This manuscript is a preprint submitted to EarthArXiv and has not been peer reviewed.

Graphical Abstract

From 2D labels to 3D structure: Scalable label transfer and benchmarking of 3D vegetation models in rangeland ecosystems

Laura N. Sotomayor, Arko Lucieer, Darren Turner, Johan Barthelemy, Teja Kattenborn

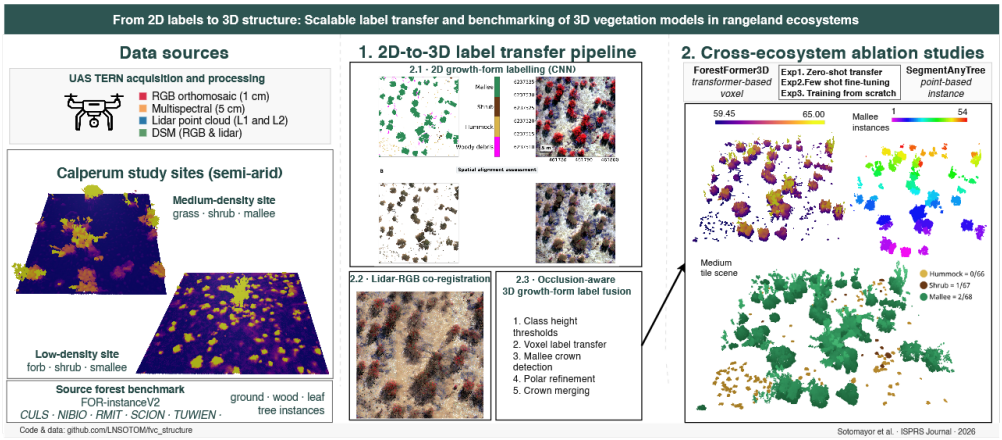

From 2D labels to 3D structure: Scalable label transfer and benchmarking of 3D vegetation models in rangeland ecosystems

Laura N. Sotomayor[a], Arko Lucieer[a], Darren Turner[a], Johan Barthelemy[b,c], Teja Kattenborn[d]

[a]*School of Geography, Planning, and Spatial Sciences, Hobart, Tasmania, Australia*
[b]*Securing Antarctica's Environmental Future (SAEF), University of Wollongong, Wollongong, , NSW, Australia*
[c]*NVIDIA NVAITC, Australia*
[d]*Sensor-based Geoinformatics (geosense), University of Freiburg, Freiburg, Breisgau, Germany*

Abstract

Three-dimensional (3D) characterisation of vegetation structure at the level of individual growth forms is critical for understanding ecosystem function and resilience, yet remains challenging in rangelands because vegetation is sparse, low-stature, and structurally heterogeneous. Recent 3D deep-learning models perform strongly in forests, but their transfer beyond closed-canopy benchmarks is limited by scarce labelled 3D data and forest-centred architectural assumptions.

We present a modular 2D-to-3D label-transfer workflow that projects high-resolution UAS multispectral labels into co-registered LiDAR point clouds, generating height-plausible and object-consistent volumetric supervision without manual point-cloud annotation. Using this supervision, we evaluated zero-shot transfer, few-shot fine-tuning, and training from scratch.

Forest-trained models declined from 85.2% mIoU in forest to 41.1% and 16.1% in medium- and low-density rangelands, while instance F1 dropped sharply. Curriculum fine-tuning produced the strongest rangeland performance, reaching 49.6% and 63.1% mIoU. Effective 3D rangeland mapping therefore requires scalable target-domain supervision and adaptation beyond forest-centred structural assumptions.

Keywords: UAS LiDAR, multispectral imagery, 2D-to-3D label transfer, 3D deep learning, vegetation structure mapping, semi-arid rangelands

1. Introduction

Understanding the three-dimensional (3D) structure of vegetation is central to ecosystem science (Ehbrecht et al., 2021; Zhang et al., 2024). Growth forms, defined by morphology, including size, shape, and vertical organisation, represent fundamental organisational units of vegetation and regulate plant function, habitat structure, and ecosystem resilience (Omasa et al., 2007; Ganz et al., 2019; Gassilloud et al., 2025). While changes in vegetation cover or productivity may be detectable before clear shifts in growth form, structural attributes provide complementary information on ecological condition that spectral metrics alone cannot resolve (Jacobs et al., 2021). Water deficit, heat, and disturbance alter plant architecture, canopy arrangement, and vertical complexity at both individual and community scales, with important consequences for photosynthesis, transpiration, habitat suitability, and ecosystem functioning (Omasa et al., 2007). Capturing these 3D characteristics is therefore essential for monitoring vegetation dynamics across heterogeneous landscapes.

Spatially explicit information on the distribution of vegetation growth forms remains limited from local to landscape scales, largely because such measurements are difficult to obtain through field surveys or conventional satellite-based Earth observation (Kellner et al., 2019). This challenge is especially pronounced in semi-arid rangeland ecosystems, where vegetation is sparse, low-stature, and highly heterogeneous. Rangelands cover approximately 45% of the Earth's terrestrial surface and nearly 70% of Australia, supporting essential ecosystem services including grazing, biodiversity conservation, and carbon sequestration (Reid et al., 2014; Harrison et al., 2021; Retallack et al., 2022). In Australia, these landscapes include extensive mallee ecosystems characterised by multi-stemmed, low-stature eucalypts and shrub-dominated vegetation, whose discontinuous canopies and fine-scale spatial variability make them particularly difficult to characterise structurally.

Field-based measurements of growth forms provide detailed structural indicators but are labour-intensive and spatially sparse (Jiang et al., 2012; Kellner et al., 2019; Zhang et al., 2021). Terrestrial laser scanning (TLS) and mobile laser scanning (MLS) systems provide highly detailed characterisation of tree architecture and understory structure at plot scale, but

their limited spatial coverage and acquisition logistics constrain scalability for landscape-level monitoring (Disney, 2019; Calders et al., 2020; D'hont et al., 2025). Satellite multispectral data have enabled broad-scale vegetation monitoring (White et al., 2009), including Landsat-based vegetation-cover mapping with shallow machine-learning approaches (Jones et al., 2018; Zhang et al., 2019). In semi-arid rangelands, Sentinel-2 and Landsat support repeated FVC and woody-cover monitoring, but mixed pixels limit their ability to resolve fine-scale vegetation components and structural heterogeneity (Sutton et al., 2022; Ruscalleda-Alvarez et al., 2025). Recent studies show that high-resolution UAS reference data can support sub-pixel vegetation-fraction estimation from Sentinel time series (Schiefer et al., 2023; Mosig et al., 2026b,a). However, satellite products still provide limited information on plant-scale height, crown extent, canopy patterns, and vertical growth-form organisation.

Unoccupied aerial systems (UAS) equipped with multispectral sensors and lidar now provide centimetre-scale 3D observations that support measurements of plant height and crown geometry of individual growth forms (Lucieer et al., 2014; Melville et al., 2019; Almeida et al., 2021; Schiefer et al., 2023; Sotomayor et al., 2025a). However, most operational workflows still reduce these data to 2D products, such as canopy height models (CHMs) derived from Structure-from-Motion (SfM) photogrammetry point clouds, which limits their capacity to represent vertical heterogeneity and multi-stemmed growth forms typical of rangelands (Kattenborn et al., 2020; Schiefer et al., 2020; Ecke et al., 2022; Harrison et al., 2024). Passive imaging also tends to underestimate vegetation height and poorly resolve mid-storey structure because of canopy occlusion and restricted viewing geometry (Dandois and Ellis, 2013; Wallace et al., 2019). Although height products such as CHMs can provide complementary structural information, their contribution in 2D CNN workflows has generally been modest or inconsistent, suggesting that some coarse structural cues are already partially encoded in RGB imagery through shadows, illumination gradients, and crown texture (Kattenborn et al., 2020; Schiefer et al., 2020; Sothe et al., 2020; Xiang et al., 2025b). More fundamentally, collapsing 3D vegetation structure into 2D surfaces obscures the vertical organisation needed for robust growth-form classification and may increase susceptibility to overfitting in sparse, structurally heterogeneous ecosystems (Kattenborn et al., 2020; Sothe et al., 2020; Sotomayor et al., 2025a).

Recent research has therefore shifted towards algorithms that operate

directly on 3D point clouds, preserving spatial context and vertical organisation essential for characterising vegetation structure (Xiang et al., 2024; Wielgosz et al., 2024; Cherlet et al., 2026). An early step in this transition involved point-based neural networks operating on unstructured point clouds, where architectures derived from PointNet++ (Qi et al., 2017) enable hierarchical feature learning from raw point sets and have since been extended for tree detection, segmentation, and structural characterisation (Krisanski et al., 2021; Wielgosz et al., 2023). Instance-level frameworks further advanced individual tree delineation, including Point2Tree, which combines hierarchical feature learning and clustering for tree segmentation (Wielgosz et al., 2023), and TreeLearn, a multi-task framework integrating semantic and instance predictions for structural interpretation (Henrich et al., 2024). More recent class-agnostic approaches such as SegmentAnyTree emphasise geometry-driven detection and aim to generalise across forest types without species-specific training (Wielgosz et al., 2024), while bottom-up frameworks such as ForAINet combine semantic and instance segmentation for object-centric ecological analysis (Xiang et al., 2024). Building on these, voxel-based sparse 3D CNN architectures such as 3D U-Net variants implemented in the Minkowski Engine (Choy et al., 2019) enable efficient processing of high-density point clouds while retaining full volumetric context, though performance remains sensitive to point density, canopy continuity, and voxel size (Choy et al., 2019; Xiang et al., 2024). Transformer-based models such as ForestFormer3D (Xiang et al., 2025b) further extend this by explicitly capturing long-range dependencies and hierarchical relationships across spatial scales, demonstrating strong performance in dense, vertically stratified forest environments. Together, these methods preserve full point-level detail, enabling representation of sub-canopy structure and vertical complexity.

Despite recent progress in 3D deep learning, a fundamental bottleneck remains: algorithmic advances have outpaced the availability of representative training data, and most labelled lidar datasets originate from dense forest ecosystems dominated by vertically stratified, single-stem trees. lidar-derived point clouds are also affected by occlusion and sampling uncertainty within lower canopy strata (Brede et al., 2022; Gassilloud et al., 2025), compounding the challenge of supervising models in structurally sparse environments. Consequently, existing frameworks including ForestFormer3D and SegmentAnyTree have been developed and validated almost exclusively in forest environments, and the structural assumptions embedded in these models, continuous canopy, single-stem architecture, vertical stratification, limit

their applicability to rangeland ecosystems where vegetation is low-statured, multi-stemmed, and discontinuous. The extent to which these assumptions affect model transferability across structurally distinct ecosystems remains largely untested, and labelled point cloud data for semi-arid rangelands are virtually nonexistent, limiting both training and objective evaluation under sparse canopy conditions.

To extend 3D vegetation modelling beyond forested benchmarks, this study builds on our previous CNN-based workflow for mapping fractional vegetation cover from high-resolution UAS multispectral imagery (Sotomayor et al., 2025a,b), adapting it to produce vegetation labels and transfer them to co-registered lidar point clouds. Our two contributions are:

1. A 2D-to-3D label-transfer workflow that adapts multispectral CNN-derived vegetation labels for scalable lidar-based volumetric supervision.

2. A set of controlled ablation studies evaluating ForestFormer3D under zero-shot transfer, few-shot fine-tuning, and training from scratch, and SegmentAnyTree as a complementary zero-shot instance-segmentation baseline, using FOR-instanceV2 (Xiang et al., 2025a) as the source forest benchmark and the Calperum semi-arid sites as target domains, isolating the effects of pretraining, adaptation strategy, and vegetation density on 3D model behaviour under structural domain shift.

2. Materials

We tested two complementary 3D deep learning approaches at two study sites, known as TERN AusPlots (White et al., 2012), each representing different vegetation types as defined by the National Vegetation Information System (NVIS) (DCCEEW, 2017). The sites sample the predominant structural formations of this semi-arid area, encompassing a range of vegetation types, soil characteristics, and structural complexities (White et al., 2012) (detailed in Table 1). The study sites covered a total area of 19 hectares (ha) (see Figs. 1, Fig.2) and are located at Calperum Station near Renmark, South Australia. This area experiences low annual rainfall (average 265 mm) and an average temperature of 17 °C.

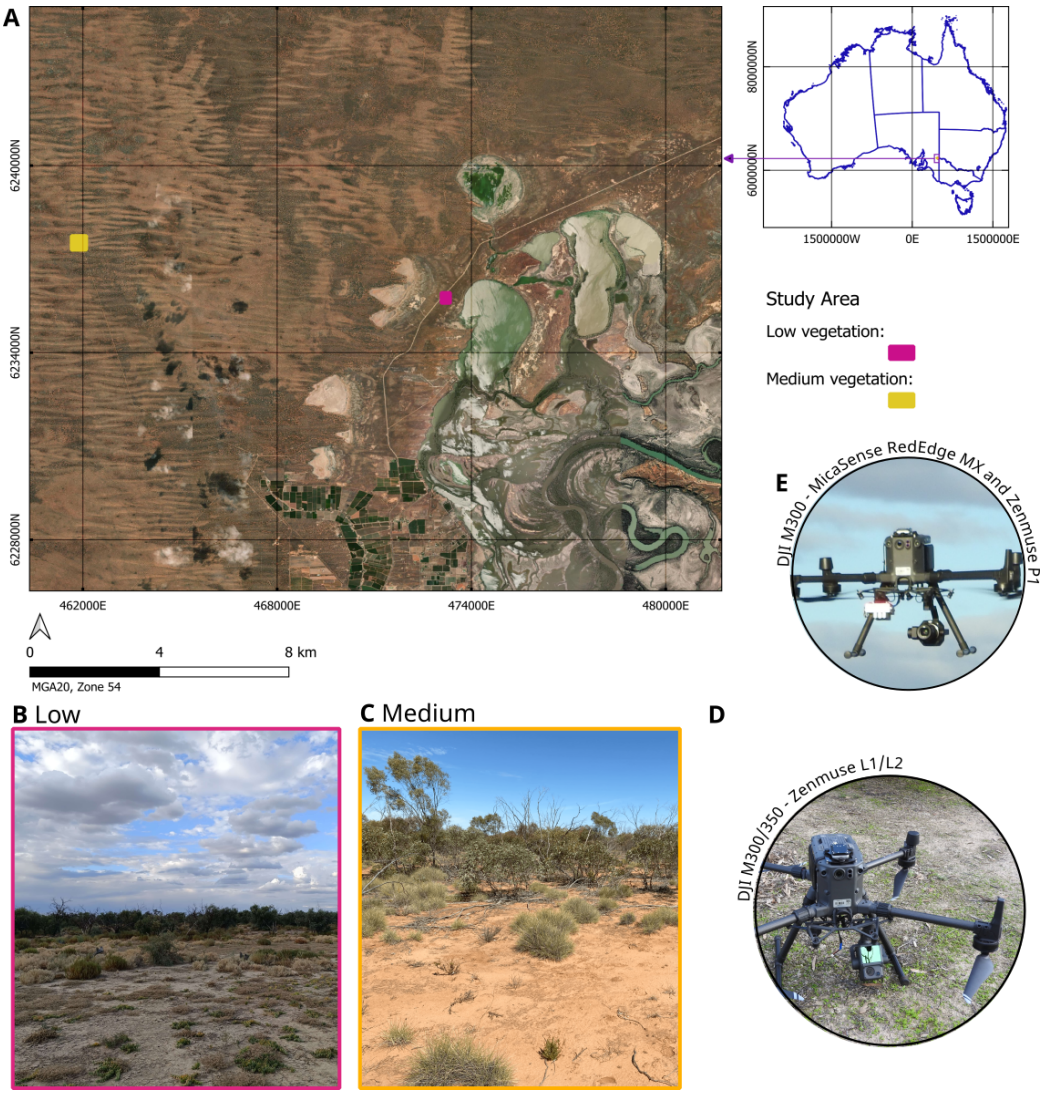

Figure 1: **A** Location of the two study sites (rectangle coloured boxes in map top-left) at Calperum Station, South Australia, each representing a different vegetation community (**B** for low, **C** for medium), and **D** the UAS platforms used for data acquisition during the survey.

The low vegetation ecosystem features open chenopod shrubland dominated by black bluebush, with sparse ground cover comprising forbs and woody debris (see Fig. 2A, B). The soil is predominantly Calcarosol. The medium-density vegetation also includes a regenerating mallee canopy, in-

fluenced by past fire events, with sparse *Triodia* hummock grasses (see Fig. 2C, D). Here, standing dead branches add to the coarse woody debris, creating structural diversity with elements such as mallee, grass hummocks, and non-photosynthetic (NPV) branches.

Table 1: Site descriptions including area, land system, dominant vegetation, height, structure and growth forms (White et al., 2012; TERN, 2021; TERN et al., 2022)

AusPlot IDs	**Site area**	**Land system and soil type**	**Dominant vegetation**	**Height of strata**	**Growth forms and structural components**
SASMDD0011 (low)	184 m x 215 m (≈ 4 ha)	Gently undulating plains; reddish-brown calcareous sandy loam	*Maireana pyramidata* (Black bluebush) low open shrubland with scattered annual grasses and forbs	Mid X m (bluebushes) and Ground Y m (herbaceous): 0.5–1.0 m	Woody chenopod shrubs with semi-succulent foliage, forbs, fallen woody branches and stems, and woody material that is part of a dead tree
SASMDD0001 (medium)	369 m x 400 m (≈ 15 ha)	Stable inland dune-swale system; reddish loamy sand	Low open woodland of multi-stemmed mallee *Eucalyptus* spp. over sparse understorey of *Triodia scariosa* hummock grasses and shrubs	Upper (Mallee trees) and Ground (mainly Triodia): 0.5–5 m	Hummock grass, multi-stemmed mallee trees, shrubs, coarse woody debris, and litter. Standing/emergent NPV branches/stem that still stand after regeneration after fire and emergent dead branches

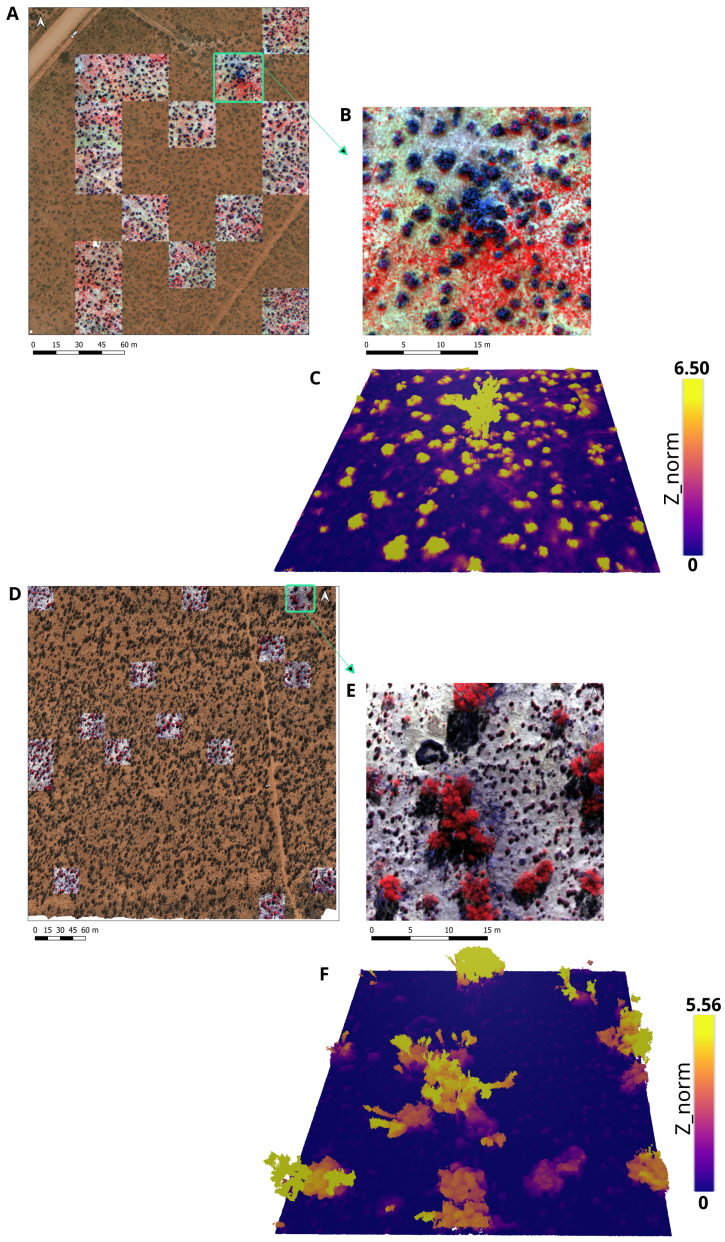

Figure 2: Multi-sensor remote sensing imagery across different vegetation types. **A** RGB orthomosaic of low vegetation, sparsely vegetated area, with considerable bare ground visible, overlaid with the 15 multispectral chunks (3070 × 3070 pixels) used for labelling. **B** Enlargement of highlighted area using multispectral imagery (band combination: NIR-842 nm, Red-668 nm, Blue-475 nm). **C** Oblique 3D visualisation of the LiDAR point cloud coloured by normalised height (Z_{norm}) in metres, with higher colour-scale values indicating taller vegetation. The scene highlights the sparse vertical structure of the low-density site, including low shrubs, isolated mallee canopies, and extensive bare ground. **D** RGB orthomosaic showing medium vegetation, with a mix of vegetation and exposed ground, with multispectral chunk overlay. **E** Enlargement of highlighted area using multispectral imagery. **F** Oblique 3D visualisation of the LiDAR point cloud coloured by normalised height (Z_{norm}), with higher colour-scale values indicating taller vegetation. The scene highlights the greater vertical structural and spatial arrangement of forbs, grass hummock, shrubs, a mallee canopy, and bare ground within the study area.

3. Methods

The proposed workflow integrates a 2D-to-3D label transfer pipeline with cross-ecosystem 3D vegetation segmentation experiments. The study area and lidar/UAS datasets are first described in Section 3.1, followed by the data acquisition and processing pipeline, the 2D-to-3D label transfer workflow, and the experimental framework used to evaluate cross-ecosystem transfer between forest and rangeland environments. The full pipeline, called *fvcStructureRangeland3D*, is publicly available at `https://github.com/LNSOTOM/fvc_structure`.

3.1. UAS acquisition and processing

3.1.1. UAS platforms and sensors

We acquired the datasets using a DJI Matrice 300 RTK UAS platform equipped with two sensors mounted simultaneously: the DJI Zenmuse P1 camera for RGB imagery and the MicaSense RedEdge-MX sensor for 5-band multispectral imagery. We used the spectral bands 'Blue-475', 'Green-560', 'Red-668, 'Red Edge-717', and 'NIR-842' for this study. A single flight was conducted at each site during the May 2022 campaign, and the flight parameters optimised for each sensor. RGB and multispectral imagery was collected at a flying height of 80 m above ground level and flying speed of 8 m/s. This resulted in ground sampling distances of approximately 1 cm/pixel for RGB imagery and 5 cm/pixel for multispectral imagery (White et al., 2012; TERN, 2021; TERN et al., 2022).

Lidar data were collected using the Zenmuse L1 sensor with Matrice 300 (M300) RTK UAS at a flying height of 50 m above ground level and a flying speed of 5 m/s. Further details on data collection and preprocessing can be found in the TERN drone data collection and imagery processing protocol (TERN, 2022). Another field campaign was carried out in 2024 using the Zenmuse L2 with the Matrice 350 (M350) RTK UAS; further details can be found at DroneScape: Collection of Penta-Return lidar for Ecosystem Plot Monitoring (Montes-Herrera et al., 2025). Data from this campaign were only used in Experiment 1 (see Supplementary materials described in Table S2) to evaluate model transferability and the potential influence of lidar point density on prediction performance. Specifically, 15 tiles per site from the Calperum study areas were processed during cross-domain inference to assess how forest-trained models performed when applied to semi-arid rangeland ecosystems.

3.1.2. UAS imagery processing

We georeferenced the RGB and multispectral imagery using the DJI D-RTK2 base station and on-board Global Navigation Satellite System Real Time Kinematic (GNSS RTK) receiver, ensuring high georeferencing accuracy without the need for ground control points (approximately 5 cm absolute accuracy). We processed the imagery using Agisoft Metashape Pro(v.2.0) to generate co-registered red-green-blue (RGB) and multispectral orthomosaics following the TERN Drone RGB and multispectral image processing protocol (TERN, 2021, 2022; TERN et al., 2022).

The RGB processing followed four main steps. First, tie points were generated through image alignment and used as a reference for the multispectral data registration alignment. Second, two dense clouds were constructed from the estimated camera positions, with quality parameters set to medium (TERN, 2022) and high, respectively. Third, the medium-quality dense cloud was used to produce a smooth mesh model (SMM) for orthorectification, while the high-quality dense cloud was used to derive a digital surface model (DSM) with interpolation disabled, preserving true elevation values by retaining the highest surface within each grid cell. The DSM was retained for later use in the lidar–RGB GeFolki co-registration step (Section 3.1.3). Fourth, the RGB orthomosaic was produced by orthorectifying the original images and projecting their pixel values onto the RGB-derived surface model (SMM) to ensure geometric consistency. More detail on this workflow can be found in (TERN, 2022) (see Supplementary materials, Table S1).

Creation of multispectral orthomosaics followed this workflow: first, the images were calibrated using a camera calibration panel. Second, tie points were generated by image alignment, and previously generated RGB tie points were used as references to align the multispectral data. Third, the dense cloud created was used to generate a digital surface model (DSM) using the Python package PDAL. Fourth, the multispectral orthomosaic was produced by orthorectifying the original images and projecting their pixel values onto the RGB-derived surface model (SMM) to ensure geometric consistency. More detail on this workflow can be found in (TERN, 2022).

3.1.3. Lidar processing and co-registration workflow

DJI Terra software (version 3.5.0) was used to process raw lidar point clouds. Further details on lidar data preprocessing can be found in the TERN Drone lidar Processing Protocol section (TERN, 2022; Montes-Herrera et al.,

2025). A Digital Surface Model (DSM) was generated from the lidar point cloud using a PDAL processing pipeline. During rasterisation, the maximum elevation value within each grid cell was retained so that each pixel in the resulting DSM represented the highest surface return within the corresponding cell. This DSM was subsequently used to support co-registration between the lidar point cloud and the RGB-derived high-quality dense point cloud.

Initial co-registration experiments were conducted using a classical point cloud-to-point cloud approach based on Iterative Closest Point (ICP). However, due to fine-scale structural differences between the lidar point cloud and the RGB-derived Structure-from-Motion (SfM) dense point cloud, particularly in vegetation structure and surface representation, ICP did not provide sufficiently stable correspondences to achieve accurate alignment. As a result, we applied a dense, non-rigid co-registration approach based on image-derived surface models to align the lidar-derived 3D point cloud with the RGB orthomosaic reference. The workflow used the Extended Flow Optical Lucas-Kanade Iterative (eFOLKI) algorithm developed by Plyer et al. (2014), an adaptation of the optical flow framework originally introduced by Le Besnerais and Champagnat (2005). This GeFolki implementation has been widely used for SAR/SAR and lidar/SAR co-registration and, more recently, for imaging spectroscopy alignment (Haynes et al., 2025).

The co-registration process consisted of three main steps:

a Grid alignment and resampling. The lidar-derived DSM was resampled onto the spatial grid of the RGB-derived DSM to ensure an identical spatial resolution, extent, and coordinate reference system (CRS). This step alignment allowed a direct pixel-to-pixel comparison of both DSMs and ensured geometric consistency for subsequent optical flow computations. Resampling used bilinear interpolation to preserve surface continuity while minimising artifacts in the continuous elevation data, with the RGB DSM defining the target grid. The resampled lidar DSM was saved as a new raster matching the dimensions of the RGB DSM, while the original datasets remained unchanged.

b The resampled lidar DSM (moving image) was co-registered to the RGB DSM (reference image) using the eFOLKI dense optical flow algorithm in Python. This method estimates continuous horizontal displacements by comparing the topographic gradients between the two DSMs, enabling the correction of non-rigid distortions and local offsets caused by sensor geometry, platform motion, or differing viewing perspectives. The analysis

was restricted to overlapping valid pixels to prevent boundary artifacts. The eFOLKI algorithm produced two dense 2D displacement fields, u and v, representing the horizontal correction vectors (in the X and Y directions) required to align the lidar DSM with the RGB surface DSM.

c Application of the 2D flow field to the original 3D lidar point cloud. The 2D displacement fields (u, v) estimated in step (b) were applied to the original full-density lidar point cloud to produce a non-rigidly co-registered 3D dataset in the RGB reference frame. This step was implemented using a custom Python script combining `laspy`, `rasterio`, and `scipy`. To reduce the pixel-scale noise in the displacement fields and prevent discontinuities ("tearing") in the warped point cloud, the u and v fields were smoothed using a Gaussian filter (`scipy.ndimage.gaussian_filter`) with $\sigma = 2$ pixels. Each lidar point's (X, Y) coordinates were mapped to the corresponding $(\text{column}, \text{row})$ pixel locations on the 2D flow grid using the inverse affine transform of the reference DSM. Points outside the flow grid were excluded. The sub-pixel displacements were then interpolated from the smoothed flow fields using scipy.ndimage.map_coordinates with a quadratic interpolation (`order = 2`) to ensure smooth and continuous displacement estimates across the flow domain. These pixel-level displacements were converted to metric corrections $(\Delta X, \Delta Y)$ using the DSM pixel size, and the resulting offsets were directly applied to update point coordinates. The resulting output was a new co-registered lidar point cloud, non-rigidly warped to align the geometric reference of the RGB orthomosaic.

3.2. 2D-to-3D label transfer pipeline

3.2.1. 2D imagery tiling and growth-form labelling

We generated 2D reference data for training the sparse 3D U-net models using a structured nomenclature with five classes and identifiers (IDs): forb (FO = 0), woody debris (WD = 1), grass hummock (GH = 2), shrub (SH = 3), and mallee (MA = 4).

Predictors from multispectral data. We adopted the orthomosaic tiling and reference-data generation workflow described by Sotomayor et al. (2025a), in which each orthomosaic was divided into regular tiles on a 5 cm raster grid. The predictor rasters used for this analysis were 614×614 pixels per tile, corresponding to approximately 30.70×30.70 m at 5 cm resolution.

Unlike the original study, where these tiles were used for CNN training, here they provided a consistent spatial framework for cropping, co-registering, and aligning the corresponding LAS point clouds.

To support visual interpretation, we randomly selected 15 chunks and generated 8-bit false-colour composites using the Blue, Red, and NIR bands. Reference labels were digitised in Labelme, an open-source annotation tool (Wada et al., 2021), with Efficient-SAM (Xiong et al., 2023, 2024) used as an AI-assisted segmentation model to support interactive annotation.

2D reference structure labels. We reused the 2D reference label schema developed for site-specific applications by Sotomayor et al. (2025a,c) and adapted it to a targeted structural schema based on vegetation growth forms. Specifically, the shadow and bare-ground classes were removed, and the remaining vegetation classes were reorganised to represent the structural components required for 3D label transfer and modelling (Table 2).

In the original workflow, reference labels were generated through manual expert interpretation of high-resolution UAS orthomosaics, supported by multispectral reflectance information and the computer-vision-assisted annotation Labelme. Image tiles were annotated to delineate vegetation components associated with fractional vegetation cover (FVC), including photosynthetic vegetation, non-photosynthetic vegetation, and background substrate fractions. For this study, the reference masks were refined and harmonised into a growth-form-based schema that better represents semi-arid vegetation structure within the lidar-supported workflow, providing the 2D semantic layer for subsequent point cloud label transfer.

We rasterised the reference labels as multi-class segmentation masks. The low site contained 3 classes (FO, WD, SH) and the medium site contained 4 classes (WD, GH, SH, and MA). Areas containing NaN values (coloured white) were excluded. The reference dataset was downsampled to 5 cm using nearest interpolation to match the resolution of the predictors or multispectral data.

Table 2: Detailed breakdown of 2D structure classes into subclasses, categorised by growth form and structural components. Grouped by class ID, this table supports reference data creation and interpretation across site-specific levels (low, medium).

ID	**Class**	**Growth form/-Structure**	**Description**	**Site Occurrence**
0/64	forb	Forb	Herbaceous or slightly woody ground vegetation.	Low
0/66	grass hummock	Hummock grass	Spiny perennial grass, initially growing as rounded hummocks, expanding to open "rings" with age.	Medium
		Hummock grass (senescent)	Aging hummock grasses.	Medium
1/65	woody debris	Standing/dead stem	Standing or emergent NPV branches/stems, often remaining after fire, likely dead, with no foliage at the tips.	Low, Medium
1/67	shrub	Shrub	Perennial plants with multiple woody stems; often height-classified (< 2 m or ≥ 2 m).	Low, Medium
		Chenopod shrub	Woody shrubs with sparse, semi-succulent foliage, and grey coatings; low IR response.	Low, Medium
		Chenopod shrub (senescent)	Aging Chenopod shrubs.	Low, Medium
2/68	mallee	Mallee	Multi-stemmed eucalypt species sprouting from an underground lignotuber that regenerates from lignotuber or epicormic buds after fire.	Low, Medium

3.2.2. 3D lidar preparation and growth-form label transfer

We implemented a multi-stage lidar processing framework to derive 3D vegetation structure at both understory and crown scales from the co-registered lidar data. The workflow integrates raster-based tiling, geometric alignment,

2D-to-3D semantic label propagation, ground normalisation, and occlusion-aware 3D label fusion to produce spatially consistent and biologically interpretable vegetation-structure classifications.

Raster-based tiling and clipping to predictor extent. For each multispectral predictor tile, we first defined a common spatial processing unit using the predictor raster footprint (5 cm grid). This step ensured that all subsequent products, rasters and point clouds, shared identical tile bounds, resolution, and CRS.

For both the RGB dense point cloud and the lidar point cloud, we applied a PDAL-based clipping workflow to generate tile-specific point clouds aligned to the predictor extent. Specifically, for each tile we:

(i) extracted the tile bounds and coordinate reference system (CRS) from the corresponding predictor raster;

(ii) transformed raster bounds into the native LAS CRS where required;

(iii) cropped the full point cloud using `filters.crop`; and

(iv) reprojected the cropped points into the predictor raster CRS using `filters.reprojection`.

This produced tile-specific RGB dense and lidar point cloud subsets spatially aligned to the predictor raster, ensuring a common tile extent and CRS without applying vegetation masks or semantic filtering (Fig. 3).

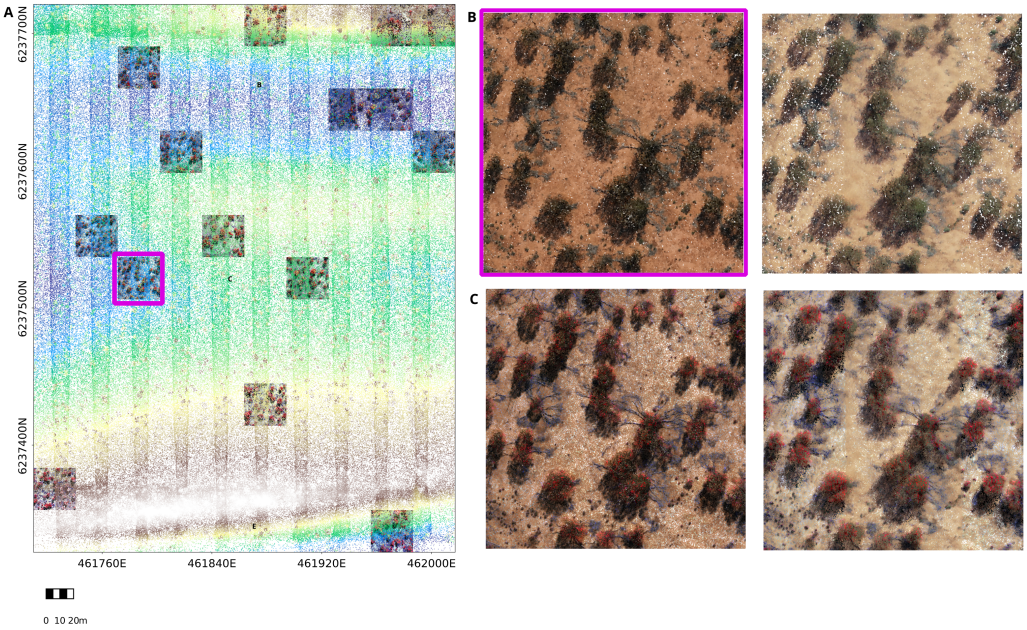

Figure 3: Multi-sensor remote sensing data for the medium-density vegetation site. A) lidar point cloud of the site, with the corresponding multispectral chunk extent overlaid. B) Enlarged view of the selected area, showing the RGB dense point cloud (left) and the lidar point cloud (right). C) Overlay of the RGB dense point cloud (left) and lidar point cloud (right) on the multispectral imagery.

Reference mask transfer to RGB dense tiles. Following raster-based tiling, 2D vegetation-structure masks were transferred to the RGB dense point cloud tiles to produce labelled 3D reference subsets for quality control and alignment diagnostics. This step was kept separate from raster-based clipping so that geometric alignment and semantic filtering were treated as distinct operations. The transferred labels corresponded to class IDs 64–68: forb (64), woody debris or standing dead stem (65), grass hummock, including senescent hummock grass (66), shrub categories, including chenopod and senescent chenopod shrub (67), and mallee (68), representing the dominant structural vegetation components across both study sites.

For each tile, the 5 cm structure mask was sampled over the footprint of the clipped RGB dense point cloud using `filters.colorization`, writing the mask values to the LAS `Classification` field. Valid vegetation-structure classes were retained using `filters.range`, while NoData or unlabelled points were removed. The resulting labelled RGB dense tiles were

used as reference data for subsequent alignment assessment.

Alignment refinement of lidar tiles using ICP. Although raster-based tiling ensured global spatial consistency, small local misalignments remained between lidar and RGB dense point clouds due to differences in sampling density and acquisition geometry. These discrepancies appeared primarily as residual vertical offsets and minor horizontal shifts at the tile level. We therefore refined the alignment of each lidar tile using Iterative Closest Point (ICP) registration against the corresponding RGB dense clouds tile (see Fig. 4). The RGB dense cloud was used as the geometric reference because it is directly tied to the orthomosaic and exhibits high internal consistency. Tile-level ICP corrected local misalignments while preserving the global RTK-based georeferencing, improving the spatial correspondence between lidar structure and image-derived geometry.

ICP registration was performed in CloudCompare v2.13 alpha (CloudCompare) using a root-mean-square (RMS) error convergence threshold of $1.0 \times 10^{-0.5}$ and a minimum overlap constraint of 90%. These parameters were selected to ensure stable point-to-point convergence while preserving fine-scale vegetation structure. Alignment quality was assessed through visual inspection of tile correspondence and evaluation of residual error distributions. The resulting rigid transformation matrix was subsequently applied to each lidar tile, producing an improved locally optimised alignment, ICP-refined lidar point clouds.

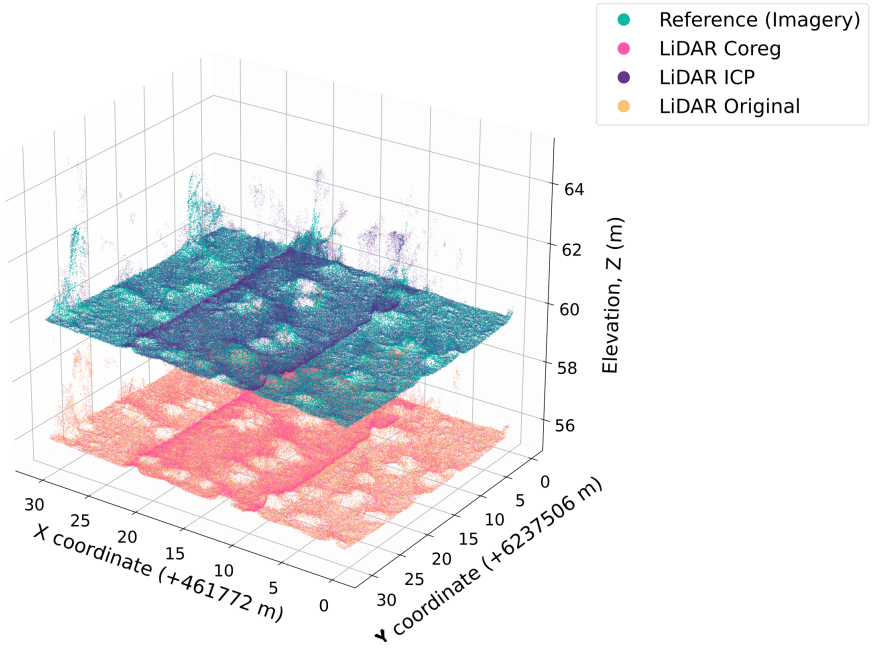

Figure 4: Alignment refinement of tile-specific lidar point clouds using ICP registration. Front-view comparison of the RGB dense reference cloud (green) and lidar tile alignments at successive stages: original before correction (cyan), raster-based co-registration (blue), and final ICP refinement (red). Vertical offsets are shown for visual separation.

Semantic label transfer to ICP-aligned lidar tiles. Following ICP refinement, 2D vegetation-structure labels were transferred from the 5 cm structure mask to the ICP-aligned lidar tiles to produce the final structure-labelled 3D point clouds used for subsequent ground classification, height normalisation, and occlusion-aware fusion (Fig. 5). Label propagation was performed in the native Cartesian coordinate system (projected MGA zone) using a projection-based top-view transfer strategy (Pellis et al., 2022), which is commonly used to assign raster-derived semantic labels to 3D point clouds (Straker et al., 2023).

The transfer was implemented using a PDAL workflow (Contributors, 2025). For each tile, the pipeline:

(i) cropped the ICP-aligned lidar points to the exact spatial extent of the multispectral predictor raster;

(ii) sampled the 5 cm vegetation-structure mask using `filters.colorization`, writing the mask values to the LAS `Classification` field; and

(iii) retained only valid vegetation-structure classes (64–68: forb, woody debris or standing dead stem, grass hummock, shrub, and mallee) using `filters.range`, thereby removing unlabelled or NoData points.

Original LAS header information, including minor version, coordinate scales, and offsets, was preserved when writing the output files to maintain full spatial precision and consistency across all ICP-aligned, mask-labelled lidar tiles.

Ground classification and height normalisation. Cloth simulation filter (CSF) for ground classification. We derived ground points for each ICP-aligned lidar tile using a PDAL-based ground-filtering workflow. Two filters were evaluated: the Simple Morphological Filter (SMRF) (Pingel et al., 2013) and the Cloth Simulation Filter (CSF) (Zhang et al., 2016). For SMRF, we used `filters.smrf` with typical parameter values tuned to our site conditions (slope = 0.15, window = 3.5 m, threshold = 0.25, scalar = 1.35), followed by `filters.range` to retain only points classified as ground (Classification = 2). For CSF, we applied `filters.csf` with a grid resolution of 0.5 m, 500 iterations, a height threshold of 0.01 m, and a rigidness parameter of 2, followed by `filters.range` to extract the ground class. In both cases, the resulting ground point clouds (Classification = 2) were extracted using `filters.range` and written to new LAS files while explicitly preserving the original LAS quantisation (scale and offset) and metadata (`forward="all"`) to maintain spatial precision and consistency with the input tiles.

Based on visual inspection and quantitative comparison of ground returns, CSF was selected for subsequent height normalisation. Although SMRF classified a larger number of ground points overall (e.g. 785,243 points for tile 106) compared to CSF (271,659 points for the same tile), CSF produced a more spatially coherent and topographically consistent ground surface, particularly in areas of low and heterogeneous vegetation. In contrast, SMRF tended to over-classify low vegetation and understory elements as ground, leading to locally inflated ground surfaces and increased artefacts in height normalisation. Consequently, CSF-derived ground points were retained for all downstream height-above-ground calculations (see Fig. 5)

We then normalised the lidar elevations to obtain the height above ground (HAG) for each tile. This was implemented in Python by combining the full ICP-aligned tile with its corresponding ground LAS file. For each tile, we built a k-nearest-neighbor search structure (cKDTree) from the ground points

in the horizontal plane (X, Y) and, for every point in the full cloud, queried the three nearest ground neighbours ($k = 3$). The ground elevation at the point location was estimated using Inverse Distance Weighted (IDW) interpolation of the neighboring ground Z values, and the HAG was computed as the difference between the original elevation and the interpolated ground level, with negative values clamped to zero. The resulting HAG values were stored as a new ExtraBytes dimension (`Z_norm`, 32-bit float) in the output LAS files, providing a normalised height metric that is independent of the absolute terrain elevation and suitable for subsequent structural analysis.

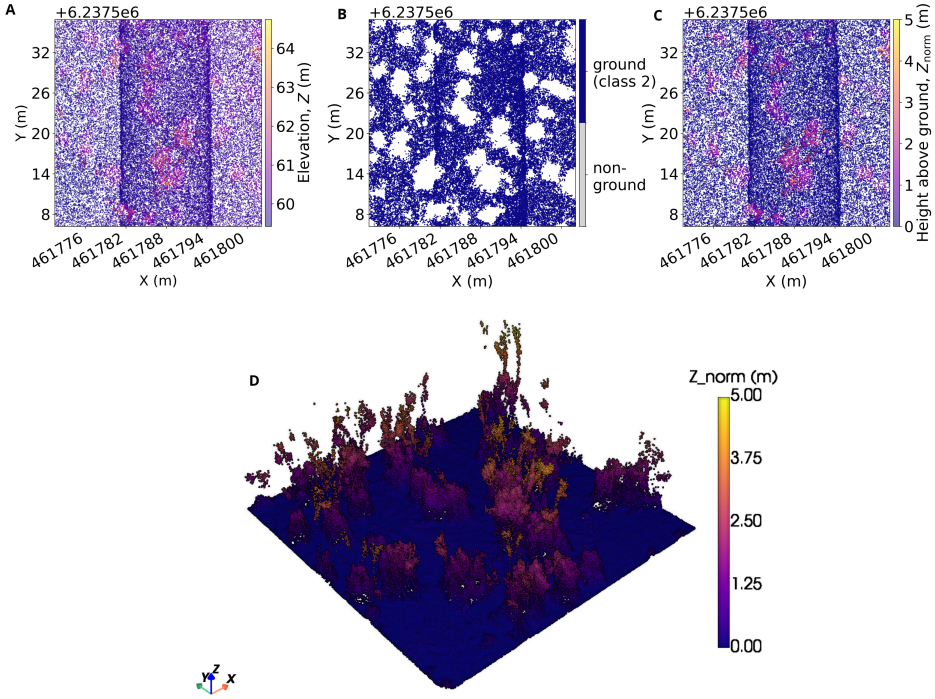

Figure 5: Workflow for lidar ground classification and terrain-normalised height derivation. **A** Raw ICP-aligned lidar point cloud coloured by absolute elevation (Z), showing the original terrain and vegetation structure before ground filtering. **B** Cloth Simulation Filter (CSF) ground classification, with ground returns (ASPRS class 2) shown in green and non-ground returns shown in grey. **C** Terrain-normalised height map derived from the CSF ground surface, where height above ground is stored as the ExtraBytes attribute `Z_norm`. **D** Oblique 3D view of the terrain-normalised point cloud, coloured by `Z_norm`, highlighting vegetation vertical structure after terrain removal. Height values are shown in metres.

Occlusion-aware 3D vegetation-structure label fusion. To obtain a complete 3D vegetation structure classification, we fused geometric lidar point clouds containing normalised heights (Z_{norm}) with structure-labelled point clouds derived from 5 cm 2D masks for Classes 64–68. The workflow was designed to mitigate irregular sampling density while explicitly accounting for occlusion and mixed understory vegetation, which are common in semi-arid Mallee ecosystems (Brede et al., 2022; Waite et al., 2025).

To characterise the vertical arrangement of these classes in structurally complex canopies, we employed a quantile-based stratification strategy. This approach prioritises stability and repeatability across varying point densities by modelling the dominant vertical profile of each class rather than relying on absolute height extrema. The resulting fused 3D analysis-ready products are directly suitable for downstream use in sparse 3D convolutional neural networks. An overview of the full fusion workflow is provided in Table 3.

Table 3: Occlusion-aware 3D vegetation-structure label fusion workflow.

Stage	**Purpose**	**Key operations / rules**	**Output**
1.Class height thresholding	Derive robust vertical priors for each vegetation class 64-68	XY nearest-neighbour matching ($\leq$0.30 m) between 2D labels and lidar points Outlier removal using P5–P95 Stratification using P15–P85 quantiles Fallback to literature-derived defaults when $n < 20$	Class-specific Z_{norm} height ranges.
2.Cartesian voxel-based label transfer	Establish a stable baseline classification under irregular sampling	Joint discretisation into a 3D voxel grid (10 cm) Majority voting within voxels Transfer of voxel labels to Z_{norm} points	Baseline 3D structure labels.
3.Mallee crown detection and instance assignment	Introduce object-level context for refinement	CHM construction from Mallee points Gaussian smoothing Watershed segmentation Assignment of per-point crown identifier ($treeID$ instances)	Detected Mallee crowns and crown instances ($treeID$s).
4.Crown-centric polar refinement with height priors	Resolve occlusion and by using local tree geometry and radial symmetry to resolve overlaps of mixed-classes	Local cylindrical coordinate transform (r, θ, Z_{norm}) Polar binning (0.10 m radial, 10° angular) Local class probability voting Height-prior modulation (P15–P85) Conservative label updating ($min\_conf = 0.3$, flip ratio=1.5)	Crown-refined structure labels.
5.Height-based consolidation and crown-level consistency enforcement	Finalise biological structural coherence	Merge woody debris (65) $\rightarrow$ Mallee (68) if $Z_{norm} > 0.5$ m; remove low woody debris Retain and promote unlabelled points if $Z_{norm} \geq 1.3$ m Crown-level majority voting (>50% Mallee $\rightarrow$ Class 68) Finalise valid Mallee $treeID$s	Final structure labels and validated $treeID$s.
6.Output generation	Produce analysis-ready 3D products	Export LAS 1.4 (Point Format 6) Generate full fused point cloud (gap-aware) Generate vegetation-only mask product	Final 3D vegetation structure datasets.

Step 1: Computing class height thresholds Because structure labels are derived from 2D imagery, they do not contain explicit vertical information. Class-specific Z_{norm} ranges were therefore estimated by spatially matching structure-labelled points to Z-normalised lidar points using a three-stage procedure:

(a) Spatial integration: structure-labelled points were linked to the lidar returns Z_{norm} values using a nearest-neighbour search in the horizontal plane (XY) with a strict distance constraint of 0.30 m (`cKDTree`).

(b) Outlier mitigation: for each class, extreme returns were suppressed by removing points outside the 5^{th}–95^{th} percentile range ($P^5 \leq Z \leq P^{95}$).

(c) Quantile stratification: final Z_{norm} ranges are determined by the 15^{th}–85^{th} percentile range ($P^{15} \leq Z \leq P^{85}$), providing robust vertical priors for subsequent refinement.

The upper quantile (P^{85}) provides a more stable estimate of canopy extent than the absolute maximum height ($Z_{\max}$), which can be inflated by isolated returns. Where sufficient matched points were available, class-specific height priors were therefore estimated directly from the data using the robust P^{15}–P^{85} interval. When insufficient points were available ($n < 20$), predefined default height ranges were applied to maintain consistency across tiles. These default fallback height ranges were informed by the structural formation terminology of the National Vegetation Information System (NVIS) and by site-specific knowledge of the observed vegetation stature, rather than by tile-level quantiles alone (DCCEEW, 2017).

Table 4: Default fallback height ranges for vegetation structure classes. Ranges were used when class-specific tile samples were insufficient for robust quantile estimation ($n < 20$).

Class ID	Vegetation Name	Default Height Range (m)
64	Forb	0.00–0.40
65	Woody Debris	0.05–5.00
66	Hummock	0.00–0.80
67	Shrub	0.40–3.00
68	Mallee	1.50–8.00

Step 2: Cartesian voxel-based semantic label transfer Following the initial threshold estimation, both point clouds were discretised into a shared

3D Cartesian voxel grid (10 cm $\times$ 10 cm in X–Y and 10 cm in Z). This abstraction standardises neighbourhood support and enables explicit identification of filled, empty, and occluded regions, consistent with established forest modelling approaches (Béland et al., 2014; Schneider et al., 2019; Brede et al., 2022; Gassilloud et al., 2025).

A voxel size of 10 cm was selected as an optimal trade-off between the effective laser beam footprint and the structural fidelity required for crown-scale analysis (Schneider et al., 2019; Gassilloud et al., 2025). Although finer resolutions (e.g. 5 cm) capture greater understory detail, preliminary testing showed that it resulted in fragmented neighbourhood support and amplified artefacts in occluded regions. The 10 cm grid preserved structural coherence while providing sufficient spatial resolution for reliable label aggregation.

Within each voxel, structure labels were aggregated using majority voting and subsequently transferred to the Z_{norm} point cloud. This procedure produces a robust baseline 3D classification across the full tile extent, serving as the initial condition for subsequent occlusion-aware, crown-centric refinement.

Step 3: Mallee crown detection and instance assignment To resolve ambiguities in complex canopy regions (e.g., woody debris versus hummocks versus Mallee), an object-oriented representation was introduced by detecting individual Mallee crowns:

(a) A canopy height model (CHM) was generated from mallee-labelled points.

(b) The CHM was Gaussian-smoothed and segmented using watershed analysis.

(c) Each detected crown was assigned a unique per-point instance identifier (*treeID*).

These instance labels provide essential spatial context for subsequent crown-centric refinement.

Step 4: Crown-centric polar coordinate refinement with height priors (if crowns detected) Residual classification ambiguities were addressed by shifting from a global Cartesian representation to a crown-centric, polar framework, which exploits local tree geometry and radial symmetry to disambiguate mixed vegetation classes under occlusion.

(a) **Local cylindrical coordinate transformation:** For points associated with detected crowns, Cartesian coordinates (X, Y) were transformed into a local cylindrical system $(r, \theta, Z_{\text{norm}})$ centred on each crown centroid. This crown-centred polar representation enables analysis along vertical columns extending outward from individual tree centres, providing structurally coherent neighbourhoods for refinement.

(b) **Polar binning and local class probability voting:** The space surrounding each crown was discretised into fine polar bins (0.10 m radial, 10° angular). Within each bin, class frequencies from structure-labelled points were aggregated to estimate local posterior class probabilities for the target Z_{norm} points. These probabilities provide a geometry-aware baseline for label inference in regions affected by occlusion or horizontal overlap.

(c) **Height priors modulation:** Posterior class probabilities were modulated using soft height priors derived from the data-driven vertical ranges defined in Step 1 (P^{15}–P^{85}). Classes whose Z_{norm} values fell within their expected height range retained full weight (1.0), while probabilities were downweighted (e.g. to 0.1) outside plausible vertical strata. This constraint prevents biologically implausible propagation, such as canopy mallee labels extending to ground-level points or low vegetation being assigned to upper canopy regions.

(d) **Conservative label updating:** Final label updates followed a conservative decision framework to avoid spurious reclassification. Previously unlabelled points were assigned a class only when posterior confidence exceeded $min\_conf = 0.3$. Existing labels were updated only if the posterior probability of the new class exceeded the current assignment by a factor of at least 1.5 ($min\_flip\_ratio = 1.5$). Column-wise refinement within polar bins further enforced neighbourhood coherence while preserving occlusion gaps.

This crown-centric refinement resolves local ambiguities (e.g., understory shrubs versus low-hanging Mallee branches) while preserving the inherent vertical structure and occlusion characteristics of the original point cloud.
Step 5: Height-based consolidation and crown-level consistency enforcement Final biological and structural consistency was enforced using knowledge-based height rules and crown-level voting:

(a) **Woody debris consolidation:** Woody debris points (Class 65) with $Z_{\text{norm}} > 0.5$ m were reclassified as mallee (Class 68), correcting for horizontal overlap with low mallee branches. Low woody debris was removed as biologically implausible noise.

(b) **Recovery of unlabelled points:** Remaining unlabelled points (Class 0) were retained only if $Z_{\text{norm}} \geq 1.3$ m. These points were promoted to woody vegetation and merged into the mallee class to preserve vertically coherent canopy structure.

(c) **Crown-level majority voting:** To enforce intra-crown label consistency, spatial majority voting was applied within detected crowns ($treeID \geq 0$). If more than 50% of points within a crown were classified as mallee, all remaining non-mallee points within that crown were reassigned to Class 68. This majority-based consolidation reduces residual mixed-class artefacts caused by occlusion or sparse sampling. An additional height-based refinement was applied within crowns: points exceeding the minimum mallee height threshold were promoted to Class 68, ensuring vertical continuity of individual tree structures.

(d) **Finalisation of tree instances (mallee crowns only):** Tree identifiers were retained only for mallee canopy points exceeding the minimum crown height threshold; all other points were assigned $treeID = 0$.

Following fusion, analysis-ready 3D products were generated to support structural assessment and downstream modelling. The pipeline produces a gap-aware fused point cloud preserving occlusion patterns, together with vegetation-only mask products that isolate structural components relevant to ecological analysis. The final 3D structure vegetation products, including the crown-centric mallee $treeID$ assignments, were written to LAS 1.4 / Point Format 6 files to ensure compatibility with modern lidar processing workflows and extended attribute storage. We exported both (a) a full fused point cloud preserving occlusion gaps for structural analysis and (b) a vegetation-only mask product for downstream modelling and volumetric assessment (Brede et al., 2022).

3.2.3. Validation of 2D-to-3D label transfer consistency

The workflow described in Section 3.2.2 and Table 3 was evaluated quantitatively on 15 representative medium-site validation tiles using three in-

ternal consistency diagnostics. Low-site validation tiles were used for visual inspection only. Validation focused on the medium-site tiles because they exhibit the canopy density, multi-stratum mixing, and occlusion patterns that the workflow is designed to resolve; low-site tiles contain sparse low-stature vegetation with minimal vertical overlap, where 2D-to-3D mask transfer is geometrically more straightforward and does not fully exercise the polar refinement or crown-level consolidation stages. The diagnostics tested voxel-level label stability, height plausibility relative to the refinement height ranges, and Mallee crown-instance retention after final *treeID* assignment. The validation is interpreted as an internal workflow consistency check, not as an independent ground-truth accuracy assessment. Quantitative diagnostics were supported by visual inspection of the 2D masks, multispectral raster tile overlays, projected 3D labels, refined structural classes, and final *treeID* instances across representative medium and low tiles.

Voxel-level label agreement.. Step 2 Cartesian voxel labels, $\hat{c}_i^{(2)}$, were compared with Step 5 consolidated labels, $\hat{c}_i^{(5)}$, at common occupied voxels. For each class k, let A_k and B_k denote the sets of voxels labelled k at Step 2 and Step 5 respectively. Overall accuracy (OA) and class-wise Intersection-over-Union (IoU) were computed as

$$\mathrm{OA} = \frac{1}{N} \sum_i \mathbb{1}\left(\hat{c}_i^{(2)} = \hat{c}_i^{(5)}\right), \qquad \mathrm{IoU}_k = \frac{|A_k \cap B_k|}{|A_k \cup B_k|}. \tag{1}$$

This diagnostic measures whether the refinement and consolidation stages preserve the Cartesian baseline labels while applying the intended corrections.

Height plausibility of refined classes.. For each tile and refined class, the median Z_{norm} and the proportion of points outside the configured class-specific height range were calculated. This diagnostic checks whether low-stature classes remain near the ground and whether taller classes occupy plausible vertical structure after 3D refinement.

Mallee crown-instance retention.. Instance consistency was assessed by comparing the number of watershed-derived Mallee crowns with the number of final non-zero *treeID*s:

$$\mathrm{Retention}_t = 100\, \frac{n_{t,\,\text{final treeID}}}{n_{t,\,\text{watershed}}}. \tag{2}$$

Retention values below 100% indicate removed or reassigned crown instances, whereas values above 100% indicate that additional valid instances were formed during final consolidation.

4. Experiments

4.1. Experimental design

The experimental design follows three standard paradigms in deep learning and domain adaptation: zero-shot transfer, few-shot fine-tuning, and training from scratch (i.e., training from randomly initialised weights). These settings are widely used to evaluate model generalisation, transferability, and data efficiency across domains (Song et al., 2022; Bafghi et al., 2024; Carós et al., 2026). In this study, they are used to examine cross-ecosystem transfer from forest to semi-arid rangeland vegetation under a consistent training and evaluation framework: zero-shot transfer assesses how forest-trained 3D frameworks behave when applied directly to structurally distinct rangeland vegetation without adaptation, few-shot fine-tuning evaluates whether limited rangeland supervision can improve cross-domain adaptation, and training from scratch provides a baseline in which models learn only from rangeland data without forest-pretrained weights.

Common semantic and instance formulation. To ensure consistency across all experiments, semantic and instance evaluation were defined within a shared structural framework. In the source forest domain, the *FOR-instanceV2* dataset uses the semantic classes *ground*, *wood*, and *leaf*. For cross-domain evaluation in the target Calperum rangeland sites, the native labels were aligned to this structural formulation, such that *hummock* and *forb* were mapped to *ground*, *shrub* to *wood*, and *mallee* to *leaf*. Instance-level evaluation was based on annotated object identifiers ($treeID > 0$) where available, representing tree instances in the *FOR-instanceV2* forest dataset and malleetree instances in the Calperum rangeland sites. This common semantic and instance formulation was maintained across zero-shot evaluation, few-shot fine-tuning, and training from scratch experiments.

Across these settings, ForestFormer3D was used as the primary semantic framework, while SegmentAnyTree was included in the zero-shot setting as a complementary instance-centric geometric reference. All experiments used identical supervision derived from the automated 2D-to-3D label-transfer pipeline, ensuring that differences in performance reflected model design and adaptation strategy rather than variation in annotation procedure.

4.1.1. Experiment 1: Zero-shot evaluation of forest-trained models

Experiment 1 assessed zero-shot transfer by applying forest-trained models directly to structurally distinct rangeland data without additional adaptation. The ForestFormer3D framework was initialised with weights pretrained on the *FOR-instanceV2* forest dataset, evaluated first on the forest test split, and then applied unchanged to the Calperum sites. SegmentAnyTree was also included in the zero-shot setting as a complementary instance-based geometric reference.

Experiment 1A: ForestFormer3D zero-shot transfer. The ForestFormer3D model pretrained on the forest *FOR-instanceV2* dataset was evaluated on both the source forest test split and the target Calperum rangeland sites without further adaptation. In this zero-shot setting, the model was applied directly to the rangeland tiles using the source-domain semantic classes. Performance was assessed using both semantic and instance-level metrics under the shared evaluation framework defined above, in order to quantify how well structural representations learned from *FOR-instanceV2* transferred to sparse semi-arid vegetation.

Experiment 1B: SegmentAnyTree geometric reference. SegmentAnyTree was evaluated under the same zero-shot conditions using instance-level metrics only. As an instance-centric geometric framework, it served as a complementary reference for examining object-grouping behaviour independently of semantic class prediction under the common instance definition described above.

4.1.2. Experiment 2: Few-shot fine-tuning for cross-domain adaptation

Experiment 2 evaluated few-shot adaptation by fine-tuning the ForestFormer3D model pretrained on the forest *FOR-instanceV2* dataset using a limited set of labelled rangeland tiles. Fine-tuning refers to end-to-end continued training of all model weights initialised from the forest-pretrained checkpoint, using the rangeland training tiles and the same loss function and optimiser settings as the original training (see Section 4.2.3). This setting was designed to examine whether target-domain supervision could improve performance under structural domain shift and better adapt forest-derived structural representations to sparse semi-arid vegetation.

Experiment 2A: Joint fine-tuning. In Experiment 2A, the pretrained model was fine-tuned using combined training data from the medium- and low-

density Calperum sites and evaluated on held-out tiles from both domains. This setting was used to assess whether simultaneous exposure to both vegetation regimes could support more generalisable cross-domain adaptation under the shared semantic formulation.

Experiment 2B: Curriculum fine-tuning. In Experiment 2B, the pretrained model was fine-tuned sequentially, first on the medium-density site and subsequently on the low-density site. Performance was evaluated after the first adaptation stage and again after the second stage on both domains. The weights obtained after fine-tuning on the medium-density site were used to initialise the second adaptation stage on the low-density site. This strategy was designed to assess whether progressive adaptation from relatively denser to sparser vegetation improved transfer across the structural gradient.

4.1.3. Experiment 3: Training from scratch on rangeland data

Experiment 3 provided a non-transfer baseline by training the ForestFormer3D framework from randomly initialised weights using only rangeland data. This setting was designed to assess whether rangeland-specific learning, in the absence of forest pretraining, could match or exceed the performance of transfer-based approaches under structural domain shift.

Three training configurations were considered: joint training across both vegetation-density domains (Exp 3A), training on the medium-density site only (Exp 3B), and training on the low-density site only (Exp 3C). All configurations used the same supervision derived from the automated 2D-to-3D label-transfer pipeline, together with the same shared semantic formulation, training protocol, and evaluation metrics as in the preceding experiments.

4.2. Experimental settings

4.2.1. Datasets, sites and sensors

The source-domain benchmark consists of the forest datasets used to train ForestFormer3D, including *FOR-instanceV2* (*CULS*, *NBIO*, *RMIT*, *SCION*, *TUWIEN*, *NBIO2*, *NIBIO_MLS*, *BlueCat*, and *Yuchen*). These datasets include multi-sensor lidar acquisitions annotated for semantic classes (ground, wood, leaf) and tree instances.

To evaluate cross-ecosystem transfer, semi-arid rangeland lidar tiles from the Calperum study area (*SASMDD0011*, *SASMDD0001*) were used. These sites represent structurally distinct environments characterised by low-stature, sparse, and multi-stemmed vegetation, providing a challenging testbed for

transfer from forest-trained models. Two rangeland conditions were considered:

(i) a medium-density site (grass, shrub, and mallee), and

(ii) a low-density site (forb, shrub, and sparse mallee).

Semantic classes correspond to vegetation growth forms, where ground represents the forb (herbaceous cover), wood corresponds to shrub structure, and leaf represents canopy elements. All primary experiments used UAS L1 lidar data, with UAS L2 used only as supplement comparisons (see results in Supplementary materials, see Table S3).

4.2.2. 3D deep learning models

This study evaluates two complementary 3D deep learning approaches for vegetation structure modelling from lidar data: a transformer-based voxel model and a point-based instance segmentation framework. Both approaches rely on sparse convolutional feature extraction implemented with Minkowski Engine (Choy et al., 2019), enabling efficient processing of large, irregular point clouds, but differ fundamentally in how spatial structure is encoded and aggregated.

Transformer-based voxel model. ForestFormer3D is employed as the primary voxel-based model (Xiang et al., 2025b), building on the OneFormer3D framework (Kolodiazhnyi et al., 2023). The architecture integrates sparse voxel feature extraction with transformer-based attention over voxel tokens, enabling global contextual reasoning across spatially distant regions and complex structural patterns in 3D vegetation. Additional components, ISA-guided query point selection, one-to-many association during training, and score-based block merging at inference, improve segmentation consistency in large and heterogeneous point clouds. This formulation is particularly suited to structurally complex environments, where local neighbourhood information alone may be insufficient to distinguish semantically similar vegetation components. Inputs are voxelised at a resolution of 0.2 m and cropped into 16 m cylindrical regions during both training and inference. Data augmentations are limited to random horizontal flipping, random rotation about the Z-axis, and random isotropic scaling in the range $[0.8, 1.2]$. The model is optimised using AdamW with an initial learning rate of 1×10^{-4} and a weight decay of 0.05. A polynomial learning-rate decay schedule with power 0.9 is used,

together with gradient clipping at a maximum norm of 10 and a batch size of 2.

Point-based instance model. For instance-level modelling, SegmentAnyTree is used as a point cloud segmentation framework (Wielgosz et al., 2024), building on the architecture of Xiang et al. (2023) and inspired by PointGroup-style clustering (Jiang et al., 2020). The model combines a sparse 3D convolutional backbone for voxel-space feature extraction with point-level instance grouping inspired by PointGroup-style clustering, followed by multiple prediction heads for semantic segmentation and instance clustering. Although features are extracted in voxel space, predictions are projected back to the original point cloud, enabling instance-level segmentation directly on unstructured data. This hybrid formulation preserves fine-grained geometric detail while leveraging efficient volumetric feature learning, emphasising local geometric grouping with lower computational demands than competing methods. Training augmentations comprise random Gaussian noise with $\sigma = 0.01$, random rotation about the Z-axis up to $\pm 180°$, anisotropic scaling in the range $[0.9, 1.1]$, random symmetry transformations, and a multi-density sparsification strategy ranging from 10 to 1000 $\mathrm{pts\,m^{-2}}$. This sparsification strategy is central to the model's sensor-agnostic design.

3D data representation and preprocessing. Both models operate on a unified point cloud representation derived from a 2D-to-3D label transfer pipeline. Each lidar tile is exported as a filtered `.ply` point cloud containing aligned semantic and instance labels:

$$P \in \mathbb{R}^{N\times 3}, \quad y \in \{0,1,2,3,4\}^N, \quad t \in \mathbb{N}^N \tag{3}$$

where P denotes the 3D spatial coordinates of N points, y represents the corresponding semantic class labels, and t denotes instance identifiers. Input features consist solely of spatial coordinates (x, y, z), recentered in the horizontal plane prior to training. No additional attributes (e.g., intensity, height-derived features) are used as model inputs. Semantic and instance labels are used exclusively as supervision targets. From this representation, two complementary 3D formulations are constructed:

(A) Sparse voxel representation (used by ForestFormer3D). The labelled point cloud is discretised into a regular 3D grid (0.20 m resolution), retaining only occupied voxels:

$$X_{\text{vox}} = \{(c_i, f_i)\}_{i=1}^{N_{\text{vox}}}, \quad c_i \in \mathbb{Z}^3,\ f_i \in \mathbb{R}^F \tag{4}$$

where c_i denotes the voxel coordinates in 3D space and f_i represents aggregated feature vectors derived from the recentered XYZ coordinates of points within each voxel. Semantic labels are assigned via majority voting and used as training targets.

(B) Point-based representation (used by SegmentAnyTree). The original point cloud is preserved, with each point retaining its spatial coordinates and associated supervision labels.

Both representations are derived from the same underlying point geometry: the voxel representation aggregates spatial information within a structured grid, whereas the point-based representation preserves it at the individual point level. Both operate on equivalent information content, ensuring that performance differences reflect model design, global contextual reasoning in voxel-based models versus local geometric grouping in point-based approaches, rather than differences in input data.

4.2.3. Training and inference protocol

Each rangeland site contained 15 tiles of 614×614 pixels, corresponding to 30.70×30.70 m at the 5 cm base resolution used in this study. For the medium-density site, 10/3/2 tiles were used for training, validation, and testing, respectively; for the low-density site, the split was 10/2/3. The same partitions were used for fine-tuning (Experiment 2) and training from scratch (Experiment 3), with the held-out test tiles reserved for evaluation. Zero-shot transfer (Experiment 1) used the full set of rangeland tiles for evaluation only, with forest performance assessed on the predefined held-out *FOR-instanceV2* test split.

To ensure comparability across experiments and with prior work, all training runs in this study inherited the original published configuration of the corresponding model (see Section 4.2.2) in full, so that the only deliberate difference between Experiments 2 and 3 was the initialisation. Experiment 2 loaded the published forest-pretrained checkpoint for few-shot adaptation, whereas Experiment 3 trained the same configuration from randomly initialised weights on the same target-domain split. This design isolates the effect of pretraining by holding the architecture, optimiser, preprocessing, and data partitioning identical within each model family. ForestFormer3D was trained for up to 3000 epochs, reflecting the long-convergence regime

typical of sparse voxel-based transformer models, while SegmentAnyTree followed the training schedule specified in its public release. Experiment-specific overrides were limited to the initialisation flag, the dataset references for the Calperum L1 splits, and the checkpoint-selection criteria. Full configuration files for both models are available in the project repository `https://github.com/LNSOTOM/fvc_structure`.

All primary results are reported using the final model at the end of each model's training schedule, following the practice of the original frameworks. Alternative checkpoints, best validation mIoU, best validation F1, and their harmonic mean (Eq. 5), were also retained and evaluated on the test set to analyse differences in optimisation behaviour; these results are reported in Supplementary Table S4 and are not used for the main comparison, to avoid selection bias.

$$H = \frac{2 \cdot \mathrm{mIoU} \cdot \mathrm{F1}}{\mathrm{mIoU} + \mathrm{F1}} \tag{5}$$

Inference was performed on the held-out test tiles using the trained or pre-trained models corresponding to each experimental setting. This consistent protocol ensured that observed performance differences were attributable to learning strategy (zero-shot, fine-tuning, or from-scratch) rather than to data splits or optimisation procedures.

4.2.4. Evaluation metrics

Semantic segmentation performance was evaluated using class-wise Intersection over Union (IoU) and mean IoU (mIoU). Instance-level performance was evaluated using precision, recall, F1-score, and coverage, defined as the average IoU between predicted and ground-truth (GT) tree or mallee instances (Wang et al., 2019; Xiang et al., 2024). In figure panels, the labels *GT semantic* and *Prediction semantic* denote the reference and predicted class maps, and *GT instance* and *Prediction instance* denote the reference and predicted instance counts per scene, allowing direct visual comparison of class confusion and under- or over-segmentation.

ForestFormer3D was evaluated using both semantic and instance-level metrics. SegmentAnyTree, which does not produce semantic class predictions, was evaluated using instance-level metrics only (precision, recall, F1-score, and coverage), providing a geometric reference for instance-level segmentation performance.

5. Results

5.1. Label transfer validation

5.1.1. Lidar-RGB co-registration

The lidar-derived DSMs were successfully standardised to the RGB/SfM DSM reference grids, enabling direct pixel-wise comparison for dense eFOLKI co-registration. Diagnostic figures showing DSM grid standardisation and before–after DSM residuals are provided in Supplementary materials, see Figure S1 and S2. The resulting displacement fields showed spatially variable lidar-RGB offsets rather than a single uniform shift, with summary statistics reported in Table 5.

Table 5: Summary of dense horizontal displacement fields estimated by eFOLKI for lidar–RGB co-registration. Values are reported in metres.

Site	Mean	Median	P_{85}	P_{95}	Max.	Mean ΔX	Mean ΔY
Low-density site	0.0455	0.0237	0.1087	0.1548	0.3856	0.0006	0.0019
Medium-density site	0.0737	0.0552	0.1522	0.2227	0.4935	0.0041	0.0060

As shown in Table 5, the medium-density site required larger local corrections than the low-density site, with higher mean, P_{95}, and maximum displacement values. In both sites, the mean signed components were small relative to the displacement magnitudes, indicating limited overall directional bias despite locally variable non-rigid corrections.

After applying the smoothed displacement fields to the full-density lidar point clouds, 97.71% of points were retained for the low-density site and 97.44% for the medium-density site. Points removed during this step fell outside the valid flow-field domain. The corresponding 2D centroid shifts were 3.126 m and 0.195 m, respectively, and are reported only as quality-control descriptors because the original and co-registered point clouds differ slightly in spatial support. The co-registered point clouds were then used for subsequent 2D–3D label transfer.

5.1.2. 2D-to-3D label transfer workflow validation

The diagnostics defined in Section 3.2.3 were applied to the 15 medium validation tiles, with visual examples also provided for the low validation tiles (see Figure 8 and Figure 9). Results are reported in the same order as the validation design: voxel-level label agreement, height plausibility of refined

classes, and Mallee crown-instance retention. The 2D reference inputs and spatial alignment underpinning the transfer are shown in Figure 6, and the representative three-stage transfer sequence is shown in Figure 7. Detailed per-tile diagnostic plots are provided in Supplementary Figures S3-S5 for the medium tiles.

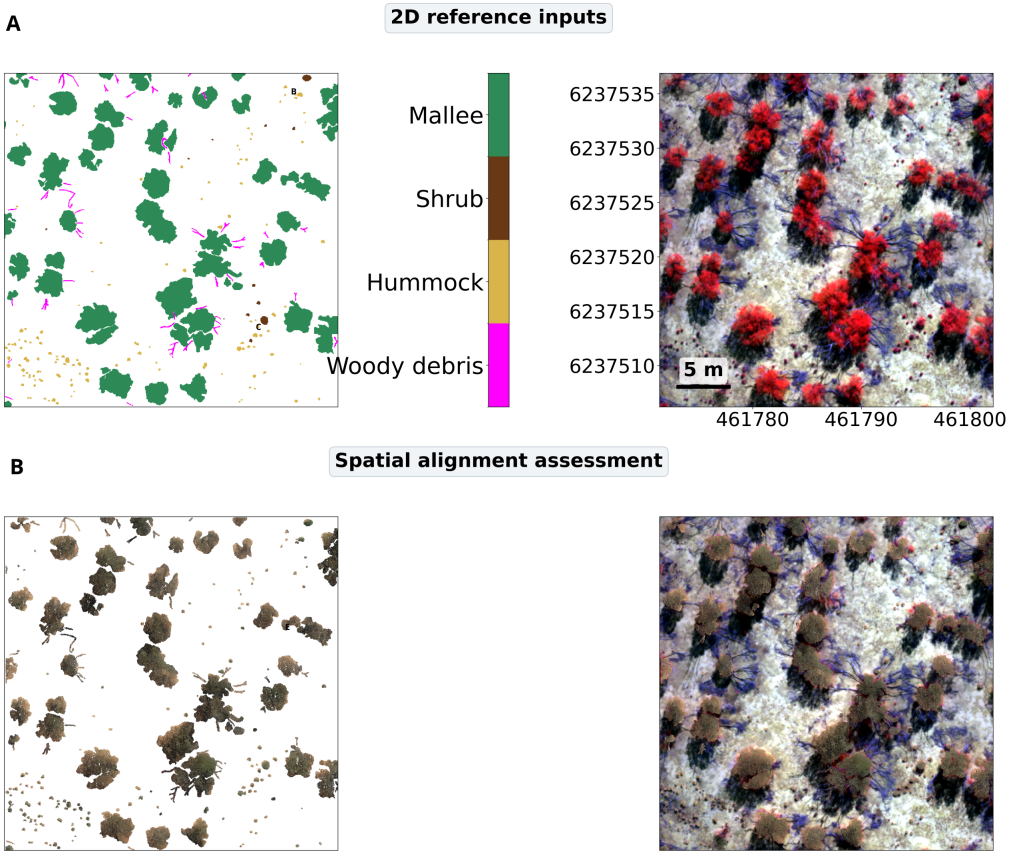

Figure 6: 2D reference inputs and spatial alignment for label transfer at the medium site. **A** 5 cm structural-class mask raster containing woody debris, hummock, shrub, and mallee, used as the source labels for transfer to the lidar point cloud, alongside the corresponding 3070 × 3070 pixel multispectral raster tile. **B** RGB-coloured lidar point cloud in plan view and the same cloud overlaid on the corresponding 5 cm multispectral raster tile, used to verify pixel-to-point alignment before Cartesian voxel merging.

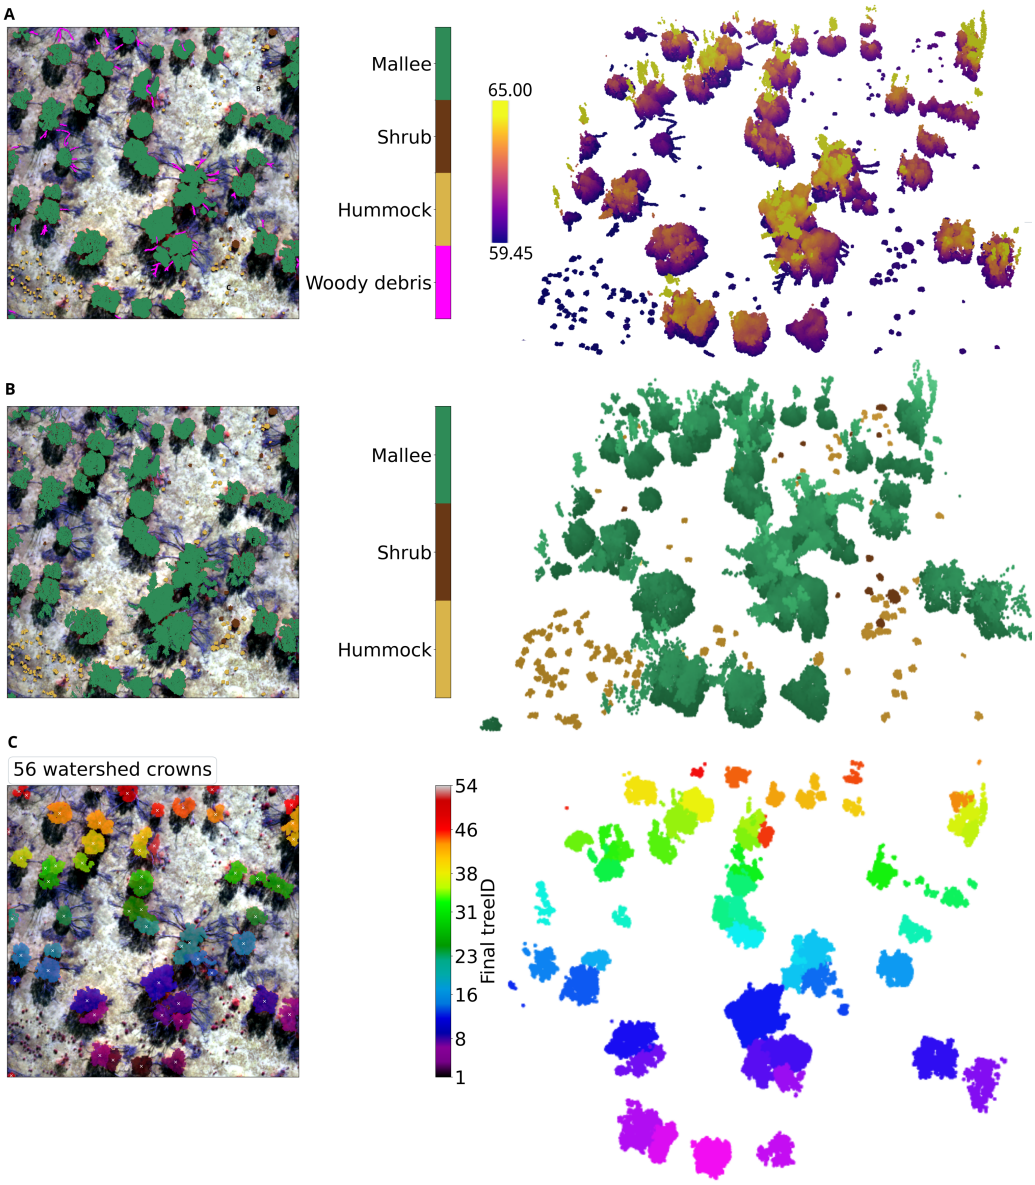

Figure 7: Three-stage label transfer from the 2D structural mask to the 3D lidar point cloud for a representative medium-site validation raster tile. **A** Initial Cartesian voxel-based label transfer, showing the 2D structural mask projected over the multispectral raster tile, including woody debris, hummock, shrub, and mallee, and the corresponding 3D lidar point cloud coloured by absolute elevation Z in metres. **B** Refined structural labels after 3D refinement, showing the updated 2D overlay and the corresponding 3D point cloud retaining hummock, shrub, and mallee classes. **C** Final mallee crown-instance delineation, showing 56 watershed crown detections on the multispectral raster tile and the corresponding 3D point cloud coloured by final *treeID*, with 54 retained final tree instances.

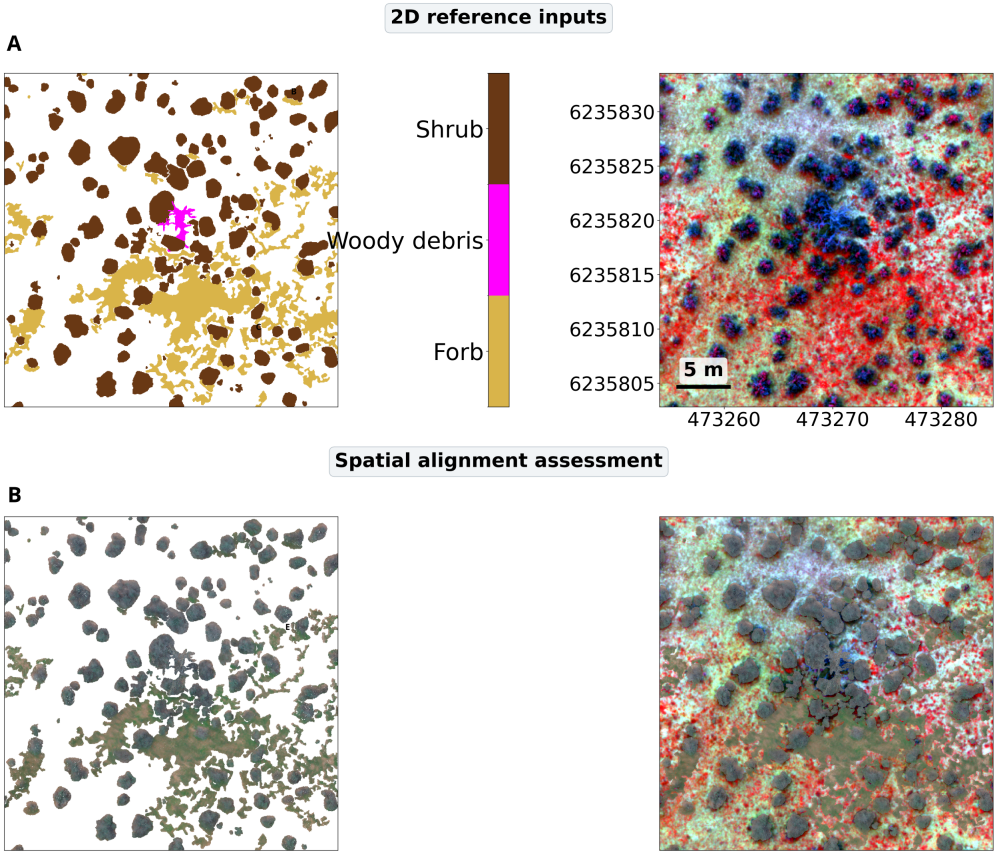

Figure 8: 2D reference inputs and spatial alignment for label transfer at the low site. **A** 5 cm structural-class mask raster containing forb, shrub and woody debris, used as the source labels for transfer to the lidar point cloud, alongside the corresponding 3070 × 3070 pixel multispectral imagery tile. **B** RGB-coloured lidar point cloud in plan view and the same cloud overlaid on the corresponding 5 cm multispectral raster tile, used to verify pixel-to-point alignment before Cartesian voxel merging.

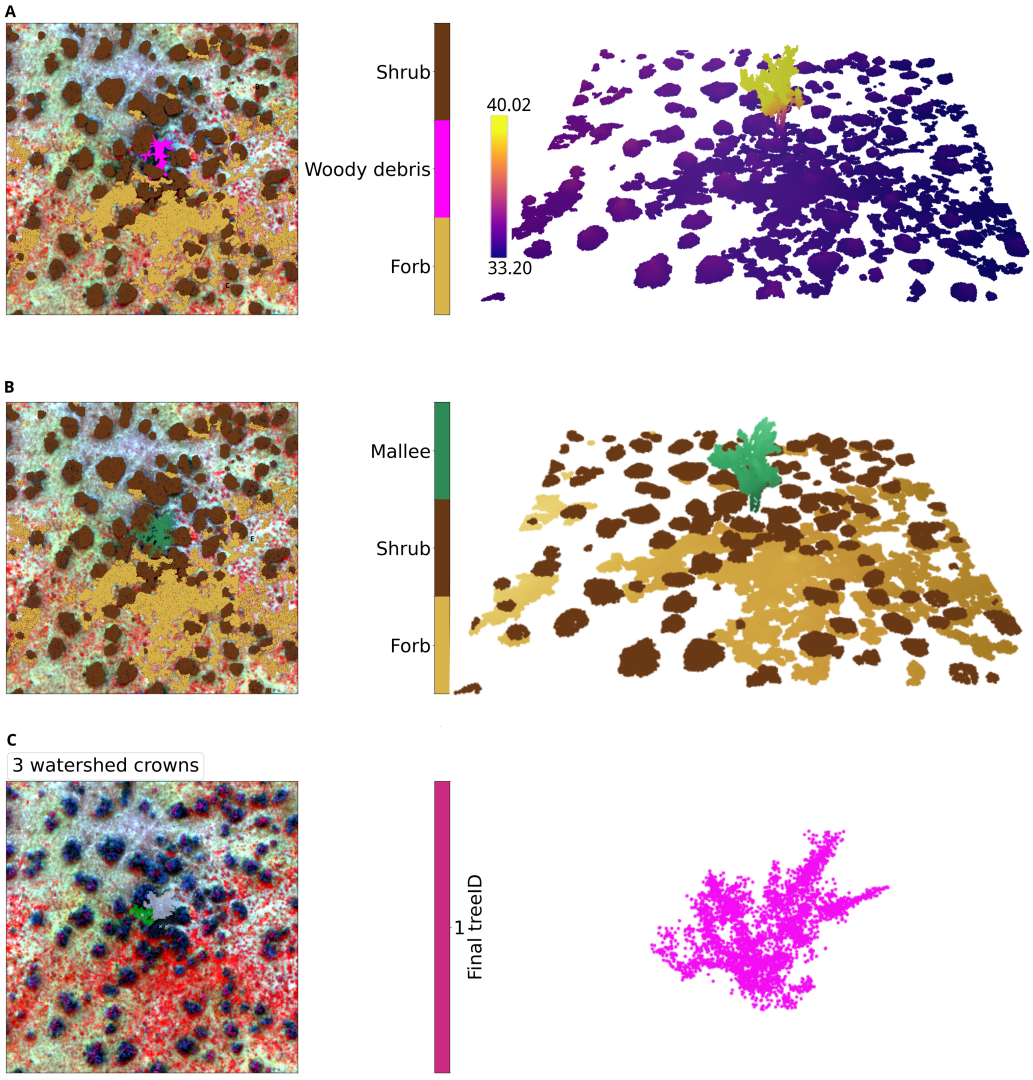

Figure 9: Three-stage label transfer from the 2D structural mask to the 3D lidar point cloud for a representative medium-site validation raster tile. **A** Initial Cartesian voxel-based label transfer, showing the 2D structural mask projected over the multispectral raster tile, including forb, woody debris, and shurb, and the corresponding 3D lidar point cloud coloured by absolute elevation Z in metres. **B** Refined structural labels after 3D refinement, showing the updated 2D overlay and the corresponding 3D point cloud retaining forb, shrub, and mallee classes. **C** Final mallee crown-instance delineation, showing 3 watershed crown detections on the multispectral raster tile and the corresponding 3D point cloud coloured by final *treeID*, with 1 retained final tree instances.

Voxel-level label agreement.. The consolidated labels showed high agreement with the Cartesian voxel baseline across the 15 validation tiles. Mean OA was 98.4%, with a range of 96.6–99.5% (Supplementary Figure S3). Mean class-wise IoU was 99.7% for Hummock, 99.9% for Shrub on tiles where Shrub remained present after consolidation, and 98.0% for Mallee. Woody debris was absent after consolidation, and Forb was absent from the validation footprint.

Height plausibility of refined classes.. Hummock medians remained within the configured [0.00, 0.80] m range across all validation tiles, with an average of 1% of points outside the configured range (Supplementary Figure S4). Shrub and Mallee showed larger outside-range fractions, averaging 95% and 72%, respectively. For these two classes, the tile-level medians were generally positioned below the configured lower bounds rather than above the upper bounds.

Mallee crown-instance retention.. Mallee crown-instance retention averaged 94.4% across the 15 validation tiles, with values ranging from 83.8% to 101.5% (Supplementary Figure S5). Tile 106 increased from 66 watershed-derived crowns to 67 final non-zero *treeID*s, producing 101.5% retention. The lowest-retention tiles were 26, 93, 192, and 205, all below 90%.

Together, these diagnostics summarise voxel-level, height-based, and crown-instance consistency across the 15 validation tiles at the medium site.

5.2. Experimental results

5.2.1. Experiment 1: Zero-shot transfer of forest-trained models

Experiment 1 evaluated how models trained on the *FOR-instanceV2* forest dataset transfer, without fine-tuning, to the structurally distinct Calperum semi-arid rangeland sites. Results are reported first for the in-domain held-out forest test split, then for the medium- and low-density Calperum UAS L1 sites. ForestFormer3D was evaluated for both semantic and instance segmentation, while SegmentAnyTree was included as an instance-centric geometric reference because it does not produce semantic classes comparable to ForestFormer3D. Table 6 summarises the zero-shot results.

Table 6: Experiment 1: Zero-shot transfer of ForestFormer3D and SegmentAnyTree from forest training data to held-out forest and Calperum rangeland test sets, without fine-tuning. UAS L1 results are reported here; UAS L2 results are in Supplementary Table S3.

	Region	Sensor	Method	Ind. Tree Seg. (%)				Semantic Seg. (%)			
				Prec	Rec	F1	Cov	Ground	Wood	Leaf	mIoU
FOR-instanceV2	Forest	ULS, MLS, TLS	ForestFormer3D	94.3	84.9	89.4	80.2	98.8	62.6	94.3	85.2
			SegmentAnyTree	78.2	76.7	76.5	71.5	–	–	–	–

	Region	Sensor	Method	Ind. Tree Seg. (%)				Semantic Seg. (%)			
				Prec	Rec	F1	Cov	Hum./Forb (Ground)	Shrub (Wood)	Mallee (Leaf)	mIoU
Calperum	Rangeland (med)	UAS L1	ForestFormer3D	7.4	3.1	4.4	12.0	42.0	2.3	78.9	41.1
			SegmentAnyTree	5.75	3.28	4.16	13.22	–	–	–	–
	Rangeland (low)	UAS L1	ForestFormer3D	5.0	40.0	8.9	52.8	33.3	0.2	14.7	16.1
			SegmentAnyTree	0.56	33.33	5.13	37.14	–	–	–	–

Experiment 1A: ForestFormer3D zero-shot transfer.. On the held-out *FOR-instanceV2* forest test split, ForestFormer3D retained strong in-domain performance, reaching 85.2% mIoU for semantic segmentation and 89.4% instance F1 with 80.2% coverage. The visual example in Figure 10 corresponds to `SCION_SCION_plot_61_annotated_test`, a *Pinus radiata* plantation plot in New Zealand, and shows close agreement between ground-truth and predicted ground, wood, and leaf structure, with only a small reduction in detected crowns.

When transferred to Calperum without adaptation, performance declined substantially. Semantic mIoU decreased to 41.1% at the medium-density site and 16.1% at the low-density site. At the medium-density site, mallee remained the strongest class (78.9% IoU), whereas hummock and shrub were much weaker (42.0% and 2.3% IoU). At the low-density site, the model degraded further across all classes, particularly shrub, which reached only 0.2% IoU.

Instance segmentation showed an even stronger domain-transfer failure. At the medium-density site, instance F1 dropped to 4.4% with 12.0% coverage, indicating poor crown-level delineation despite partial recovery of mallee structure. At the low-density site, recall increased to 40.0%, but precision remained low at 5.0%, yielding an F1 score of 8.9%. Figure 11 illustrates these failure modes: reduced mallee instance recovery in the medium-density tile, shrub-to-forb collapse in the low-density tiles, and sparse false mallee predictions where mallee is absent.

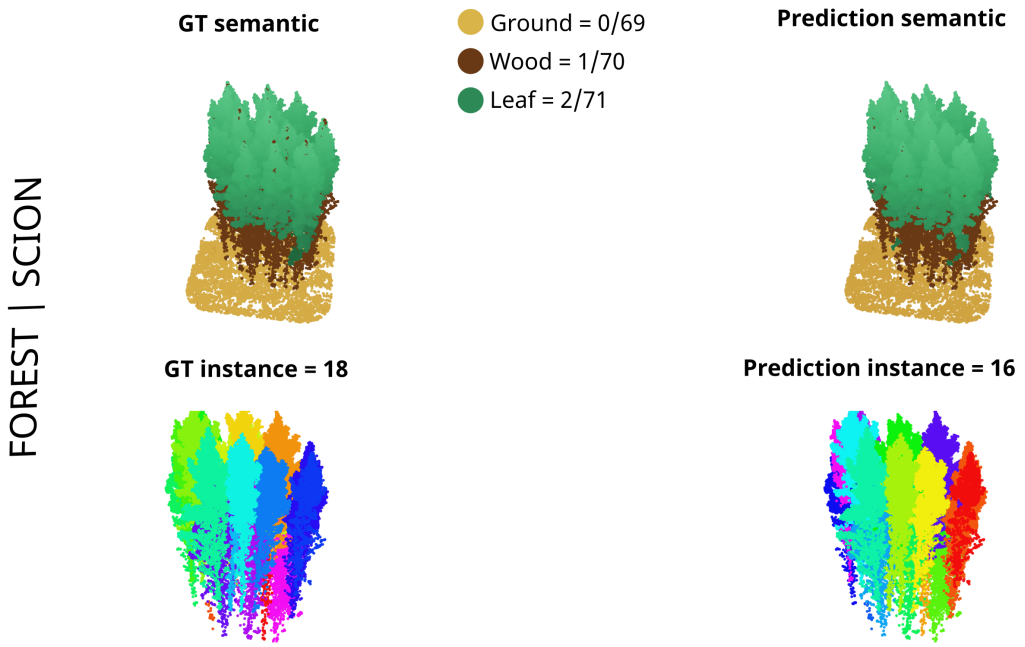

Figure 10: Experiment 1A zero-shot ForestFormer3D predictions on the held-out *FOR-instanceV2* forest test scene `SCION_SCION_plot_61_annotated_test`. The figure compares ground-truth and predicted semantic labels for ground, wood, and leaf, together with tree instance segmentation. The model reproduces the layered semantic structure closely and recovers most individual crowns, with only a small reduction in detected instances (18 to 16). Semantic colours follow the legend; instance colours distinguish individual trees.

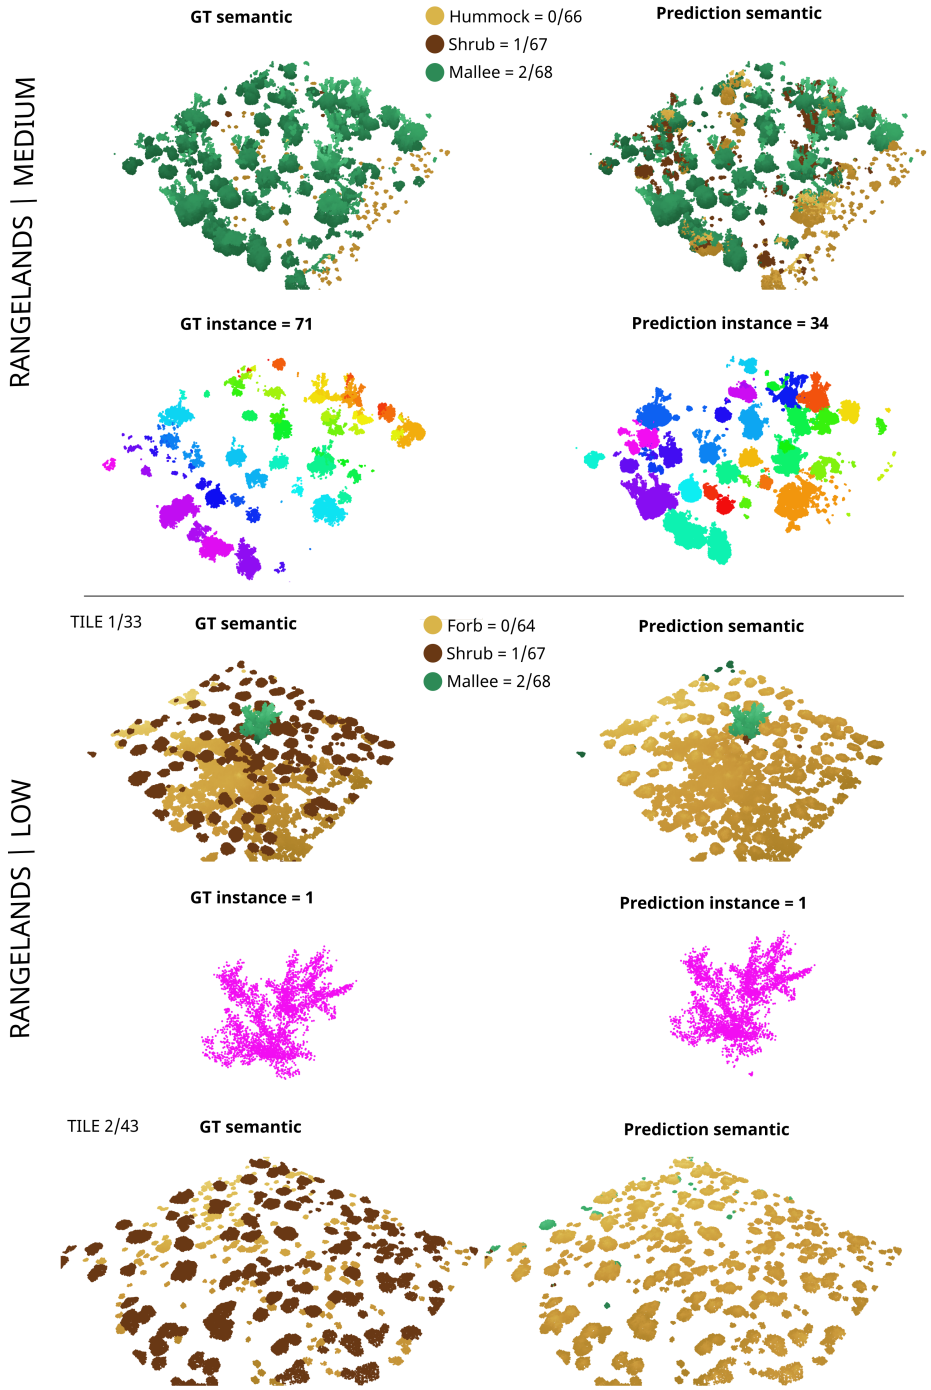

Figure 11: Experiment 1A zero-shot ForestFormer3D predictions on representative Calperum UAS L1 rangeland tiles. The medium-density example shows tile `znorm_3Dlabels_m_L1_93` with semantic predictions for hummock, shrub, and mallee, and mallee instance predictions. The low-density examples show `znorm_3Dlabels_l_L1_33`, containing a single mallee tree, and `znorm_3Dlabels_l_L1_43`, where mallee is absent. The model partially recovers mallee structure in the medium site but reduces the number of detected instances from 71 to 34. In the low-density tiles, shrub points largely collapse into forb, with sparse false mallee fragments where mallee is absent. Semantic colours follow the legend; instance colours distinguish individual mallee plants.

Experiment 1B: SegmentAnyTree zero-shot instance transfer.. SegmentAnyTree was evaluated under the same zero-shot protocol as an instance-level geometric reference. On the held-out forest test split, it reached 76.5% instance F1 with 71.5% coverage, confirming that the model retained effective crown grouping within the source forest domain.

The same model transferred poorly to the Calperum rangeland sites. At the medium-density site, instance F1 declined to 4.16% with 13.22% coverage, indicating strong under-detection and weak spatial correspondence with reference mallee crowns. At the low-density site, recall reached 33.33%, but precision dropped to 0.56%, producing an F1 score of 5.13%. Figure 12 shows that SegmentAnyTree preserved most forest crowns but failed to provide reliable rangeland instance delineation, particularly under sparse low-density conditions.

Overall, Experiment 1 establishes a clear zero-shot baseline: both forest-trained models generalised well within the forest source domain but showed limited transferability to semi-arid rangeland vegetation, especially for shrub discrimination and coherent mallee instance segmentation. These results define the reference point for the adapted models in Experiment 2 and the from-scratch models in Experiment 3.

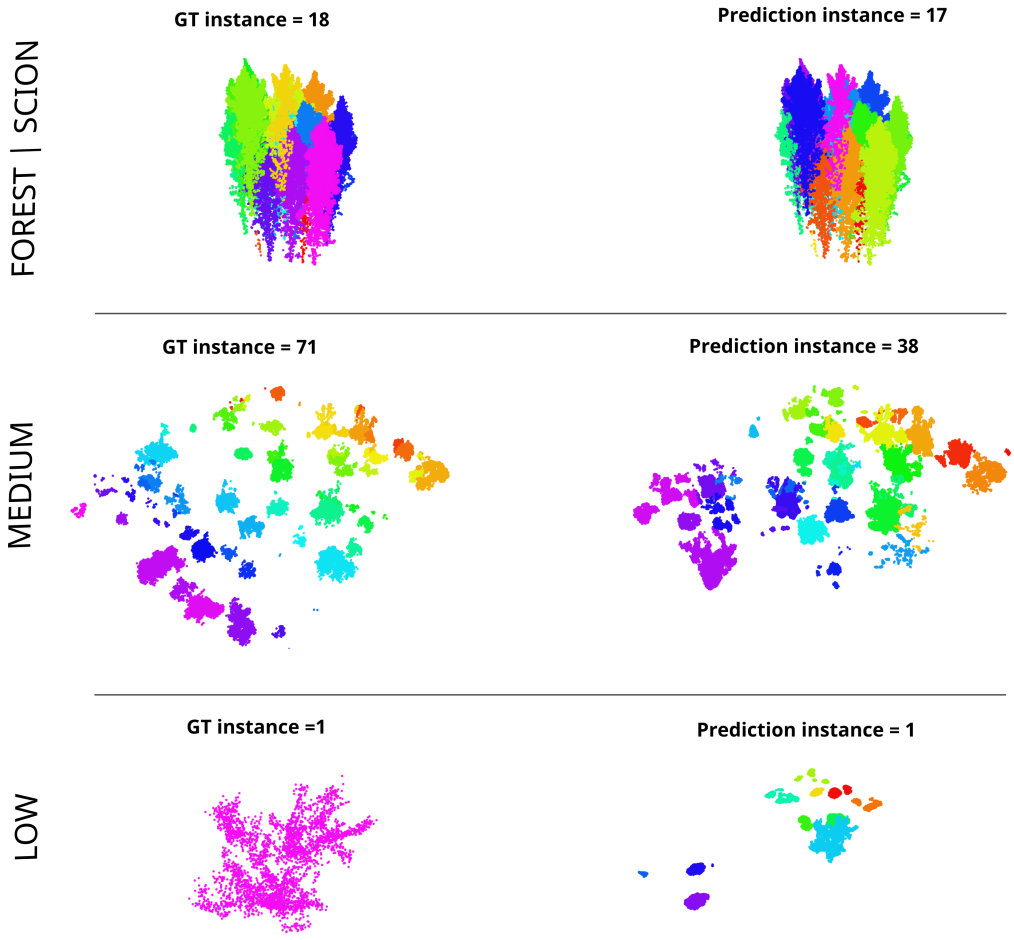

Figure 12: Experiment 1B zero-shot SegmentAnyTree instance predictions on the held-out *FOR-instanceV2* forest test scene `SCION_SCION_plot_61_annotated_test`, the medium-density Calperum tile `znorm_3Dlabels_m_L1_93`, and the low-density tile `znorm_3Dlabels_l_L1_33`. Each row compares ground-truth and predicted instances. The model preserves most crowns in the forest scene (18 to 17), but detects fewer instances in the medium-density rangeland tile (71 to 38) and shows weaker spatial agreement in the low-density example. Instance colours distinguish individual trees or mallee crowns.

5.2.2. Experiment 2: Few-shot fine-tuning for cross-domain adaptation

Experiment 2 evaluated whether limited rangeland supervision could improve the transfer of the forest-pretrained ForestFormer3D model to the Calperum UAS L1 sites. Two adaptation strategies were tested: joint fine-tuning using combined medium- and low-density rangeland samples (Exper-

iment 2A), and curriculum fine-tuning from the medium-density site to the low-density site (Experiment 2B). Tables 7 and 8 summarise the results.

Table 7: Experiment 2A joint few-shot fine-tuning of the forest-pretrained ForestFormer3D model using combined medium- and low-density Calperum rangeland samples.

	Region	Sensor	Method	Ind. Tree Seg. (%)				Semantic Seg. (%)			
				Prec	Rec	F1	Cov	Hum./Forb (Ground)	Shrub (Wood)	Mallee (Leaf)	mIoU
Calperum	Rangeland (med)	UAS L1	ForestFormer3D	58.9	29.2	39.1	26.2	56.3	4.1	86.3	48.9
	Rangeland (low)	UAS L1	ForestFormer3D	0.0	0.0	0.0	0.0	74.3	26.7	–	50.5

Table 8: Experiment 2B curriculum few-shot fine-tuning of the forest-pretrained ForestFormer3D model. Stage 1 was adapted and evaluated on the medium-density site, and the resulting weights were then used to initialise Stage 2 adaptation on the low-density site.

	Region	Sensor	Method	Ind. Tree Seg. (%)				Semantic Seg. (%)			
				Prec	Rec	F1	Cov	Hum./Forb (Ground)	Shrub (Wood)	Mallee (Leaf)	mIoU
Calperum	Rangeland (med)	UAS L1	ForestFormer3D (Stage 1, med.)	54.2	46.0	49.8	44.8	49.2	7.1	92.5	49.6
	Rangeland (low)	UAS L1	ForestFormer3D (Stage 2, low)	0.0	0.0	0.0	0.0	80.1	46.2	–	63.1

Experiment 2A: Joint fine-tuning. Joint fine-tuning improved semantic transfer relative to the zero-shot baseline, increasing mIoU to 48.9% at the medium-density site and 50.5% at the low-density site. At the medium-density site, mallee remained the strongest class, reaching 86.3% IoU, while hummock reached 56.3% IoU and shrub remained weak at 4.1% IoU. At the low-density site, forb and shrub reached 74.3% and 26.7% IoU, respectively.

Instance segmentation improved only for the medium-density site, where precision, recall, and F1-score reached 58.9%, 29.2%, and 39.1%, with 26.2% coverage. In contrast, instance-level metrics remained zero at the low-density site. Thus, joint fine-tuning substantially improved semantic adaptation,

but the gains remained class-dependent and did not translate into reliable instance delineation across both rangeland conditions.

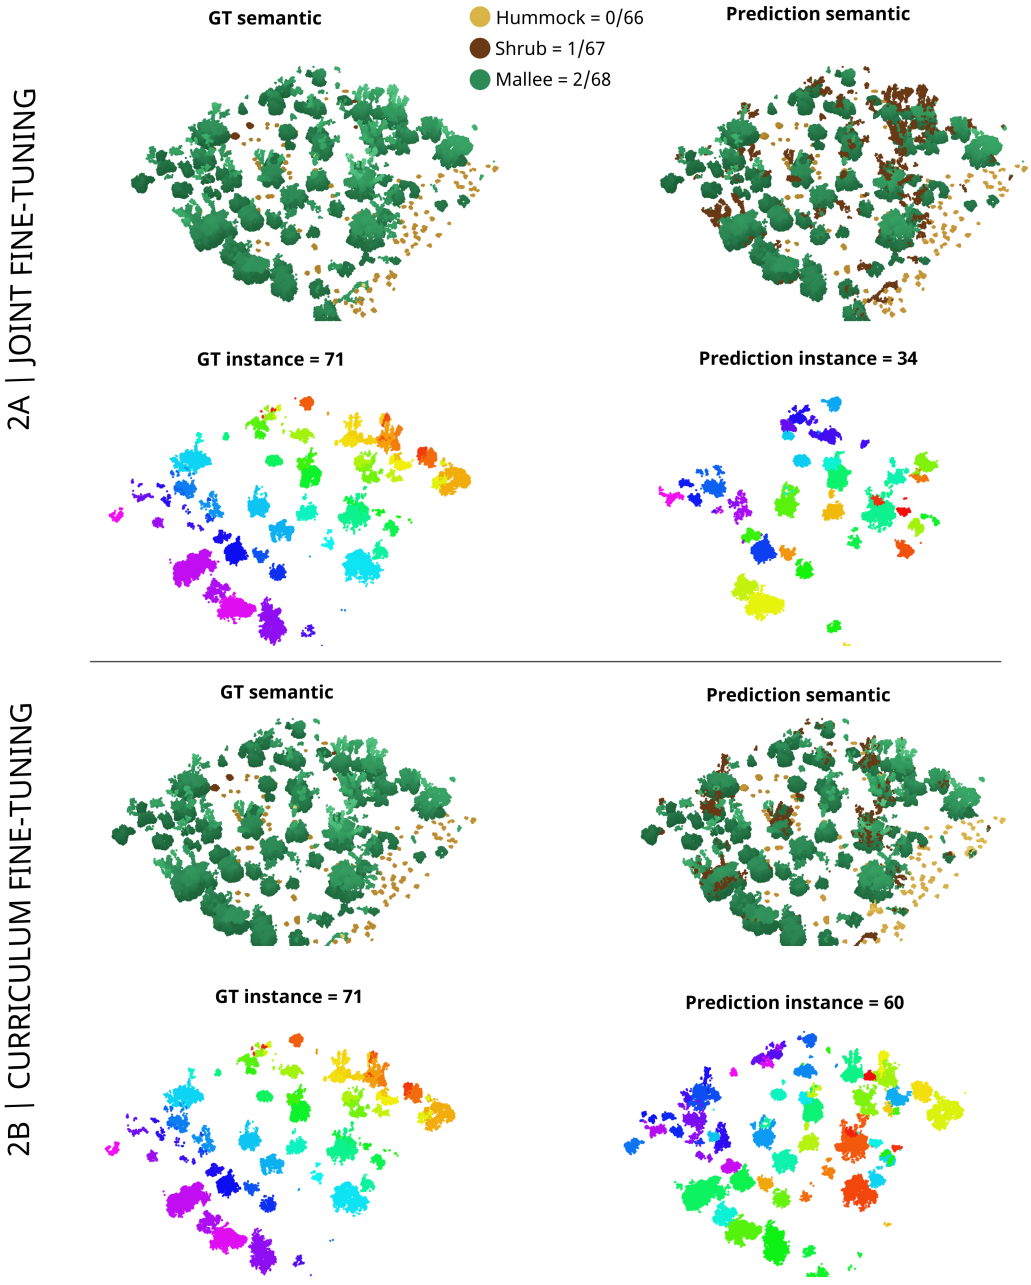

Figure 13: Experiment 2A and Experiment 2B ForestFormer3D predictions on the medium-density Calperum tile `znorm_3Dlabels_m_L1_93`. Experiment 2A shows joint fine-tuning with combined medium- and low-density samples, while Experiment 2B shows curriculum fine-tuning after medium-site adaptation. Each block compares ground-truth and predicted semantic labels for hummock, shrub, and mallee, together with mallee instance segmentation. Curriculum fine-tuning increased the number of detected mallee instances from 34 to 60 relative to joint fine-tuning, closer to the 71 ground-truth instances. Semantic colours follow the legend; instance colours distinguish individual mallee plants.

Experiment 2B: Curriculum fine-tuning. Curriculum fine-tuning produced the strongest Experiment 2 results, but at different stages for each site. Stage 1 was initialised from the forest-pretrained weights and produced the strongest medium-site semantic and instance performance after adaptation to the medium-density site. Stage 2 was initialised from the Stage 1 medium-adapted weights and produced the strongest low-site semantic performance after sequential medium-to-low adaptation.

In Stage 1, where the forest-pretrained model was adapted and evaluated on the medium-density site, semantic performance reached 49.6% mIoU, slightly higher than joint fine-tuning at the same site. Class-wise IoU was highest for mallee at 92.5%, followed by hummock at 49.2% and shrub at 7.1%. Instance performance also improved relative to joint fine-tuning, with precision, recall, and F1-score reaching 54.2%, 46.0%, and 49.8%, respectively, and coverage increasing to 44.8%.

In Stage 2, the medium-adapted weights from Stage 1 were used to initialise fine-tuning on the low-density site. This produced the strongest semantic performance across the low-density experiments, reaching 63.1% mIoU, with 80.1% IoU for forb and 46.2% IoU for shrub. Mallee IoU was not reported because mallee was absent from the low-density test set. For the same reason, instance-level precision, recall, F1-score, and coverage were zero: there were no mallee tree instances available for instance-level evaluation in this test set.

Overall, Experiment 2 shows that few-shot fine-tuning improves cross-domain semantic transfer compared with zero-shot inference. Curriculum fine-tuning provided the strongest semantic adaptation for the medium-density site in Stage 1 and for the low-density site in Stage 2. It also produced the strongest instance delineation for the medium-density site. However, shrub discrimination remained weak at the medium-density site, and instance segmentation remained dependent on the presence of coherent mallee crown structure in the evaluation data.

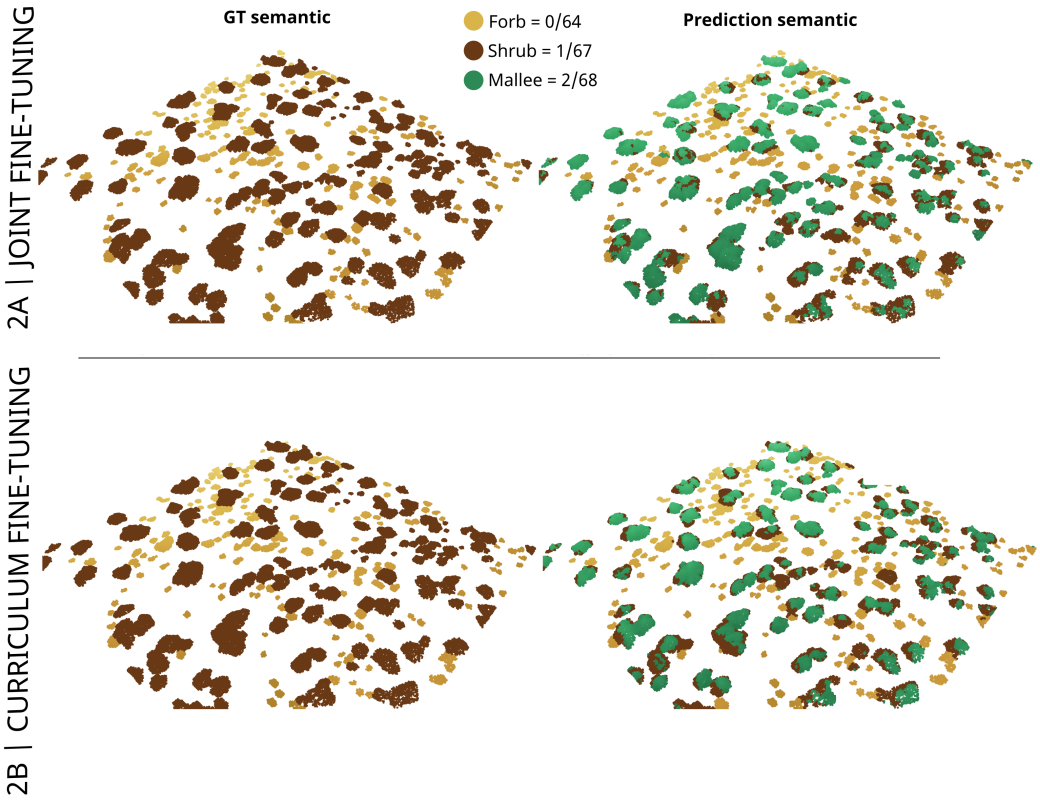

Figure 14: Experiment 2A and Experiment 2B ForestFormer3D semantic predictions on the low-density Calperum tile `znorm_3Dlabels_l_L1_43`. Experiment 2A shows joint fine-tuning with combined medium- and low-density samples, while Experiment 2B shows curriculum fine-tuning after sequential medium-to-low adaptation. Curriculum fine-tuning improves the separation of forb and shrub relative to joint fine-tuning, but sparse false mallee predictions remain. Semantic colours follow the legend.

5.2.3. Experiment 3: Training from scratch on rangeland data

Experiment 3 tested whether rangeland supervision alone could replace forest pretraining. ForestFormer3D was trained from randomly initialised weights using only Calperum UAS L1 rangeland data. Three configurations were evaluated: joint training on medium- and low-density sites (Experiment 3A), training on the medium-density site only (Experiment 3B), and training on the low-density site only (Experiment 3C). Table 9 summarises the results.

Table 9: Experiment 3 training from scratch on Calperum rangeland data. ForestFormer3D was trained using randomly initialised weights and only rangeland samples, without forest-pretrained initialisation.

	Region	Sensor	Method	Ind. Tree Seg. (%)				Semantic Seg. (%)			
				Prec	Rec	F1	Cov	Hum./Forb (Ground)	Shrub (Wood)	Mallee (Leaf)	mIoU
Calperum	Rangeland (med)	UAS L1	ForestFormer3D (3A, joint)	0.0	0.0	0.0	0.0	0.0	0.6	75.1	25.2
	Rangeland (low)	UAS L1	ForestFormer3D (3A, joint)	0.0	0.0	0.0	0.0	0.1	16.6	–	8.3
	Rangeland (med)	UAS L1	ForestFormer3D (3B, med. only)	0.0	0.0	0.0	0.0	25.8	0.0	88.42	38.07
	Rangeland (low)	UAS L1	ForestFormer3D (3C, low only)	0.0	0.0	0.0	0.0	38.19	22.23	–	30.21

Note: Mallee IoU is not reported for the low-density site because mallee was absent from the corresponding test set.

Experiment 3A: Joint training from scratch. Joint training from scratch on combined medium- and low-density rangeland data produced weak semantic performance compared with the transfer-based settings. The model reached 25.2% mIoU on the medium-density site and 8.3% mIoU on the low-density site. At the medium-density site, performance was concentrated almost entirely in the mallee class, which reached 75.1% IoU, while hummock and shrub were nearly absent from the prediction, with IoU values of 0.0% and 0.6%, respectively. At the low-density site, forb and shrub reached only 0.1% and 16.6% IoU, respectively; mallee IoU was not reported because mallee was absent from the test set.

Instance-level precision, recall, F1-score, and coverage were zero. For the medium-density site, this indicates that the model failed to produce coherent mallee crown instances despite the presence of 71 reference instances in the evaluated tile. For the low-density site, instance-level metrics are not informative because no mallee tree instances were present in the test set. These results show that joint training from scratch did not provide a viable alternative to transfer-based initialisation under the available rangeland sample size.

Experiment 3B: Training on the medium-density site only. Training from scratch on the medium-density site only improved semantic performance relative to the joint scratch setting, increasing mIoU from 25.2% to 38.07%. Class-wise performance again remained strongly unbalanced: mallee reached 88.42% IoU and hummock increased to 25.8% IoU, but shrub was not recovered, with 0.0% IoU. Thus, the improvement was driven mainly by better recovery of the dominant mallee structure rather than balanced class separation.

Instance-level precision, recall, F1-score, and coverage all remained zero. Compared with Experiment 2, medium-only training from scratch improved over the joint scratch baseline but remained below the fine-tuned models and did not recover usable mallee instance segmentation.

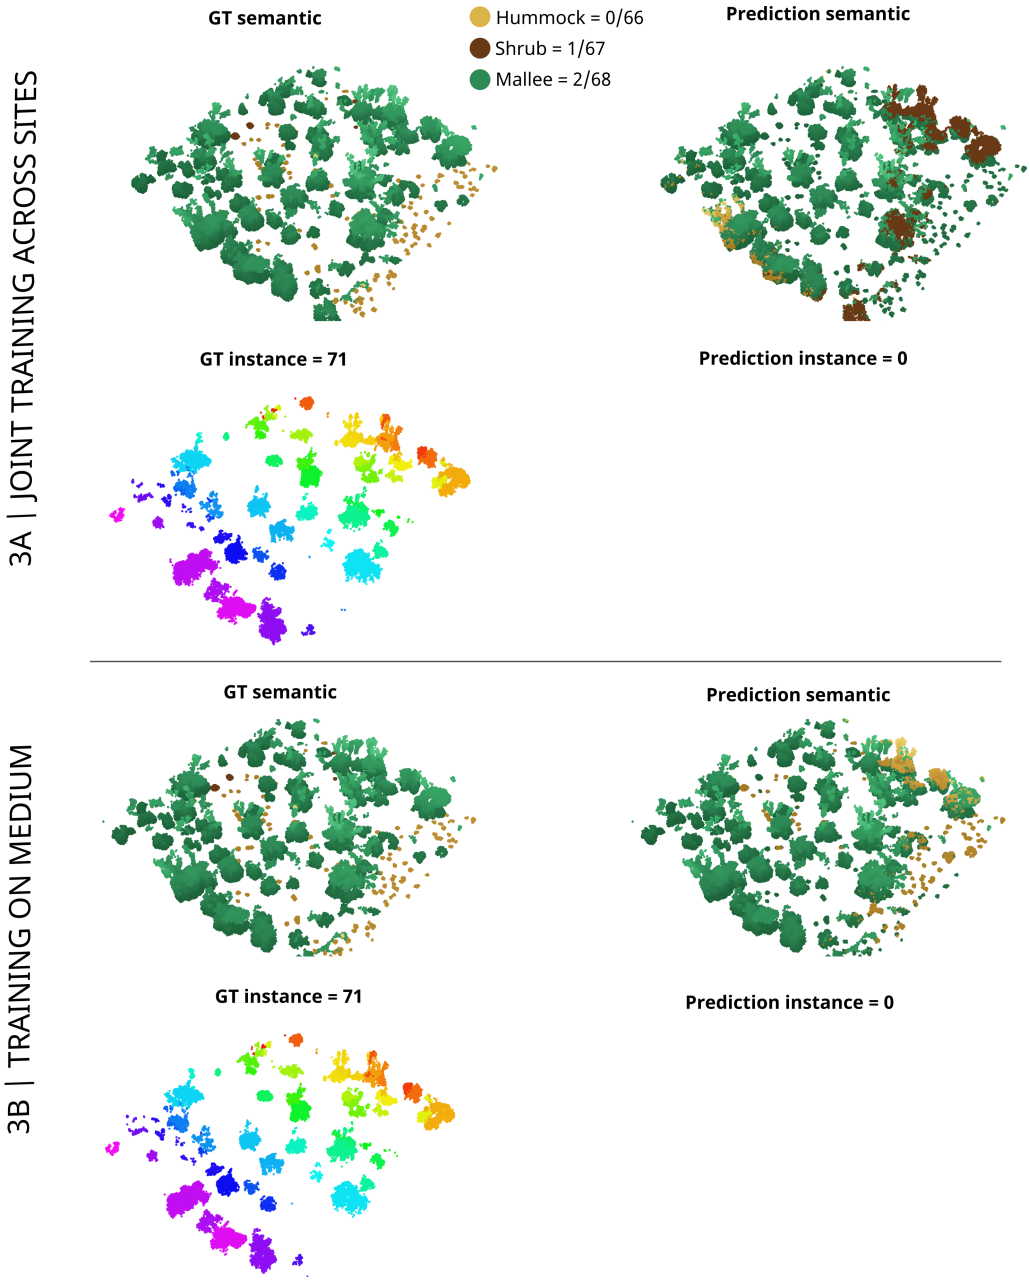

Figure 15: Experiment 3A and Experiment 3B ForestFormer3D predictions on the medium-density Calperum tile `znorm_3Dlabels_m_L1_93`. Experiment 3A shows joint training from scratch using combined medium- and low-density samples, while Experiment 3B shows training from scratch on the medium-density site only. Each block compares ground-truth and predicted semantic labels for hummock, shrub, and mallee, together with mallee instance segmentation. Both scratch-trained models fail to predict mallee instances on this tile (71 to 0), although medium-only training improves semantic recovery of mallee. Semantic colours follow the legend; instance colours distinguish individual mallee plants.

Experiment 3C: Training on the low-density site only. Training from scratch on the low-density site only also improved over the corresponding joint scratch result. Semantic mIoU increased from 8.3% under joint training to 30.21% under low-only training. Class-wise IoU values were 38.19% for forb and 22.23% for shrub. Mallee IoU was not reported because mallee was absent from the low-density test set. This indicates that site-specific training improved recovery of the non-mallee classes compared with joint training from scratch, but performance still remained below the few-shot fine-tuned models.

Instance-level precision, recall, F1-score, and coverage again remained zero because no mallee tree instances were available for instance-level evaluation in the low-density test set. Overall, site-specific scratch training improved semantic performance relative to joint scratch training for the corresponding site, but it did not match the transfer-based models and did not provide evidence of reliable instance-level recovery.

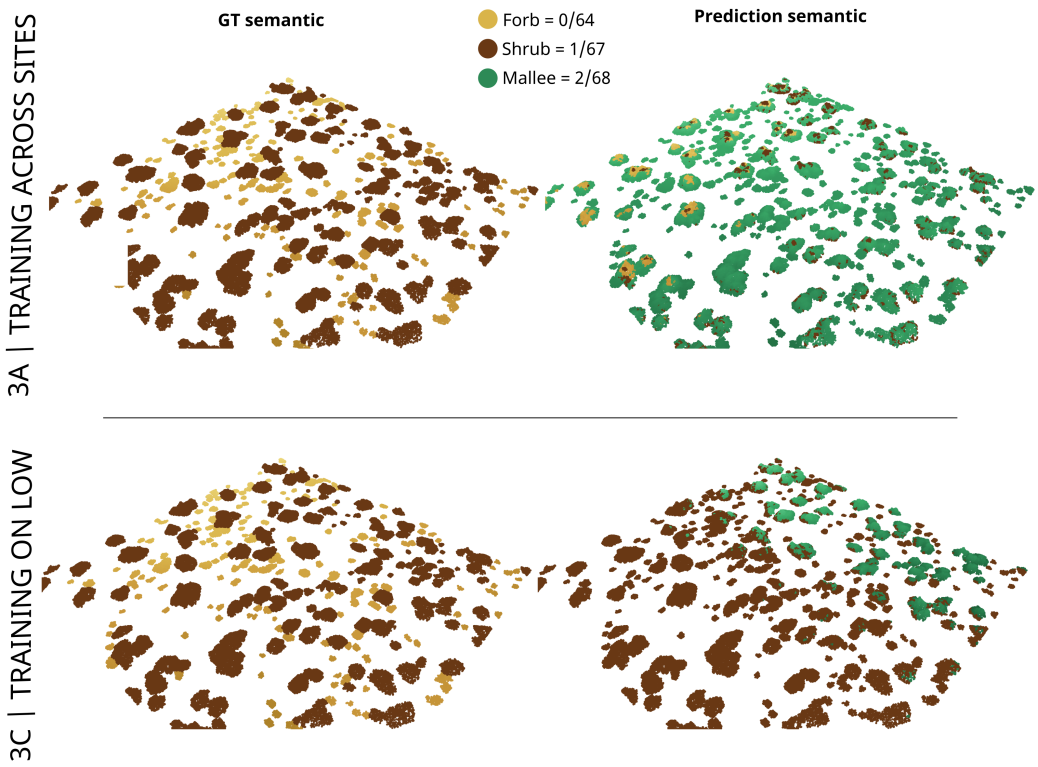

Figure 16: Experiment 3A and Experiment 3C ForestFormer3D semantic predictions on the low-density Calperum tile `znorm_3Dlabels_1_L1_43`. Experiment 3A shows joint training from scratch using combined medium- and low-density samples, while Experiment 3C shows training from scratch on the low-density site only. Low-only training improves forb and shrub separation relative to joint training, but predictions remain noisy and include sparse false mallee fragments despite mallee being absent from the test tile. Semantic colours follow the legend.

Overall, Experiment 3 shows that training ForestFormer3D from scratch on the available rangeland samples was consistently weaker than few-shot adaptation from forest-pretrained weights. Site-specific scratch training improved semantic performance relative to joint scratch training, but the gains were class-dependent and did not produce usable instance segmentation. These results indicate that forest-pretrained initialisation provided a more effective starting point for cross-domain adaptation than random initialisation under limited rangeland supervision.

6. Discussion

6.1. 2D-3D label transfer

We demonstrated a workflow for generating 3D semantic supervision for semi-arid rangeland vegetation by transferring high-resolution 2D labels from multispectral imagery to co-registered UAS lidar point clouds. This addresses a central limitation in vegetation remote sensing: growth forms are defined by morphology, vertical organisation, and spatial arrangement, yet many operational workflows still rely on 2D spectral or surface-based products (Omasa et al., 2007; Ganz et al., 2019; Jacobs et al., 2021; Gassilloud et al., 2025). By retaining point-level structure, the proposed workflow provides a pathway for mapping growth forms in ecosystems where field surveys, TLS, and MLS provide detailed structural information but remain difficult to scale across landscapes (Kellner et al., 2019; Disney, 2019; Calders et al., 2020; D'hont et al., 2025).

The co-registration results show that dense non-rigid alignment can connect RGB-derived and multispectral-derived semantic information with lidar geometry at a scale suitable for point cloud learning. This is important because small geometric offsets between orthomosaics, DSMs, and lidar point clouds can propagate directly into incorrect point labels. Similar to recent UAS studies that emphasise the importance of image resolution, orthorectification quality, sensor characteristics, and acquisition conditions for model transferability (Kattenborn et al., 2020; Schiefer et al., 2020, 2023), our results show that geometric consistency is a prerequisite for reliable multi-sensor supervision.

The validation diagnostics indicate that the 2D-to-3D transfer behaved as a robust refinement of the Cartesian baseline under occlusion and irregular sampling, rather than a wholesale reclassification. These diagnostics were applied quantitatively to the medium-site validation tiles, where canopy density, multi-stratum vegetation, and local occlusion provided the strongest test of the refinement workflow; low-site tiles were used as visual checks because their sparse, low-stature vegetation made the transfer geometrically more straightforward. Z_{norm} was used as a structural constraint during fusion to keep low-stature and canopy classes in their expected vertical strata, but it is not propagated into the downstream learning model; the diagnostics therefore aim to confirm that the labels are stable and structurally coherent, not to estimate absolute classification accuracy. High voxel-level agreement between the Step 2 and Step 5 labels shows that the 3D refinement preserved

the dominant 2D mask-derived class structure while applying only targeted corrections, indicating that the effects of occlusion, mixed understory vegetation, and irregular point density (Brede et al., 2022; Gassilloud et al., 2025) were controlled rather than amplified. Hummock remained tightly constrained within the near-ground height range, while the larger outside-range fractions for shrub and mallee reflect the deliberate transition from foliage-oriented priors to whole-object 3D labelling, in which mallee labels include lower stems, trunks, and within-crown returns, consistent with structural-modelling approaches that treat canopy returns as part of an integrated tree object (Schneider et al., 2019; Brede et al., 2022). At the object level, the high mean retention of mallee *treeID*s shows that most watershed-derived crowns survived the polar refinement and consolidation stages, with lower-retention tiles representing marginal instances correctly reassigned by the crown-level voting rule.

Because the validation compares sequential outputs within the same workflow and Z_{norm} is not used by the downstream model, the diagnostics should be read as evidence of label robustness rather than independent accuracy. The workflow produces stable, height-plausible, and object-consistent 3D labels under occlusion conditions that are typical of semi-arid mallee canopies, providing supervision suitable for downstream sparse 3D learning. Independent field validation or manually labelled 3D reference data would still be required to quantify absolute accuracy across the full range of vegetation density, illumination, and canopy conditions.

6.2. Cross-ecosystem ablation studies

6.2.1. Zero-shot transfer of forest-trained models (Experiment 1)

The zero-shot experiments show that forest-trained 3D models have limited direct transferability to sparse semi-arid rangeland vegetation. ForestFormer3D retained strong performance on the held-out *FOR-instanceV2* forest test data, confirming that the pretrained model remained effective within its source domain. However, performance declined sharply when the same model was applied to the Calperum rangeland sites. SegmentAnyTree showed a similar pattern at the instance level, preserving reasonable crown delineation in the forest scene but failing to generalise reliably to rangeland growth forms.

This behaviour is consistent with the broader remote-sensing literature showing that deep learning models often specialise to the visual, structural, and sensor conditions represented in their training data (Jones et al., 2018;

Zhang et al., 2019; Kattenborn et al., 2020). It also reflects the forest-centred design of many current 3D vegetation benchmarks and algorithms, which are typically developed around individual tree delineation in vertically stratified forest stands (Wielgosz et al., 2023; Henrich et al., 2024; Xiang et al., 2024; Wielgosz et al., 2024; Xiang et al., 2025b). These assumptions do not transfer cleanly to mallee- and shrub-dominated rangelands, where vegetation is lower, more discontinuous, more horizontally heterogeneous, and often multi-stemmed. As a result, forest-trained models can retain useful geometric priors, but those priors are not sufficient for direct zero-shot rangeland mapping.

6.2.2. Few-shot fine-tuning for cross-domain adaptation (Experiment 2)

Few-shot fine-tuning partly reduced this domain shift. Compared with zero-shot transfer, both joint and curriculum fine-tuning improved semantic segmentation, and the strongest results were obtained under curriculum adaptation. In Stage 1, the model was initialised from forest-pretrained weights and adapted to the medium-density site, producing the strongest medium-site semantic and instance performance in Experiment 2. In Stage 2, the medium-adapted weights were then used to initialise fine-tuning on the low-density site, producing the strongest low-site semantic performance. This suggests that progressive adaptation from denser to sparser rangeland vegetation can help the model adjust to changing structural context.

The improvement after fine-tuning indicates that forest-pretrained models contain useful geometric representations, but that these representations require target-domain supervision before they can support rangeland growth-form mapping. This is consistent with recent point-cloud studies showing that fine-tuning can substantially improve segmentation models when adapted with higher-quality or more domain-relevant data (Henrich et al., 2024). In this study, target-domain supervision improved not only semantic growth-form classification but also object-level structure in the medium-density site. ForestFormer3D instance F1 increased from 4.4% under zero-shot transfer to 39.1% after joint fine-tuning and 49.8% after curriculum fine-tuning. Recall and coverage also improved, indicating that fine-tuning helped recover more coherent mallee crown structure.

However, the instance-level benefit did not generalise equally to the low-density site. Instance metrics remained zero because the low-density evaluation data contained very few mallee tree instances, with only two trees present across the low site and no coherent mallee instances available in some

evaluated test tiles. Therefore, the low-density instance results should be interpreted cautiously: they indicate that the model did not recover reliable mallee instances under sparse conditions, but they also reflect the limited availability of tree-like objects for instance-level evaluation. This highlights a broader limitation of tree-centric instance metrics in rangelands, where vegetation may be dominated by shrubs, hummocks, forbs, or sparse clumps rather than distinct tree crowns.

The class-wise results further show that adaptation was not equally effective for all growth forms. Mallee was generally recovered more successfully in the medium-density site, whereas shrubs remained difficult to identify across experiments. This likely reflects the intermediate and variable structure of shrubs, which can overlap with both low vegetation and taller woody growth forms. Lidar occlusion and uneven sampling of lower canopy strata may further reduce separability, particularly in clumped or multi-layered vegetation (Brede et al., 2022; Gassilloud et al., 2025). These results support the argument that 3D structure provides information beyond spectral cover alone (Jacobs et al., 2021), but also show that point-cloud geometry does not remove all ambiguity when growth forms are structurally transitional.

6.2.3. Training from scratch on rangeland data (Experiment 3)

The from-scratch experiments highlight the value of forest-pretrained initialisation under limited rangeland supervision. Models trained only on the available rangeland data performed worse than the fine-tuned models, indicating that the target dataset was too small to train a complex 3D architecture effectively from random initialisation. Joint training from scratch produced weak semantic performance in both sites and failed to recover usable instance segmentation. Site-specific scratch training improved semantic performance relative to joint scratch training, particularly for the dominant classes in each site, but it still remained below the fine-tuned transfer models.

These results show that neither zero-shot transfer nor random initialisation was sufficient on its own. Forest pretraining provided a useful starting point, but direct application to rangelands was unreliable. Conversely, rangeland-only training provided domain-specific supervision, but the available sample size was insufficient to train the full 3D model effectively from scratch. The best performance was obtained by combining both: using forest-pretrained weights as an initial representation and then adapting the model with labelled rangeland data. This balance is important for future 3D vegetation mapping in non-forest ecosystems, where large manually la-

belled point-cloud datasets remain rare and where vegetation structure differs substantially from the tree-centric assumptions embedded in many existing benchmarks.

6.3. Implications for rangeland monitoring and 3D vegetation research

Together, these findings support the main contributions of this study. The proposed 2D-to-3D label-transfer workflow provides a scalable pathway for generating volumetric supervision from existing high-resolution semantic masks, avoiding the need for fully manual point-cloud annotation. By combining dense co-registration, height-normalised LiDAR, crown-aware refinement, and class-specific structural rules, the workflow produced labels that were spatially consistent, height-plausible, and more object-consistent under the canopy density and occlusion conditions typical of semi-arid mallee ecosystems. The cross-ecosystem ablation studies further show that forest-pretrained 3D models contain useful structural priors, but these priors require target-domain adaptation before they can support reliable rangeland growth-form mapping.

For operational rangeland monitoring, the results indicate that scalable 3D growth-form mapping is feasible, but only when each stage of the pipeline is designed for the target ecosystem. Co-registration quality constrains the reliability of transferred labels; semantic label quality determines the ecological meaning of the 3D supervision; and the modelling results show that direct transfer from forest benchmarks is insufficient under strong structural domain shift. However, few-shot adaptation substantially improved cross-ecosystem performance, suggesting that modest amounts of well-aligned rangeland supervision can make pretrained 3D models more useful for semi-arid vegetation mapping. This has practical relevance for biodiversity monitoring, vegetation-structure assessment, and biomass or carbon accounting in landscapes where field surveys are difficult to scale and 2D spectral products often miss vertical organisation (Omasa et al., 2007; Jacobs et al., 2021; Gassilloud et al., 2025).

The results also highlight the need to move beyond tree-centric assumptions in 3D vegetation research. Current benchmarks and models are largely designed around forest structure, where individual trees form relatively distinct vertical objects. In contrast, semi-arid rangelands contain shrubs, hummocks, forbs, woody debris, and multi-stemmed mallee forms that are structurally discontinuous and often horizontally mixed. The weak zero-shot results and the persistent difficulty of shrub discrimination show that rangeland

vegetation cannot be treated simply as a lower-density version of forest. Instead, future 3D models and evaluation protocols should represent growth forms, clump structure, canopy openness, and vertical layering in ways that are ecologically meaningful for non-forest ecosystems.

Several limitations remain. First, the label-transfer validation reported here is an internal consistency assessment rather than an independent accuracy estimate. It was evaluated on both medium- and low-density condition tiles, allowing the workflow to be tested under contrasting rangeland structures, including denser mallee canopies, sparse vegetation, occlusion, and multi-stratum mixing. However, broader validation across additional sites, vegetation densities, seasons, and acquisition conditions is still needed to quantify generalisation. Second, the transferred 3D labels inherit errors from the 2D semantic model, meaning that uncertainty-aware label transfer and confidence-weighted supervision should be explored in future work. Third, UAS LiDAR point clouds remain affected by occlusion and variable point density, particularly in the lower vegetation layers most relevant to rangeland growth-form discrimination (Brede et al., 2022; Gassilloud et al., 2025). Finally, the evaluation still partly depends on forest-derived semantic and instance assumptions. Developing rangeland-specific 3D benchmarks will be essential for testing whether models genuinely learn growth-form structure rather than adapting forest concepts to a structurally different ecosystem.

7. Conclusions

This study shows that the main challenge to applying 3D deep-learning models in semi-arid rangelands is not only model capacity, but the mismatch between forest-centred structural assumptions and the fragmented, low-stature, and horizontally heterogeneous structure of dryland vegetation. We addressed this by combining an automated 2D-to-3D label-transfer workflow with a controlled evaluation of zero-shot transfer, few-shot adaptation, and training from scratch. The proposed workflow generated height-plausible and object-consistent volumetric supervision from existing semantic masks, avoiding the need for fully manual point-cloud annotation.

The cross-ecosystem experiments showed that forest pretraining provides useful structural priors, but direct transfer is insufficient. Under zero-shot transfer, ForestFormer3D declined from 85.2% mIoU on the held-out forest test data to 41.1% and 16.1% mIoU on the medium- and low-density rangeland sites, with instance F1 decreasing from 89.4% to 4.4% and 8.9%.

Few-shot fine-tuning substantially improved semantic performance. Joint fine-tuning reached 48.9% and 50.5% mIoU on the medium- and low-density sites, while curriculum fine-tuning produced the strongest results, reaching 49.6% and 63.1% mIoU, respectively. In contrast, training from scratch on rangeland data produced weaker and class-biased predictions, indicating that the available rangeland supervision was insufficient to replace pretrained structural representations.

Overall, effective 3D rangeland mapping requires both scalable target-domain supervision and deliberate adaptation of pretrained models. The 2D-to-3D label-transfer framework makes this possible by converting existing high-resolution semantic labels into usable volumetric training data. More broadly, the results highlight the need for 3D vegetation models and benchmarks that move beyond tree-centric assumptions and better represent sparsity, structural discontinuity, and growth-form variability across non-forest ecosystems.

8. Acknowledgements

We would like to acknowledge the Terrestrial Ecosystem Research Network (TERN) for support, in particular the TERN Ecosystem Surveillance team and TERN Landscapes program. We would like to thank Ben Sparrow, Poornima Sivanandam and Juan Carlos Montes-Herrera for drone operation and data processing. We thank the managers of Calperum Station (Australian Landscape Trust), South Australia, for providing access to the study sites. Laura Sotomayor is supported by a Research Training Program (RTP) scholarship through the University of Tasmania. The Australian Research Council (ARC) supported some of the initial fieldwork and infrastructure under Discovery grant DP180103460. The authors gratefully acknowledge NVIDIA for selecting the project Mapping Ecosystem Resilience with High-Resolution Remote Sensing Data for inclusion in the NVIDIA Academic Grant Program. This six-month award provided cloud access through Brev to one reserved instance of 8× NVIDIA A100 80GB GPUs from October 1, 2025 to March 31, 2026. The instance comprised 80 GiB VRAM per GPU, 2 TiB RAM, 128 CPUs, 60 TiB of storage, and 560 GiB NVMe. Multiple team members were permitted to use the resource, although only one instance was allocated to the project.

9. Code and Data Availability

The 2D-to-3D label transfer pipeline with cross-ecosystem 3D vegetation segmentation experiments, (see Section 3 called *fvcStructureRangeland3D*, is publicly available at `https://github.com/LNSOTOM/fvc_structure`. The version used in this paper is 2026.28.4. The code to reproduce the transfer labelling and results will be made public after acceptance of the manuscript. The entire dataset of manual labels and UAS imagery that were used in this study are available as tiled dataset and shapefiles on `https://zenodo.org/records/15036860`.

10. Declaration of generative AI and AI-assisted technologies in the manuscript preparation process.

During the preparation of this work the author(s) used ChatGPT for language editing. During analysis, the authors employed GitHub Copilot for programming assistance. After using these tools, the author(s) reviewed and edited the content and code as needed and take(s) full responsibility for the content of the published article and code.

References

Almeida, D.R.A.d., Broadbent, E.N., Ferreira, M.P., Meli, P., Zambrano, A.M.A., Gorgens, E.B., Resende, A.F., de Almeida, C.T., do Amaral, C.H., Corte, A.P.D., Silva, C.A., Romanelli, J.P., Prata, G.A., de Almeida Papa, D., Stark, S.C., Valbuena, R., Nelson, B.W., Guillemot, J., Féret, J.B., Chazdon, R., Brancalion, P.H.S., 2021. Monitoring restored tropical forest diversity and structure through UAV-borne hyperspectral and lidar fusion. Remote Sensing of Environment 264, 112582. URL: `https://www.sciencedirect.com/science/article/pii/S0034425721003023`, doi:`10.1016/j.rse.2021.112582`.

Bafghi, R.A., Harilal, N., Monteleoni, C., Raissi, M., 2024. Parameter Efficient Fine-tuning of Self-supervised ViTs without Catastrophic Forgetting. URL: `http://arxiv.org/abs/2404.17245`, doi:`10.48550/arXiv.2404.17245`. arXiv:2404.17245 [cs].

Brede, B., Bartholomeus, H.M., Barbier, N., Pimont, F., Vincent, G., Herold, M., 2022. Peering through the thicket: Effects of UAV LiDAR scanner settings and flight planning on canopy volume discovery. International Journal of Applied Earth Observation and Geoinformation 114, 103056. URL: https://www.sciencedirect.com/science/article/pii/S1569843222002448, doi:10.1016/j.jag.2022.103056.

Béland, M., Baldocchi, D.D., Widlowski, J.L., Fournier, R.A., Verstraete, M.M., 2014. On seeing the wood from the leaves and the role of voxel size in determining leaf area distribution of forests with terrestrial LiDAR. Agricultural and Forest Meteorology 184, 82–97. URL: https://www.sciencedirect.com/science/article/pii/S0168192313002608, doi:10.1016/j.agrformet.2013.09.005.

Calders, K., Adams, J., Armston, J., Bartholomeus, H., Bauwens, S., Bentley, L.P., Chave, J., Danson, F.M., Demol, M., Disney, M., Gaulton, R., Krishna Moorthy, S.M., Levick, S.R., Saarinen, N., Schaaf, C., Stovall, A., Terryn, L., Wilkes, P., Verbeeck, H., 2020. Terrestrial laser scanning in forest ecology: Expanding the horizon. Remote Sensing of Environment 251, 112102. URL: https://www.sciencedirect.com/science/article/pii/S0034425720304752, doi:10.1016/j.rse.2020.112102.

Carós, M., Just, A., Seguí, S., Vitrià, J., 2026. Enhancing point cloud semantic segmentation via scalable domain adaptation with LoRA-enabled PointNet++. ISPRS Open Journal of Photogrammetry and Remote Sensing 19, 100119. URL: https://www.sciencedirect.com/science/article/pii/S2667393226000050, doi:10.1016/j.ophoto.2026.100119.

Cherlet, W., Dayal, K., Chen, S., Cooper, Z., Disney, M., Hanzl, A., Levick, S., Nightingale, J., Origo, N., Senf, C., Soenens, L., Terryn, L., Van den Broeck, W.A.J., Calders, K., 2026. Benchmarking tree instance segmentation of terrestrial laser scanning point clouds. ISPRS Journal of Photogrammetry and Remote Sensing 231, 230–247. URL: https://www.sciencedirect.com/science/article/pii/S092427162500423X, doi:10.1016/j.isprsjprs.2025.10.033.

Choy, C., Gwak, J., Savarese, S., 2019. 4D Spatio-Temporal ConvNets: Minkowski Convolutional Neural Networks. URL: http://arxiv.org/abs/1904.08755, doi:10.48550/arXiv.1904.08755. arXiv:1904.08755 [cs].

CloudCompare, . Releases · CloudCompare/CloudCompare. URL: https://github.com/CloudCompare/CloudCompare/releases.

Contributors, P., 2025. PDAL point data abstraction library. URL: https://doi.org/10.5281/zenodo.10884408, doi:10.5281/zenodo.10884408.

Dandois, J.P., Ellis, E.C., 2013. High spatial resolution three-dimensional mapping of vegetation spectral dynamics using computer vision. Remote Sensing of Environment 136, 259–276. URL: https://www.sciencedirect.com/science/article/pii/S0034425713001326, doi:10.1016/j.rse.2013.04.005.

DCCEEW, 2017. Australian Vegetation Attribute Manual Version 7.0 - DCCEEW. URL: https://www.dcceew.gov.au/environment/land/publications/australian-vegetation-attribute-manual-version-7.

D'hont, B., Calders, K., Antonelli, A., Berg, T., Cherlet, W., Dayal, K., Fitzpatrick, O.J., Hambrecht, L., Leponce, M., Lucieer, A., Pascal, O., Raumonen, P., Verbeeck, H., 2025. Integrating terrestrial and canopy laser scanning for comprehensive analysis of large old trees: Implications for single tree and biodiversity research. Remote Sensing in Ecology and Conservation n/a. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/rse2.70021, doi:10.1002/rse2.70021. _eprint: https://zslpublications.onlinelibrary.wiley.com/doi/pdf/10.1002/rse2.70021.

Disney, M., 2019. Terrestrial LiDAR: a three-dimensional revolution in how we look at trees. New Phytologist 222, 1736–1741. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/nph.15517, doi:10.1111/nph.15517. _eprint: https://nph.onlinelibrary.wiley.com/doi/pdf/10.1111/nph.15517.

Ecke, S., Dempewolf, J., Frey, J., Schwaller, A., Endres, E., Klemmt, H.J., Tiede, D., Seifert, T., 2022. UAV-Based Forest Health Monitoring: A Systematic Review. Remote Sensing 14. URL: https://www.mdpi.com/2072-4292/14/13/3205. publisher: Multidisciplinary Digital Publishing Institute.

Ehbrecht, M., Seidel, D., Annighöfer, P., Kreft, H., Köhler, M., Zemp, D.C., Puettmann, K., Nilus, R., Babweteera, F., Willim, K., Stiers,

M., Soto, D., Boehmer, H.J., Fisichelli, N., Burnett, M., Juday, G., Stephens, S.L., Ammer, C., 2021. Global patterns and climatic controls of forest structural complexity. Nature Communications 12, 519. URL: https://www.nature.com/articles/s41467-020-20767-z, doi:10.1038/s41467-020-20767-z. publisher: Nature Publishing Group.

Ganz, S., Käber, Y., Adler, P., 2019. Measuring Tree Height with Remote Sensing—A Comparison of Photogrammetric and LiDAR Data with Different Field Measurements. Forests 10, 694. URL: https://www.mdpi.com/1999-4907/10/8/694, doi:10.3390/f10080694. publisher: Multidisciplinary Digital Publishing Institute.

Gassilloud, M., Koch, B., Göritz, A., 2025. Occlusion mapping reveals the impact of flight and sensing parameters on vertical forest structure exploration with cost-effective UAV based laser scanning. International Journal of Applied Earth Observation and Geoinformation 139, 104493. URL: https://www.sciencedirect.com/science/article/pii/S1569843225001402, doi:10.1016/j.jag.2025.104493.

Harrison, B.A., Gibson, R., Bastin, G., Thackway, R., Huete, A., Donald, G., Lyons, M., Sparks, T., Byrne, G., Lewis, M.M., Xie, Q., 2021. Earth Observation: Data, Processing and Applications. Volume 3A: Applications—Terrestrial Vegetation. Australia and New Zealand CRC for Spatial Information. URL: https://opus.lib.uts.edu.au/handle/10453/155085. accepted: 2022-03-09T00:53:51Z.

Harrison, G.R., Shrestha, A., Strand, E.K., Karl, J.W., 2024. A comparison and development of methods for estimating shrub volume using drone-imagery-derived point clouds. Ecosphere 15, e4877. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/ecs2.4877, doi:10.1002/ecs2.4877. _eprint: https://esajournals.onlinelibrary.wiley.com/doi/pdf/10.1002/ecs2.4877.

Haynes, R.S., Lucieer, A., Turner, D., Cimoli, E., 2025. Co-Registration of Multi-Modal UAS Pushbroom Imaging Spectroscopy and RGB Imagery Using Optical Flow. Drones 9, 132. URL: https://www.mdpi.com/2504-446X/9/2/132, doi:10.3390/drones9020132. publisher: Multidisciplinary Digital Publishing Institute.

Henrich, J., Delden, J.v., Seidel, D., Kneib, T., Ecker, A., 2024. TreeLearn: A deep learning method for segmenting individual trees from ground-based LiDAR forest point clouds. Ecological Informatics 84, 102888. URL: http://arxiv.org/abs/2309.08471, doi:10.1016/j.ecoinf.2024.102888. arXiv:2309.08471 [cs].

Jacobs, M., Rais, A., Pretzsch, H., 2021. How drought stress becomes visible upon detecting tree shape using terrestrial laser scanning (TLS). Forest Ecology and Management 489, 118975. URL: https://www.sciencedirect.com/science/article/pii/S0378112721000645, doi:10.1016/j.foreco.2021.118975.

Jiang, L., Zhao, H., Shi, S., Liu, S., Fu, C.W., Jia, J., 2020. PointGroup: Dual-Set Point Grouping for 3D Instance Segmentation. URL: http://arxiv.org/abs/2004.01658, doi:10.48550/arXiv.2004.01658. arXiv:2004.01658 [cs].

Jiang, N., Zhu, W., Mou, M., Wang, L., Zhang, J., 2012. A phenology-preserving filtering method to reduce noise in NDVI time series, in: 2012 IEEE International Geoscience and Remote Sensing Symposium, pp. 2384–2387. URL: https://ieeexplore.ieee.org/document/6351013, doi:10.1109/IGARSS.2012.6351013. iSSN: 2153-7003.

Jones, M.O., Allred, B.W., Naugle, D.E., Maestas, J.D., Donnelly, P., Metz, L.J., Karl, J., Smith, R., Bestelmeyer, B., Boyd, C., Kerby, J.D., McIver, J.D., 2018. Innovation in rangeland monitoring: annual, 30 m, plant functional type percent cover maps for U.S. rangelands, 1984–2017. Ecosphere 9, e02430. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/ecs2.2430, doi:10.1002/ecs2.2430. _eprint: https://esajournals.onlinelibrary.wiley.com/doi/pdf/10.1002/ecs2.2430.

Kattenborn, T., Eichel, J., Wiser, S., Burrows, L., Fassnacht, F.E., Schmidtlein, S., 2020. Convolutional Neural Networks accurately predict cover fractions of plant species and communities in Unmanned Aerial Vehicle imagery. Remote Sensing in Ecology and Conservation 6, 472–486. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/rse2.146, doi:10.1002/rse2.146. _eprint: https://zslpublications.onlinelibrary.wiley.com/doi/pdf/10.1002/rse2.146.

Kellner, J.R., Albert, L.P., Burley, J.T., Cushman, K.C., 2019. The case for remote sensing of individual plants. American Journal of Botany 106, 1139–1142. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/ajb2.1347, doi:10.1002/ajb2.1347. _eprint: https://bsapubs.onlinelibrary.wiley.com/doi/pdf/10.1002/ajb2.1347.

Kolodiazhnyi, M., Vorontsova, A., Konushin, A., Rukhovich, D., 2023. OneFormer3D: One Transformer for Unified Point Cloud Segmentation. URL: http://arxiv.org/abs/2311.14405, doi:10.48550/arXiv.2311.14405. arXiv:2311.14405 [cs].

Krisanski, S., Taskhiri, M.S., Aracil, S.G., Herries, D., Muneri, A., Gurung, M.B., Montgomery, J., Turner, P., Krisanski, S., Taskhiri, M.S., Aracil, S.G., Herries, D., Muneri, A., Gurung, M.B., Montgomery, J., Turner, P., 2021. Forest Structural Complexity Tool—An Open Source, Fully-Automated Tool for Measuring Forest Point Clouds. Remote Sensing 13. URL: https://www.mdpi.com/2072-4292/13/22/4677. publisher: Multidisciplinary Digital Publishing Institute.

Le Besnerais, G., Champagnat, F., 2005. Dense optical flow by iterative local window registration, in: IEEE International Conference on Image Processing 2005, pp. I–137. URL: https://ieeexplore.ieee.org/abstract/document/1529706, doi:10.1109/ICIP.2005.1529706. iSSN: 2381-8549.

Lucieer, A., Turner, D., King, D.H., Robinson, S.A., 2014. Using an Unmanned Aerial Vehicle (UAV) to capture micro-topography of Antarctic moss beds. International Journal of Applied Earth Observation and Geoinformation 27, 53–62. URL: https://www.sciencedirect.com/science/article/pii/S0303243413000603, doi:10.1016/j.jag.2013.05.011.

Melville, B., Fisher, A., Lucieer, A., 2019. Ultra-high spatial resolution fractional vegetation cover from unmanned aerial multispectral imagery. International Journal of Applied Earth Observation and Geoinformation 78, 14–24. URL: https://www.sciencedirect.com/science/article/pii/S0303243418311693, doi:10.1016/j.jag.2019.01.013.

Montes-Herrera, J.C., Robbins, A., Franson, C., Bektas, T., Sparrow, B., Lucieer, A., 2025. DroneScape: Collection of Penta-Return LiDAR for Ecosystem Plot Monitoring URL: https://www.protocols.io/print/dronescape-collection-of-penta-return-lidar-for-ec-g3c3byiyp.

Mosig, C., Kattenborn, T., Loaiza, D.M., Vanja-Jehle, J., Brandt, J., Jacobs, N., Khanal, S., Xing, E., Schwartz, M., Muller-Landau, H.C., Beloiu, M., Bozzini, A., Cheng, Y., Ganz, K., Grüning, B., Hartmann, H., Hempel, J., Horion, S., Junttila, S., Korznikov, K., Kraemer, G., Mönks, M., Nardi, D., Neumeier, P., Schmid, J., Soltani, S., Therese-Schmehl, M., Veitch-Michaelis, J., Mahecha, M., 2026a. Sub-pixel mapping of disturbance and tree mortality dynamics from Sentinel-2 time series around the globe URL: `https://eartharxiv.org/repository/view/11912/`. publisher: EarthArXiv.

Mosig, C., Vajna-Jehle, J., Mahecha, M.D., Cheng, Y., Hartmann, H., Montero, D., Junttila, S., Horion, S., Schwenke, M.B., Koontz, M.J., Maulud, K.N.A., Adu-Bredu, S., Al-Halbouni, D., Ali, M., Allen, M., Altman, J., Amorós, L., Angiolini, C., Astrup, R., Awada, H., Barrasso, C., Bartholomeus, H., Beck, P.S.A., Bozzini, A., Braun-Wimmer, J., Brede, B., Breunig, F.M., Brugnaro, S., Buras, A., Burchard-Levine, V., Camarero, J.J., Candotti, A., Capuder, L., Carrieri, E., Centritto, M., Chirici, G., Cloutier, M., Conciani, D., Cushman, K., Dalling, J.W., Dao, P.D., Dempewolf, J., Denter, M., Dogotari, M., Díaz-Delgado, R., Ecke, S., Eichel, J., Eltner, A., Fabbri, A., Fabi, M., Fassnacht, F., Ferreira, M.P., Fischer, F.J., Frey, J., Frick, A., Fuentes, J., Ganz, S., Garbarino, M., García, M., Gassilloud, M., Gazol, A., Gea-Izquierdo, G., Gerberding, K., Ghasemi, M., Giannetti, F., Gillan, J., Gonzalez, R., Gosper, C., Greene, T., Greinwald, K., Grieve, S., Große-Stoltenberg, A., Gutierrez, J.A., Göritz, A., Hajek, P., Hedding, D., Hempel, J., Heremans, S., Hernández, M., Heurich, M., Honkavaara, E., Höfle, B., Jackisch, R., Jucker, T., Kalwij, J.M., Kepfer-Rojas, S., Khatri-Chhetri, P., Kleinebecker, T., Klemmt, H.J., Klouček, T., Koivumäki, N., Kolagani, N., Komárek, J., Korznikov, K., Kraszewski, B., Kruse, S., Krüger, R., Kuechly, H., Kwong, I.H.Y., Laliberté, E., Langan, L., Latifi, H., Leal-Medina, C., Lehmann, J.R.K., Li, L., Lines, E., Lisiewicz, M., Lopatin, J., Lucieer, A., Ludwig, A., Ludwig, M., Lyytikäinen-Saarenmaa, P., Ma, Q., Mansuy, N., Peña, J.M., Marino, G., Maroschek, M., Martín, M.P., Martín-Benito, D., Matham, P., Mazzoni, S., Meloni, F., Menzel, A., Meyer, H., Miraki, M., Moreno, G., Moreno-Fernández, D., Muller-Landau, H.C., Mälicke, M., Möhring, J., Müllerova, J., Naidu, S.S., Nardi, D., Neumeier, P., Nita, M.D., Näsi, R., Oppgenoorth, L., Orunbaev, S., Palmer, M., Paul, T., Pfenning, M., Potts, A., Prasanna, G.L., Prober,

S., Puliti, S., Pérez-Luque, A.J., Pérez-Priego, O., Reudenbach, C., Revuelto, J., Rivas-Torres, G., Roberge, P., Roggero, P.P., Rossi, C., Ruehr, N.K., Ruiz-Benito, P., Runge, C.M., Satta, G.G.A., Scanu, B., Scherer-Lorenzen, M., Schiefer, F., Schiller, C., Schladebach, J., Schmehl, M.T., Schmid, J., Schmidt, T.A., Schwarz, S., Seidl, R., Seifert, T., Barba, A.S., Shafeian, E., Shapiro, A., de Simone, L., Sohrabi, H., Soltani, S., Sotomayor, L., Sparrow, B., Steer, B.S.C., Stenson, M., Stöckigt, B., Su, Y., Suomalainen, J., Tamudo, E., Barbieri, M.J.T., Tomelleri, E., Torresani, M., Trepekli, K., Ullah, S., Ullah, S., Umlauft, J., Vargas-Ramírez, N., Vatandaslar, C., Visacki, V., Volpi, M., Vásquez, V., Wallis, C., Weinstein, B., Weiser, H., Wich, S., Ximena, T.C., Zarco-Tejada, P.J., Zdunic, K., Zielewska-Büttner, K., de Oliveira, R.A., van Wagtendonk, L., von Dosky, V., Kattenborn, T., 2026b. deadtrees.earth — An open-access and interactive database for centimeter-scale aerial imagery to uncover global tree mortality dynamics. Remote Sensing of Environment 332, 115027. URL: https://www.sciencedirect.com/science/article/pii/S0034425725004316, doi:10.1016/j.rse.2025.115027.

Omasa, K., Hosoi, F., Konishi, A., 2007. 3D lidar imaging for detecting and understanding plant responses and canopy structure. Journal of Experimental Botany 58, 881–898. URL: https://doi.org/10.1093/jxb/erl142, doi:10.1093/jxb/erl142.

Pellis, E., Murtiyoso, A., Masiero, A., Tucci, G., Betti, M., Grussenmeyer, P., 2022. 2D to 3D Label Propagation for the Semantic Segmentation of Heritage Building Point Clouds. The International Archives of the Photogrammetry Remote Sensing and Spatial Information Sciences XLIII-B2-2022, 861–867. doi:10.5194/isprs-archives-XLIII-B2-2022-861-2022.

Pingel, T.J., Clarke, K.C., McBride, W.A., 2013. An improved simple morphological filter for the terrain classification of airborne LIDAR data. ISPRS Journal of Photogrammetry and Remote Sensing 77, 21–30. URL: https://www.sciencedirect.com/science/article/pii/S0924271613000026, doi:10.1016/j.isprsjprs.2012.12.002.

Plyer, A., Le Besnerais, G., Champagnat, F., 2014. Massively parallel Lucas Kanade optical flow for real-time video processing applications. Journal of Real-Time Image Processing 11, 713–730. URL: https:

//link.springer.com/article/10.1007/s11554-014-0423-0, doi:10.1007/s11554-014-0423-0. publisher: Springer.

Qi, C.R., Yi, L., Su, H., Guibas, L.J., 2017. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space, in: Advances in Neural Information Processing Systems, Curran Associates, Inc. URL: https://proceedings.neurips.cc/paper/2017/hash/d8bf84be3800d12f74d8b05e9b89836f-Abstract.html.

Reid, R.S., Fernández-Giménez, M.E., Galvin, K.A., 2014. Dynamics and Resilience of Rangelands and Pastoral Peoples Around the Globe. Annual Review of Environment and Resources 39, 217–242. URL: https://www.annualreviews.org/content/journals/10.1146/annurev-environ-020713-163329, doi:10.1146/annurev-environ-020713-163329. publisher: Annual Reviews.

Retallack, A., Finlayson, G., Ostendorf, B., Lewis, M., 2022. Using deep learning to detect an indicator arid shrub in ultra-high-resolution UAV imagery. Ecological Indicators 145, 109698. URL: https://www.sciencedirect.com/science/article/pii/S1470160X22011712, doi:10.1016/j.ecolind.2022.109698.

Ruscalleda-Alvarez, J., Page, G.F.M., Zdunic, K., Prober, S.M., 2025. Estimating woody vegetation cover in arid and semi-arid rangelands. Ecological Indicators 177, 113741. URL: https://www.sciencedirect.com/science/article/pii/S1470160X25006715, doi:10.1016/j.ecolind.2025.113741.

Schiefer, F., Kattenborn, T., Frick, A., Frey, J., Schall, P., Koch, B., Schmidtlein, S., 2020. Mapping forest tree species in high resolution UAV-based RGB-imagery by means of convolutional neural networks. ISPRS Journal of Photogrammetry and Remote Sensing 170, 205–215. URL: https://www.sciencedirect.com/science/article/pii/S0924271620302938, doi:10.1016/j.isprsjprs.2020.10.015.

Schiefer, F., Schmidtlein, S., Frick, A., Frey, J., Klinke, R., Zielewska-Büttner, K., Junttila, S., Uhl, A., Kattenborn, T.,

2023. UAV-based reference data for the prediction of fractional cover of standing deadwood from Sentinel time series. ISPRS Open Journal of Photogrammetry and Remote Sensing 8, 100034. URL: https://www.sciencedirect.com/science/article/pii/S2667393223000054, doi:10.1016/j.ophoto.2023.100034.

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. URL: https://www.sciencedirect.com/science/article/pii/S0168192319300267, doi:10.1016/j.agrformet.2019.01.033.

Song, Y., Wang, T., Mondal, S.K., Sahoo, J.P., 2022. A Comprehensive Survey of Few-shot Learning: Evolution, Applications, Challenges, and Opportunities. URL: http://arxiv.org/abs/2205.06743, doi:10.48550/arXiv.2205.06743. arXiv:2205.06743 [cs].

Sothe, C., De Almeida, C.M., Schimalski, M.B., La Rosa, L.E.C., Castro, J.D.B., Feitosa, R.Q., Dalponte, M., Lima, C.L., Liesenberg, V., Miyoshi, G.T., Tommaselli, A.M.G., 2020. Comparative performance of convolutional neural network, weighted and conventional support vector machine and random forest for classifying tree species using hyperspectral and photogrammetric data. GIScience & Remote Sensing 57, 369–394. URL: https://doi.org/10.1080/15481603.2020.1712102, doi:10.1080/15481603.2020.1712102. publisher: Taylor & Francis _eprint: https://doi.org/10.1080/15481603.2020.1712102.

Sotomayor, L.N., Lucieer, A., Turner, D., Lewis, M., Kattenborn, T., 2025a. Mapping fractional vegetation cover in UAS RGB and multispectral imagery in semi-arid Australian ecosystems using CNN-based semantic segmentation. Landscape Ecology 40, 169. URL: https://doi.org/10.1007/s10980-025-02193-y, doi:10.1007/s10980-025-02193-y.

Sotomayor, L.N., Megan, L., Krishna, L., Hoyer, S., Molly, M., Arko, L., 2025b. Fractional Vegetation Cover Mapping - UAS RGB and Multispectral Imagery, CNN algorithms, Semi-Arid Australian Ecosystems Coverage URL: https://zenodo.org/records/15036860.

Sotomayor, L.N., Megan, L., Krishna, L., Sophia, H., Molly, M., Arko, L., 2025c. Fractional vegetation cover mapping - UAS RGB and multispectral imagery, CNN algorithms, semi-arid australian ecosystems coverage. URL: https://doi.org/10.5281/zenodo.15036860, doi:10.5281/zenodo.15036860.

Straker, A., Puliti, S., Breidenbach, J., Kleinn, C., Pearse, G., Astrup, R., Magdon, P., 2023. Instance segmentation of individual tree crowns with YOLOv5: A comparison of approaches using the ForInstance benchmark LiDAR dataset. ISPRS Open Journal of Photogrammetry and Remote Sensing 9, 100045. URL: https://www.sciencedirect.com/science/article/pii/S2667393223000169, doi:10.1016/j.ophoto.2023.100045.

Sutton, A., Fisher, A., Metternicht, G., 2022. Assessing the Accuracy of Landsat Vegetation Fractional Cover for Monitoring Australian Drylands. Remote Sensing 14, 6322. URL: https://www.mdpi.com/2072-4292/14/24/6322, doi:10.3390/rs14246322. publisher: Multidisciplinary Digital Publishing Institute.

TERN, 2021. Summary of TERN plots on Calperum, South Australia.

TERN, L., 2022. Field Survey Apps & Protocols | TERN Land Ecosystem Research Network. URL: https://www.tern.org.au/field-survey-protocols/.

TERN, L., Monitoring, T.S., Stenson, M., Sparrow, B., Lucieer, A., 2022. Drone RGB and Multispectral Imagery from TERN plots across Australia. URL: https://portal.tern.org.au/metadata/TERN/39de90f5-49e3-4567-917c-cf3e3bc93086.

Wada, K., mpitid, Buijs, M., N, Z.C., , Kubovčík, B.M., Myczko, A., latentix, Zhu, L., Yamaguchi, N., Fujii, S., iamgd67, IlyaOvodov, Patel, A., Clauss, C., Kuroiwa, E., Iyengar, R., Shilin, S., Malygina, T., Kawaharazuka, K., Engelberts, J., J, A., AlexMa, Song, C., Charlie, Rose, D., Livingstone, D., Doug, Erik, Toft, H., 2021. wkentaro/labelme: v4.6.0. URL: https://zenodo.org/records/5711226, doi:10.5281/zenodo.5711226.

Waite, O.J.M., Coops, N.C., Grubinger, S., Isaac-Renton, M., Degner, J., King, J., Liu, A., 2025. Responses of spectral indices to heat and

drought differ by tree size in Douglas-fir. Science of Remote Sensing 11, 100193. URL: https://www.sciencedirect.com/science/article/pii/S2666017224000774, doi:10.1016/j.srs.2024.100193.

Wallace, L., Bellman, C., Hally, B., Hernandez, J., Jones, S., Hillman, S., 2019. Assessing the Ability of Image Based Point Clouds Captured from a UAV to Measure the Terrain in the Presence of Canopy Cover. Forests 10, 284. URL: https://www.mdpi.com/1999-4907/10/3/284, doi:10.3390/f10030284. number: 3 Publisher: Multidisciplinary Digital Publishing Institute.

Wang, X., Liu, S., Shen, X., Shen, C., Jia, J., 2019. Associatively Segmenting Instances and Semantics in Point Clouds. URL: http://arxiv.org/abs/1902.09852, doi:10.48550/arXiv.1902.09852. arXiv:1902.09852 [cs].

White, A., Sparrow, B., Leitch, E., Foulkes, J., Flitton, R., Lowe, A.J., Caddy-Retalic, S., 2012. AusPlots Rangelands Survey Protocols Manual. University of Adelaide Press. URL: https://digital.library.adelaide.edu.au/dspace/handle/2440/101955. accepted: 2016-10-24T01:29:28Z Publication Title: http://www.tern.org.au/AusPlots-Rangelands-Survey-Protocols-Manual-pg23944.html.

White, M.A., De BEURS, K.M., Didan, K., Inouye, D.W., Richardson, A.D., Jensen, O.P., O'keefe, J., Zhang, G., Nemani, R.R., Van LEEUWEN, W.J.D., Brown, J.F., De WIT, A., Schaepman, M., Lin, X., Dettinger, M., Bailey, A.S., Kimball, J., Schwartz, M.D., Baldocchi, D.D., Lee, J.T., Lauenroth, W.K., 2009. Intercomparison, interpretation, and assessment of spring phenology in North America estimated from remote sensing for 1982–2006. Global Change Biology 15, 2335–2359. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1365-2486.2009.01910.x, doi:10.1111/j.1365-2486.2009.01910.x. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1365-2486.2009.01910.x.

Wielgosz, M., Puliti, S., Wilkes, P., Astrup, R., Wielgosz, M., Puliti, S., Wilkes, P., Astrup, R., 2023. Point2Tree(P2T)—Framework for Parameter Tuning of Semantic and Instance Segmentation Used with Mobile Laser Scanning Data in Coniferous Forest. Remote Sensing 15. URL: https://www.mdpi.com/2072-4292/15/15/3737. publisher: Multidisciplinary Digital Publishing Institute.

Wielgosz, M., Puliti, S., Xiang, B., Schindler, K., Astrup, R., 2024. SegmentAnyTree: A sensor and platform agnostic deep learning model for tree segmentation using laser scanning data. Remote Sensing of Environment 313, 114367. URL: https://www.sciencedirect.com/science/article/pii/S0034425724003936, doi:10.1016/j.rse.2024.114367.

Xiang, B., Peters, T., Kontogianni, T., Vetterli, F., Puliti, S., Astrup, R., Schindler, K., 2023. Towards accurate instance segmentation in large-scale LiDAR point clouds. URL: http://arxiv.org/abs/2307.02877, doi:10.48550/arXiv.2307.02877. arXiv:2307.02877 [cs].

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 Sensing of Environment 305, 114078. URL: https://www.sciencedirect.com/science/article/pii/S0034425724000890, doi:10.1016/j.rse.2024.114078.

Xiang, B., Wielgosz, M., PULITI, S., Král, K., Krucek, M., Missarov, A., Astrup, R., 2025a. FOR-instanceV2 dataset and pre-trained model in paper entitled "ForestFormer3D: A Unified Framework for End-to-End Segmentation of Forest LiDAR 3D Point Clouds". URL: https://zenodo.org/records/16742708, doi:10.5281/zenodo.16742708.

Xiang, B., Wielgosz, M., Puliti, S., Král, K., Krůček, M., Missarov, A., Astrup, R., 2025b. ForestFormer3D: A Unified Framework for End-to-End Segmentation of Forest LiDAR 3D Point Clouds. URL: http://arxiv.org/abs/2506.16991, doi:10.48550/arXiv.2506.16991. arXiv:2506.16991 [cs].

Xiong, Y., Varadarajan, B., Wu, L., Xiang, X., Xiao, F., Zhu, C., Dai, X., Wang, D., Sun, F., Iandola, F., Krishnamoorthi, R., Chandra, V., 2023. EfficientSAM: Leveraged Masked Image Pretraining for Efficient Segment Anything. URL: http://arxiv.org/abs/2312.00863, doi:10.48550/arXiv.2312.00863. arXiv:2312.00863 [cs].

Xiong, Y., Zhou, C., Xiang, X., Wu, L., Zhu, C., Liu, Z., Suri, S., Varadarajan, B., Akula, R., Iandola, F., Krishnamoorthi, R., Soran, B., Chandra, V., 2024. Efficient Track Anything. URL: http://arxiv.org/abs/2411.18933, doi:10.48550/arXiv.2411.18933. arXiv:2411.18933 [cs].

Zhang, B., Fischer, F.J., Prober, S.M., Yeoh, P.B., Gosper, C.R., Zdunic, K., Jucker, T., 2024. Robust retrieval of forest canopy structural attributes using multi-platform airborne LiDAR. Remote Sensing in Ecology and Conservation 10, 725–742. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/rse2.398, doi:10.1002/rse2.398. _eprint: https://zslpublications.onlinelibrary.wiley.com/doi/pdf/10.1002/rse2.398.

Zhang, J., Okin, G.S., Zhou, B., 2019. Assimilating optical satellite remote sensing images and field data to predict surface indicators in the Western U.S.: Assessing error in satellite predictions based on large geographical datasets with the use of machine learning. Remote Sensing of Environment 233, 111382. URL: https://www.sciencedirect.com/science/article/pii/S0034425719304018, doi:10.1016/j.rse.2019.111382.

Zhang, J., Okin, G.S., Zhou, B., Karl, J.W., 2021. UAV-derived imagery for vegetation structure estimation in rangelands: validation and application. Ecosphere 12, e03830. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/ecs2.3830, doi:10.1002/ecs2.3830. _eprint: https://esajournals.onlinelibrary.wiley.com/doi/pdf/10.1002/ecs2.3830.

Zhang, L., Zhang, L., Du, B., 2016. Deep Learning for Remote Sensing Data: A Technical Tutorial on the State of the Art. IEEE Geoscience and Remote Sensing Magazine 4, 22–40. doi:10.1109/MGRS.2016.2540798. conference Name: IEEE Geoscience and Remote Sensing Magazine.

References

Leslie Lamport, *LaTeX: a document preparation system*, Addison Wesley, Massachusetts, 2nd edition, 1994.