Synthetic Forest: A UAV Laser Scanning Benchmark Dataset for Individual Tree Segmentation, Classification, and Wood Volume Estimation

Habib Pourdelan[1], Martin Tomko[1], Kourosh Khoshelham[1]

[1] The University of Melbourne
(hpourdelan@student.unimelb.edu.au, tomkom@unimelb.edu.au, k.khoshelham@unimelb.edu.au)

Keywords: UAV LiDAR simulation, Synthetic point clouds, 3D forest modeling, Instance segmentation, Biomass estimation

Abstract

Accurate tree-level analysis in forests via LiDAR scanning is essential for biomass estimation, canopy structure assessment, and carbon monitoring, yet remains constrained by the scarcity of large-scale annotated LiDAR datasets and the high cost of manual annotation. To address this, we present a novel approach that integrates 3D tree models with UAV-borne LiDAR simulation to generate synthetic forest point clouds with comprehensive annotations. Our approach generates diverse woodland, open, and closed forest structures, producing Synthetic Forest, a benchmark dataset of three 1 ha scenes containing 38–47 million points each, with densities of 3302–3862 points/m² and average spacing of 2 cm. Each scene contains between 70 and 216 individual trees, along with understory vegetation, deadwood, stumps, rocks, and bushes, all automatically annotated with semantic classification IDs, instance IDs, and tree IDs for volume estimation. The proposed pipeline provides automated, error-free ground truth for leaf-wood classification, instance segmentation, and wood volume estimation. We provide guidelines for generating forest plots and utilizing the datasets for diverse forestry tasks. By eliminating the need for costly field data collection, our pipeline offers scalable, customizable synthetic datasets that accelerate forest inventory. The Synthetic Forest dataset is publicly released via Zenodo (DOI: 10.5281/zenodo.17568131), enabling reproducible research and supporting further developments in forest monitoring and management.

1. Introduction

Forests play a pivotal role in global ecosystems, contributing to biodiversity conservation, carbon sequestration, and climate regulation (Bohn and Huth, 2017; Sanquetta et al., 2011). Light Detection and Ranging (LiDAR) technology has transformed forest monitoring by providing high-resolution, three-dimensional mapping of canopy structure, biomass, and individual tree attributes (Borsah et al., 2023; Tamiminia et al., 2021).

Recent advances in deep learning have shown great promise for improving the accuracy of forest mapping and enabling large-scale identification of multiple tree species. However, the development of robust deep learning models requires large benchmark dataset with high-quality annotations, which remain scarce in forestry applications. In particular, the use of LiDAR for tree-level analysis, such as leaf–wood classification, instance segmentation, and volume estimation, is limited by the availability of annotated data (Gaydon and Roche, 2025; Liang et al., 2025). Real-world LiDAR acquisition is resource-intensive and frequently limited by the substantial cost of manual annotation (Weiser et al., 2022; Puliti et al., 2023).

Utilizing synthetic data generation emerges as a promising solution, as demonstrated in works on simulated point clouds for model validation (Feng et al., 2025; Xiang et al., 2024; Zhao et al., 2024). Several studies have employed synthetic data to improve forest analysis. Song et al. (2023) generated 3D tree structures using Arbaro and simulated LiDAR point clouds across different forest scenarios. These synthetic datasets were then used to train machine learning models for above-ground biomass estimation, demonstrating improved accuracy through data augmentation. Similarly, Wang (2020) employed SpeedTree to construct synthetic forest environments and used HELIOS++ to simulate corresponding point clouds, enabling effective single-tree segmentation and classification with high accuracy in structural measurements relevant to AGB estimation. In a related study, Bryson et al. (2023) introduced a tree simulation framework for producing synthetic LiDAR data, showing that models such as PointNet++ trained on these data can achieve performance comparable to, or out-perform, models trained on limited real-world datasets, particularly for stem segmentation tasks across varied forest conditions. A recent study demonstrated the effectiveness of using synthetic LiDAR point clouds with known ground truth to train deep regression networks for wood volume and above-ground biomass estimation. By leveraging synthetic data for model training, the approach reduces reliance on costly field measurements while enabling scalable learning across diverse forest conditions. The trained models were successfully transferred to real-world datasets, showing improved accuracy compared to traditional allometric methods and highlighting the potential of synthetic-to-real learning in forest analysis (Pourdelan et al., 2026). However, most existing synthetic forest datasets are limited in scale and structural diversity, highlighting the need for a flexible method to generate synthetic datasets tailored to every forest type for broader applicability and improved accuracy. She et al. (2025) introduced a synthetic data generation pipeline using game-engine environments and physics-based LiDAR simulation to produce large-scale annotated forest datasets for tree segmentation. Their experiments demonstrated that synthetic data can substantially reduce the requirement for labeled real-world data through effective pre-training. Specifically, they showed that after fine-tuning on just a single real forest plot of less than 0.1 hectare, a pre-trained model achieves segmentation performance competitive with models trained on full-scale real datasets. While their study identified physics, diversity, and scale as the critical factors for robust 3D forest vision systems, their dataset lacks pre-computed wood volume annotations. This is an essential metric for monitoring wood supply and forest carbon balance, highlighting the need for a large-scale dataset with volume ground truth for every individual tree (Abegg et al., 2023).

In this study, we introduce a novel pipeline for generating synthetic LiDAR point clouds, focusing on Australian Eucalyptus forests. By combining 3D tree models, physics-based LiDAR simulation, and automated annotation, our approach produces large-scale, fully labeled datasets for tree-level analysis. Importantly, we calculate the volume for each individual tree and provide this ground truth alongside the dataset for every tree in the scene. This work addresses the limitations of real-world LiDAR data by enabling reproducible, cost-efficient generation of realistic, annotated forest scenes.

The key contributions of this paper are as follows:

- Synthetic data generation pipeline: A fully automated framework integrating 3D modeling, LiDAR simulation, and Python post-processing to produce annotated point clouds in .pcd format.
- Benchmark datasets (Synthetic Forest): Three 1 ha Eucalyptus forest scenes, such as woodland, open, and closed forest, with detailed annotations for tree components, ground, deadwood, stumps, rocks, and bushes.
- Applications: The datasets enable diverse forestry tasks, including leaf–wood classification, individual tree instance segmentation, and wood volume estimation, serving as a benchmark for training and validating machine learning models.

This paper is organized as follows: Section 2 details the methodology, including scene generation, ground-truth volume computation, and LiDAR simulation; Section 3 presents results on point cloud characteristics and tree-level analysis approaches; and Section 4 concludes with key implications and future directions.

2. Methodology

The dataset generation process begins with the creation of a synthetic 1 ha Eucalyptus forest scene, incorporating detailed tree models and understory elements such as bushes, rocks, and deadwood. Prior to simulation, wood volumes are computed for each tree based on their mesh components, providing ground-truth metrics integrated into the scene description. This is followed by LiDAR scanning via Helios++ (Esmorís et al., 2022) to produce annotated point clouds, with all elements automatically labeled to support applications such as instance segmentation and volume estimation. Figure 1 shows an overview of the synthetic data generation process.

2.1 Scene generation

To create synthetic tree representations in our simulated forest scene, realistic 3D mesh models are essential. While various methods exist for generating lifelike 3D tree meshes, such as the Blender (Blender Foundation, 2023) Sapling Tree add-on for procedural creation or SpeedTree software for high-fidelity modeling, we acquired 3D mesh models of Eucalyptus trees developed by Yankobe (2023). We developed a Python-based pipeline to construct the forest scene, enabling the simulation of a realistic 1 ha Eucalyptus plantation spanning a flat 100 m

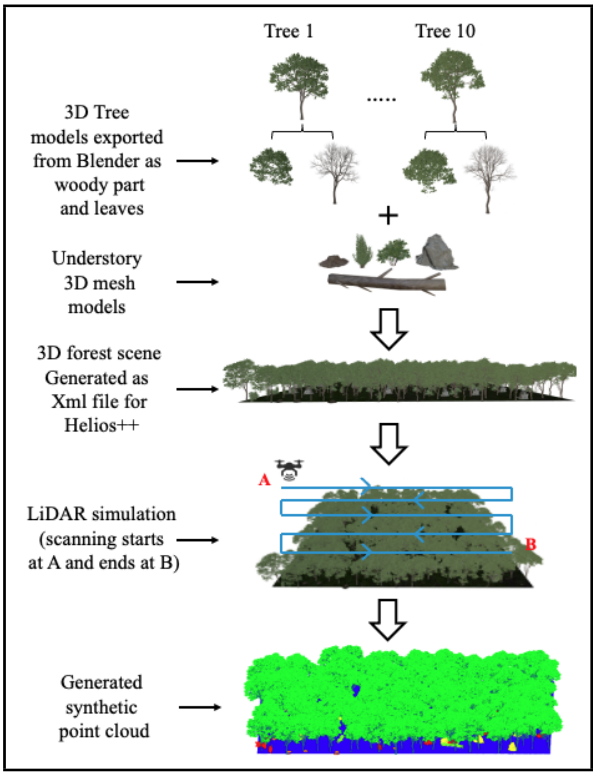

Figure 1. Generating synthetic forest point cloud pipeline

by 100 m ground area. This pipeline generates an XML configuration file that defines the placement and orientation of all scene elements. For the Eucalyptus trees, each 3D model is first processed to separate leaves from woody parts, with these components exported as individual .obj files into one of Helios++'s designated folders. Similarly, understory and ground features, including bushes, rocks, deadwood, and terrain surfaces, are exported as .obj files to the same folder. To enhance realism, trees are randomly positioned across the flat ground, with rotations applied around the Z-axis to vary canopy patterns when viewed from above, ensuring that repeated tree models (e.g., multiple instances of the same tree type, potentially up to 20 for a single variant) do not exhibit uniform canopy profiles when viewed from overhead. Additional small-angle rotations (±5 degrees) around the X and Y axes are introduced to tilt trees slightly, avoiding unnaturally straight alignments and mimicking natural growth variations. Understory elements are also randomly distributed throughout the scene to create more realistic understory vegetation. We generated three types of Eucalyptus forest scenes to align with Australia's national classification of native forests based on crown cover: woodland (20-50% canopy coverage), open (50-80%), and closed (80-100%) (National Forest Inventory, 2003). This classification reflects the proportion of ground area covered by tree canopies, enabling our synthetic datasets to represent diverse ecological conditions observed in real Eucalyptus plantations. Each object and component in the scene is assigned a unique identifier, which persists through the simulation process, allowing for precise labeling in the resulting point cloud data. The XML file references these .obj meshes from the folder, and Helios++ assembles them into a unified 3D environment, ready for LiDAR simulation and point cloud generation. All object positions and orientations within the scene are generated and managed entirely by the Python pipeline, ensuring precise and reproducible placement control.

https://doi.org/10.5194/isprs-annals-XI-3-2026-245-2026 |

2.2 Ground truth volume computation

Wood volumes are computed through a structured mesh processing pipeline applied to 10 distinct types of Eucalyptus 3D meshes. For each tree model, the triangulated mesh is first separated into distinct components, such as stems, branches, and twigs, using Blender's native "separate by loose parts" function. These components are then exported from Blender and processed with PyMeshFix (Attene, 2010), a Python library for repairing mesh defects and ensuring watertight integrity. Following this, the Trimesh (Dawson-Haggerty et al., 2019) library is used to calculate the volume of each watertight component, with the total wood volume for a tree derived by summing the volumes of all its parts. This process is repeated for each of the 10 tree models, and all woody elements along with additional objects like deadwood are similarly annotated to facilitate ground truth data for volume estimation in subsequent algorithms.

2.3 LiDAR simulation

Helios++ simulates laser scanning of the assembled forest scenes using a virtual Riegl VUX-1 UAV 22 sensor model mounted on an unmanned aerial vehicle (UAV). Table 1 shows the key scanner settings used in the simulation to replicate realistic UAV-borne LiDAR acquisition. As illustrated in Figure 1, a scanner trajectory is generated to cover the entire 1 ha area, ensuring comprehensive data capture across all trees and understory components in the simulated forest scene. The scene is scanned along multiple legs of this trajectory, producing separate point cloud segments for each pass. We developed a Python script to merge these individual point clouds into a single cohesive .pcd file. This combined point cloud integrates geometric data with unique identifiers derived from the XML scene description, preserving exact annotations for each object and tree component. To incorporate wood volume calculations, tree IDs are assigned to each tree, allowing volumes to be read directly from a corresponding .txt file generated during the mesh processing stage. Additionally, unique IDs for every scene part are stored in a text file, which can be referenced for tasks such as classification and instance segmentation.

Parameter	Value
Flight altitude	30 m
Flight speed	2 m/s
Scan frequency	50 Hz
Pulse frequency	200 kHz

Table 1. Key scanner settings for the HELIOS++ simulation to mimic UAV-borne LiDAR.

3. Results

3.1 Wood volume computation from 3D meshes

By separating the woody components of each tree, we employed PyMeshFix to ensure all mesh components were watertight, addressing potential geometric inconsistencies. Subsequently, the Trimesh library was utilized to calculate the volume of each individual part, enabling accurate estimation of total wood volume per tree. Figure 2 illustrates a branch from Tree 1 that was repaired using our method, where (a) shows the original branch with surface holes and (b) presents the fixed version after mesh correction. This process was applied to the 10 distinct Eucalyptus tree models and deadwood 3d mesh, providing ground-truth data for subsequent volume estimation tasks. Table 2 presents the number of components per tree and the estimated wood volume for each object, derived from the mesh computations.

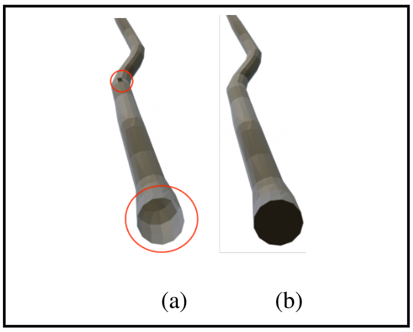

Figure 2. Comparison between (a) not fixed and (b) fixed tree component, where red circles in (a) indicate holes and flaws in the branch mesh.

Object name	Wood volume (m^3)
Tree 1	0.666
Tree 2	0.914
Tree 3	0.975
Tree 4	0.299
Tree 5	1.364
Tree 6	1.150
Tree 7	0.666
Tree 8	0.711
Tree 9	0.398
Tree 10	0.837
Dead wood	0.760

Table 2. Calculated wood volume for each object

3.2 Characteristics of the point clouds

The synthetic datasets generated through the pipeline described in Section 2 provide a diverse set of point clouds representing Eucalyptus-dominated forest environments at varying densities. Figure 3 illustrates the generated point clouds for these three distinct forest types (woodland, open, and closed forest) each spanning a 1 ha area (100 m $\times$ 100 m) on flat terrain. Key characteristics of these synthetic scenes are summarized in Table 3, including forest type, area, number of trees, point density, sensor type, and comparison with selected real-world datasets from the FOR-instance benchmark. The woodland scene includes 70 trees with 38% canopy cover, promoting sparser point distributions and greater visibility of understory elements. The open forest features 121 trees and 63% coverage, balancing canopy density with occasional gaps. The closed forest, with 216 trees and 84% coverage, exhibits the highest structural complexity, resulting in denser occlusions and richer vertical layering.

3.3 Point cloud structure

Following scene assembly and LiDAR simulation, the resulting point clouds (.pcd format) were post-processed to integrate annotations for downstream analysis. Each file contains 38-47 million points, with an average point spacing of 2 cm across all forest types, and overall density ranging from 3302 to 3862 points/m², influenced by canopy interactions and scan geometry. The point clouds are structured with 10 fields per point, as detailed in Table 4, encompassing spatial coordinates, semantic labels, and LiDAR metadata. These fields support a range of

https://doi.org/10.5194/isprs-annals-XI-3-2026-245-2026

Dataset	Country	Forest Type	Area (ha)	Number of Trees	Point density (pts/m^2)	Sensor
FOR-instance						
RMIT	Australia	Eucalyptus	0.37	223	498	Riegl miniVUX-1 UAV
NIBIO	Norway	Coniferous boreal	1.21	575	9529	Riegl miniVUX-1 UAV
CULS	Czech Republic	Coniferous temperate	0.33	47	2585	Riegl VUX-1 UAV
TU_WIEN	Austria	Deciduous alluvial	0.55	150	1717	Riegl VUX-1 UAV
SCION	New Zealand	Coniferous temperate	0.33	135	4576	Riegl miniVUX-1 UAV
Synthetic Forest						
Woodland	Australia	Eucalyptus	1.0	70	3302	Virtual Riegl VUX-1 UAV 22
Open forest	Australia	Eucalyptus	1.0	121	3593	Virtual Riegl VUX-1 UAV 22
Closed forest	Australia	Eucalyptus	1.0	216	3862	Virtual Riegl VUX-1 UAV 22

Table 3. Comparison of point cloud characteristics between real-world forest datasets (For-Instance) and Synthetic Forest

analytical tasks, such as semantic classification (e.g., leaf-wood separation), instance segmentation (e.g., individual tree segmentation), and tree volume estimation.

Field index	Field name
1	treeid
2	x
3	y
4	z
5	classification
6	instance
7	object_id
8	Number_of_Returns
9	Return_Number
10	Intensity

Table 4. Point cloud file column structure

3.4 Approaches for tree-level analysis in synthetic forest scenes

The Synthetic Forest, with its detailed annotations, enables advanced tree-level analyses that are challenging in real-world data due to annotation costs and variability. This section outlines key approaches for leveraging this dataset in tasks such as semantic classification of tree components, wood volume estimation, and instance segmentation at the individual tree scale.

3.4.1 Classification of tree woody and leaf components

Table 5 lists the classification IDs used in the point cloud structure for the synthetic forest scenes. These IDs describe the type of each point in the dataset, including ground, woody parts of trees, leaves, deadwood, and other objects. Figure 4 illustrates a small representative tile (20 m × 20 m) from the closed forest scene, visualizing the distribution of these elements. For the purpose of tree part classification, the synthetic forest scenes can be used to train and evaluate models that distinguish woody parts from leaves. In this context:

- Points with `classification` = 2 correspond to tree woody parts.
- Points with `classification` = 3 correspond to tree leaves.

By filtering points based on these classification IDs, it is possible to separate woody and leafy components for further analysis, such as biomass estimation or tree structure reconstruction.

Class ID	Description
1	ground
2	tree_wood
3	tree_leaves
4	deadwood
5	rock
6	stump
7	bush

Table 5. Classification IDs used in the point cloud structure.

3.4.2 Wood volume estimation

Our synthetic forest scenes can be used to estimate the wood volumes of individual trees and deadwood. Each 1 ha scene contains multiple tree instances,

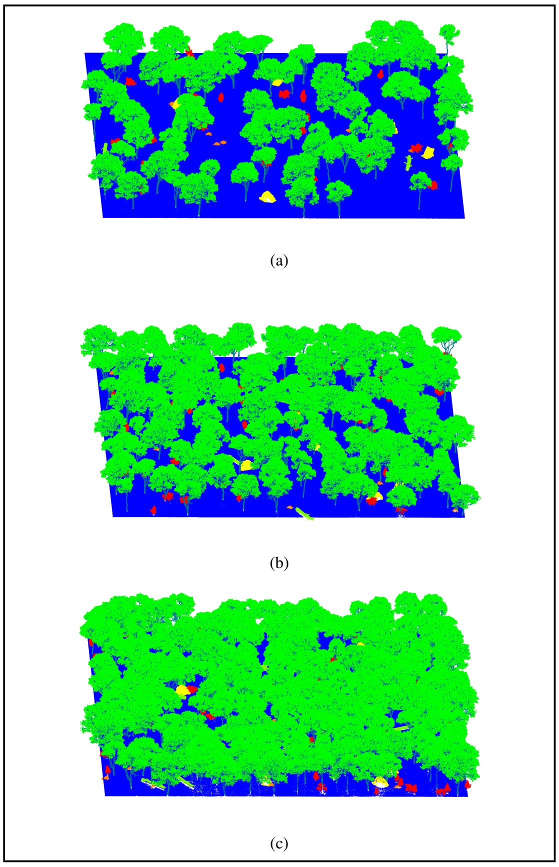

Figure 3. Synthetic Forest dataset scenes (100 m × 100 m): (a) Woodland; (b) Open forest; (c) Closed forest.

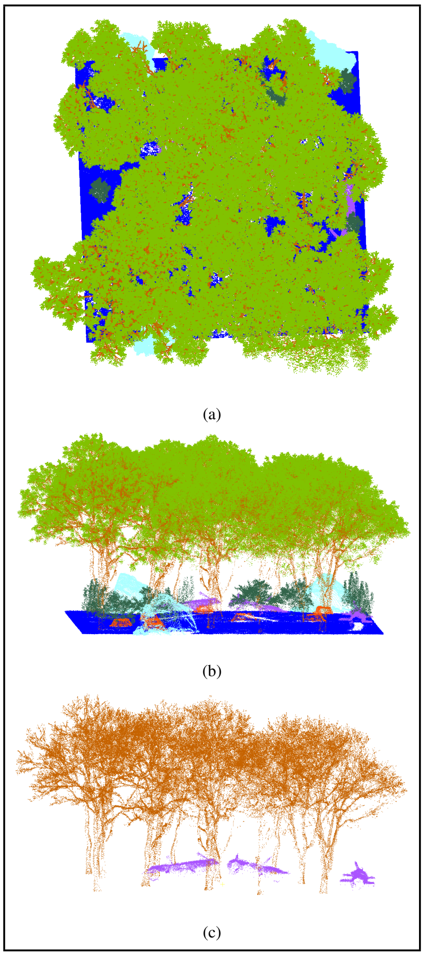

Figure 4. Visualization of a small tile (20 m × 20 m) from the closed forest scene: (a) Top view showing tree distribution; (b) Front view with all components (ground (blue), tree stem (brown), leaves (light green), deadwood (purple), rocks (cyan), stumps (red), bushes (dark green)) visible; (c) Woody components (trees and deadwood only).

but only 10 distinct Eucalyptus trees, which may be repeated across the scene. These trees are uniquely identified by tree IDs 1–10 in the `treeid` field. Deadwood elements are grouped under tree ID 11, while all other components (e.g., ground, rocks, stumps, bushes) are assigned tree ID 0 and should be ignored as noise during estimation. The datasets enable training deep learning models to predict volumes based on point groupings, with ground-truth values available from mesh computations (see Section 2 for details). For woody-specific volume (excluding leaves), combine the `treeid` and `classification` fields to filter relevant points and only consider tree wood points for this purpose, as visualized in Figure 4(c), which shows the woody components.

3.4.3 Tree-level instance segmentation For instance segmentation, the `instance` field assigns a unique identifier to each object in the scene. For trees, both woody parts and leaves are grouped under the same instance ID, ensuring that each individual tree is represented as a single instance. Other objects, such as deadwood or understory vegetation, are also assigned their own unique instance IDs. Figure 5 depicts the instance labeling of individual trees with trees highlighted in distinctive colors and other elements visible for context. To use the dataset for individual tree segmentation tasks, the `instance` field should be combined with the `treeid` field. This allows filtering of non-relevant objects, where ground points are colored blue and understory elements (deadwood, bushes, rocks, stumps) are colored cyan, while retaining only tree-related instances for analysis, as shown in Figure 5(c).

4. Conclusion

This work presents a novel approach for generating synthetic LiDAR point clouds of forests, addressing key challenges in tree-level analysis where real datasets are limited and manual annotation is labor-intensive and time-consuming. By combining 3D tree models, Blender graphics, LiDAR simulation, and Python-based post-processing, this pipeline generates forest datasets optimized for machine learning applications with automatic annotations in .pcd format. By simulating diverse woodland, open, and closed forest structures, we generated Synthetic Forest, a benchmark dataset of Australian Eucalyptus forests. Specifically, this comprises three 1 ha scenes with high point densities of 3302–3862 points/m² and an average point spacing of 2 cm, enabling fine-grained tree-level detail. These detailed synthetic point clouds are enriched with comprehensive annotations, including classification IDs for semantic separation, instance IDs for per-tree segmentation, and tree IDs for volume estimation. This pipeline enables the generation of any forest type in any size, solving the persistent problem of limited annotated data by providing scalable, error-free ground truth without field collection costs. By providing a large-scale annotated source, this dataset can support pre-training of models which can then be fine-tuned on smaller, real-world datasets via domain adapta-

https://doi.org/10.5194/isprs-annals-XI-3-2026-245-2026 | © Author(s) 2026. CC BY 4.0 License.

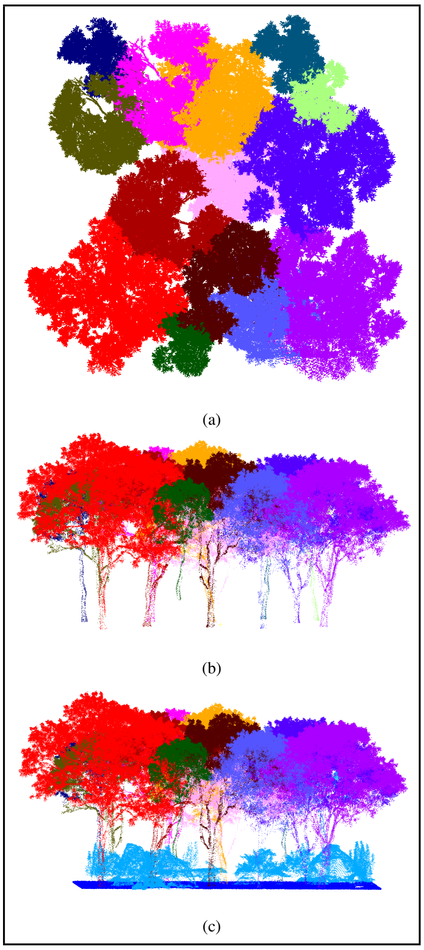

Figure 5. The instance labeling of individual tree in a small tile (20 m × 20 m) from the closed forest scene: (a)Top view; (b) Front view showing tree only; (c) Front view including all components, with labeled trees in distinctive colors and other elements (ground, deadwood, rocks, stumps, and bushes).

tion, effectively reducing the domain gap inherent in simulated LiDAR data. Future work will extend this approach to multi-species forests and incorporate additional LiDAR sensors to enhance simulation realism, paving the way for more robust forest analysis and modeling.

5. Acknowledgments

The first author acknowledges the financial support provided by the University of Melbourne through the Melbourne Research Scholarship.

References

Abegg, M., Bösch, R., Kükenbrink, D., Morsdorf, F., 2023. Tree volume estimation with terrestrial laser scanning—Testing for bias in a 3D virtual environment. *Agricultural and Forest Meteorology*, 331, 109348.

Attene, M., 2010. A lightweight approach to repairing digitized polygon meshes. *The visual computer*, 26, 1393–1406.

Blender Foundation, 2023. Blender.

Bohn, F. J., Huth, A., 2017. The importance of forest structure to biodiversity–productivity relationships. *Royal Society open science*, 4(1), 160521.

Borsah, A. A., Nazeer, M., Wong, M. S., 2023. LIDAR-Based Forest Biomass Remote Sensing: A Review of Metrics, Methods, and Assessment Criteria for the Selection of Allometric Equations. *Forests*, 14(10).

Bryson, M., Wang, F., Allworth, J., 2023. Using synthetic tree data in deep learning-based tree segmentation using LiDAR point clouds. *Remote Sensing*, 15(9), 2380.

Dawson-Haggerty et al., 2019. trimesh.

Esmorís, A. M., Yermo, M., Weiser, H., Winiwarter, L., Höfle, B., Rivera, F. F., 2022. Virtual LiDAR Simulation as a High Performance Computing Challenge: Toward HPC HELIOS++. *IEEE Access*, 10, 105052–105073.

Feng, Z., She, Y., Keshav, S., 2025. SPREAD: A large-scale, high-fidelity synthetic dataset for multiple forest vision tasks. *Ecological Informatics*, 87, 103085.

Gaydon, C., Roche, F., 2025. Pureforest: A large-scale aerial lidar and aerial imagery dataset for tree species classification in monospecific forests. *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, IEEE, 5895–5904.

Liang, X., Qi, H., Deng, X., Chen, J., Cai, S., Zhang, Q., Wang, Y., Kukko, A., Hyyppä, J., 2025. ForestSemantic: a dataset for semantic learning of forest from close-range sensing. *Geo-spatial Information Science*, 28(1), 185–211.

National Forest Inventory, 2003. Australia’s state of the forests report 2003. Technical report, Bureau of Rural Sciences, Canberra, Australia.

Pourdelan, H., Xiang, Z., Stewart, H., Nicholson, C., Tomko, M., Khoshelham, K., 2026. Direct Estimation of Tree Volume and Aboveground Biomass Using Deep Regression with Synthetic Lidar Data. *arXiv preprint arXiv:2603.04683*.

Puliti, S., Pearse, G., Surovỳ, P., Wallace, L., Hollaus, M., Wielgosz, M., Astrup, R., 2023. For-instance: a uav laser scanning benchmark dataset for semantic and instance segmentation of individual trees. *arXiv preprint arXiv:2309.01279*.

Sanquetta, C., Dalla Corte, A., Maas, G. B., 2011. The role of forests in climate change. *Quebracho-Revista de Ciencias Forestales*, 19(1-2), 84–96.

She, Y., Blake, A., Coomes, D., Keshav, S., 2025. Scaling Up Forest Vision with Synthetic Data. *arXiv preprint arXiv:2509.11201*.

Song, Q., Wang, Y., Zhu, X. X., 2023. 3d point cloud simulation for above-ground forest biomass estimation. *IGARSS 2023-2023 IEEE International Geoscience and Remote Sensing Symposium*, IEEE, 1357–1360.

Tamiminia, H., Salehi, B., Mahdianpari, M., Beier, C. M., Johnson, L., Phoenix, D. B., 2021. A comparison of decision tree-based models for forest above-ground biomass estimation using a combination of airborne lidar and landsat data. *ISPRS annals of the photogrammetry, remote sensing and spatial information sciences*, 3, 235–241.

Wang, D., 2020. Unsupervised semantic and instance segmentation of forest point clouds. *ISPRS Journal of Photogrammetry and Remote Sensing*, 165, 86–97.

Weiser, H., Schäfer, J., Winiwarter, L., Krašovec, N., Fassnacht, F. E., Höfle, B., 2022. Individual tree point clouds and tree measurements from multi-platform laser scanning in German forests. *Earth System Science Data*, 14(7), 2989–3012.

Xiang, Z., Huang, Z., Khoshelham, K., 2024. Synthetic lidar point cloud generation using deep generative models for improved driving scene object recognition. *Image and Vision Computing*, 150, 105207.

Yankobe, 2023. Eucalypt pack. Sketchfab. Licensed under Creative Commons Attribution 4.0 International (CC-BY-4.0).

Zhao, H., Zhao, Y., Tomko, M., Khoshelham, K., 2024. MoLiPoseNet: Model-based indoor relocalization using deep pose regression from synthetic LiDAR scans. *IEEE Sensors Journal*.