

Article

Pinus pinaster Seedling Detection in Coastal Dune Plantations Using a UAS Multispectral Point Cloud and Point Transformer V3

Tiago van der Worp da Silva ¹  and Luísa Gomes Pereira ^{2,*} 

¹ Center for Environmental and Marine Studies (CESAM), Department of Environment and Planning, University of Aveiro, Campus Universitário de Santiago 7, 3810-193 Aveiro, Portugal; tworp@ua.pt

² Águeda School of Technology and Management, Aveiro University, R. Cmte. Pinho e Freitas 28, 3750-127 Águeda, Portugal

* Correspondence: luisapereira@ua.pt

Highlights

What are the main findings?

- The baseline model retrieved 88% of field-mapped plantation-line seedlings, while the hard-negative refined model retained 84% recall and reduced off-buffer candidate detections by 83%.
- Under the adopted buffer-based retrieval framework, omission was most consistently associated with smaller crown expression and the absence of recorded branches.

What are the implications of the main findings?

- Multispectral photogrammetric point clouds and Point Transformer V3 can support object-level monitoring of early-stage pine seedlings in heterogeneous plantation environments.
- Model configuration should reflect the monitoring objective: prioritising either maximum seedling retrieval or a lower false-positive verification burden.

Abstract

Early detection of tree-seedling establishment is essential for monitoring regeneration success in coastal-dune plantations, where conventional field assessments remain labour-intensive and spatially limited. This study presents a deep-learning workflow for detecting early-stage *Pinus pinaster* seedlings using multispectral UAS-derived point clouds. Field surveys in the Quiaios National Forest, Portugal, mapped approximately 1500 seedlings using RTK GNSS positioning, biometric measurements, and field photographs. Multispectral imagery acquired with a DJI Mavic 3 Multispectral platform was processed through Structure-from-Motion to generate calibrated orthomosaics, terrain products, and dense point clouds. Training-data preparation combined pine-centred buffers, spectral conditioning, manual refinement and point-cloud class assignment. Point Transformer V3 models were trained in ArcGIS Pro and evaluated using field-mapped buffers withheld from model training within plantation-line areas. The Baseline high-recall model achieved 88% object-level recall at the operational threshold of at least three classified Pine-Seedling points per buffer. The refined hard-negative model retained 84% recall while reducing off-buffer detections from 243 to 41. False-negative analysis showed that omissions were associated with reduced crown diameter and limited branch development under the adopted buffer-based retrieval framework. These results support transformer-based multispectral point-cloud classification for scalable monitoring of early-stage pine regeneration in heterogeneous coastal environments.



Academic Editor: Michael Sprintsin

Received: 30 June 2026

Revised: 28 August 2026

Accepted: 31 August 2026

Published: 4 September 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: UAS-based monitoring; multispectral point clouds; seedling detection; *Pinus pinaster*; Point Transformer V3; deep learning

1. Introduction

Unmanned Aerial Systems (UASs), high-resolution remote sensing, and deep learning are increasingly used to support forest monitoring, species discrimination, forest-health assessment, and vegetation-structure analysis at spatial scales that are difficult to achieve through conventional field surveys alone [1–5]. These technologies are particularly relevant in forest regeneration and restoration contexts, where spatially explicit information on establishment patterns, survival, vegetation condition, and management needs is required across entire planting areas [6,7]. In Mediterranean and Atlantic coastal landscapes, where field monitoring may be labour-intensive, spatially discontinuous, and affected by heterogeneous vegetation structure, UAS-based workflows offer a practical route for improving the scale, repeatability, and spatial detail of regeneration assessment. Nevertheless, operational monitoring of early regeneration remains constrained by site-specific vegetation structure, sensor configuration, radiometric consistency, ground-reference quality, model-training design, and the limited transferability of models trained under complex field conditions [7].

Seedling detection is a distinct and more demanding problem than the detection of mature trees. Whereas mature crowns often form spectrally and structurally coherent objects, recently planted seedlings may occupy only a few centimetres of horizontal crown area, present sparse or stressed foliage, and occur within herbaceous or shrubby backgrounds with similar spectral responses. Previous UAS-based studies have demonstrated the feasibility of conifer seedling mapping and large-scale seedling detection, but have also emphasised the persistent influence of small target size, occlusions, background vegetation, and site variability on detection accuracy [6,8]. These limitations are especially relevant for pine seedlings, whose spectral–spatial signatures may strongly overlap with those of surrounding ground vegetation and whose weak canopy expression reduces their separability in image-based products [9,10].

In *Pinus pinaster* Aiton plantations, reliable information on seedling establishment is important for both ecological and silvicultural decision-making. Early-stage seedling density, spatial distribution, and condition influence subsequent stand structure, gap formation, competition dynamics, and the need for adaptive interventions such as replanting, vegetation control, targeted field verification, or localised management adjustments. In even-aged plantation systems, where uniform establishment is often a central objective, early identification of missing, weak, or poorly distributed seedlings can support more efficient monitoring and improve the timing of corrective actions. However, the same characteristics that make early monitoring valuable also make it technically difficult: seedlings are small, structurally variable, frequently stressed after planting, and often embedded in complex vegetation mosaics. In this context, conventional field assessment remains essential but is time-consuming and labour-intensive, particularly when survival, spatial gaps, and weak seedlings need to be assessed across entire plantation areas. Studies using UAS imagery and deep learning in pine plantations have shown that spectral confusion and background complexity can reduce segmentation reliability even for larger trees [9]. Reviews of UAS-assisted reforestation monitoring similarly highlight that sensor resolution, model-training strategy, and ecological heterogeneity remain major barriers to robust operational deployment [11].

Multispectral UAS data offer a practical way to address part of this challenge because they combine centimetre-scale spatial detail with reflectance information with visible, red-

edge, and near-infrared wavelengths. When imagery is radiometrically and geometrically calibrated, multispectral products can support vegetation discrimination, phytosanitary assessment, and the calculation of vegetation indices such as NDVI, which are sensitive to live vegetation and canopy vigour [12,13]. Repeated or seasonally timed UAS acquisitions can also improve species or vegetation-type discrimination by exploiting phenological differences, although this requires careful survey design, radiometric consistency, and representative sampling strategies [14,15]. For early-stage pine monitoring, the spectral information is useful but not sufficient on its own, because the target objects may be represented by fragmented, sparse, or mixed pixels and may remain difficult to separate from surrounding vegetation using raster imagery alone.

Most existing UAS-based seedling-monitoring workflows rely primarily on RGB or multispectral raster imagery, object-based image analysis, convolutional neural networks, or spectral-spatial image processing [6,8,10,16]. These approaches have demonstrated the potential of UAS data for seedling mapping, but they generally treat the detection problem through two-dimensional image objects, pixels, or image-derived features. In contrast, photogrammetric point clouds derived from UAS Structure-from-Motion processing provide an alternative representation that combines three-dimensional point geometry with multispectral attributes [17]. This representation is potentially useful for seedling detection because it preserves local spatial neighbourhoods, fine-scale surface structure, and point-level reflectance information, rather than forcing the detection task to rely exclusively on raster texture, pixel values, or image-object boundaries.

Recent advances in point-cloud deep learning, including transformer-based architectures, provide new opportunities for modelling irregular point distributions and local spatial-spectral relationships. However, the application of multispectral UAS-derived point clouds and transformer-based point-cloud classification for early-stage tree-seedling detection remains limited, particularly when the input data consist of high-density multispectral photogrammetric point clouds rather than airborne laser scanning or raster-derived image products. For small, planted seedlings, this gap is relevant because the target signal may be expressed as sparse, fragmented, and spatially localised point clusters embedded within a dominant background class.

The contribution of the present study is therefore primarily methodological and applied. Rather than proposing a new deep learning architecture, the study develops an integrated workflow for translating field-mapped early-stage seedlings into point-cloud training labels, applying Point Transformer V3 [18] to multispectral photogrammetric point clouds, and evaluating seedling retrieval at the object level. The novelty lies in the combination of: (i) field-derived pine-centred buffers used for supervised point-cloud labelling; (ii) multispectral point-cloud classification using Point Transformer V3; (iii) evaluation of field-mapped seedlings withheld from model training; (iv) hard-negative refinement to reduce shrub-related off-buffer responses; and (v) explicit interpretation of the operational trade-off between omission errors and commission-related verification burden.

The present study applies this workflow in a heterogeneous coastal-dune plantation located in the Quiaios National Forest (Matas Nacionais de Quiaios), central Portugal, where *Pinus pinaster* seedlings occur within sandy dune terrain and a complex herbaceous-shrub vegetation matrix typical of Atlantic dune-pine ecosystems [19,20]. The site represents a realistic operational monitoring context: seedlings are arranged mainly along plantation lines, but they are small, variably developed, and surrounded by spectrally similar vegetation, residual biomass, and occasional natural regeneration. This setting provides a demanding test case for evaluating whether high-density multispectral UAS-derived point clouds can support individual seedling retrieval under field conditions relevant to post-planting assessment.

To this end, this study integrates centimetre-resolution multispectral UAS imagery, RTK GNSS field mapping, seedling-centred reference buffers, spectral conditioning, manual label refinement, and deep-learning-based point-cloud classification. The workflow was designed to move from field-mapped individual seedlings to point-level semantic labelling and then to object-level evaluation, allowing model performance to be interpreted in terms of individual seedling retrieval rather than pointwise classification alone. This distinction is important because the operational question is not only whether individual points are labelled correctly, but also whether known planted seedlings can be reliably detected to support monitoring decisions.

The specific objectives were to: (i) construct a field-validated multispectral point-cloud dataset for early-stage *Pinus pinaster* seedling detection; (ii) train and compare a Baseline high-recall model and a Refined hard-negative model; (iii) quantify object-level retrieval of field-mapped seedlings withheld from model training within plantation-line areas; (iv) screen off-buffer detections to characterise commission-error behaviour; and (v) identify the seedling morphological attributes associated with omission errors. By focusing on small, planted seedlings in a heterogeneous coastal plantation, this study contributes an applied UAS-based workflow for early regeneration monitoring and provides evidence on the practical detectability limits of transformer-based multispectral point-cloud classification.

2. Materials and Methods

The following section describes the materials and the methodology adopted in this study, structured according to four main phases: (i) data acquisition, (ii) image processing, (iii) dataset curation for model training and evaluation; and (iv) point-cloud data preparation, model training, and evaluation, as illustrated in Figure 1.




Figure 1. Workflow for UAS-based detection of *Pinus pinaster* seedling using multispectral point-cloud data and deep learning.

The methodological pipeline integrates field and multispectral data acquisition, photogrammetric reconstruction, training-model dataset preparation, deep-learning-based point-cloud classification, and performance assessment. Field data acquisition combined RTK GNSS surveying of ground control points with the multispectral imagery of the survey area to ensure spatial accuracy. The individual *Pinus pinaster* seedlings were also georeferenced, and biometric attributes were recorded together with two geotagged photographs (I).

Photogrammetric processing included radiometric calibration, image alignment, dense point-cloud generation, and RGB and multispectral orthomosaics production (II).

Dataset curation for model training and evaluation was implemented using tree-centred buffers generated from measured seedling positions and crown dimensions, which served as object-based units for supervised training and subsequent detection validation. Spectral conditioning was performed within these buffer areas and further refined manually using orthomosaics and field imagery, resulting in a labelled point-cloud dataset (III).

Deep-learning classification was conducted using the Point Transformer V3 architecture [18,21], integrating multispectral and structural point-cloud attributes within a recall-oriented optimisation framework [22,23]. The final stage involved model evaluation and performance assessment for detecting *Pinus pinaster* seedlings in heterogeneous coastal-dune environments (IV).

A summary of the main resources (hardware and software) used throughout the workflow is provided in Table 1.

Table 1. Summary of tools, platforms, and devices used across the methodological workflow.

Component	Resources	Role in Workflow
GNSS positioning	Trimble R780 (Trimble Inc., Westminster, CO, USA); Emlid Reach RX (Emlid Tech Kft., Budapest, Hungary)	RTK-based georeferencing of individual seedlings and GCPs
Biometric data collection	iPhone 16 Pro (Apple Inc., Cupertino, CA, USA) + ArcGIS Field Maps (Esri, Redlands, CA, USA) version 25.2.0	Attribute recording, geotagged field photographs, and real-time data submission
UAS platform	DJI Mavic 3 Multispectral (M3M) (SZ DJI Technology Co., Ltd., Shenzhen, China)	Acquisition of high-resolution multispectral images
Radiometric calibration	Spectralon [®] reflectance panels (Labsphere, Inc., North Sutton, NH, USA)	Radiometric calibration of multispectral imagery to reflectance values
Spectral measurement	USB2000 Fibre Optic Spectrometer (Ocean Optics, Inc., Dunedin, FL, USA); HL120W halogen lamp (ELRO Europe/CoSa Connects B.V., Zundert, The Netherlands)	Laboratory measurement of reflectance-panel values
Photogrammetry	Agisoft Metashape Pro (Agisoft LLC, St. Petersburg, Russia) version 2.2.3	Radiometric calibration, image alignment, dense point-cloud processing, orthomosaic production
GIS processing	ArcGIS Pro (Esri, Redlands, CA, USA) version 3.6	Data curation, point-cloud labelling, model training, and evaluation
Deep learning	Point Transformer V3 implemented in ArcGIS Pro	Point-cloud classification and seedling detection
Processing Platform	Workstation equipped with Intel i9-13900k (Intel Corporation, Santa Clara, CA, USA), 64 GB RAM, NVIDIA GeForce RTX 4080 GPU (NVIDIA Corporation, Santa Clara, CA, USA)	Model training and GPU-accelerated computation

The study area description, as well as the workflow components of each phase, including materials, inputs, and outputs, are described in detail in the following sections.

2.1. Study Area and Ecological Context

The study was conducted in a coastal pine forest area in central Portugal, Mata Nacional de Quiaios, near Praia da Tocha, characterised by sandy coastal soils and vegetation assemblages typical of Atlantic dune–pine ecosystems [19,20] (Figure 2).




Figure 2. Survey-area context in the Matas Nacionais de Quiaios, Portugal.

The survey site is dominated by herbaceous and shrubby vegetation, including *Corynephorus canescens*, *Acacia longifolia*, *Helichrysum italicum*, *H. halimifolium*, *Cistus salvifolius*, *Ulex* spp., *Carex arenaria*, *Corema album*, along with other smaller seasonal herbaceous species. The area also contains residual biomass and debris originating from a 2017 wildfire, followed by mechanical clearing and site preparation in 2021 for pine-tree reforestation [24]. The area was replanted with *Pinus pinaster* in 2023 using a tractor-mounted soil ripper and currently supports predominantly planted individuals, with limited natural regeneration [25]. There are 27 plantation lines in this study area.

The UAS survey covers 1.82 ha. To minimise edge effects and ensure sufficient overlap for photogrammetric reconstruction, a 15-m internal buffer was applied, resulting in a final study area of 1.23 ha. The terrain is gently undulating, with elevations ranging from approximately 46 to 50.5 m, typical of stabilised coastal-dune systems (Figure 3).

2.2. Data Acquisition

The first methodological phase consisted of the acquisition of spatially explicit field and UAS data (2.2.1 and 2.2.2 respectively) required to construct the reference dataset and generate the remote-sensing products used for model development (Figure 4).




Figure 3. (a) Detailed RGB orthomosaic of the survey area containing ground-control points (GCPs) locations and (b) a DTM.




Figure 4. Overview of the methodological workflow in the data acquisition phase.

2.2.1. Field Data Acquisition

Field data consists of georeferenced positions of five ground-control points collected using a Trimble R780 RTK-enabled GNSS receiver from the nearest reference station in the national permanent network, ReNEP (Rede Nacional de Estações Permanentes GNSS). This is a public geopositioning service provided by the Portuguese Directorate General for Territory (DGT), consisting of GNSS stations with continuous observation that broadcast observations in ETRS89 (mainland) and ITRF93 (autonomous regions), supporting real-time RTK positioning and post-processing using RINEX files [26]. The nearest station, identified by the code MVEL, located at 14 km from the survey area, provided real-time corrections.

Individual *Pinus pinaster* seedlings were systematically mapped in situ within the cropped area using GNSS receivers, with positional data acquired as close as operationally feasible to the trunk centre or seedling base at ground level to minimise horizontal offset errors. Since all individuals were still at an early developmental stage and did not exceed knee height, the interference effects of vegetation, such as multipath propagation and signal attenuation, were minimal. Consequently, positional data acquisition was conducted under

effectively open-sky conditions, which are known to enhance GNSS positional accuracy and stability relative to closed-canopy forest environments [27,28]. All spatial data were referenced to the ETRS89/Portugal TM06 coordinate system with Cascais Height (GeodPT08), which was applied consistently across field surveys, photogrammetric processing, and GIS analyses.

For each mapped seedling, biometric attributes were collected through direct field measurement. Tree height and two orthogonal crown diameters were recorded to characterise vertical and horizontal structure, ensuring consistency and supporting individual-level analysis.

In addition to those geometric attributes, semantic information was recorded, including health status indicators such as crown condition (trunk or severe defoliation) and visible signs of stress (low or heavy stress). Two geotagged photographs per pine were also acquired to serve as visual references for post-field validation and interpretation, as illustrated in Figure 5.



Figure 5. Field picture from a *Pinus pinaster* seedling in two perspectives: nadir on the **left** and profile on the **right**.

This seedling field-data collection was carried out over six consecutive days immediately following image acquisition, ensuring strong temporal correspondence between ground observations and remotely sensed data. This temporal alignment minimises discrepancies associated with short-term environmental variability and supports the use of the dataset as reliable ground truth. The resulting dataset comprised 1541 mapped seedlings.

All data capture and attribute management were implemented using the ArcGIS Field Maps mobile application, enabling real-time recording of spatial and descriptive information within a structured geodatabase. Biometric data collection was conducted using an iPhone 16, which served as an integrated platform for attribute entry, geotagged-photograph acquisition, and direct-data submission. The device was connected to the Emlid Reach RX RTK-enabled GNSS receiver via the iOS environment, allowing high-

precision positioning data to be recorded and automatically linked to each seedling in real time. This configuration provided centimetre-level positional accuracy while maintaining a streamlined data acquisition workflow.

The integration of positioning, attribute collection, and data synchronisation within a single mobile platform significantly reduced the time required per sampled individual, thereby improving field efficiency and optimising the time-to-information ratio. Data were synchronised directly with the desktop ArcGIS project, ensuring immediate integration, minimising transcription errors, and supporting consistent data management across devices and operators.

2.2.2. UAS Data Acquisition

Multispectral UAS imagery was acquired on 3 September 2025, using a DJI Mavic 3 Multispectral (M3M), between 12:21 and 13:11 local time, with RTK enabled throughout the mission, following recent applications of RTK-equipped multispectral DJI platforms for high-accuracy mapping and land-cover classification [29,30].

The M3M integrates two complementary imaging systems: a 20 MP RGB camera with a 4/3 CMOS sensor and four 5-MP multispectral cameras with 1/2.8-inch CMOS sensors. The multispectral module acquires green, red, red-edge, and near-infrared bands centred at 560 ± 16 nm, 650 ± 16 nm, 730 ± 16 nm, and 860 ± 26 nm, respectively, with each band recorded at 2592×1944 pixels, while the RGB camera records images at a maximum resolution of 5280×3956 pixels.

The flight was planned for a multispectral ground sampling distance of approximately 1 cm pixel^{-1} , based on the spatial resolution of the M3M multispectral sensors, to ensure visibility of fine vegetation details such as pine seedlings, consistent with UAV multispectral studies that exploit centimetre-scale resolution for vegetation monitoring and modelling [31,32]. Under the same flight configuration, the higher native resolution of the RGB sensor was enabled to produce a finer-resolution RGB orthomosaic. Image overlap was set to 83% forward and 75% side, in line with recommended high-overlap configurations for Structure-from-Motion photogrammetry in coastal and vegetated environments to improve point-cloud robustness [33–36].

Meteorological conditions during the UAS acquisition were characterised using station-based observations. Observations were obtained from the Figueira da Foz weather station (≈ 12.5 km from the study area) using data accessed via Meteostat [37]. The proximity and comparable coastal exposure support the assumption that the recorded conditions are representative of the study site, although local wind variability cannot be ruled out.

The UAS acquisition was conducted under thermally stable, precipitation-free atmospheric conditions, with air temperature remaining nearly constant throughout the acquisition window (Table 2). These conditions indicate a stable boundary layer, which is generally favourable for radiometric consistency in multispectral data acquisition [35,38,39].

Table 2. Meteorological conditions during UAS image acquisition.

Time	Temperature (°C)	Wind Speed ($\text{m}\cdot\text{s}^{-1}$)	Wind Direction (°)	Wind Gust ($\text{m}\cdot\text{s}^{-1}$)	Precipitation ($\text{mm}\cdot\text{h}^{-1}$)	Relative Humidity (%)
12:00	23.01	3.41	262	18.5	0.017	71
13:00	23.11	3.97	263	23.0	0.007	57

Wind conditions were characterised by a coastal sea-breeze regime, with westerly to north-westerly flow and moderate background wind speeds. However, station observations revealed elevated gust intensity. Such conditions are typical of Atlantic coastal environments and therefore representative of the study context. The presence of gusty

winds may have induced minor vegetation movement, particularly in low-height pine seedlings and shrub components. This effect can introduce small-scale geometric variability in photogrammetric reconstruction and increase spectral variability within vegetation classes [26,38].

Five ground control points (GCPs) were deployed and surveyed using the RTK GNSS receiver Trimble R780 (Figure 3) and distributed approximately homogeneously across the site to capture topographic variability of the dune surface, with four GCPs placed near the corners of the sampling area and used in the bundle block adjustment, while a centrally located GCP was withheld and used exclusively as an independent checkpoint, following best practice on GCP number, spatial configuration, and separation of control and checkpoints in UAV-SfM accuracy assessment. This distribution is in line with recommendations and provides sufficient geometric constraints to minimise distortions and improve horizontal and vertical accuracy of the derived digital-terrain model and dense point cloud [12,40–43].

In addition, four 60×60 cm Spectralon® reflectance panels were deployed within the survey area to support radiometric calibration of the multispectral imagery. Panel reflectance values were measured under laboratory conditions in a dark room using an Ocean Optics USB2000 Fibre Optic Spectrometer, with illumination provided exclusively by an HL120W Elro Halogen lamp [12]. Corresponding M3M reflectance bands were derived from these measurements to ensure accurate calibration of the collected images. Panels were positioned adjacent to the GCPs located near the corners of the survey area, ensuring clear visibility across multiple flight lines [38,44].

2.3. Image Processing

The second methodological phase consisted of photogrammetric image processing, including radiometric calibration (2.3.1), point-cloud generation (2.3.2), and orthomosaic production (2.3.3) (Figure 6).




Figure 6. Overview of the methodological workflow in the image processing phase.

2.3.1. Photogrammetric Processing and Radiometric Calibration

Photogrammetric processing was carried out in Agisoft Metashape Professional 2.3.2 (Agisoft LLC, St. Petersburg, Russia) following a workflow designed to preserve both geometric accuracy and spectral consistency. Radiometric calibration was performed using the Spectralon® reflectance panels (2.2.2) with known reference values [12]. To account for

potential illumination variability, the image dataset was divided into two temporal blocks, each calibrated independently and cross-validated using a secondary reflectance panel. For each block, radiometric calibration was anchored to one panel, while a secondary panel within the same block was used for independent verification [12,38,44,45]. Radiometric consistency was quantified in ArcGIS Pro by computing mean reflectance values for each band from panel regions of interest using spectral profile analysis, thereby enabling assessment of internal radiometric stability across panels and flight segments [12].

Image alignment was performed at high accuracy, followed by bundle block adjustment incorporating the four corner GCPs (2.2.2). Dense point-cloud reconstruction was performed by generating ultra-high quality depth maps with mild depth filtering to preserve fine-scale structure [12,40,46]. The resulting dense point cloud constituted the primary dataset for terrain modelling and GIS-based and deep-learning processing.

2.3.2. Point-Cloud Generation

Dense point-cloud generation was performed in Agisoft Metashape Professional after image alignment and bundle-block adjustment. Depth maps were generated at ultra-high quality using mild depth filtering, as this configuration was intended to preserve fine-scale vegetation structure while limiting excessive removal of small or fragmented seedling features. The resulting dense point cloud constituted the principal three-dimensional product of the photogrammetric workflow. It served two methodological functions: first, it provided the surface information required to generate a DSM and to support orthorectification of the RGB and multispectral orthomosaics; second, it formed the analytical point-cloud dataset used for subsequent spectral attribution, semantic labelling, deep-learning classification, and object-level seedling evaluation. The point cloud retained three-dimensional point coordinates and multispectral reflectance attributes derived from calibrated M3M imagery, corresponding to the green, red, red-edge, and near-infrared bands. This structure allowed each point to carry both geometric position and spectral information, essential for representing *Pinus pinaster* seedlings.

The dense point cloud was exported in LAS format for subsequent processing in ArcGIS Pro. LAS class codes were used as the semantic labelling framework throughout the following stages. The assignment of *Pinus pinaster* seedling points to a dedicated user-defined LAS class was performed only after spectral conditioning and manual refinement of pine-centred buffers, as described in Section 2.4.3.

2.3.3. Orthomosaic Production

Orthomosaic products were generated in Agisoft Metashape Professional 2.3.2 after the photogrammetric processing. A multispectral orthomosaic was processed using the green, red, red-edge, and near-infrared reflectance values, while an RGB orthomosaic was produced separately to support visual interpretation of the study area and manual verification of detection outputs. Orthorectification was performed using the Digital Surface Model (DSM) produced from the point cloud (2.3.2), ensuring that the orthomosaic geometry was corrected for relief displacement and perspective-related distortions. The resulting orthomosaics, multispectral and RGB, were exported in the ETRS89/PT-TM06 coordinate reference system with ground sampling distance set to 0.5 cm pixel⁻¹ for the RGB and 1 cm pixel⁻¹ for the multispectral orthomosaics. These orthomosaics were subsequently used as reference layers for NDVI production, interpretation of seedling and shrub distribution (2.4.2), and visual assessment of model detections, particularly in the analysis of false-positive and false-negative cases (2.5.3).

2.4. Dataset Curation for Model Training and Evaluation

The third methodological phase converted the georeferenced field inventory (2.2.1) and photogrammetric products (2.3.3) into spatially explicit reference objects and labelled point-cloud inputs for model training and evaluation (Figure 7).




Figure 7. Overview of the methodological workflow in the dataset curation for model training and evaluation phase.

This stage used the field-mapped *Pinus pinaster* locations, crown-diameter measurements (2.2.1), multispectral point cloud (2.3.2), vegetation-index products, and RGB/multispectral orthomosaics (2.3.3) to define pine-centred buffers (2.4.1) and refine seedling-related vegetation signals (2.4.2). Spectral conditioning (2.4.2) was applied only within pine-centred buffers to retain vegetation with reflectance characteristics consistent with live pine foliage, after which manual refinement was performed using the orthomosaics and field photographs (2.4.3). The resulting outputs were buffer-based seedling objects and a labelled point cloud in which background points remained assigned to Class 0 and pine-seedling points were assigned to Class 64 (2.4.3). These outputs formed the basis for subsequent point-cloud training-data preparation, model training, and evaluation (2.5).

2.4.1. Pine-Centred Buffer Construction

Pine-centred buffers were used to translate the field inventory into spatial reference units, linking individual surveyed seedlings to the corresponding photogrammetric point-cloud data. They were constructed individually for each surveyed *Pinus pinaster* seedling to provide weak-supervision spatial units and to support subsequent model training

and performance assessment. Buffer radii were calculated as half of the mean of the two measured crown diameters, with an additional 4-cm radial margin to account for the expected geometric uncertainty associated with RTK GNSS positioning, photogrammetric reconstruction, and point-cloud alignment [27,41,42]. Buffers were centred on the RTK GNSS coordinates of the tagged pine. Figure 8 represents an example of a pine buffer, exhibiting its position in the field, on top of the radiometrically calibrated multispectral orthomosaic, followed by the nadir picture taken during field-data acquisition (2.2.1).



Figure 8. Example of pine-centred buffer construction for an individual *Pinus pinaster* seedling. The GNSS seedling position is shown over the multispectral orthomosaic, together with the corresponding buffer derived from the measured crown dimensions and additional radial margin. The inset shows the field photograph of the same seedling.

The buffers were designed to encompass the spatial extent of each seedling rather than to delineate exact crown boundaries, thereby reducing the influence of positional uncertainty and irregular crown geometry during point classification. Additionally, the vegetation fraction for each buffer was calculated using the spectral conditioning applied (2.4.2). This output measures the vegetation density within buffers.

2.4.2. Spectral Conditioning and Pine-Seedling Data Refinement

Prior to the preparation of the model training and evaluation data, spectral conditioning was applied to multispectral raster data within the pine-centred buffers to reduce the influence of non-vegetated and spectrally ambiguous areas before raster-derived information was transferred to the point cloud. A series of NDVI and near-infrared (NIR) (2.3.3) reflectance thresholds was evaluated during exploratory analysis and was used as empirical, site-specific filters [10,47]. The most permissive and stable configuration applied

an NDVI threshold of at least 0.30 and a minimum NIR reflectance threshold of 0.25. These thresholds were applied exclusively within the pine-centred buffers and were not used to define class labels. NDVI values as low as 0.30 were retained to account for variability in pine health and illumination conditions, while the additional NIR constraint reduced the influence of shadowed surfaces and spectrally inconsistent returns that could otherwise inflate NDVI values. This spectral conditioning step served solely to suppress obvious non-vegetation signals.

Following this raster-based spectral filtering, a manual inspection and refinement procedure was conducted at the buffer level. Each pine-centred buffer was reviewed individually, and localised adjustments to the spectral masks were made using the Pixel Editor tool in ArcGIS Pro. This process was supported by visual cross-referencing with high-resolution RGB orthomosaics (2.3.3) and field-acquired imagery of the corresponding pine seedlings (2.2.1), enabling contextual interpretation of canopy structure, surrounding vegetation, and illumination conditions. Manual corrections were applied conservatively to address evident misclassifications, such as incorporating non-vegetated surfaces, removing valid vegetation points, and, critically, correcting the misrepresentation of surrounding non-pine vegetation within pine-designated buffers. This step aimed not only to reduce spectral noise ambiguity but also to improve the separability between *Pinus pinaster* seedlings and adjacent vegetation, ensuring that the resulting point cloud more accurately reflected the true presence of pine while minimising ambiguity in the training data.

2.4.3. Point-Cloud Labelling and Class Assignment

Point-cloud labelling was performed following raster-based spectral conditioning and manual refinement within the pine-centred buffers. The refined raster masks were used to assign semantic class labels to the dense point cloud through a raster-to-point classification procedure implemented in ArcGIS Pro. Specifically, the Set LAS Class Codes Using Raster tool was applied to identify points that spatially intersected raster cells classified as pine vegetation, and to assign them to a dedicated class code (Class 64), defined in this study as pine seedling. All remaining points were treated as background (Class 0) during subsequent processing.

This procedure ensured that point-level class labels were derived directly from the refined raster interpretation, maintaining consistency between multispectral image analysis and point-cloud representation.

2.5. Point-Cloud Data Preparation, Model Training, and Evaluation

The fourth methodological phase converted the labelled point cloud (2.4.3) into model-ready inputs for deep-learning-based seedling detection (2.5.1), training the Point Transformer V3 (2.5.2), and evaluation of the trained model's outputs (2.5.3) (Figure 9).

First, the labelled point cloud was partitioned into spatially structured blocks compatible with the input requirements of the deep-learning architecture, while preserving the high-density spectral and geometric information of the UAS-derived photogrammetric reconstruction. Second, these blocks were used to train a binary point-cloud classification model in which Class 64 represented *Pinus pinaster* seedlings and Class 0 represented background points. Third, the trained model was applied to the complete point cloud. Fourth, the model's output was evaluated using object-level detection, allowing performance to be assessed at the object-level rather than only through point-wise classification metrics.

2.5.1. Point-Cloud Training-Data Preparation for Model Ingestion

Training-data preparation was conducted in ArcGIS Pro using the labelled point cloud derived from the previous processing steps (2.4.3).




Figure 9. Overview of the methodological workflow in the point-cloud data preparation, model training, and evaluation phase.

Model-ready point-cloud samples were generated using the Prepare Point Cloud Training Data tool in ArcGIS Pro. This tool partitions a labelled point cloud into fixed-size spatial blocks suitable for ingestion by deep learning models, and it requires defining training and validation areas, class-labelled input points, block dimensions, point-count limits, and optional elevation-normalisation inputs. Consequently, before executing the tool, the field reference dataset and labelled point cloud were organised to define: (i) the valid seedling population, (ii) the subset of seedlings used for semantic labelling and model development, (iii) the subset withheld from positive supervision for object-level evaluation, and (iv) the spatial training and internal validation regions used during model training.

Prior to the preparation of the training and evaluation data, the field-mapped seedling dataset was quality-controlled to retain a valid dataset. The initial field inventory comprised 1541 mapped *Pinus pinaster* seedlings. Of these, 50 records were excluded because they fell outside the plantation line buffer, leaving 1491 mapped seedlings. A further 75 individuals recorded as dead during field surveys were excluded, resulting in 1416 living pine-centred buffers. This final living-buffer dataset was divided into two subsets before point-cloud labelling (2.4.3). A subset of 889 seedlings was used for model development, including training and internal validation (2.5.2). The remaining 527 seedlings were withheld from model training and reserved for object-level model evaluation within the plantation domain (2.5.3). These seedlings remained physically present in the dense point cloud as Class 0, but they were not included in the model-training samples. For the baseline model (2.5.2), Class 64 was selected as the class of interest during model training. For the refined model (2.5.2), in which the class of interest setting was not available, training-data preparation was configured with filter blocks by class code set to Class 64. Consequently, only blocks containing labelled pine-seedling points were exported for refined-model training, and Class-0-only withheld seedlings were not introduced as negative seedling examples. This filtering was necessary to counter the severe class imbalance between pine seedling and background points.

The model development subset was further divided spatially into training and internal validation datasets using polygon boundaries defined by plantation lines (Figure 10).

The study area comprised twenty-seven plantation lines, of which the first nineteen, from north to south, were allocated to the training region, and the subsequent eight to the internal validation region, corresponding approximately to a 70/30% spatial partition, respectively [48]. Within the model development subset, 634 seedlings were in the

training region and 255 in the validation region. The use of complete plantation lines as partitioning units was intended to reduce spatial leakage between training and validation samples [49]. In geospatial datasets, neighbouring observations frequently share similar spectral, structural, and environmental characteristics, which may inflate model performance if spatially adjacent samples are included in both the training and validation sets [50,51]. By assigning complete plantation lines exclusively to either subset, the training and validation regions remained spatially separated while retaining comparable plantation structure and background vegetation conditions.




Figure 10. Spatial partitioning of the analytical study area into model-training and internal validation regions. Plantation lines are overlaid on the RGB orthomosaic, with the first nineteen lines assigned to the training region and the remaining eight lines assigned to the internal validation region. Pine-centred buffers used for model development and withheld buffers reserved for object-level evaluation are shown separately.

Model-ready point-cloud blocks were generated using the Prepare Point Cloud Training Data tool in ArcGIS Pro, which partitions the point cloud into fixed-size spatial blocks suitable for ingestion by deep-learning models. Block-based sampling is commonly adopted in point-cloud deep learning because it enables efficient processing of large datasets while preserving local neighbourhood relationships required for feature learning.

The model development point cloud was generated using spatial blocks of 0.6 m. This block size was selected based on the characteristic dimensions of the target objects, as the vast majority of *Pinus pinaster* seedlings observed in the field had crown diameters of less than 0.5 m. A 0.6 m block extent was therefore considered sufficient to capture most individual seedling structures while limiting excessive inclusion of adjacent vegetation that could increase class ambiguity. At the same time, the selected block size preserved the

immediate local context required for attention-based feature aggregation in transformer-based architectures [18,21,52].

The block-point limit was defined empirically after inspecting the point-count distributions of the generated training and internal validation blocks. This ensured that the selected limit was compatible with the native point density of the photogrammetric reconstruction, while avoiding unnecessary subsampling of dense blocks and maintaining feasible GPU processing.

The resulting output consisted of spatially structured point-cloud blocks formatted for direct ingestion into the deep-learning pipeline. These blocks preserved the native geometric density and spectral variability of the photogrammetric reconstruction while maintaining the training, internal validation, and withheld object-level evaluation structure required for subsequent model development and assessment.

2.5.2. Deep-Learning Model Training

Deep-learning model training was conducted in ArcGIS Pro using the Train Point Cloud Classification Model tool, implementing the Point Transformer V3 architecture [18,53]. This architecture was selected because attention-based mechanisms are well-suited to modelling local neighbourhood relationships in irregular point-cloud data [18,21,54]. In this study, *Pinus pinaster* seedlings are small targets embedded within spectrally similar surrounding vegetation; therefore, the ability to capture fine-scale local spatial–spectral relationships was considered important for class discrimination [22,55].

Model input consisted of the three-dimensional point geometry of each training block, defined by the x, y, and z coordinates of the photogrammetric point cloud, along with selected point attributes. The selected non-geometric attribute channels consisted of multispectral radiometric values from the M3M sensor, including those of the green, red, red-edge, and near-infrared bands.

The classification scheme (2.4.3) defined Class 64 as pine seedling, while all remaining points were treated as background or unclassified points under Class 0. The baseline model was trained with the class-of-interest parameter explicitly set to Class 64. This configuration prioritised learning for the minority class over the background class. In this study, pine seedlings occupied a small proportion of the total point cloud and were embedded within spectrally similar vegetation; therefore, the use of Class 64 as the class of interest was important for stable convergence.

Training followed a recall-oriented optimisation strategy [53], reflecting the operational priority of minimising omission errors during seedling detection. This was considered appropriate because undetected seedlings directly limit the workflow's usefulness for field monitoring and plantation assessment. Models were trained for up to 35 epochs, with this upper bound serving as the training budget rather than a fixed optimisation target. This limit allowed validation metrics to stabilise while avoiding unnecessarily prolonged training. Early stopping was enabled to terminate training when no further improvement in validation was observed, thereby reducing the risk of overfitting to the labelled training blocks. The retained model was selected primarily according to validation recall for the pine-seedling class, but validation precision and F1 score were also examined to avoid selecting a model that achieved high recall only by overpredicting pine-seedling points in the background class. Therefore, model selection prioritised seedling retrieval while maintaining a defensible balance between omission and commission errors [55,56].

Model training used cross-entropy loss, consistent with binary semantic classification, in which each point was assigned either to the pine-seedling class (Class 64) or to the background class (Class 0) [56–59]. The learning rate was automatically managed using a

one-cycle policy, enabling dynamic adjustment during training and reducing the need for manual learning-rate tuning.

Model development followed an iterative parameter-refinement procedure. After each model was trained, it was applied to classify the entire point cloud, and object-level detection metrics were calculated using the withheld plantation-line buffers described in Section 2.5.3. These repeated evaluations were used to compare the effect of training parameters, including block size, batch size, and class-of-interest configuration, on seedling retrieval. The primary selection criterion was object-level recall, reflecting the objective of maximising detection of known field-mapped seedlings. However, configurations that achieved high recall by excessively classifying background vegetation as pine seedling were not retained. Therefore, parameter refinement was guided by the objective of achieving the highest practical recall while maintaining a defensible level of commission-related responses.

Following an evaluation of the baseline model (baseline) (2.5.3), a refined model (refined) was developed through additional training, using the trained baseline model as a starting checkpoint, to reduce residual false-positive detections of shrub vegetation. False-positive clusters from the baseline classification were visually inspected using the high-resolution RGB orthomosaic (2.3.3) and field imagery (2.3.1). Areas confidently associated with shrub vegetation were manually digitised as hard-negative polygons and retained as background Class 0, preserving the binary classification structure [60]. The refined model retained the same model development and withheld seedling subsets, but incorporated layers defined by plantation line buffers (width 0.6 m) and hard negatives for training and internal validation (Figure 11).




Figure 11. Spatial distribution of the training and internal validation plantation-line buffers and hard-negative shrub/background areas used during refined model development, displayed over the RGB orthomosaic.

Unlike the baseline configuration, the class-of-interest constraint was not applied during refinement, allowing the model to incorporate a broader representation of background vegetation during training, specifically within plantation lines. Consistent with the strategy of increasing background representation during model training, the block size was increased to 1.0 m (2.5.1).

2.5.3. Model Evaluation

Model evaluation was designed to quantify object-level retrieval of individual *Pinus pinaster* records from the trained point-cloud classifications, rather than relying solely on point-wise semantic classification. The evaluation subset consisted of 527 pine-centred buffers withheld from model training and retained for post-classification object-level assessment. As described in Section 2.5.1, Class-0-only withheld seedlings were excluded from model-training samples through the class-of-interest setting for the baseline model and through filter blocks by class code set to Class 64 during refined-model training-data preparation. Of the 527 withheld buffers, 442 were located within the spatial-training region and 85 within the spatial-internal-validation region.

After each trained model was applied to classify the complete point cloud, object-level evaluation was performed by intersecting the classified output with the withheld pine-centred buffers. For each buffer, the number of points predicted as pine seedling (Class 64) was quantified using the Spatial Join tool in ArcGIS Pro. A pine record was considered successfully detected when the number of predicted Class 64 points within its corresponding buffer met or exceeded a predefined point-count threshold. Multiple thresholds were evaluated, including 1, 3, 5, and 10 classified points per buffer, to assess the sensitivity of seedling retrieval to the required number of predicted pine points. Lower thresholds favour sensitivity to small or weakly represented seedlings, whereas higher thresholds require stronger within-buffer evidence and may therefore increase omission of small, partially occluded, or weakly reconstructed individuals. A threshold of at least three Class 64 points per buffer was retained as the operational threshold for the main object-level assessment, balancing sensitivity to small seedlings with reduced influence of isolated one- or two-point responses.

Object-level recall was calculated at the level of individual withheld pine-centred buffers rather than at the level of individual points. At a given detection threshold (T), a buffer was classified as a true positive (TP) when it contained at least (T) predicted Class 64 points, and as a false negative (FN) when it contained fewer than (T) predicted Class 64 points. Recall was calculated following Equation (1):

$$Recall_{(T)} = \frac{TP_{(T)}}{TP_{(T)} + FN_{(T)}} \quad (1)$$

where $TP_{(T)}$ is the number of withheld pine-centred buffers containing at least T predicted pine-seedling points, and $FN_{(T)}$ is the number of withheld pine-centred buffers containing fewer than T predicted pine-seedling points. Therefore, at the operational threshold of ($T = 3$), recall represented the proportion of known field-mapped seedlings that were retrieved by the model with at least three predicted Class 64 points inside their corresponding reference buffer. Because this assessment was conducted only on known pine-centred buffers, it quantified seedling retrieval and omission error, but did not directly estimate landscape-wide commission error. This evaluation should therefore be interpreted as a retrieval-oriented object-level assessment of known field-mapped seedlings, rather than as a complete object-detection evaluation across all pine and non-pine objects in the landscape. Conventional object-detection precision and F1 score were not calculated because a complete set of explicit non-pine reference objects was not available. Commission-related

behaviour was instead assessed separately through off-buffer candidate screening within the plantation-line domain.

Given the intended operational application of the workflow, model evaluation was primarily assessed within plantation-line buffers. This restriction reflected the main management objective of detecting planted seedlings within plantation rows, where seedling presence is expected and where field-based monitoring is operationally relevant. It also reduced the influence of ambiguous inter-row vegetation outside the model's primary application domain.

Commission-error behaviour was assessed through a complementary post-classification spatial-screening procedure implemented in ArcGIS Pro. Accordingly, predicted Class 64 points were first extracted from the classified point cloud and compared spatially with the reference pine-centred buffers. Class 64 points falling inside known reference buffers were excluded from the commission-screening workflow, while Class 64 points located outside the buffers but within the 1.2-m plantation-line buffers were retained as candidate off-buffer detections. To reduce the influence of isolated point detections, the retained off-buffer Class 64 points were processed using density-based clustering in ArcGIS Pro. The point features were clustered using DBSCAN, with a search distance of 0.10 m and a minimum of three points per cluster, corresponding to the operational threshold definition ($T = 3$). The 0.10 m distance approximately corresponded to the mean field-measured crown size of the mapped pine seedlings, while the minimum cluster size of three points prevented isolated one- or two-point responses from being interpreted as candidate objects. Points classified by DBSCAN as noise or unclustered were not considered, and only points belonging to valid clusters were retained for subsequent screening.

Each retained cluster was assigned a unique cluster identifier and treated as a pine seedling candidate off-buffer detection. Cluster centroids were generated using the cluster identifier as the case field, producing one representative point per candidate cluster. These centroids were used as the primary inspection layer because they provided a clear spatial reference for visual interpretation and, where required, field verification. Cluster hulls were also generated to support visual interpretation of each candidate's spatial extent, although final interpretation was based on clustered points, centroid locations, and ancillary imagery rather than on hull geometry alone.

The retained candidate clusters were visually screened in ArcGIS Pro using the RGB and multispectral orthomosaics, original pine-centred buffers and 1.2-m plantation-line buffers. Each candidate detection was inspected to determine whether it corresponded to a likely non-pine vegetation response, a valid pine individual not represented by the buffers, or an ambiguous object. Candidate clusters clearly associated with shrubs, herbaceous vegetation, and residual biomass were classified as false positives. Candidate clusters corresponding to visually identifiable pine individuals absent from, or not correctly represented by, the buffer dataset were classified as valid pine detections in the commission-screening context. Candidate clusters that could not be confidently interpreted from the orthomosaic and ancillary layers were retained as uncertain and flagged for field verification based on their centroid coordinates. These detections were not included in the calculation of object-level recall because recall was defined exclusively from the withheld reference buffers. Instead, the screening provided a complementary assessment of commission-error behaviour and allowed off-buffer detections to be separated into false positives, valid off-buffer pine detections, and uncertain cases.

To characterise the detectability limits of the models, detected and omitted seedlings were compared within the withheld plantation-line evaluation subset at the operational threshold of ($T = 3$). Mann–Whitney U tests were used for continuous variables, including height, mean crown diameter and vegetation fraction within the buffer, while Fisher's exact

tests were used for binary categorical traits. Effect sizes for continuous comparisons were reported using Cliff's delta. Paired differences in detection outcome between the baseline and refined models were assessed using McNemar's exact test, because both models were evaluated on the same withheld buffer set.

Model comparison was conducted using the same withheld plantation-line evaluation subset across configurations. The baseline model represented the high-recall configuration for seedling retrieval within plantation-line buffers. The refined model was used to assess whether targeted inclusion of hard-negative shrub samples and broader background exposure could reduce shrub-related commission detections. This comparison allowed model behaviour to be interpreted as a trade-off between omission and commission errors rather than as a single absolute performance outcome.

3. Results

The results are presented in the order outlined in Figure 4, progressing from intermediate outputs generated during data acquisition and photogrammetric processing and dataset preparation to the final evaluation of the trained models. First, the UAS-derived products are described, including the reconstructed point cloud, orthomosaics, and terrain products. Second, the field reference dataset is characterised in terms of positional consistency, living-seedling filtering and allocation to model-development and withheld evaluation subsets. Third, the composition and morphological characteristics of the seedling buffers are reported, with emphasis on the complete model-development subset and the withheld subset used for model evaluation. Fourth, the point-cloud training data and model configurations are summarised. Finally, model performance is assessed through object-level seedling retrieval within withheld plantation-line buffers, complementary off-buffer commission screening and false-negative characterisation.

3.1. UAS-Derived Products and Analytical Dataset

The UAS-SfM workflow generated spatial, spectral, and point-cloud products. A total of 3636 band-specific multispectral images were aligned in Agisoft Metashape, covering a reconstructed area of 1.66 ha. The imagery was acquired at an average flying altitude of 22.1 m, yielding a ground resolution of approximately 1 cm pixel⁻¹. Image alignment generated 3,276,710 tie points from 19,858,277 projections, with a mean reprojection error of 0.566 pixels.

After applying the internal edge-exclusion buffer (2.2), the final analytical area used for reference-buffer construction, model preparation, point-cloud classification, and object-level evaluation corresponded to 12,314.66 m², equivalent to 1.23 ha. Table 3 summarises the UAS-derived products and photogrammetric outputs used in the analytical workflow.

The dense photogrammetric point cloud comprised 89,928,442 points, with a coordinate precision of 4.81 mm. The reported point-cloud coordinate precision refers to the internal coordinate precision estimate associated with the photogrammetric dense cloud and should not be interpreted as millimetre-level absolute positional accuracy. Absolute spatial correspondence between field and photogrammetric data was constrained by the RTK-enabled UAS acquisition, the four GCPs used in bundle block adjustment, the independent checkpoint, and the positional agreement between the Emlid Reach RX and Trimble R780 receivers. The independent checkpoint confirmed centimetre-level agreement between the photogrammetric products and the field-surveyed reference point. Given the 1-cm multispectral GSD, the 2.6–3.0 cm mean horizontal difference between GNSS receivers, and the small size of the target seedlings, positional uncertainty remained meaningful relative to the approximately 7.5–10.5 cm crown-diameter scale of many individuals. The additional 4-cm radial margin used in pine-centred buffer construction was therefore

included to reduce the risk that valid seedling points were excluded from reference buffers due to combined GNSS, photogrammetric, and alignment uncertainties. Point attributes included multispectral information corresponding to the green, red, red-edge, and near-infrared bands, together with photogrammetric descriptors such as normal vectors and confidence values. A digital terrain model was generated at 1-cm pixel^{−1} resolution (2.3.2), with elevations ranging from approximately 46 to 51 m across the reconstructed area. The multispectral orthomosaic provided the basis for vegetation-index generation and spectral conditioning, while the RGB orthomosaic served as an ancillary high-detail visual layer for manual inspection, interpretation of seedling and shrub context, and post-classification screening of model detections.

Table 3. Summary of UAS-derived products and photogrammetric outputs used in the analytical workflow.

Product or Metric	Value	Role in Workflow
Aligned band-specific multispectral images	3636	Photogrammetric reconstruction
Average flying altitude	22 m	Acquisition geometry
Final analytical area	12,314.66 m ² /1.23 ha	Cropped area used for analysis
Tie points	3,276,710	Image alignment
Projections	19,858,277	Bundle adjustment support
Mean reprojection error	0.57 pixels	Alignment quality indicator
Dense point-cloud points	89,928,442	Analytical point cloud
Mean dense-cloud point density	7303 points m ^{−2}	Point-cloud density context for seedling representation
Approximate mean point spacing	1.17 cm	Spatial context for fine-scale seedling detectability
Point-cloud coordinate precision	4.81 mm	Internal dense-cloud coordinate precision
Independent checkpoint horizontal error	0.85 cm	Absolute horizontal-accuracy check
Independent checkpoint vertical error	0.08 cm	Absolute vertical-accuracy check
Independent checkpoint 3D error	0.85 cm	Absolute 3D-accuracy check
Point attributes	Four-band multispectral attributes, normals, point confidence	Point-cloud structure and quality descriptors
Digital surface model	1-cm pixel ^{−1}	Orthorectification and surface representation
Multispectral orthomosaic	1-cm pixel ^{−1}	Spectral conditioning, visual refinement, and vegetation index generation
RGB orthomosaic	0.5-cm pixel ^{−1}	Visual interpretation and manual inspection

3.2. Field Dataset and Positional Consistency

The field inventory produced a spatially explicit reference dataset of mapped *Pinus pinaster* seedlings. After excluding individuals recorded as dead, the living pine-centred buffer dataset comprised 1416 records. Of these, 889 buffers were assigned to the model-development subset and retained the broader spatial and morphological vari-

ability of the mapped seedlings across the study area. For object-level model evaluation, the withheld subset was restricted to the 527 buffers located within plantation-line areas, corresponding to the operational domain in which planted-seedling retrieval was assessed.

The positional consistency between the Emlid Reach RX receiver used for seedling mapping and the Trimble R780 receiver used as the higher-grade reference instrument was assessed at 31 common control points. Instantaneous RTK measurements showed a mean horizontal difference of 2.6 cm between receivers, with a 3D RMSE of 4.54 cm. The 3-min averaged positions showed a mean horizontal difference of 3.0 cm and a 3D RMSE of 4.95 cm. The observed differences indicate centimetre-level agreement between receivers. Because pine-centred buffers were generated from horizontal field positions, the mean horizontal differences of 2.6–3.0 cm support the suitability of the Emlid Reach RX for seedling-centred buffer generation and field reference data acquisition.

3.3. Dataset Composition and Seedling Characteristics

The complete living buffer dataset was dominated by small, weakly developed pine individuals. Across all 1416 living buffers, median seedling height was 17.0 cm, with an interquartile range of 13.0–22.0 cm. Median mean crown diameter was 10.5 cm, with an interquartile range of 8.0–13.0 cm, and the median number of recorded branches was 0.0. The median vegetation fraction within pine-centred buffers was 0.136, indicating that spectrally retained vegetation occupied only a limited proportion of each buffer after the applied conditioning procedure.

The model-development subset showed very similar central tendencies to the full living buffer dataset, with a median height of 17.0 cm and median mean-crown diameter of 10.5 cm. The withheld evaluation subset also presented a median height of 16.0 cm and a median mean-crown diameter of 10.5 cm. These similarities indicate that the evaluation subset retained the main morphological characteristics of the plantation-line model-development domain. These similarities were consistent with the broad spatial distribution of both subsets across the study area, as shown in Figure 10. Heavy stress was recorded in 36.7% of the model-development subset and 38.1% of the withheld subset, while trunk or severe defoliation was recorded in 82.7% and 82.7% of records, respectively. These similarities, listed in Table 4, indicate that the withheld subset retained the general morphological and condition variability observed in the model-development subset.

Table 4. Morphological and condition characteristics of the living pine-centred buffer dataset and analytical subsets. Continuous variables are reported as median [interquartile range].

Variable	Full Living-Buffer Dataset	Model-Development Subset	Withheld-Evaluation Subset
<i>n</i>	1416	889	527
Height [cm]	17.0 [13.0–22.0]	17.0 [13.0–22.0]	16.0 [12.0–21.0]
Mean crown diameter [cm]	10.5 [8.0–13.0]	10.5 [8.0–13.0]	10.5 [8.0–13.5]
Number of branches	0.0 [0.0–2.0]	0.0 [0.0–2.0]	0.0 [0.0–2.0]
Pine pixels within buffer	34.0 [18.0–70.0]	35.0 [19.0–69.8]	33.0 [16.0–71.0]
Pine fraction within buffer	0.14 [0.08–0.21]	0.14 [0.09–0.21]	0.13 [0.08–0.20]
Heavy stress	523 (36.9%)	326 (36.7%)	201 (38.1%)
Trunk or severe defoliation	1159 (81.9%)	735 (82.7%)	430 (81.6%)

The comparison between the model development subset and the withheld evaluation subset showed only minor differences in the analysed variables. The withheld subset

had a slightly lower median height, median number of pine pixels within buffers, and median pine fraction, but height, crown diameter, branch number, heavy stress, and severe defoliation indicators remained highly comparable. This supports the use of the withheld subset as an operationally relevant evaluation set for object-level seedling retrieval.

3.4. Point-Cloud Training Data and Model Configurations

The labelled point cloud was converted into model-ready training data using block-based sampling. Model development followed an iterative configuration-testing procedure in which trained models produced under different parameter settings were applied to the complete point cloud and subsequently evaluated using the object-level workflow described in Section 2.5.3. This procedure was used to assess how training parameters affected seedling retrieval and commission-related responses. The objective was not to identify a universally optimal parameter set, but to retain the configuration that provided the highest practical recall for the pine-seedling class while avoiding excessive classification of spectrally similar background vegetation.

For the baseline configuration, the 0.6 m block size provided the most suitable spatial unit for representing the small seedling structures observed in the field while retaining immediate local context. The observed point-count distributions showed that the maximum number of points per block remained below approximately 12,000, with most training and internal validation blocks concentrated between approximately 3000 and 8000 points (Figure 12). Based on this empirical distribution, the block point limit was set to 16,000 points. This value provided a margin above the observed maximum and avoided unnecessary subsampling of dense blocks, thereby preserving the native point density of the photogrammetric reconstruction within each training sample.




Figure 12. Point-count distributions of the 0.6-m point-cloud blocks generated during training-data preparation: (a) training subset and (b) internal validation subset. The x-axis represents the number of points per block, and the y-axis represents the corresponding number of blocks.

Exploratory testing of batch size showed that this parameter influenced model behaviour beyond its role in controlling GPU memory use. While hardware constraints defined the feasible range of batch sizes, the tested values also affected the model's sensitivity to weak seedling signals and the degree of shrub-related misclassification. Smaller batch sizes tended to increase sensitivity to weak or marginal seedling responses but also produced more frequent misclassification of spectrally similar shrub and herbaceous vegetation as pine seedling. Larger batch sizes produced more conservative and stable classifications but tended to suppress weaker seedling detections. The batch size of six blocks, therefore, represented the best practical balance observed during model development, maintaining

high object-level seedling retrieval while limiting shrub-related overprediction relative to more sensitive configurations.

Model input consisted of three-dimensional point coordinates and the selected radiometric point attributes for the green, red, red-edge, and near-infrared bands. During model development, internal validation recall for the pine-seedling class was evaluated alongside validation precision and F1 score, while object-level recall and off-buffer detections were assessed after completing point-cloud classification. This combination of internal validation metrics and object-level evaluation enabled comparisons of configurations in terms of both point-level learning behaviour and operational seedling retrieval.

Following baseline classification and evaluation, a refined model configuration was developed to reduce residual shrub-related commission responses. Rather than being trained as an entirely independent model, the refined model used the trained baseline model as the starting checkpoint for further training. This retraining step incorporated visually identified hard-negative shrub/background areas and broader background exposure within the plantation-line domain. The refined configuration retained the same general model-development structure and was evaluated using the same withheld plantation-line buffers as the baseline model. Its purpose was to test whether targeted retraining from the high-recall baseline model could reduce off-buffer false-positive detections while preserving as much object-level seedling recall as possible. The main differences between the baseline and refined model configurations are summarised in Table 5.

Table 5. Summary of baseline and refined model configurations.

Component	Baseline Model	Refined Model
Main objective	Maximise seedling retrieval	Reduce off-buffer commission detections
Positive class	Class 64, pine seedling	Class 64, pine seedling
Background class	Class 0	Class 0, with a targeted hard-negative background
Model input	XYZ coordinates + green, red, red-edge, and NIR attributes	XYZ coordinates + green, red, red-edge, and NIR attributes
Training domain	Model-development subset	Model-development subset with plantation-line/hard-negative refinement
Block size	0.6 m	1.0 m
Model selection emphasis	Validation recall, with precision/F1-score inspection	Omission–commission balance after hard-negative refinement
Evaluation subset	527 withheld plantation-line buffers	Same 527 withheld plantation-line buffers

3.5. Object-Level Seedling Retrieval Within Plantation-Line Buffers

Object-level seedling detection was assessed using the 527 withheld pine-centred buffers located within plantation-line areas. A seedling was considered detected when at least three points classified as pine seedling (Class 64) were found inside its corresponding buffer. Buffers containing fewer than three Class 64 points were treated as omission errors (2.5.3).

At the operational threshold of ≥ 3 classified points per buffer, the baseline model detected 463 of the 527 withheld plantation-line buffers, corresponding to a recall of 0.88. The refined model detected 444 buffers at the same threshold, corresponding to a recall of 0.84. Detection performance decreased progressively as the point-count threshold became stricter. For the baseline model, recall decreased from 0.89 at the ≥ 1 -point threshold to 0.83 at the ≥ 10 -point threshold. For the refined model, recall decreased from 0.86 to 0.80 across the same threshold range. This response indicates that a proportion of the detected

seedlings was represented by relatively sparse Class-64 point counts within the reference buffers, consistent with the small size and limited structural expression of the target-plantation seedlings. Object-level detection performance across the evaluated point-count thresholds is summarised in Table 6.

Table 6. Object-level detection performance of the baseline and refined models within the plantation-line withheld buffers across point-count thresholds.

Model	Detection Threshold	Reference Buffers	TP	FN	Recall
Baseline	≥ 1 point	527	471	56	0.89
Baseline	≥ 3 point		463	64	0.88
Baseline	≥ 5 point		458	69	0.87
Baseline	≥ 10 point		438	89	0.83
Refined	≥ 1 point	527	454	73	0.86
Refined	≥ 3 point		444	83	0.84
Refined	≥ 5 point		438	89	0.83
Refined	≥ 10 point		422	105	0.80

The threshold comparison shows that the baseline model consistently retained more detections than the refined model at the lower point-count thresholds. However, the difference between models narrowed as the threshold increased, with both models converging at the ≥ 5 point threshold. This pattern indicates that the refined model primarily suppressed weaker or more marginal detections, whereas seedlings, as represented by stronger Class-64 point counts, were more consistently retrieved by both models. Commission-screening results are reported separately because off-buffer detections were evaluated using visual interpretation rather than the pine-centred buffer retrieval framework.

3.6. Off-Buffer Commission Screening

Commission-error behaviour was assessed by visually screening Class 64 detections located outside the known pine-centred reference buffers but within the 1.2 m-wide plantation-line evaluation domain. This analysis was performed separately from the object-level retrieval assessment because off-buffer Class 64 detections may correspond either to false-positive vegetation responses or to pine individuals not represented by the reference-buffer dataset. Figure 13 shows the location of candidate off-buffer detections across the analytical area and provides a close-up example of overlapping and model-specific Class 64 responses within a plantation-line buffer.

Baseline model produced 243 off-buffer detections within plantation-line areas. Of these, 218 were classified as false positives, 14 were classified as true positives in the commission-screening context, and 11 remained uncertain after visual interpretation. False positives represented 89.7% of the screened baseline-model candidate detections, while visually confirmed pine detections outside the reference-buffer dataset represented 5.8%. The remaining 4.5% were retained as uncertain cases.

The refined model produced 41 off-buffer detections within the same plantation-line evaluation domain. Of these, 37 were classified as false positives, two were classified as true positives in the commission-screening context, and two remained uncertain. False positives, therefore, accounted for 90.2% of the screened refined-model candidates, while true positives and uncertain detections each accounted for 4.9%.

Overall, the refined model reduced the total number of off-buffer detections from 243 to 41, corresponding to an 83.1% reduction relative to the baseline model. The number of visually confirmed false positives decreased from 218 to 37, corresponding to

an 83.0% reduction. Of the 41 refined model candidate detections, 29 overlapped with those identified in the baseline model, indicating that most of the remaining refined model off-buffer detections corresponded to persistent candidate responses shared by both models. These results indicate that the refined model did not represent a universal improvement over the baseline model. Instead, it shifted model behaviour towards a more conservative response, reducing off-buffer candidate detections and visually confirmed false positives while increasing omission errors among known withheld seedlings. The relative changes in off-buffer screening outcomes between the baseline and refined models are summarised in Table 7.




Figure 13. Spatial distribution of off-buffer candidate detections within the plantation-line evaluation domain, displayed over the multispectral orthomosaic. (a) Overview of baseline and refined model off-buffer detections across the analytical area. (b) Close-up example showing baseline and refined model Class 64 points outside pine buffers and within the plantation-line buffer.

Table 7. Relative change in candidate off-buffer detections from baseline to refined model.

Screening Outcome	Baseline Model	Refined Model	Absolute Change	Relative Reduction
False Positives	218	37	−181	83%
Valid pine detections in commission-screening context	14	2	−12	86%
Uncertain detections	11	2	−9	82%
Total Detections	243	41	−202	83%

3.7. False-Negative Characterisation and Detectability Limits

At the operational threshold of ($T = 3$), the baseline model omitted 64 of the 527 withheld plantation-line buffers, while the refined model omitted 83. Most false negatives were complete omissions rather than marginal detections below the operational threshold. In the baseline model, 56 of the 64 false-negative buffers contained no Class 64 points, while eight contained only one or two predicted Class 64 points. In the refined model, 73 of the 83 false-negative buffers contained no Class 64 points, while 10 contained one or two predicted Class 64 points.

To identify the seedling characteristics associated with omission errors, detected and omitted records were compared within the withheld plantation-line evaluation subset at the operational threshold of ($T = 3$). This analysis was performed separately for the

baseline and refined models, enabling comparison of their detectability patterns within the same reference-buffer framework. Continuous variables—seedling height, mean crown diameter, number of branches, and vegetation fraction within the buffer—were summarised using medians and interquartile ranges (IQRs) and compared using Mann–Whitney U tests. Binary traits (height ≤ 20 cm, mean crown diameter ≤ 10 cm, absence of recorded branches, vegetation fraction ≤ 0.15 , heavy stress, and trunk or severe defoliation) were analysed using Fisher’s exact tests. The statistical comparison between detected and omitted seedlings for both model configurations is summarised in Table 8.

Table 8. Statistical comparison between detected and omitted seedlings within the withheld plantation-line evaluation subset at the operational threshold of $T = 3$. Continuous variables are reported as median [interquartile range]. Cliff’s delta (δ) is reported with positive values, indicating higher values in detected seedlings. For binary variables, the odds ratio (OR) represents the odds of omission associated with the presence of the trait.

Variable	Baseline Detected ($n = 463$)	Baseline Omitted ($n = 64$)	Baseline Test Result	Refined Detected ($n = 444$)	Refined Omitted ($n = 83$)	Refined Test Result
Height [cm]	17.0 [12.0–21.5]	15.5 [11.0–18.0]	$p = 0.057$; $\delta = 0.147$	17.0 [12.0–22.0]	16.0 [11.0–18.0]	$p = 0.011$; $\delta = 0.175$
Mean crown diameter [cm]	11.0 [8.5–13.5]	7.5 [6.0–9.1]	$p < 0.001$; $\delta = 0.543$	11.0 [8.5–14.0]	7.5 [6.0–9.5]	$p < 0.001$; $\delta = 0.599$
Number of branches	0.0 [0.0–2.0]	0.0 [0.0–0.0]	$p < 0.001$; $\delta = 0.264$	0.0 [0.0–2.0]	0.0 [0.0–0.0]	$p < 0.001$; $\delta = 0.257$
Vegetation fraction within buffer	0.133 [0.078–0.203]	0.114 [0.067–0.175]	$p = 0.274$; $\delta = 0.084$	0.131 [0.077–0.197]	0.139 [0.066–0.197]	$p = 0.983$; $\delta = 0.001$
Height ≤ 20 cm	324/463 (70.0%)	53/64 (82.8%)	OR = 2.07; $p = 0.038$	306/444 (68.9%)	71/83 (85.5%)	OR = 2.67; $p = 0.002$
Mean crown diameter ≤ 10 cm	193/463 (41.7%)	51/64 (79.7%)	OR = 5.49; $p < 0.001$	177/444 (39.9%)	67/83 (80.7%)	OR = 6.32; $p < 0.001$
No recorded branches	290/463 (62.6%)	57/64 (89.1%)	OR = 4.86; $p < 0.001$	275/444 (61.9%)	72/83 (86.7%)	OR = 4.02; $p < 0.001$
Low vegetation fraction ≤ 0.15	264/463 (57.0%)	39/64 (60.9%)	OR = 1.18; $p = 0.591$	259/444 (58.3%)	44/83 (53.0%)	OR = 0.81; $p = 0.398$
Heavy stress	174/463 (37.6%)	27/64 (42.2%)	OR = 1.21; $p = 0.494$	170/444 (38.3%)	31/83 (37.3%)	OR = 0.96; $p = 0.903$
Trunk or severe defoliation	391/463 (84.4%)	39/64 (60.9%)	OR = 0.29; $p < 0.001$	372/444 (83.8%)	58/83 (69.9%)	OR = 0.45; $p = 0.005$

The statistical comparison between detected and omitted seedlings showed that omission errors were primarily associated with reduced horizontal crown development and limited branch structure. For both models, omitted seedlings had smaller mean crown diameters than detected seedlings. The median crown diameter of omitted seedlings was 7.5 cm in both the baseline and refined models, compared with 11.0 cm for detected seedlings. This difference was associated with large effect sizes and strong statistical support in both configurations. Seedlings with a mean crown diameter ≤ 10 cm had 5.49 times higher odds of omission in the baseline model and 6.32 times higher odds in the refined model.

Branch development also influenced detectability. Branchless seedlings are nearly five times more likely to be omitted by the baseline model than those with branches (OR = 4.86) and four times in the refined model (OR = 4.02). Height showed a weaker, less-consistent relationship with omission: it was not statistically significant in the baseline model and was significant in the refined model, with a small effect size. Vegetation fraction within the buffer and heavy-stress status did not differ significantly between detected and omitted seedlings in either model. Although trunk or severe defoliation differed significantly, the relationship was inverse: seedlings assigned to this category had lower omission odds.

Therefore, this variable was not interpreted as a primary cause of missed detections. Overall, the most consistent detectability constraints were small crown diameter and absence of recorded branch development.

4. Discussion

This study addresses a particularly demanding remote-sensing problem: the detection of recently planted *Pinus pinaster* seedlings that are small, structurally ill-defined and embedded within spectrally similar coastal-dune vegetation. Rather than evaluating tree detection in the context of well-developed crowns, the workflow tested whether centimetre-resolution multispectral photogrammetric point clouds could support object-level retrieval of early-stage seedlings under realistic plantation-monitoring conditions. The results indicate that transformer-based point-cloud classification can retrieve a high proportion of field-mapped seedlings within plantation-line buffers, but they also show that model performance is governed by a clear trade-off between sensitivity to weak seedling signals and suppression of shrub-related false-positive responses.

The main contribution of the study is therefore not only the achieved recall, but the characterisation of the operational limits of seedling detectability. The baseline model demonstrated the workflow's capacity to prioritise seedling retrieval, whereas the refined model showed that hard-negative retraining from the baseline checkpoint can substantially reduce off-buffer candidate detections, with increased omission of less evident seedlings. False-negative analysis further showed that missed detections were not randomly distributed across the seedling population but were mainly associated with reduced crown diameter and absence of recorded branch development. These findings indicate that, for very young *P. pinaster* seedlings, horizontal crown expression and local structural complexity are more important for UAS point-cloud detectability than height alone or buffer-level vegetation abundance.

The following discussion interprets these findings in relation to previous UAS-based seedling-detection studies, the quality and limitations of field-derived reference buffers, the omission–commission trade-off between the baseline and refined models, and the operational implications of using multispectral point-cloud deep learning for early-stage plantation monitoring.

4.1. Detection of Early-Stage Seedlings in Structurally Complex Coastal Vegetation

The detection task addressed in this study differs from many tree-detection applications focused on mature or well-developed crowns. The target population consisted mainly of recently planted seedlings, most of which were below 20 cm in height and had limited crown diameter and branch development. This represents a particularly demanding object-detection problem because the target individuals occupy only a small spatial footprint and may share spectral and structural characteristics with surrounding herbaceous and shrubby vegetation.

These results are consistent with previous studies identifying seedling detection as a persistent challenge in UAS-based forest monitoring. Feduck et al. [6], Jayathunga et al. [8], and Poudel and Bevilacqua [16] reported the value of UAS imagery for seedling mapping, but also highlighted limitations associated with small target size, background vegetation, and site-specific variability. Finn et al. [10] specifically emphasised the difficulty of identifying pine seedlings because of their small canopy footprint and spectral similarity to surrounding ground vegetation. The present study extends this line of work by showing that a multispectral point-cloud workflow can retrieve a high proportion of field-mapped seedlings under real plantation conditions, while also explicitly characterising the types of individuals that remain difficult to detect.

The use of multispectral point-cloud attributes is relevant in this context because the model can exploit both spatial neighbourhood relationships and reflectance information. Point Transformer V3 is designed to process irregular point-cloud data and model local spatial relationships using attention-based mechanisms [18,61]. In this study, that capability was relevant because seedlings were not represented as large, continuous objects but as small, sometimes fragmented clusters of vegetation points embedded in a complex background. However, considering the image pixel size used to produce the high-density point cloud, the results also show that even high-density multispectral point clouds have a practical detectability limit when the target crown is very small or poorly developed.

4.2. Field Reference Quality and Implications for Seedling-Centred Evaluation

Accurate field data were central to this workflow because model training and evaluation depended on pine-centred buffers derived from GNSS positions and field-measured crown dimensions. The positional uncertainty remains meaningful relative to the size of the target seedlings. Many target seedlings had mean crown diameters close to 10 cm or less. Therefore, even centimetre-level horizontal differences may influence the number of classified points falling inside a small buffer. The buffer design, which incorporated crown dimensions and an additional radial margin, was therefore essential to reducing the risk of valid seedling points being excluded solely due to the GNSS or photogrammetric accuracy. This is particularly important in early-stage seedling detection, where the difference between a true positive and a false negative may depend on only a few classified points inside a small object-centred buffer.

The object-level evaluation also provided a more operationally meaningful assessment than point-wise accuracy alone. Point-wise semantic accuracy can be misleading in strongly imbalanced point-cloud datasets, where the background class dominates the total number of points [57,59]. By evaluating whether each field-mapped seedling buffer contained a minimum number of predicted pine-seedling points, the analysis focused directly on whether individual seedlings could be retrieved. This aligns with the intended monitoring objective, in which managers need to know whether planted individuals are present or absent, rather than only whether individual points are correctly classified.

At the same time, the buffer-retrieval assessment primarily quantified omission error because all reference objects in the calculation were known seedling locations. For this reason, conventional precision and F1 score were not treated as primary performance indicators in the object-level buffer evaluation. Commission-error behaviour was instead assessed through a separate off-buffer screening procedure. This separation is methodologically important because off-buffer predictions can include both false-positive vegetation responses and valid pine detections that were absent from, or not correctly represented by, the field survey dataset.

4.3. Baseline and Refined Model Behaviour

The comparison between the baseline and refined models should be interpreted in the context of the iterative model-development procedure used in this study. The baseline model was not the result of a single predefined training run, but of repeated configuration testing in which trained models were applied to the complete point cloud and evaluated using the same object-level buffer-retrieval workflow. This procedure was necessary because the detection problem was highly imbalanced: pine-seedling points constituted a small minority class embedded within a dominant, spectrally complex background. In such conditions, model assessment based only on global point-wise accuracy can be misleading, and recall-oriented evaluation is often more informative for assessing rare-target retrieval [22,23].

The selected baseline configuration, therefore, reflected the best practical compromise observed during this iterative process for maximising seedling retrieval. The 0.6-m block size was retained because it was large enough to capture most seedling structures and their immediate spatial context, yet small enough to avoid excessive inclusion of adjacent background vegetation.

Batch size should therefore be not only a hardware-constrained parameter, but also a methodological factor influencing optimisation behaviour and model sensitivity. Although GPU memory defined the maximum feasible number of blocks processed per iteration, the selected batch size affected the stochasticity of gradient estimation, the stability of model convergence, and the balance between sensitivity to the minority pine seedling class and suppression of background responses [57–59]. In the present point-cloud workflow, smaller batch sizes appeared to increase sensitivity to weak or marginal seedling signals but also increased misclassification of spectrally similar shrub and herbaceous vegetation as pine seedling. Larger batch sizes produced more stable and conservative classifications but tended to suppress weaker seedling responses. A batch size of six blocks was therefore retained not simply because it was computationally feasible, but because it provided the most favourable practical balance observed during iterative model development between seedling retrieval and control of shrub-related overprediction.

The class-of-interest configuration further contributed to the baseline model behaviour. By explicitly defining Class 64 as the class of interest, the baseline model prioritised learning the minority pine-seedling class. This was appropriate and necessary for the first modelling stage because the operational priority was to minimise omission errors and to determine whether the multispectral point-cloud representation could retrieve field-mapped seedlings within plantation lines. However, this sensitivity also increased the risk that spectrally similar vegetation would be classified as pine seedling, particularly where shrub or herbaceous vegetation presented local reflectance and structural patterns similar to small pine crowns. The baseline model, therefore, achieved the highest object-level recall, detecting 463 of the 527 withheld plantation-line buffers at the operational threshold, but also produced a larger number of off-buffer candidate detections.

The refined model was developed in response to this residual commission-related behaviour. Rather than representing an independent alternative model, it was generated through additional training from the trained baseline checkpoint. This strategy preserved the seedling-detection capability learned by the baseline model while introducing additional information about difficult background cases. Visually identified shrub/background regions from the baseline output were used as hard-negative examples, and broader background exposure was introduced within the plantation-line domain. Hard-negative mining and related refinement strategies are commonly used in detection workflows to expose the model to difficult non-target examples that were previously confused with the target class [60]. In this study, the same principle was applied to reduce shrub-related pine-seedling responses in the point-cloud classification output.

The behavioural change between models is therefore methodologically coherent. The refined model reduced total off-buffer candidate detections from 243 to 41 and visually confirmed false positives from 218 to 37, indicating that hard-negative retraining substantially constrained the model response to shrub-like background vegetation. However, this also reduced object-level recall from 88% to 84%, resulting in 19 additional omitted seedlings compared with the baseline model. This indicates that some weak but valid seedling detections shared spectral–structural characteristics with the hard-negative background examples introduced during refinement. Consequently, the refined model became more conservative: it suppressed many false-positive shrub responses, but also removed some marginal seedling responses that the baseline model had retained.

This reduction in off-buffer false-positive burden should not be described as conventional precision in the strict object-detection sense, because precision was not calculated from a complete set of explicit non-pine reference objects. Instead, the refined model improved operational specificity within the plantation-line domain by reducing the number of candidate detections requiring visual inspection or field verification. The comparison between models reflects an omission–commission trade-off rather than a simple improvement or degradation in model quality. If the operational objective is to maximise seedling retrieval and minimise missed seedlings, the baseline model is preferable. If the objective is to reduce false-positive screening effort while retaining most known seedlings, the refined model is more appropriate.

The commission-screening results also show why off-buffer detections should not automatically be interpreted as false positives. Both models detected visually identifiable pine individuals outside the field buffer dataset, suggesting that some off-buffer Class 64 clusters may represent unreferenced seedlings or natural regeneration. This reinforces the methodological decision to separate object-level recall assessment from off-buffer commission screening. It also shows that model outputs can be used not only as final classifications, but also as diagnostic layers for identifying potential gaps in the field inventory and for guiding targeted verification in complex plantation environments.

4.4. Detectability Limits and Characteristics of False-Negative Seedlings

The false-negative analysis provided insight into the field conditions under which the workflow failed to retrieve known seedlings. Across both baseline and refined models, omission errors were primarily associated with reduced horizontal crown development and limited branch structure. Omitted seedlings had a median mean crown diameter of 7.5 cm in both models, compared with 11.0 cm for detected seedlings. Seedlings with a mean crown diameter ≤ 10 cm had substantially higher odds of omission in both the baseline and refined models. Similarly, branchless seedlings had markedly higher omission odds in both configurations. This result indicates that the practical detectability limit was driven mainly by the degree of crown expression represented in the multispectral point cloud. Seedlings with small crowns and no recorded branches probably generated fewer distinctive point clusters and less consistent spectral–spatial patterns, making them harder to distinguish from sand, litter, herbaceous vegetation, or low shrub components. This is consistent with the ecological and remote-sensing context of young *Pinus pinaster* plantations in coastal dunes, where seedlings may be sparse, partially stressed, and embedded in spectrally similar ground vegetation.

The association between crown diameter and omission should be interpreted with caution because the evaluation buffer radius was derived from measured crown diameter (2.4.1). Seedlings with larger crowns were represented by larger buffer areas and therefore had a greater spatial opportunity to contain the minimum number of predicted Class 64 points required for detection. However, because buffer radii were derived from field-measured crown dimensions, this evaluation effect is intrinsically linked to actual seedling size and to the weaker structural and spectral–spatial representation expected for smaller seedlings in a photogrammetric point cloud. Consequently, the crown-diameter effect partly reflects the adopted buffer-based retrieval rule and cannot be interpreted solely as an independent biophysical limitation of the model. Nevertheless, the consistent association between omission and both smaller crown-derived buffer areas and absence of recorded branch development indicates that weak above-ground structural expression constrained seedling retrieval in this dataset.

Seedling height exhibited a weaker and less-consistent relationship with omission errors. It was not statistically significant in the baseline model but significant in the

refined model. This suggests that height alone was less informative than horizontal crown development and branch structure. In very young seedlings, vertical expression may be too limited to produce a reliable photogrammetric or point-cloud signal, especially when the surrounding vegetation and microtopography introduce similar local height variability. Vegetation fraction within the buffer and heavy-stress status did not differ significantly between detected and omitted seedlings in either model, indicating that overall buffer-level vegetation abundance and broad stress classification were not the main drivers of omission. Although trunk exposure or severe defoliation differed significantly, the relationship was inverse: highly defoliated seedlings had lower omission odds. Therefore, this condition was associated with enhanced detection rather than missed records, meaning this trait cannot be considered a driver of missed records within this dataset.

The comparison between models further supports the interpretation that omission was not random. Many false negatives were complete omissions, with no predicted Class-64 points within the corresponding buffer, rather than marginal detections that fell just below the operational threshold. Lowering the detection threshold would therefore recover only a small proportion of omitted seedlings. The additional omissions introduced by the refined model are likely due to its more conservative response after hard-negative retraining, which suppressed some weaker detections and shrub-related false positives. This reinforces the interpretation that detectability depends on the interaction between seedling morphology, background complexity, and model sensitivity.

4.5. Contribution of Multispectral Point-Cloud Classification

The workflow presented here contributes to seedling-monitoring research by combining centimetre-resolution multispectral UAS data, field-derived object buffers, point-cloud semantic labelling, and transformer-based classification. Most UAS-based seedling-detection studies rely primarily on RGB or multispectral raster imagery, object-based image analysis, convolutional neural networks, or spectral-spatial image processing [6,8,10,16]. In contrast, this study used radiometric-calibrated multispectral point cloud as the deep-learning input, allowing the model to operate on irregular three-dimensional point structures with associated spectral attributes.

This representation is relevant because early-stage seedlings may not be adequately represented as continuous image objects. By preserving point-level geometry and spectral attributes, the multispectral point cloud provides a way to model small local structures and their surrounding context. The results suggest that this representation can support high seedling retrieval under operational plantation conditions, particularly when evaluation is performed at the object-level using field-derived reference buffers.

The iterative model-development process also provides useful methodological insight. The retained baseline configuration resulted from repeated evaluation of trained models generated under different parameter settings, including block size, batch size, and class-of-interest configuration. The 0.6-m block size and batch size of six represented the best practical baseline configuration among the tested alternatives, balancing seedling sensitivity with control of shrub-related misclassification. The refined model then used the baseline model as a starting checkpoint and introduced hard-negative background samples to reduce commission-related responses. This sequence is relevant for operational remote-sensing applications because it shows that parameter tuning and model refinement should be interpreted in relation to the monitoring objective, rather than as a search for a universally optimal model.

The results also highlight the importance of spatially constrained evaluation. By focusing object-level recall on withheld buffers within plantation-line areas, the evaluation addressed the operational question of whether planted seedlings could be retrieved where

they were expected to occur. At the same time, off-buffer screening within the plantation-line domain provided a complementary assessment of commission-related responses. This distinction is particularly important in managed plantations, where planting geometry provides strong prior information, but natural regeneration, shrubs, and residual vegetation may still produce ambiguous detections.

4.6. Operational Implications for Plantation Monitoring

From an operational perspective, the workflow is directly relevant to early-stage plantation assessment. A baseline recall of 0.88 indicates that the model can retrieve a high proportion of field-mapped seedlings within plantation-line buffers. Such a configuration may be useful when the cost of missing seedlings is high, for example when estimating survival, identifying planting gaps, or planning targeted field verification. However, the higher number of off-buffer candidate detections produced by the baseline model would require additional visual filtering or field inspection.

The refined model offers a more conservative alternative. Although object-level recall decreased to 0.84, the number of off-buffer candidate detections decreased by more than 80%. This reduction may be valuable in operational contexts where the number of false alerts must be controlled and where field verification resources are limited. In practical terms, the refined model may support more efficient field verification by reducing the number of candidate locations requiring inspection, while still retrieving most known plantation-line seedlings.

The choice between models should therefore depend on the monitoring objective. A high-recall configuration is more appropriate for broad candidate identification and omission-sensitive monitoring. A refined, more conservative configuration is more appropriate when the main objective is to reduce false-positive burden and prioritise a smaller set of candidate detections for verification. In both cases, the workflow should be interpreted as a decision-support approach rather than a complete replacement for field assessment. Its greatest operational value may lie in directing field effort towards likely seedling locations, suspected omissions, ambiguous detections and areas where plantation establishment is spatially uneven.

4.7. Limitations and Future Work

Several limitations should be considered when interpreting these results.

First, the workflow was evaluated within a single coastal plantation and under the acquisition conditions of one UAS survey. Although the dataset included substantial local heterogeneity, broader validation across sites, seasons, plantation ages, and background vegetation conditions is required to assess transferability. Seasonal variation may be particularly important because pine seedlings and surrounding vegetation can show different phenological trajectories, potentially improving or reducing spectral separability depending on acquisition timing [14,15].

Second, the model-development process involved iterative evaluation of alternative trained configurations using the withheld plantation-line buffer framework. This was appropriate for identifying the most useful configuration for the present operational objective, but the resulting performance should be interpreted as an object-level evaluation within this study domain rather than as a fully independent external validation. The 527 evaluation seedlings were withheld from model training and were not introduced as either positive or negative training examples. Although they remained physically present in the dense point cloud as Class 0, they were excluded from the exported model-training samples through the class-of-interest setting in the baseline model and through filter blocks by class code set to Class 64 during refined-model training-data preparation. Future work

should include an additional independent test area or a separate acquisition date to more robustly assess generalisation.

Third, the evaluation relied on pine-centred buffers derived from field measurements and GNSS positions. Although receiver comparison, independent checkpoint assessment, and photogrammetric control support centimetre-level accuracy, the target seedlings were extremely small. Minor positional offsets, irregular crown geometry, or imperfect buffer placement may still influence point-count-based detection outcomes. In addition, because buffer radii were derived from measured crown diameter, seedlings with larger crowns were evaluated using larger search areas. The association between crown diameter and omission may therefore partly reflect the adopted buffer-based retrieval rule, in addition to true differences in seedling detectability. However, because buffer radius was based on field-measured crown dimensions, this effect is also linked to actual seedling size and to the reduced number of reconstructed vegetation points expected for smaller, weakly developed seedlings in a photogrammetric point cloud. This limitation is inherent to object-level evaluation of very small vegetation targets and should be addressed in future work through repeated field positioning, refined crown delineation, fixed-radius sensitivity analysis, centre-distance matching, probabilistic buffer designs or direct object-level crown annotations.

Fourth, the commission-error assessment was based on off-buffer candidate screening within the plantation-line domain, rather than on a complete set of explicit non-pine reference objects across the landscape. This approach was suitable for characterising operational false-positive responses, but does not provide a conventional landscape-wide precision estimate. Future studies should incorporate explicit negative-reference samples, systematic object annotations of non-pine objects, or independent field verification of candidate detections to more fully quantify commission error.

Fifth, although the refined model substantially reduced off-buffer detections, it also increased omission errors among known seedlings. Future refinements should explore approaches that reduce false positives without suppressing weak but valid seedling detections. Potential strategies include active learning using visually verified commission detections, more targeted hard-negative sampling, multi-date inputs to exploit phenological differences, integration of higher-resolution RGB-derived texture and object-level post-processing based on plantation-line geometry, and local crown morphology.

Finally, the workflow was developed in a commercial GIS and deep-learning environment. This improves accessibility for practitioners using ArcGIS Pro but may limit direct reproducibility in open-source environments. Future work could replicate the workflow using open-source point-cloud deep-learning frameworks and compare model behaviour across point-based and hybrid raster–point-cloud approaches.

Taken together, these limitations indicate that the proposed workflow should be interpreted as an operationally promising but context-dependent approach for early-stage seedling monitoring. Its performance depends not only on model architecture, but also on field reference quality, seedling morphology, plantation geometry, background vegetation complexity, point-cloud density, reconstruction quality, buffer geometry, and the selected balance between omission and commission errors.

5. Conclusions

This study demonstrated that UAS-derived radiometrically calibrated multispectral photogrammetric point clouds combined with Point Transformer V3 can support object-level detection of early-stage *Pinus pinaster* seedlings in a heterogeneous coastal-dune plantation. The proposed workflow integrated RTK-based field mapping, seedling-centred buffer construction, spectral conditioning, point-cloud semantic labelling, iterative

model development, and object-level evaluation of withheld field-mapped seedlings within plantation-line areas.

The baseline model achieved high seedling retrieval, detecting 463 of 527 withheld plantation-line buffers at the operational threshold of at least three predicted pine-seedling points per buffer, corresponding to a recall of 88%. This confirmed the multispectral point-cloud approach's capacity to retrieve small, recently planted seedlings under complex field conditions. However, the baseline model also produced an elevated rate of off-buffer detections, reflecting its greater sensitivity to weak seedling-like signals and spectrally similar background vegetation.

The refined model, produced through additional training on the baseline model, introduced hard-negative shrub/background examples and broader background exposure, reducing off-buffer candidate detections from 243 to 41 and visually confirmed false positives from 218 to 37. This represented an 83% reduction in the operational false-positive burden within the plantation-line domain. The reduction was achieved at the cost of lower object-level recall, with the refined model detecting 444 of 527 withheld buffers, corresponding to 84% recall. These results indicate that hard-negative refinement can effectively constrain shrub-related commission responses but may also suppress marginal seedling detections.

False-negative analysis showed that omission errors were mainly associated with reduced crown-level structural expression. Omitted seedlings had smaller mean crown diameters and were more frequently branchless than detected seedlings, whereas height, buffer-level vegetation fraction, and heavy-stress status were not consistent explanatory factors. This indicates that, for very young *P. pinaster* seedlings, horizontal crown development and branch structure are more relevant to UAS point-cloud detectability than height, surrounding vegetation density, or health status. This suggests that the model relies primarily on the horizontal spatial footprint and branching architecture to identify targets, making its detection capability robust against variations in seedling height, background vegetation density, and stress-induced changes in foliage.

Overall, the study supports the use of multispectral point-cloud deep learning as a decision-support tool for early-stage plantation monitoring. The workflow is particularly relevant where planted seedlings are expected along defined plantation lines and where model outputs can be used to prioritise field verification. The results also show that model configuration should be selected according to the monitoring objective: the baseline model is preferable when minimising omission errors is the priority, whereas the refined model is preferable when reducing false-positive verification effort is more important. Future work should test the workflow across additional sites, acquisition dates, seedling age classes and seasonal vegetation conditions, and should include independent validation areas and explicit negative reference samples to further quantify generalisation and commission-error behaviour.

Author Contributions: Conceptualisation, T.v.d.W.d.S. and L.G.P.; methodology, T.v.d.W.d.S. and L.G.P.; software, T.v.d.W.d.S.; validation, T.v.d.W.d.S. and L.G.P.; formal analysis, T.v.d.W.d.S.; investigation, T.v.d.W.d.S. and L.G.P.; resources, T.v.d.W.d.S. and L.G.P.; data curation, T.v.d.W.d.S.; writing—original draft preparation, T.v.d.W.d.S.; writing—review and editing, T.v.d.W.d.S. and L.G.P.; visualisation, T.v.d.W.d.S.; supervision, L.G.P.; project administration, L.G.P.; funding acquisition, L.G.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by national funds (OE) through FCT—Fundação para a Ciência e Tecnologia, I.P. under the PhD studentship 2023.01767.BDANA awarded to Tiago van der Worp da Silva. This work was also funded by national funds through FCT—Fundação para a Ciência e a Tecnologia, I.P., under the project CESAM—Centre for Environmental and Marine Studies,

references UID/50017/2025 (<https://doi.org/10.54499/UID/50017/2025>) and LA/P/0094/2020 (<https://doi.org/10.54499/LA/P/0094/2020>).

Data Availability Statement: Additional data, including the full raw UAS image dataset, the full dense point-cloud dataset, and photogrammetric project files, may be made available from the corresponding author upon reasonable request, subject to data-use and storage limitations.

Acknowledgments: The authors acknowledge the institutional and technical support provided by the University of Aveiro and CESAM (Centre for Environmental and Marine Studies). The use of the ReNEP GNSS infrastructure provided by the Portuguese Directorate-General for Territory is also acknowledged. The authors further acknowledge GeoLayer Geoengenharia e Serviços, Lda., for providing access to and operational support for the DJI Mavic 3 Multispectral platform used for UAS data acquisition. During the preparation of this manuscript, the authors used ChatGPT, GPT-5.5 Thinking, developed by OpenAI, for AI-assisted copyediting and language refinement. The authors reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
CMOS	Complementary Metal–Oxide–Semiconductor
DBSCAN	Density-Based Spatial Clustering of Applications with Noise
DGT	Directorate-General for Territory
DSM	Digital Surface Model
DTM	Digital Terrain Model
ETRS89	European Terrestrial Reference System 1989
F1 score	Harmonic mean of precision and recall
FN	False Negative
GCP	Ground Control Point
GIS	Geographic Information System
GNSS	Global Navigation Satellite System
GPU	Graphics Processing Unit
ITRF93	International Terrestrial Reference Frame 1993
LAS	LAS point-cloud file format
M3M	DJI Mavic 3 Multispectral
MP	Megapixel
NDVI	Normalised Difference Vegetation Index
NIR	Near-Infrared
OR	Odds Ratio
PT-TM06	Portugal Transverse Mercator 2006
ReNEP	Rede Nacional de Estações Permanentes GNSS
RGB	Red–Green–Blue
RINEX	Receiver Independent Exchange Format
RMSE	Root Mean Square Error
RTK	Real-Time Kinematic
SfM	Structure-from-Motion
TP	True Positive
UAS	Unmanned Aerial System
UAV	Unmanned Aerial Vehicle
XYZ	Three-Dimensional Cartesian Coordinates

References

1. Ecke, S.; Stehr, F.; Dempewolf, J.; Frey, J.; Klemmt, H.-J.; Seifert, T.; Tiede, D. Species-Specific Machine Learning Models for UAV-Based Forest Health Monitoring: Revealing the Importance of the BNDVI. *Int. J. Appl. Earth Obs. Geoinf.* **2024**, *135*, 104257. [\[CrossRef\]](#)
2. Abreu-Dias, R.; Santos-Gago, J.M.; Martín-Rodríguez, F.; Álvarez-Sabucedo, L.M. Advances in the Automated Identification of Individual Tree Species: A Systematic Review of Drone- and AI-Based Methods in Forest Environments. *Technologies* **2025**, *13*, 187. [\[CrossRef\]](#)
3. Estrada, J.S.; Fuentes, A.; Reszka, P.; Auat Cheein, F. Machine Learning Assisted Remote Forestry Health Assessment: A Comprehensive State of the Art Review. *Front. Plant Sci.* **2023**, *14*, 1139232. [\[CrossRef\]](#) [\[PubMed\]](#)
4. Sandino, J.; Pegg, G.S.; Gonzalez, F.; Smith, G.R. Aerial Mapping of Forests Affected by Pathogens Using UAVs, Hyperspectral Sensors, and Artificial Intelligence. *Sensors* **2018**, *18*, 944. [\[CrossRef\]](#) [\[PubMed\]](#)
5. Guimarães, N.; Pádua, L.; Marques, P.; Silva, N.; Peres, E.; Sousa, J.J. Forestry Remote Sensing from Unmanned Aerial Vehicles: A Review Focusing on the Data, Processing and Potentialities. *Remote Sens.* **2020**, *12*, 1046. [\[CrossRef\]](#)
6. Feduck, C.; McDermid, G.J.; Castilla, G. Detection of Coniferous Seedlings in UAV Imagery. *Forests* **2018**, *9*, 432. [\[CrossRef\]](#)
7. Mohan, M.; Richardson, G.; Gopan, G.; Aghai, M.M.; Bajaj, S.; Galgamuwa, G.A.P.; Vastaranta, M.; Arachchige, P.S.P.; Amorós, L.; Corte, A.P.D.; et al. UAV-Supported Forest Regeneration: Current Trends, Challenges and Implications. *Remote Sens.* **2021**, *13*, 2596. [\[CrossRef\]](#)
8. Jayathunga, S.; Pearse, G.D.; Watt, M.S. Unsupervised Methodology for Large-Scale Tree Seedling Mapping in Diverse Forestry Settings Using UAV-Based RGB Imagery. *Remote Sens.* **2023**, *15*, 5276. [\[CrossRef\]](#)
9. Pérez-Carrasco, M.; Karelavic, B.; Molina, R.; Saavedra, R.; Cerulo, P.; Cabrera-Vives, G. Precision Silviculture: Use of UAVs and Comparison of Deep Learning Models for the Identification and Segmentation of Tree Crowns in Pine Crops. *Int. J. Digit. Earth* **2022**, *15*, 2223–2238. [\[CrossRef\]](#)
10. Finn, A.; Kumar, P.; Peters, S.; O'Hehir, J. Unsupervised Spectral-Spatial Processing of Drone Imagery for Identification of Pine Seedlings. *ISPRS J. Photogramm. Remote Sens.* **2022**, *183*, 363–388. [\[CrossRef\]](#)
11. Stamatopoulos, I.; Le, T.C.; Daver, F. UAV-Assisted Seeding and Monitoring of Reforestation Sites: A Review. *Aust. For.* **2024**, *87*, 90–98. [\[CrossRef\]](#)
12. Da Silva, T.V.D.W.; Gomes Pereira, L.; Oliveira, B.R.F. Assessing Geometric and Radiometric Accuracy of DJI P4 MS Imagery Processed with Agisoft Metashape for Shrubland Mapping. *Remote Sens.* **2024**, *16*, 4633. [\[CrossRef\]](#)
13. Huo, L.; Matsiakh, I.; Bohlin, J.; Cleary, M. Estimation of Tree Vitality Reduced by Pine Needle Disease Using Multispectral Drone Images. *Remote Sens.* **2025**, *17*, 271. [\[CrossRef\]](#)
14. Kleinsmann, J.; Verbesselt, J.; Kooistra, L. Monitoring Individual Tree Phenology in a Multi-Species Forest Using High Resolution UAV Images. *Remote Sens.* **2023**, *15*, 3599. [\[CrossRef\]](#)
15. Singh, K.K.; Surasinghe, T.D.; Frazier, A.E. Systematic Review and Best Practices for Drone Remote Sensing of Invasive Plants. *Methods Ecol. Evol.* **2024**, *15*, 998–1015. [\[CrossRef\]](#)
16. Poudel, A.; Bevilacqua, E. Young White Pine Detection Using UAV Imagery and Deep Learning Object Detection Models. *Sensors* **2026**, *26*, 1284. [\[CrossRef\]](#) [\[PubMed\]](#)
17. Poudel, A.; Bevilacqua, E. Assessing Red Pine Seedlings Using Uav Point Clouds and Field-Verified Data. *Int. Arch. Photogramm. Remote Sens. Spat. Inf. Sci.* **2022**, XLVI-M-2–2022, 173–176. [\[CrossRef\]](#)
18. Wu, X.; Jiang, L.; Wang, P.-S.; Liu, Z.; Liu, X.; Qiao, Y.; Ouyang, W.; He, T.; Zhao, H. Point Transformer V3: Simpler, Faster, Stronger. *arXiv* **2024**, arXiv:2312.10035.
19. Rodrigues, D.; Corticeiro, S.; Maia, P. Short Communication: Maritime Pine Natural Regeneration in Coastal Central Portugal: Effects of the Understory Composition. *For. Syst.* **2022**, *31*, eSC06. [\[CrossRef\]](#)
20. Danielsen, R. Palaeoecological Development of the Quiaios–Mira Dunes, Northern-Central Littoral Portugal. *Rev. Palaeobot. Palynol.* **2008**, *152*, 74–99. [\[CrossRef\]](#)
21. Zhao, H.; Jiang, L.; Jia, J.; Torr, P.H.S.; Koltun, V. Point Transformer. *arXiv* **2021**, arXiv:2012.09164.
22. He, H.; Garcia, E.A. Learning from Imbalanced Data. *IEEE Trans. Knowl. Data Eng.* **2009**, *21*, 1263–1284. [\[CrossRef\]](#)
23. Saito, T.; Rehmsmeier, M. The Precision-Recall Plot Is More Informative than the ROC Plot When Evaluating Binary Classifiers on Imbalanced Datasets. *PLoS ONE* **2015**, *10*, e0118432. [\[CrossRef\]](#) [\[PubMed\]](#)
24. Rodrigues, D.; Da Silva, T.V.D.W.; Maia, P.; Oliveira, B.R.F. Comparative Analysis of Native and Exotic Invasive Species in Post-Fire Ecosystem Recovery and Carbon Cycling. *Forests* **2026**, *17*, 535. [\[CrossRef\]](#)
25. Instituto da Conservação da Natureza e das Florestas (ICNF). *Plano de Gestão Florestal—Mata Nacional das Dunas de Quiaios 2020–2038*; ICNF: Lisbon, Portugal, 2020.
26. ReNEP. Available online: <https://resep.dgterritorio.gov.pt/> (accessed on 6 February 2026).
27. Lee, T.; Bettinger, P.; Merry, K.; Cieszewski, C. The Effects of Nearby Trees on the Positional Accuracy of GNSS Receivers in a Forest Environment. *PLoS ONE* **2023**, *18*, e0283090. [\[CrossRef\]](#) [\[PubMed\]](#)

28. Tomašík, J.; Mokroš, M.; Saloň, Š.; Chudý, F.; Tunák, D. Accuracy of Photogrammetric UAV-Based Point Clouds under Conditions of Partially-Open Forest Canopy. *Forests* **2017**, *8*, 151. [\[CrossRef\]](#)
29. Sefercik, U.G.; Kavzoglu, T.; Colkesen, I.; Adali, S.; Dinc, S.; Nazar, M.; Ozturk, M.Y. Land Cover Classification Performance of Multispectral Rtk UAVs. *Int. Arch. Photogramm. Remote Sens. Spat. Inf. Sci.* **2021**, *XLVI-4-W5-2021*, 489–492. [\[CrossRef\]](#)
30. Dlamini, S.M.; Ouma, Y.O. Large-Scale Topographic Mapping Using RTK-GNSS and Multispectral UAV Drone Photogrammetric Surveys: Comparative Evaluation of Experimental Results. *Geomatics* **2025**, *5*, 25. [\[CrossRef\]](#)
31. Dash, J.P.; Pearse, G.D.; Watt, M.S. UAV Multispectral Imagery Can Complement Satellite Data for Monitoring Forest Health. *Remote Sens.* **2018**, *10*, 1216. [\[CrossRef\]](#)
32. Pan, W.; Wang, X.; Sun, Y.; Wang, J.; Li, Y.; Li, S. Karst Vegetation Coverage Detection Using UAV Multispectral Vegetation Indices and Machine Learning Algorithm. *Plant Methods* **2023**, *19*, 7. [\[CrossRef\]](#) [\[PubMed\]](#)
33. Duffy, J.P.; Shutler, J.D.; Witt, M.J.; DeBell, L.; Anderson, K. Tracking Fine-Scale Structural Changes in Coastal Dune Morphology Using Kite Aerial Photography and Uncertainty-Assessed Structure-from-Motion Photogrammetry. *Remote Sens.* **2018**, *10*, 1494. [\[CrossRef\]](#)
34. Mancini, F.; Dubbini, M.; Gattelli, M.; Stecchi, F.; Fabbri, S.; Gabbianelli, G. Using Unmanned Aerial Vehicles (UAV) for High-Resolution Reconstruction of Topography: The Structure from Motion Approach on Coastal Environments. *Remote Sens.* **2013**, *5*, 6880–6898. [\[CrossRef\]](#)
35. Taddia, Y.; Corbau, C.; Zambello, E.; Pellegrinelli, A. UAVs for Structure-From-Motion Coastal Monitoring: A Case Study to Assess the Evolution of Embryo Dunes over a Two-Year Time Frame in the Po River Delta, Italy. *Sensors* **2019**, *19*, 1717. [\[CrossRef\]](#) [\[PubMed\]](#)
36. Wouter, H. Maes Practical Guidelines for Performing UAV Mapping Flights with Snapshot Sensors. *Remote Sens.* **2025**, *17*, 606. [\[CrossRef\]](#)
37. Meteostat Meteostat Data Access. Available online: <https://dev.meteostat.net/data> (accessed on 24 February 2026).
38. Aasen, H.; Honkavaara, E.; Lucieer, A.; Zarco-Tejada, P.J. Quantitative Remote Sensing at Ultra-High Resolution with UAV Spectroscopy: A Review of Sensor Technology, Measurement Procedures, and Data Correction Workflows. *Remote Sens.* **2018**, *10*, 1091. [\[CrossRef\]](#)
39. Wang, C. At-Sensor Radiometric Correction of a Multispectral Camera (RedEdge) for sUAS Vegetation Mapping. *Sensors* **2021**, *21*, 8224. [\[CrossRef\]](#) [\[PubMed\]](#)
40. Iglhaut, J.; Cabo, C.; Puliti, S.; Piermattei, L.; O'Connor, J.; Rosette, J. Structure from Motion Photogrammetry in Forestry: A Review. *Curr. For. Rep.* **2019**, *5*, 155–168. [\[CrossRef\]](#)
41. Tonkin, T.N.; Midgley, N.G. Ground-Control Networks for Image Based Surface Reconstruction: An Investigation of Optimum Survey Designs Using UAV Derived Imagery and Structure-from-Motion Photogrammetry. *Remote Sens.* **2016**, *8*, 786. [\[CrossRef\]](#)
42. Sanz-Ablanedo, E.; Chandler, J.H.; Rodríguez-Pérez, J.R.; Ordóñez, C. Accuracy of Unmanned Aerial Vehicle (UAV) and SfM Photogrammetry Survey as a Function of the Number and Location of Ground Control Points Used. *Remote Sens.* **2018**, *10*, 1606. [\[CrossRef\]](#)
43. Stott, E.; Williams, R.D.; Hoey, T.B. Ground Control Point Distribution for Accurate Kilometre-Scale Topographic Mapping Using an RTK-GNSS Unmanned Aerial Vehicle and SfM Photogrammetry. *Drones* **2020**, *4*, 55. [\[CrossRef\]](#)
44. Mamaghani, B.; Salvaggio, C. Multispectral Sensor Calibration and Characterization for sUAS Remote Sensing. *Sensors* **2019**, *19*, 4453. [\[CrossRef\]](#) [\[PubMed\]](#)
45. Žižala, D.; Minařík, R.; Zádorová, T. Soil Organic Carbon Mapping Using Multispectral Remote Sensing Data: Prediction Ability of Data with Different Spatial and Spectral Resolutions. *Remote Sens.* **2019**, *11*, 2947. [\[CrossRef\]](#)
46. Dandois, J.P.; Ellis, E.C. High Spatial Resolution Three-Dimensional Mapping of Vegetation Spectral Dynamics Using Computer Vision. *Remote Sens. Environ.* **2013**, *136*, 259–276. [\[CrossRef\]](#)
47. Tucker, C.J. Red and Photographic Infrared Linear Combinations for Monitoring Vegetation. *Remote Sens. Environ.* **1979**, *8*, 127–150. [\[CrossRef\]](#)
48. Rommel, E.; Giese, L.; Fricke, K.; Kathöfer, F.; Heuner, M.; Mölter, T.; Deffert, P.; Asgari, M.; Nätke, P.; Dzunic, F.; et al. Very High-Resolution Imagery and Machine Learning for Detailed Mapping of Riparian Vegetation and Substrate Types. *Remote Sens.* **2022**, *14*, 954. [\[CrossRef\]](#)
49. Kuhn, M.; Johnson, K. *Applied Predictive Modeling*; Springer Science & Business Media: Berlin/Heidelberg, Germany, 2013; ISBN 978-1-4614-6849-3.
50. Roberts, D.R.; Bahn, V.; Ciuti, S.; Boyce, M.S.; Elith, J.; Guillera-Arroita, G.; Hauenstein, S.; Lahoz-Monfort, J.J.; Schröder, B.; Thuiller, W.; et al. Cross-Validation Strategies for Data with Temporal, Spatial, Hierarchical, or Phylogenetic Structure. *Ecography* **2017**, *40*, 913–929. [\[CrossRef\]](#)
51. Wang, Y.; Khodadadzadeh, M.; Zurita-Milla, R. Spatial+: A New Cross-Validation Method to Evaluate Geospatial Machine Learning Models. *Int. J. Appl. Earth Obs. Geoinf.* **2023**, *121*, 103364. [\[CrossRef\]](#)

52. ArcGIS Pro Documentation Prepare Point Cloud Training Data (3D Analyst Tools) | ArcGIS Pro Documentation. Available online: <https://doc.esri.com/en/arcgis-pro/latest/tool-reference/3d-analyst/prepare-point-cloud-training-data.html> (accessed on 24 February 2026).
53. ArcGIS Pro Documentation Train Point Cloud Classification Model (3D Analyst Tools) | ArcGIS Pro Documentation. Available online: <https://doc.esri.com/en/arcgis-pro/latest/tool-reference/3d-analyst/train-point-cloud-classification-model.html?tabs=dialog> (accessed on 24 June 2026).
54. Guo, Y.; Wang, H.; Hu, Q.; Liu, H.; Liu, L.; Bennamoun, M. Deep Learning for 3D Point Clouds: A Survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **2021**, *43*, 4338–4364. [[CrossRef](#)] [[PubMed](#)]
55. Buda, M.; Maki, A.; Mazurowski, M.A. A Systematic Study of the Class Imbalance Problem in Convolutional Neural Networks. *Neural Netw.* **2018**, *106*, 249–259. [[CrossRef](#)] [[PubMed](#)]
56. Goodfellow, I.; Bengio, Y.; Courville, A. *Deep Learning*; MIT Press: Cambridge, MA, USA, 2016.
57. Keskar, N.S.; Mudigere, D.; Nocedal, J.; Smelyanskiy, M.; Tang, P.T.P. On Large-Batch Training for Deep Learning: Generalization Gap and Sharp Minima. *arXiv* **2017**, arXiv:1609.04836.
58. Smith, S.L.; Le, Q.V. A Bayesian Perspective on Generalization and Stochastic Gradient Descent. *arXiv* **2018**, arXiv:1710.06451.
59. Smith, L.N. A Disciplined Approach to Neural Network Hyper-Parameters: Part 1—Learning Rate, Batch Size, Momentum, and Weight Decay. *arXiv* **2018**, arXiv:1803.09820.
60. Tang, T.; Zhou, S.; Deng, Z.; Zou, H.; Lei, L. Vehicle Detection in Aerial Images Based on Region Convolutional Neural Networks and Hard Negative Example Mining. *Sensors* **2017**, *17*, 336. [[CrossRef](#)] [[PubMed](#)]
61. Zhu, X.X.; Tuia, D.; Mou, L.; Xia, G.-S.; Zhang, L.; Xu, F.; Fraundorfer, F. Deep Learning in Remote Sensing: A Comprehensive Review and List of Resources. *IEEE Geosci. Remote Sens. Mag.* **2017**, *5*, 8–36. [[CrossRef](#)]

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