstruction of Tree 1 generated by ODGNet. However, for Tree 3, ODGNet produces a reconstruction that more closely resembles the reference tree compared to ODGNet*, particularly in the lower canopy where some dispersion is visible in the ODGNet* reconstruction. Overall, the overlaid visualizations confirm that Both ODGNet and ODGNet* successfully reconstruct trees with strong overall structural similarity to the dense UAV lidar references from sparse ALS inputs (100 points per tree). However, ODGNet* produces slightly more detailed branching patterns and a better representation of the lower parts of the canopy.

3.2 Structural Evaluation of Experiments

Structural attributes (tree height, crown area, and tree volume) are derived from sparse inputs, reconstructed point clouds (ODGNet and ODGNet*), and ground-truth point clouds to evaluate the impact of upsampling on structural attribute estimation. Table 3 and Figure 3 present the evaluation results by comparing estimated attributes with their corresponding ground-truth values at 16×, 32×, and 64× upsampling rates. As shown in Table 3, both ODGNet and ODGNet* improve the accuracy of structural metric estimation compared to sparse inputs across all upsampling rates. The improvement is greater for crown area and tree volume, which are more structurally complex attributes than tree height. For example, at 128 input points, upsampling reduces the RMSE of crown area from 3.52 to 1.38 m^2 (ODGNet) and 1.33 m^2 (ODGNet*), and tree volume from 44.1 to 15.7 m^3 (ODGNet) and 16.3 m^3 (ODGNet*), while tree height improves more modestly from 0.364 to 0.227 m (ODGNet) and 0.209 m (ODGNet*). As the input resolution increases, performance improves across all metrics; however, the differences between sparse inputs and reconstructed trees remain noticeable for crown area and tree volume even at the highest input resolution (512 points).

A comparison between ODGNet and ODGNet* shows that synthetic pretraining generally improves the estimation of tree height and crown area, although the improvement is not consistent across all input resolutions. In contrast, ODGNet provides slightly more accurate tree volume estimates. This difference may be because tree volume is more sensitive to detailed crown structure, and models trained exclusively on real UAV lidar data may better adapt to the structural variability and irregularity present in natural tree crowns. Overall, tree volume remains the most challenging metric, as it depends on the complete three-dimensional structure of the tree crown, and even minor structural deviations can significantly affect volumetric estimates.

The scatter plots in Figure 3 further confirm the improvements reported in Table 3. Estimates derived from reconstructed point clouds (both ODGNet and ODGNet*) align more closely with the ground truth than those derived from sparse inputs. Even sparse point clouds generated under ALS-like distributions often capture the top of the tree and ground points, allowing relatively reliable estimation of tree height. In contrast, the larger dispersion of sparse-input estimates demonstrates that crown area and tree volume are more sensitive to missing structural information, which is improved through upsampling.

4. Conclusion

The proposed upsampling framework demonstrates promising results in addressing the challenge of reconstructing detailed tree structures from sparse ALS point clouds. Experimental results based on geometric and structural metrics, together with qualitative analyses, show that the upsampling network can generate realistic tree point clouds from low-density inputs that closely resemble high-density UAV lidar data in terms of both spatial geometry and forest structural attributes. Geometric evaluation confirms that the reconstructed trees exhibit strong spatial similarity to the ground truth. These metrics also indicate that synthetic pretraining produces reconstructions that more closely resemble the ground truth compared to training solely on real data. Qualitative results further demonstrate the applicability of the trained framework to real low-density ALS point clouds. The upsampling network trained on paired sparse–dense data derived from UAV lidar is applied to real ALS data. The results

show that both synthetic-pretrained and real-data-trained models produce slightly different reconstructions, but both successfully recover the overall tree structure and generate trees that more closely resemble dense UAV lidar observations than the original sparse ALS inputs.

Structural metrics further demonstrate that the reconstructed trees improve the estimation of forest structural attributes compared with sparse inputs. In particular, synthetic pretraining tends to improve the estimation of tree height and crown area across most input resolutions, although this trend is not consistent in all cases. In contrast, tree volume estimation is slightly more accurate when the model is trained solely on real data.

Overall, these findings demonstrate that the proposed upsampling framework can generate dense and realistic tree point clouds from sparse ALS data, providing a practical approach for bridging the wide spatial coverage of ALS with the detailed structural information typically obtained from UAV lidar.

5. References

- Kükenbrink, D., Schneider, F. D., Leiterer, R., Schaepman, M. E., & Morsdorf, F. (2017). Quantification of hidden canopy volume of airborne laser scanning data using a voxel traversal algorithm. *Remote Sensing of Environment*, 194, 424–436. <https://doi.org/10.1016/j.rse.2016.10.023>
- Li, R., Li, X., Fu, C.-W., Cohen-Or, D., & Heng, P.-A. (2019). *PU-GAN: a Point Cloud Upsampling Adversarial Network*. <http://arxiv.org/abs/1907.10844>
- Liu, H., Yang, M., Xi, B., Wang, X., Huang, Q., Xu, C., & Meng, W. (2025). MFCPopulus: A Point Cloud Completion Network Based on Multi-Feature Fusion for the 3D Reconstruction of Individual Populus Tomentosa in Planted Forests. *Forests*, 16(4), 635. <https://doi.org/10.3390/f16040635>
- Liu, H., Zhong, H., Lin, W., & Wu, J. (2024). Tree species classification based on PointNet++ deep learning and true-colour point cloud. *International Journal of Remote Sensing*, 45(16), 5577–5604. <https://doi.org/10.1080/01431161.2024.2377837>
- Liu, H., Zhong, H., Xie, G., & Zhang, P. (2025a). Research on Tree Point Cloud Enhancement Based on Deep Learning. *Forests*, 16(6), 915. <https://doi.org/10.3390/f16060915>
- Liu, H., Zhong, H., Xie, G., & Zhang, P. (2025b). Tree Species Classification Based on Point Cloud Completion. *Forests*, 16(2), 280. <https://doi.org/10.3390/f16020280>
- Qian, G., Abualshour, A., Li, G., Thabet, A., & Ghanem, B. (2021). *PU-GCN: Point Cloud Upsampling using Graph Convolutional Networks*. <http://arxiv.org/abs/1912.03264>
- Schneider, F. D., Kükenbrink, D., Schaepman, M. E., Schimel, D. S., & Morsdorf, F. (2019). Quantifying 3D structure and occlusion in dense tropical and temperate forests using close-range LiDAR. *Agricultural and Forest Meteorology*, 268, 249–257. <https://doi.org/10.1016/j.agrformet.2019.01.033>
- Su, Y., Chen, Z., & Xue, X. (2025). TreeDBH: Dual Enhancement Strategies for Tree Point Cloud Completion in Medium–Low Density UAV Data. *Forests*, 16(4), 667. <https://doi.org/10.3390/f16040667>
- Abdollahi, A., Yebra, M., 2025. Challenges and Opportunities in Remote Sensing-Based Fuel Load Estimation for Wildfire Behavior and Management: A Comprehensive Review. *Remote Sens (Basel)* 17, 415. <https://doi.org/10.3390/rs17030415>
- Yu, L., Li, X., Fu, C.-W., Cohen-Or, D., & Heng, P.-A. (2018). *PU-Net: Point Cloud Upsampling Network*. <http://arxiv.org/abs/1801.06761>
- Balestra, M., Marselis, S., Sankey, T.T., Cabo, C., Liang, X., Mokroš, M., Peng, X., Singh, A., Stereńczak, K., Vega, C., Vincent, G., Hollaus, M., 2024. LiDAR Data Fusion to Improve Forest Attribute Estimates: A Review. *Current Forestry Reports* 10, 281–297. <https://doi.org/10.1007/s40725-024-00223-7>
- Besl, P.J., McKay, N.D., 1992. A method for registration of 3-D shapes. *IEEE Trans Pattern Anal Mach Intell* 14, 239–256. <https://doi.org/10.1109/34.121791>
- Cai, P., Scott, D., Li, X., Wang, S., 2024. Orthogonal Dictionary Guided Shape Completion Network for Point Cloud. *Proceedings of the AAAI Conference on Artificial Intelligence* 38, 864–872. <https://doi.org/10.1609/aaai.v38i2.27845>
- Hamraz, H., Contreras, M.A., Zhang, J., 2017. Forest understory trees can be segmented accurately within sufficiently dense airborne laser scanning point clouds. *Sci Rep* 7. <https://doi.org/10.1038/s41598-017-07200-0>
- Moran, C.J., Kane, V.R., Seielstad, C.A., 2020. Mapping Forest Canopy Fuels in the Western United States with LiDAR–Landsat Covariance. *Remote Sens (Basel)* 12, 1000. <https://doi.org/10.3390/rs12061000>
- Terryn, L., Calders, K., Bartholomeus, H., Bartolo, R.E., Brede, B., D’hont, B., Disney, M., Herold, M., Lau, A., Shenkin, A., Whiteside, T.G., Wilkes, P., Verbeeck, H., 2022. Quantifying tropical forest structure through terrestrial and UAV laser scanning fusion in Australian rainforests. *Remote Sens Environ* 271, 112912. <https://doi.org/10.1016/j.rse.2022.112912>
- Toivonen, J., Kangas, A., Maltamo, M., Kukkonen, M., Packalen, P., 2023. Assessing biodiversity using forest structure indicators based on airborne laser scanning data. *For Ecol Manage* 546, 121376. <https://doi.org/10.1016/j.foreco.2023.121376>
- Weber, J., Penn, J., 1995. Creation and rendering of realistic trees, in: *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques - SIGGRAPH '95*. ACM Press, New York, New York, USA, pp. 119–128. <https://doi.org/10.1145/218380.218427>
- Winiwarter, L., Esmoris Pena, A.M., Weiser, H., Anders, K., Martínez Sánchez, J., Searle, M., Höfle, B., 2022. Virtual laser scanning with HELIOS++: A novel take on ray tracing-based simulation of topographic full-waveform 3D laser scanning. *Remote Sens Environ* 269, 112772. <https://doi.org/10.1016/j.rse.2021.112772>
- Xiang, B., Wielgosz, M., Kontogianni, T., Peters, T., Puliti, S., Astrup, R., Schindler, K., 2024. Automated forest inventory: Analysis of high-density airborne LiDAR point clouds with 3D deep learning. *Remote Sens Environ* 305, 114078. <https://doi.org/10.1016/j.rse.2024.114078>
- Zhang, D., Král, K., Krůček, M., Cushman, K.C., Kellner, J.R., 2024. Near-Complete Sampling of Forest Structure from High-

Density Drone Lidar Demonstrated by Ray Tracing. *Remote Sens (Basel)* 16, 2774. <https://doi.org/10.3390/rs16152774>

Loshchilov, I., & Hutter, F. (2019). *Decoupled Weight Decay Regularization*. <http://arxiv.org/abs/1711.05101>

Zhang, W., Qi, J., Wan, P., Wang, H., Xie, D., Wang, X., & Yan, G. (2016). An Easy-to-Use Airborne LiDAR Data Filtering Method Based on Cloth Simulation. *Remote Sensing*, 8(6), 501. <https://doi.org/10.3390/rs8060501>

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., ... Vázquez-Baeza, Y. (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>

Petras, V., Petrasova, A., McCarter, J. B., Mitasova, H., & Meentemeyer, R. K. (2023). Point Density Variations in Airborne Lidar Point Clouds. *Sensors*, 23(3), 1593. <https://doi.org/10.3390/s23031593>

Zhou, H., Cao, Y., Chu, W., Zhu, J., Lu, T., Tai, Y., & Wang, C. (2022). *SeedFormer: Patch Seeds based Point Cloud Completion with Upsample Transformer*. <http://arxiv.org/abs/2207.10315>