

Article

ForestSplat: Proof-of-Concept for a Scalable and High-Fidelity Forestry Mapping Tool Using 3D Gaussian Splatting

Belal Shaheen¹, Matthew David Zane¹, Bach-Thuan Bui^{1,2} , Shubham¹, Tianyuan Huang³ , Manuel Merello⁴, Ben Scheelk⁴, Steve Crooks^{1,5}  and Michael Wu^{1,*}

¹ Coolant, San Francisco, CA 94111, USA

² Graduate School of Information Science and Engineering, Ritsumeikan University, 2-150 Iwakuracho, Ibaraki 567-8570, Osaka, Japan

³ Department of Civil and Environmental Engineering, Stanford University, 473 Via Ortega, Stanford, CA 94305, USA

⁴ The Ocean Foundation, Washington, DC 20036, USA

⁵ Silvestrum Climate Associates, Sausalito, CA 94965, USA

* Correspondence: michael@coolant.earth

Abstract: Accurate, scalable forestry insights are critical for implementing carbon credit-based reforestation initiatives and data-driven ecosystem management. However, existing forest quantification methods face significant challenges: hand measurement is labor-intensive, time-consuming, and difficult to trust; satellite imagery is not accurate enough; and airborne LiDAR remains prohibitively expensive at scale. In this work, we introduce ForestSplat: an accurate and scalable reforestation monitoring, reporting, and verification (MRV) system built from consumer-grade drone footage and 3D Gaussian Splatting. To evaluate the performance of our approach, we map and reconstruct a 200-acre mangrove restoration project in the Jobos Bay National Estuarine Research Reserve. ForestSplat produces an average mean absolute error (MAE) of 0.17 m and mean error (ME) of 0.007 m compared to canopy height maps derived from airborne LiDAR scans, using 100× cheaper hardware. We hope that our proposed framework can support the advancement of accurate and scalable forestry modeling with consumer-grade drones and computer vision, facilitating a new gold standard for reforestation MRV.



check for updates

Academic Editor: Brenden E. McNeil

Received: 1 February 2025

Revised: 1 March 2025

Accepted: 10 March 2025

Published: 12 March 2025

Citation: Shaheen, B.; Zane, M.D.; Bui, B.-T.; Shubham; Huang, T.; Merello, M.; Scheelk, B.; Crooks, S.; Wu, M. ForestSplat: Proof-of-Concept for a Scalable and High-Fidelity Forestry Mapping Tool Using 3D Gaussian Splatting. *Remote Sens.* **2025**, *17*, 993. <https://doi.org/10.3390/rs17060993>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Keywords: canopy height; Gaussian Splatting; MRV; ForestSplat

1. Introduction

Reforestation is widely recognized as one of the most important tools in combating climate change [1], offering numerous benefits such as carbon sequestration and restoring biodiversity for native species. However, legacy tools for quantifying reforestation efforts face significant challenges in terms of accuracy, cost and scalability, holding back capital investments and preventing reforestation from playing a larger role in regulating our climate. Specifically, traditional field work for forest inventories to obtain tree positions and diameter at breast height (DBH) measurements is labor-intensive and time-consuming [2–5]. Moreover, field samples today are measured using simple tools like tape measures, calipers, and clinometers, all of which are susceptible to human error and lead to inaccuracies in collected data [6]. Traditional fieldwork, being labor-intensive, is also constrained in its ability to cover large areas. This limitation forces forestry projects to rely on sampling small regions, which increases the likelihood of measurement errors [7].

A recent NASA project, the Global Ecosystem Dynamics Investigation (GEDI), provides a repeated canopy height map for the first time, collecting relative heights and

canopy heights at a resolution of 25 m [8]. The Ice, Cloud, and land Elevation Satellite-2 (ICESat-2) collects canopy height measurements with a 13×100 m footprint [9]. To improve the accuracy of spaceborne LiDAR data from GEDI or ICESat-2, several projects have proposed using multi-sensor fusion to enhance aboveground biomass mapping [10], including integration with optical imagery from Sentinel-2 [11]. Additionally, two recent global vegetation height maps are freely available at 10-m and 30-m spatial resolutions, derived from Sentinel-2 [12] and Landsat [13], respectively. However, while these maps are useful for scientific applications, they remain limited for local forestry applications that require very high accuracy and resolution due to their relatively coarse resolution and large uncertainties.

On the other hand, studies that apply state-of-the-art deep network models directly to satellite imagery to estimate forest details, such as canopy height and above-ground carbon stocks in forest ecosystems [14–17], are becoming increasingly popular. Mugabowindekwe et al. applied deep learning techniques to nationwide tree-level carbon stock estimation in Rwanda, leveraging high-resolution imagery for precise mapping [14]. Other studies have explored different methodologies for carbon stock estimation, including the integration of multi-modal satellite time-series data [16] and the use of self-supervised vision transformers trained on aerial LiDAR to generate high-resolution canopy height maps from RGB satellite images [17]. However, these models are still limited by satellite resolution. Even high-resolution commercial satellites such as WorldView-3 typically provide ground sample distances (GSDs) of about 0.3–3 m. In addition, publicly available datasets, including Landsat and Sentinel-2, have even lower spatial resolutions, with GSDs of 30 and 20 m, respectively, [15]. Satellite imagery is also predominantly captured from near-vertical perspectives, which limits the accuracy of canopy height estimation due to occlusion and the overlap of tree crowns [14]. Moreover, weather conditions, such as cloud cover and atmospheric interference, pose significant challenges for optical satellite imagery [18,19]. These limitations hinder the ability to capture detailed forest features such as crown size and canopy height [20]. The issues are particularly critical in reforestation settings, where measuring year-over-year changes in carbon stock is essential for generating carbon credits, yet satellite-based models often struggle to perform accurately.

Airborne LiDAR is considered the gold standard and ground truth for many forestry applications [19]. Previous studies demonstrate the effectiveness of airborne LiDAR for 3D reconstruction in forestry mapping, canopy height calculations, and Digital Terrain Maps (DTMs)—all of which are highly important for reforestation projects [21–23]. However, one key challenge is the high operational cost associated with LiDAR surveys, which require specialized aircraft and sensors. This makes frequent large-scale monitoring financially unfeasible, particularly for developing regions and conservation projects with limited budgets. Additionally, the high costs of equipment and logistics further constrain the scalability of LiDAR deployment, rendering it impractical for most forestry projects [24].

Three-dimensional Gaussian Splatting [25] and its scene reconstruction capabilities have shown promise in various applications such as robotics and autonomous systems, reproducing texture more faithfully compared to traditional methods like LiDAR and point cloud techniques [26–29]. In this project, we present a novel reforestation MRV system integrating large-scale 3D Gaussian Splatting. Our framework relies solely on camera footage from consumer-grade inexpensive hardware such as DJI Mini drones, bridging the current gap between inaccurate hand measurements and expensive LiDAR scans, enabling large-scale, cost-effective, and accurate measurement for forestry.

Technically, we propose a comprehensive and cost-effective pipeline for accurate and scalable forestry modeling using only high-resolution images captured by a consumer-grade drone. For Gaussian Splatting, we introduce a simple yet effective method to address

training challenges in large-scale environments by integrating a neural-agnostic scaffold densification strategy with a lightweight partitioning process. Additionally, we present an efficient approach for estimating canopy height maps from Multi-View Stereo point clouds and 3D Gaussian models.

Furthermore, we open-source a large-scale forestry dataset covering 200 acres, which includes LiDAR data and 13,657 drone images. We believe this dataset will serve as a valuable benchmark and comparison baseline for advancing MRV methodologies using Gaussian Splatting or other computer vision-based approaches.

2. Data

In partnership with the Ocean Foundation and the Jobos Bay National Estuarine Research Reserve in Puerto Rico, we conducted three months of field data collection starting in late 2023. To serve as ground truth for our study, we collected 200 acres of airborne LiDAR scans using a DJI Zenmuse L1 and DJI Matrice 300 RTK. The UAV was flown two months after the LiDAR data collection at a height of 30 m and a speed of 2 m/s, parameters which were determined to be optimal in prior experiments. To build our Gaussian Splatting models, we photographed our sites with RGB camera footage from a consumer-grade DJI Mini 2. The Mini 2 was flown at a height of 60 m and a speed of 5 m/s, with footage captured at two angles: 90° and 60°. This resulted in 13,657 drone images for the survey area with an image-side overlap of about 30%. It is also worth mentioning that the proposed dataset includes images captured at a 60° pitch angle, which potentially provides a better view of tree crowns and helps mitigate occlusions caused by near-vertical perspectives commonly found in satellite imagery. We show a satellite visualization of this area in Figure 1.



Figure 1. Two-hundred-acre reforestation study plot within the Jobos Bay Mangrove Forest, Puerto Rico. The red dashed line outlines the region of interest (ROI). Data were collected in November 2023.

3. ForestSplat Pipeline

We introduce ForestSplat, a large-scale Gaussian Splatting-based pipeline designed for mapping vegetation in reforestation projects. ForestSplat leverages a novel combination of Structure from Motion (SfM), Multi-View Stereo (MVS), and 3D Gaussian Splatting (GS) to enhance the fidelity of modeled trees and successfully generate canopy height maps (CHMs) from GS models. Figure 2 illustrates the overall design of our approach.

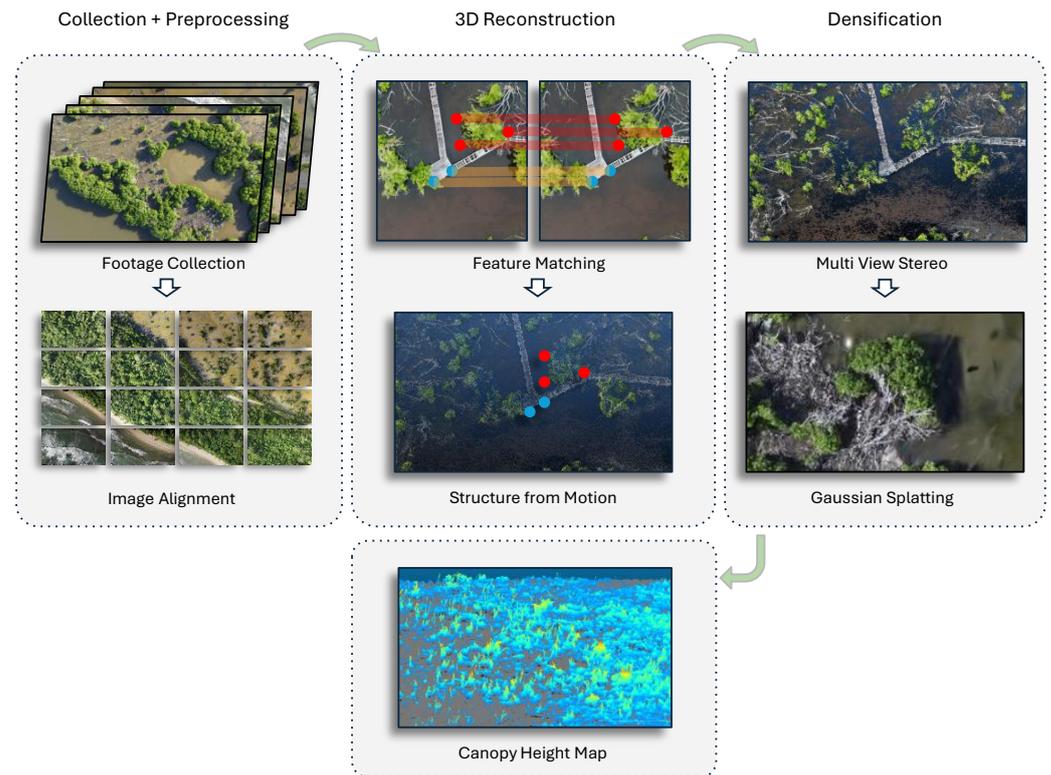


Figure 2. ForestSplat Pipeline Overview. *Collection and Preprocessing:* Sample drone imagery is captured by the DJI Mini 2 during aerial surveys, followed by preprocessing steps including image rotation correction and alignment. *Three-dimensional Reconstruction:* Feature detection and matching are performed to establish correspondences, forming the foundation for a 3D sparse point cloud via Structure from Motion. *Densification:* The sparse 3D model is further densified using Multi-View Stereo and enhanced with 3D Gaussian Splatting. *Canopy Height Map Generation:* Finally, a canopy height map is generated, where gray represents the ground level, and colors indicate varying canopy heights.

3.1. Problem Statement

Given a set of N aerial images $\{I_1, I_2, \dots, I_N\}$, the objective is to derive a high-resolution canopy height map representing relative vegetation height to the ground. This process involves reconstructing a sparse 3D model (Section 3.3) by leveraging SfM to estimate sparse 3D points $\{X_j\}$ and camera poses $\{P_i\}$, densifying the reconstruction using MVS (Section 3.5), and fitting Gaussian distributions to the dense point cloud through GS to obtain \mathcal{G}_j map (Section 3.6). The final step involves extracting height values from Gaussian splats and differentiating canopy and ground levels to produce the CHM (Section 3.7). The overall pipeline is summarized as follows:

$$I_1, I_2, \dots, I_N \xrightarrow{\text{SfM}} \{X_j\} \xrightarrow{\text{MVS}} \{X_j^{\text{dense}}\} \xrightarrow{\text{GS}} \mathcal{G}_j \xrightarrow{\text{CHM}} \text{CHM}(u, v) \quad (1)$$

3.2. Preprocessing

Image Pairs Generation

Initially, SfM requires a set of image pairs to guide feature matching. However, generating image pairs exhaustively would be inefficient, as it creates $\frac{N(N-1)}{2}$ image pairs. This is impractical for thousands of images. Instead, we leverage GNSS coordinates for more efficient image pair generation, focusing on images with high overlap. This approach

reduces unnecessary computations while still maintaining a highly accurate 3D sparse model. Specifically, we calculate the ground footprint size of each image as follows:

$$G_w = \frac{AS_w}{f}; G_h = \frac{AS_h}{f} \quad (2)$$

where A is the altitude, and S_w and S_h are the sensor width and height of the camera, respectively. From this, we derive the 2D corner coordinates of the initial ground footprint F_{init} as:

$$F_{init} = \left\{ \left(-\frac{G_w}{2}, \frac{G_h}{2} \right), \left(\frac{G_w}{2}, \frac{G_h}{2} \right), \left(\frac{G_w}{2}, -\frac{G_h}{2} \right), \left(-\frac{G_w}{2}, -\frac{G_h}{2} \right) \right\}. \quad (3)$$

We then calculate the exact ground coordinates for each frame by $F'_{final} = R \cdot F'_{init} + t'$ where $'$ indicates the homogeneous coordinates, $R = R_z(yaw)R_y(pitch)R_x(roll)$ and $t' = (E, N, 0)^T$ is the UTM coordinates. Since the ground footprints are represented as polygons, we calculate the intersection over union as $IoU = \frac{Area_{inter}}{Area_{union}}$ and use this to represent the score value of each image pair.

3.3. Structure from Motion

Now, given a limited number of robust image pairs $M \ll \frac{N(N-1)}{2}$ generated using the technique presented in the previous section, we aim to determine the camera extrinsics $[R_i, t_i]$ for each image i , where $R_i \in SO(3)$ and $t_i \in \mathbb{R}^3$, as well as a set of sparse 3D coordinates $\{X_j\} \in \mathbb{R}^3$ corresponding to feature j .

For the forestry-based dataset, we observed that sparse keypoints present significant challenges for achieving accurate matching between image pairs. To address this issue, we leverage semi-dense features for local feature matching of each pair. Specifically, we use TopicFM [30] to extract 2D–2D correspondences. Due to the snake-pattern flight path, consecutive flight legs result in alternate images being flipped by 180° or rotated by 90° . To improve feature matching accuracy during the Structure from Motion (SfM) process, we rotate these images to maintain a consistent orientation.

For global applicability to any rotation angle present in the image pairs, we calculate a global yaw angle θ that appears most frequently across the dataset. The rotation needed for each image i is then determined as $\Delta\theta_i = \theta - \theta_i$.

After completing all matching steps for the aligned images, we reapply the inverse transformation to the matched pixels j using $-\Delta\theta_i$ to restore the original orientation.

Additionally, we observed that the traditional incremental SfM pipeline [31] requires a significant amount of time to fully reconstruct 13K images. This is because incremental methods attempt to reconstruct the scene starting from two views and then sequentially register new camera images along with their associated 3D structures. To reduce processing time, we instead leverage a global SfM pipeline for this sparse 3D reconstruction. Specifically, we use Glomap [32], a method that combines camera positioning, bundle adjustment, and structure refinement into a single global positioning step. This approach reduces the 3D sparse reconstruction time on our large dataset from several days to just a couple of hours.

3.4. Transformation to World Coordinates

Sparse 3D models reconstructed using (SfM) are inherently unscaled [32], as the camera poses and 3D scene coordinates do not align with a real-world coordinate system. To address this limitation, we propose a simple method for estimating a transformation model that maps SfM-derived coordinates to real-world coordinates using noisy GNSS data.

Given $\mathbf{x}_i \in \mathbb{R}^3$, the camera position in the original coordinate system of the SfM model, and $\mathbf{x}'_i \in \mathbb{R}^3$, the corresponding camera position in the real-world coordinate

system derived from GNSS data, we aim to estimate a similarity transformation $T \in \mathbb{R}^{4 \times 4}$. The transformation T is represented as:

$$T = \begin{bmatrix} sR & \mathbf{t} \\ \mathbf{0}^\top & 1 \end{bmatrix}, \quad (4)$$

where s is a scaling factor, $R \in \mathbb{R}^{3 \times 3}$ is a rotation matrix, and $\mathbf{t} \in \mathbb{R}^3$ is the translation vector. The transformation from \mathbf{x}_i to \mathbf{x}'_i can be expressed as:

$$\mathbf{x}'_i = T(\mathbf{x}_i) = sR\mathbf{x}_i + \mathbf{t}. \quad (5)$$

To obtain a robust estimation of T , we leverage the RANSAC algorithm [33] to minimize the following objective function:

$$\arg \min_{s,R,\mathbf{t}} \sum_{i \in \text{inliers}} \|\mathbf{x}'_i - (sR\mathbf{x}_i + \mathbf{t})\|^2. \quad (6)$$

In comparison to traditional photogrammetry methods, such as those employed by Pix4D software [34], which typically require multiple ground control points (GCPs) to establish a proper scale and coordinate alignment, our approach eliminates this dependency. This significantly enhances the level of automation in the reconstruction pipeline.

3.5. Multi-View Stereo

Following SfM, we densify sparse 3D SfM point clouds using Multi-View Stereo (MVS) algorithms. Specifically, we use ET-MVSNet (Enhanced Texture Multi-View Stereo), which leverages enhanced texture information to generate high-fidelity, dense point clouds [35]. This method significantly improves the detail and provides a more comprehensive representation of the surveyed environment, forming a better foundation for subsequent splatting.

In detail, given the camera poses $\{R_i, t_i\}$, intrinsic matrix K , and images $\{I_i\}$, as well as sparse 3D coordinates $\{X_j\}^S$ obtained from a previous SfM step, the goal is to densify the sparse model from $\{X_j\}^S$ to a denser representation $\{X_j^{\text{dense}}\}^D$, where $D \gg S$.

Specifically, given a reference image i and $N - 1$ source images, the method first extracts feature representations for all images using a Feature Pyramid Network (FPN) integrated with an Epipolar Transformer module at the coarsest resolution. These enhanced features are then propagated through subsequent layers of the pipeline.

Next, the cost volumes are constructed by warping source image pixels into the reference camera frustum and measuring their similarity. These feature volumes are aggregated to construct a 3D cost volume for each depth hypothesis of the reference image. Finally, a 3D CNN is applied to the cost volume for regularization, enabling the inference of the most likely depth hypothesis for each pixel in the reference image.

Leveraging this, each reference image i now can have a depth map D_i . In the subsequent Gaussian Splatting step, this depth map is further utilized as a robust loss function for training a dense and robust 3D model.

3.6. Gaussian Splatting

We develop a custom large-scale Gaussian Splatting model called ForestSplat that is built from the *gsplat* [36] framework and incorporates elements from the implementation of Level of Gaussians [36,37]. Gaussian splats are especially adept at modeling tree features like leaves and branches, and allows us to further enrich the models developed from MVS with pixel-perfect 3D reconstructions. Our approach captures fine-grained details and accurate texture representations, which are crucial for precise carbon stock estimation

and tree dimension measurements. It also ensures efficient processing and rendering of large-scale datasets, making it suitable for extensive reforestation projects.

Preliminary, 3D-GS [36] represents a scene using a set of anisotropic 3D Gaussians, denoted as $\mathcal{G}_j = (X_j, \Sigma_j, \alpha_j)$, where $X_j \in \mathbb{R}^3$ is the 3D position, typically initialized from Structure from Motion (SfM) models. $\Sigma_j \in \mathbb{R}^{3 \times 3}$ is the covariance matrix of the 3D Gaussian, encoding the scale and orientation in 3D space, while α_j represents its opacity, which is used for rendering and pruning. The Gaussian splats are projected onto the image plane during rendering, where gradients in 2D image space are leveraged for optimization.

In this paper, instead of initializing X_j from $\{X_j\}^{\text{SfM}}$, we initialize the Gaussian means X_j from the results of a prior coarse densification process, denoted as $\{X_j^{\text{dense}}\}^{\text{MVS}}$. This approach significantly enhances the density and quality of the surface Gaussian model \mathcal{G}_j , which is critical for generating a high-accuracy and high-resolution CHM model. In addition, we also transform the initial 3D points and camera poses to a world coordinate system using T estimated in Section 3.4.

Inspired from GS-scaffold [38], our work introduces a neural-agnostic scaffold densification strategy that enhances scene representation without directly relying on neural features. Specifically, we aim to address challenges in dense scene coverage by introducing gradient-driven anchor growing and structured pruning methods. This later avoids reliance on computationally expensive neural predictions while maintaining high-quality scene representation. We present these techniques in the following sub-sections.

3.6.1. Gradient-Driven Anchor Growing

In the proposed approach, new Gaussians are added adaptively based on 2D image-plane gradients, allowing for a more geometrically informed densification process. Each Gaussian splat \mathcal{G}_j contributes a normalized 2D image-plane gradient, denoted as $\nabla I_{2D}(\mathcal{G}_j)$. Anchor candidates are selected based on a threshold τ_g , forming the growth mask:

$$\mathcal{M}_{\text{grow}} = \{j \mid \|\nabla I_{2D}(\mathcal{G}_j)\| \geq \tau_g\}. \quad (7)$$

Similar to [38], an anchor is treated as the center of each voxel that represent M point cloud $\mathbf{X} \in \mathbb{R}^{M \times 3}$ as:

$$\mathbf{V} = \left\{ \left\lfloor \frac{\mathbf{X}}{\epsilon} \right\rfloor \right\} \cdot \epsilon, \quad (8)$$

where $\mathbf{V} \in \mathbb{R}^{N \times 3}$ denotes voxel centers, and ϵ is the voxel size.

To prevent over-densification, candidate positions are quantized into a voxel grid of size ϵ , ensuring spatial consistency. Unique positions within the voxel grid are retained, and new anchors are initialized with appropriate scales and opacities. This geometric-only approach allows efficient expansion of the Gaussian set, even in scenarios where neural features are unavailable.

3.6.2. Structured Pruning

To maintain computational efficiency and avoid overgrowth, low-opacity and geometrically inconsistent Gaussians are pruned periodically. The pruning process uses three criteria:

1. **Opacity Threshold:** Gaussians with opacity α_j below a threshold τ_o are removed:

$$\mathcal{M}_{\text{prune, opa}} = \{j \mid \alpha_j < \tau_o\}. \quad (9)$$

2. **Scale Constraint:** Gaussians with overly large scales (as determined by the eigenvalues of Σ_j) are pruned:

$$\mathcal{M}_{\text{prune, scale}} = \{j \mid \lambda_{\max}(\Sigma_j) > \tau_s\}. \quad (10)$$

3. Geometric Bounds: Gaussians outside the scene's vertical bounds $[z_{\min}, z_{\max}]$ are removed:

$$\mathcal{M}_{\text{prune, geo}} = \{j \mid Z_j < z_{\min} \text{ or } Z_j > z_{\max}\}. \quad (11)$$

The final pruning mask is the union of these individual masks:

$$\mathcal{M}_{\text{prune}} = \mathcal{M}_{\text{prune, opa}} \cup \mathcal{M}_{\text{prune, scale}} \cup \mathcal{M}_{\text{prune, geo}}. \quad (12)$$

3.6.3. Gradient-Driven Updates

To refine the representation over time, we track the running average of the gradient norms for each Gaussian:

$$\text{grad2d}_j = \frac{\sum_{t=1}^T \|\nabla I_{2D}(\mathcal{G}_j)\|}{\max(1, \text{count}_j)},$$

where T is the current epoch, and count_j is the number of frames in which the Gaussian \mathcal{G}_j is visible. This allows us to prioritize highly relevant regions for future densification or refinement.

3.6.4. Loss Function

To train the 3D Gaussian models we use the following loss function:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{recon}} + \lambda_{\text{scale}} \mathcal{L}_{\text{scale}} + \lambda_{\text{opa}} \mathcal{L}_{\text{opacity}} + \lambda_{\text{depth}} \mathcal{L}_{\text{depth}}, \quad (13)$$

where λ is the balance coefficient for each loss function. In detail, these loss functions are defined as:

1. Reconstruction Loss: The reconstruction loss ensures that the rendered colors I_{rendered} match the ground truth I_{gt} :

$$\mathcal{L}_{\text{recon}} = (1 - \lambda_{\text{ssim}}) \|I_{\text{rendered}} - I_{\text{gt}}\|_1 + \lambda_{\text{ssim}} \cdot (1 - \text{SSIM}(I_{\text{rendered}}, I_{\text{gt}})), \quad (14)$$

where λ_{ssim} balances the L1 loss and the Structural Similarity Index Measure (SSIM).

2. Scale Regularization: The scale regularization penalizes excessively large scales of Gaussians:

$$\mathcal{L}_{\text{scale}} = \frac{1}{N} \sum_{j=1}^N \|\exp(\mathbf{s}_j)\|, \quad (15)$$

where \mathbf{s}_j represents the logarithmic scale parameters of Gaussian j .

3. Opacity Regularization: The opacity regularization ensures meaningful opacity values:

$$\mathcal{L}_{\text{opacity}} = \frac{1}{N} \sum_{j=1}^N \|\sigma(\mathbf{a}_j)\|, \quad (16)$$

where σ is the sigmoid function applied to the opacity parameters \mathbf{a}_j .

4. Depth Loss: The depth loss enforces consistency between rendered depth and ground-truth depth:

$$\mathcal{L}_{\text{depth}} = \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \|D_{\text{rendered}, i} - D_{\text{gt}, i}\|_1, \quad (17)$$

where \mathcal{M} is a mask selecting pixels with valid depth measurements, and $D_{\text{gt}, i}$ is derived from MVS in Section 3.5.

3.6.5. Partitioning for Training Large-Scale GS Model

Since the reconstructed 3D sparse model from SfM is too large to train a single 3D Gaussian Splatting (GS) model, we propose a simple partitioning process to divide the SfM

model into multiple smaller partitions. Each partition is then trained individually using the GS settings described above. Figure 3 illustrates the process of dividing a large SfM model into multiple sub-models.

Specifically, we calculate the origin point of the SfM model by averaging the latitude and longitude of all images. Using this origin point as a reference, we define the position of each partition within a grid of boxes, each measuring 300 m in width and height. This approach results in multiple boxes, as illustrated in Figure 3. To better accommodate large-scale datasets, we further employ a Level of Detail (LoD) strategy [39] to enhance rendering quality and train multiple detail levels for each partition.

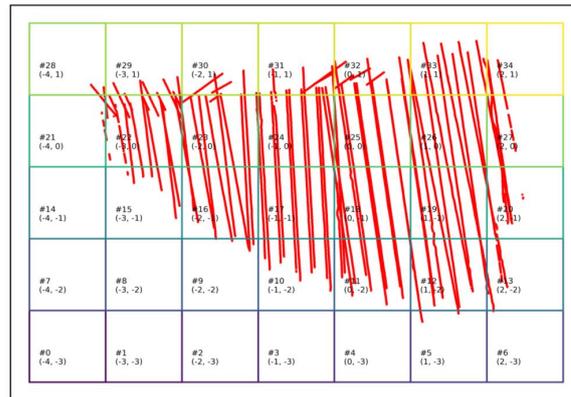


Figure 3. Partitioning the SfM model into multiple partitions to handle large-scale Gaussian Splatting. The red trajectories represent camera positions reconstructed by SfM.

3.7. Estimate Canopy Height Map from Gaussian Models

In this section, we describe the process of estimating the canopy height map (CHM) from the dense Gaussian models introduced earlier. The CHM construction begins with calculating the Digital Terrain Model (DTM) using the Cloth Simulation Filter (CSF) [40], which estimates the ground surface from the given point cloud data. The CSF algorithm simulates a cloth draped over the inverted point cloud, classifying the points touched by the cloth as ground points.

Starting with the dense point cloud $\{X_j^{\text{dense}}\}$, derived from ET-MVS [35] (see Section 3.5), which represents a combination of surfaces including the ground, vegetation, and other objects, we extract a set of cloth nodes $\{X_{\text{cloth}}\}$ from $\{X_j^{\text{dense}}\}$ using CSF [40]. Note that the cloth nodes are not directly extracted from the MVS point cloud; rather, they are points from the simulated cloth. Specifically, the z positions of the cloth nodes represent the simulated cloth's height at certain (x, y) positions, which are arranged in a grid formation. These cloth nodes are subsequently used to interpolate the DTM as a continuous surface, $H_{\text{ground}}(u, v)$.

Once the DTM is generated, the CHM is computed by estimating the height of the Gaussian splats relative to the ground surface $H_{\text{ground}}(u, v)$. Using the 3D Gaussian model \mathcal{G} obtained in Section 3.6, we render a dense depth map $D_{\text{rendered}}(u, v)$ via orthographic projection [25], ensuring a high-resolution and geometrically accurate representation of the canopy. The initial CHM is then computed as:

$$\text{CHM}(u, v) = H_{\text{camera}} - D_{\text{rendered}}(u, v) - H_{\text{ground}}(u, v), \quad (18)$$

where H_{camera} is the camera height above the reference plane. To mitigate noise and interpolation errors, we use an external vegetation filter based on a U-Net [41] model $\mathcal{F}_{\theta}(\cdot)$,

which predicts a vegetation mask $M_{\text{veg}}(u, v)$ from the orthographic RGB image rendered from \mathcal{G} :

$$M_{\text{veg}} = \mathcal{F}_{\theta}(I_{\text{rendered}}), \quad (19)$$

where $\mathcal{F}_{\theta} : \mathbb{R}^{H \times W \times C} \rightarrow [0, 1]^{H \times W}$, and $M_{\text{veg}}(u, v) \in [0, 1]$ represents the probability that pixel (u, v) corresponds to vegetation. This filter refines the CHM as:

$$\text{CHM}_{\text{final}}(u, v) = \begin{cases} \text{CHM}(u, v), & \text{if } M_{\text{veg}}(u, v) > \tau_{\text{th}}, \\ 0, & \text{otherwise.} \end{cases} \quad (20)$$

However, the cloth nodes $\{X_{\text{cloth}}\}$ obtained from CSF may include unreliable points, especially in areas with canopy cover. To address this, we refine the cloth nodes using the vegetation filter $\mathcal{F}_{\theta}(\cdot)$. Specifically, we create a binary mask $\hat{M}_{\text{veg}}(u, v)$ for the cloth nodes based on a threshold τ_{veg} :

$$\hat{M}_{\text{veg}}(u, v) = \begin{cases} 1, & \text{if } M_{\text{veg}}(u, v) \geq \tau_{\text{veg}}, \\ 0, & \text{otherwise.} \end{cases} \quad (21)$$

For each cloth node $X_{\text{cloth},i} = (x_i, y_i, z_i)$, we compute a vegetation score by extracting a local region of the vegetation mask within a radius r :

$$S_{\text{veg},i} = \frac{\sum_{(u,v) \in \text{Region}_i} M_{\text{veg},i}(u, v)}{\text{total pixels in Region}_i}, \quad (22)$$

where Region_i is defined as the set of pixels (u, v) within a circular region centered at (x_i, y_i) with radius r . Cloth nodes with $S_{\text{veg},i} < \tau_{\text{node}}$ are considered invalid and removed. Heights for these invalid nodes are interpolated linearly using their nearest valid neighbors.

Finally, the refined ground model $H_{\text{ground}}(u, v)$ is used to compute the final CHM using the earlier equations, ensuring accurate canopy height estimation.

3.8. Evaluation Metrics and Baselines

Evaluation metrics: Following previous works [12,17], we use several metrics to evaluate the CHM estimation results of the proposed method. These metrics include the RMSE (Equation (23)), which emphasizes higher errors in height estimation, the MAE (Equation (24)), which provides an equal average of all height errors, and the ME Equation (25), which quantifies height bias, with a negative bias indicating that the estimated height \hat{h}_i is lower than the LiDAR reference ground truth h_i .

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (h_i - \hat{h}_i)^2} \quad (23)$$

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |h_i - \hat{h}_i| \quad (24)$$

$$\text{ME} = \frac{1}{N} \sum_{i=1}^N (h_i - \hat{h}_i) \quad (25)$$

We additionally report the R^2 -block (R^2), which is defined as follows:

$$R^2 = 1 - \frac{\sum_{i=1}^N (h_i - \hat{h}_i)^2}{\sum_{i=1}^N (h_i - \bar{h})^2}, \quad (26)$$

where \bar{h} is the mean of actual target values. The R^2 , also known as the coefficient of determination, is a statistical measure that indicates how well a regression model fits the observed data. In this paper, we use it to evaluate how well a block of 150×150 pixels ($\sim 15 \text{ m} \times 15 \text{ m}$) corresponds to the same block size of LiDAR results. Unlike [17], we select a block size of 15 m instead of 30 m to achieve more reliable results, ensuring higher accuracy for future applications such as biomass measurement. Furthermore, we also report the performance of the proposed method using smaller, more challenging block sizes of 10 m and 5 m.

Baselines: To assess the modeling fidelity of ForestSplat, we compare its canopy height models (CHMs) with those derived from airborne LiDAR data. Canopy height is a critical proxy for biomass and serves as the sole metric where airborne LiDAR is unequivocally considered the ground truth. The CHM for airborne LiDAR is generated using ArcGIS [42], a widely used tool for processing and analyzing LiDAR data in geospatial research.

To benchmark ForestSplat, we select two baseline methods for comparison, ensuring an objective comparison of our technique against industry gold standards. The first is SSL-Satellite [17], a state-of-the-art approach for producing very high-resolution canopy height maps from RGB imagery using a self-supervised vision transformer and a convolutional decoder. This method can predict CHMs with high accuracy directly from satellite images.

The second baseline is a traditional photogrammetry-based approach [34], which employs Pix4D software to generate CHMs. This method relies on well-established photogrammetric techniques to reconstruct canopy height models from imagery.

4. Experiments

This section presents a detailed explanation of the experimental setups and the corresponding results, benchmarked against the baselines, for reconstructing high-resolution and highly accurate CHMs. The data collection methodology, including drone flight parameters and image acquisition details, is described in Section 2.

4.1. Experimental Setups

SSL-Satellite [17]: Using Google Earth Pro, we downloaded satellite imagery of the survey area corresponding to the same time in November 2023 as the LiDAR data collection. The satellite image is shown in Figure 1, where the region of interest, enclosed by the red polygon, represents the study area initially used for reforestation research and serves as the primary benchmark dataset in this work. Specifically, the satellite image was downloaded at a high resolution of 8192×4437 pixels (approximately 0.34 m GSD per pixel). Following the settings in [17], we cropped the satellite image into a set of 256×256 pixel tiles, resulting in 544 non-overlapping crops. We then applied the same configurations proposed in [17] to generate CHMs for each crop, where each loaded RGB satellite image was normalized using a mean of (0.420, 0.411, 0.296) and a standard deviation of (0.213, 0.156, 0.143). Note that the GSD and image normalization configurations were recommended by the authors [17]. Since the SSL-Satellite provides several pre-trained models for predicting CHMs, we selected the same evaluation model as [17], namely SSLhuge-Satellite.

After obtaining all CHMs for 544 non-overlapping crops through SSLhuge-Satellite model, we merged them back in their correct order to reconstruct an image of the original size, 8192×4437 pixels, and converted it to the GeoTIFF format for accurate comparison with our method and other baselines.

Photogrammetry-based [34]: For this baseline, we used 13K images captured in this site in January 2024. We then divided the site into four non-overlapping areas for processing. This step was necessary because the entire survey area was too large to be processed at once using Pix4D. Additionally, since the proposed dataset lacks the recommended 70% side

overlap for Pix4D (having only 30% side overlap, as mentioned in Section 2, which makes the task highly challenging), we had to manually label matching pixels between images to ensure acceptable results. Without this manual intervention, the CHM's performance would have been significantly worse. This process required at least 24 h of manual labor. Finally, after exporting the point cloud from Pix4D, we used CloudCompare, which employs the CSF algorithm [40], to estimate the ground surface and calculate the CHMs. The four obtained CHM parts from this model were also merged into a unified one and registered as a GeoTIFF for later comparison.

The Proposed ForestSplat Pipeline: We used the following configurations. We set $\text{IoU} = 0.2$ for image pair generation, resulting in approximately 500K image pairs compared to 93M pairs generated using an exhaustive approach. For Gaussian Splatting, we used a partition size of $300 \text{ m} \times 300 \text{ m}$, which created 18 non-overlapping partitions. Each partition contains a different number of valid regions, as some partitions are located near the borders of the survey area.

To train each Gaussian Splatting partition model, we set the growing anchor threshold $\tau_g = 6 \times 10^{-4}$, pruning opacity threshold $\tau_o = 0.04$, and scale $\tau_s = 0.4$. The weighting coefficients in the total loss function were configured as follows: $\lambda_{\text{scale}} = 0.01$ to regulate the Gaussian scales, $\lambda_{\text{opa}} = 0.001$ to penalize opacity values, $\lambda_{\text{depth}} = 0.01$ to enforce depth consistency, and $\lambda_{\text{ssim}} = 0.2$ to balance the reconstruction loss between the L1 term and the SSIM term.

For CHM generation, we used $\tau_{\text{node}} = 0.1$, $\tau_h = 0.5$, and $\tau_{\text{veg}} = 0.25$ for all partitions. For training the vegetation filter $\mathcal{F}_\theta(\cdot)$, we randomly selected 90 tiles (each covering $50 \text{ m} \times 50 \text{ m}$ of ground area) from nearby regions of the proposed dataset and manually labeled them with binary masks to distinguish tree areas from non-tree areas. To ensure that the model does not encounter out-of-distribution data when applied to the proposed dataset, we included several randomly labeled tiles selected from the proposed dataset of the 200-acre area. Note that the total number of labeled tiles used for training and evaluation remained at 90. The filter was then trained for 50 epochs on 75% of the training data (67 tile images) using a learning rate of 0.001. The trained filter $\mathcal{F}_\theta(\cdot)$ was subsequently applied to the entire 200-acre dataset. Interestingly, we found that training $\mathcal{F}_\theta(\cdot)$ on this small amount of data was sufficient to achieve good results for filtering in the final CHM calculation.

Since the images were captured two months after the LiDAR data, we adjusted the CHMs by reducing the height of all values greater than 1 m by 10 cm to enhance the reliability of comparisons. This adjustment assumes a tree growth rate of 5 cm per month and was applied to the photogrammetry-based method [34] and the proposed method. In contrast, the SSL-Satellite results [17] were left unchanged, as the satellite images were obtained at the same time as the LiDAR data.

4.2. Results

4.2.1. Results at $15 \text{ m} \times 15 \text{ m}$ Block Size

We present the obtained results in Tables 1 and 2. Specifically, Table 1 summarizes the comparison results of the proposed pipeline against SSL-Satellite [17] and the photogrammetry-based method [34] across all partitions in terms of RMSE, MAE, and ME metrics. The proposed pipeline consistently outperformed the other methods, achieving the lowest weighted average RMSE (0.272 m), MAE (0.172 m), and ME (0.007 m) across all partitions. In contrast, the SSL-Satellite method exhibited the highest average errors, with RMSE, MAE, and ME values of 0.854 m, 0.522 m, and 0.090 m, respectively. The photogrammetry-based method demonstrated intermediate performance, with average errors of 0.653 m (RMSE), 0.433 m (MAE), and 0.280 m (ME). These results highlight the robustness and

reliability of the proposed approach in reconstructing CHMs with high precision. In Table 2, we provide a detailed number of valid blocks, R^2 scores, and an arbitrary accuracy at a threshold of 0.5 m. In the weighted average of all partitions, the proposed ForestSplat achieved the highest accuracy at 90.9% and demonstrated an R^2 score of 0.79, compared to -1.21 and -0.99 for the photogrammetry-based and SSL-Satellite methods, respectively. This high R^2 score indicates that ForestSplat’s estimations align closely with the ground truth, outperforming the baseline methods that fail to fit the ground truth effectively.

Table 1. Comparison results of the proposed pipeline against SSL-Satellite [17] and the photogrammetry-based method [34] in terms of RMSE, MAE, and ME metrics for all partitions. All results are reported using a block size of $15\text{ m} \times 15\text{ m}$. The final row shows the weighted average metrics over different valid block sizes. The best results are in bold.

Partition		SSL-Satellite [17] RMSE/MAE/ME	Photogrammetry [34] RMSE/MAE/ME	ForestSplat (Ours) RMSE/MAE/ME
#10	-01_-02	0.897/0.567/0.528	0.133/0.100/−0.014	0.146/0.116/−0.074
#11	000_-02	1.149/0.641/0.530	0.520/0.389/0.349	0.202/0.128/−0.010
#12	001_-02	1.755/0.960/0.644	0.645/0.451/0.384	0.503/0.297/0.148
#13	002_-02	1.357/0.937/0.478	1.192/1.064/1.046	0.715/0.578/0.458
#16	-02_-01	0.332/0.187/−0.048	0.834/0.508/0.507	0.152/0.092/−0.024
#17	-01_-01	0.625/0.330/0.146	0.287/0.138/0.115	0.178/0.115/−0.068
#18	000_-01	0.693/0.350/0.262	0.256/0.153/0.129	0.167/0.115/−0.113
#19	001_-01	1.325/0.781/0.679	0.485/0.263/0.227	0.230/0.133/0.014
#20	002_-01	1.952/1.350/1.122	1.823/1.610/1.608	0.393/0.272/0.073
#22	-03_000	1.889/1.228/1.154	1.844/1.663/1.435	0.809/0.587/0.537
#23	-02_000	0.595/0.306/ 0.024	0.969/0.720/0.718	0.315/0.164/0.080
#24	-01_000	0.406/0.160/0.044	0.567/0.442/0.442	0.160/0.073/0.022
#25	000_000	0.297/0.136/−0.101	0.473/0.164/0.125	0.147/0.072/−0.010
#26	001_000	1.199/0.969/−0.820	0.972/0.662/−0.217	0.489/0.357/−0.021
#27	002_000	2.200/1.817/1.473	1.530/1.220/1.190	1.167/0.805/0.535
#31	-01_001	0.075/0.029/−0.019	1.498/1.441/1.441	0.057/0.027/−0.006
#32	000_001	0.996/0.744/−0.589	1.283/0.665/0.371	0.363/0.247/−0.093
#33	001_001	1.340/1.201/−0.977	1.062/0.773/0.161	0.519/0.404/0.112
E(X)		0.854/0.522/0.090	0.653/0.433/0.280	0.272/0.172/0.007

Table 2. Comparison of additional metrics across different methods. Detailed results include the number of valid blocks for each partition, the R^2 metric, and the accuracy metric for the number of blocks with an average error lower than 50 cm. The best results are in bold.

Partition		SSL-Satellite [17] #Blocks/R2/Acc.	Photogrammetry [34] #Blocks/R2/Acc.	ForestSplat (Ours) #Blocks/R2/Acc.
#10	-01_-02	12/−7.83/58.3	12/ 0.81 / 100.0	12/0.77/100.0
#11	000_-02	100/−1.10/70.0	97/0.64/74.2	100/ 0.94 / 96.0
#12	001_-02	177/−1.24/49.7	177/0.74/73.5	177/ 0.81 / 78.5
#13	002_-02	24/−1.69/41.7	24/−0.81/20.8	24/ 0.25 / 50.0
#16	-02_-01	134/0.04/95.5	122/−4.18/75.4	134/ 0.80 / 98.5
#17	-01_-01	371/−1.70/83.6	352/0.52/93.8	371/ 0.78 / 98.7
#18	000_-01	400/−1.73/78.2	399/0.67/93.5	400/ 0.84 / 99.2
#19	001_-01	398/−2.25/61.8	398/0.63/87.4	398/ 0.90 / 95.2
#20	002_-01	83/−0.95/33.7	83/−0.55/10.8	83/ 0.92 / 81.9
#22	-03_000	22/−1.67/31.8	22/−1.35/9.1	22/ 0.51 / 54.5
#23	-02_000	330/0.33/85.5	248/−0.37/52.4	330/ 0.81 / 90.6
#24	-01_000	396/0.15/89.9	264/−0.15/65.5	396/ 0.87 / 97.0
#25	000_000	400/0.20/91.2	365/−0.74/94.5	400/ 0.81 / 97.8
#26	001_000	381/−1.23/36.0	381/−0.44/49.3	381/ 0.63 / 73.8
#27	002_000	22/−4.99/13.6	22/−1.64/31.8	22/− 0.69 / 50.0
#31	-01_001	14/0.34/ 100.0	14/−232.02/0.0	14/ 0.63 /100.0
#32	000_001	61/−1.10/47.5	61/−2.31/67.2	61/ 0.72 / 83.6
#33	001_001	88/−2.30/14.8	88/−0.98/45.5	88/ 0.51 / 67.0
E(X)		3413/−0.99/70.5	3129/−1.21/73.4	3413/ 0.79 / 90.9

Figures 4 and 5 show the comparative results of canopy height models (CHMs) generated by different methods, using two example partitions (#19 and #24), evaluated against

LiDAR as the ground truth. For both partitions, the proposed ForestSplat pipeline consistently demonstrates superior performance compared to SSL-Satellite and photogrammetry-based methods. The photogrammetry-based method completely failed to produce CHMs in areas with water, leaving these regions blank. Similarly, the SSL-Satellite method shows poor performance when merging crop regions, as evidenced by the frequent appearance of sharp edges in its CHMs, disrupting the continuity of the height maps. Additionally, the block-wise CHMs and scatter plots visually highlight ForestSplat’s ability to deliver consistent and accurate height estimations, further underscoring its effectiveness in producing accurate and reliable canopy height maps.

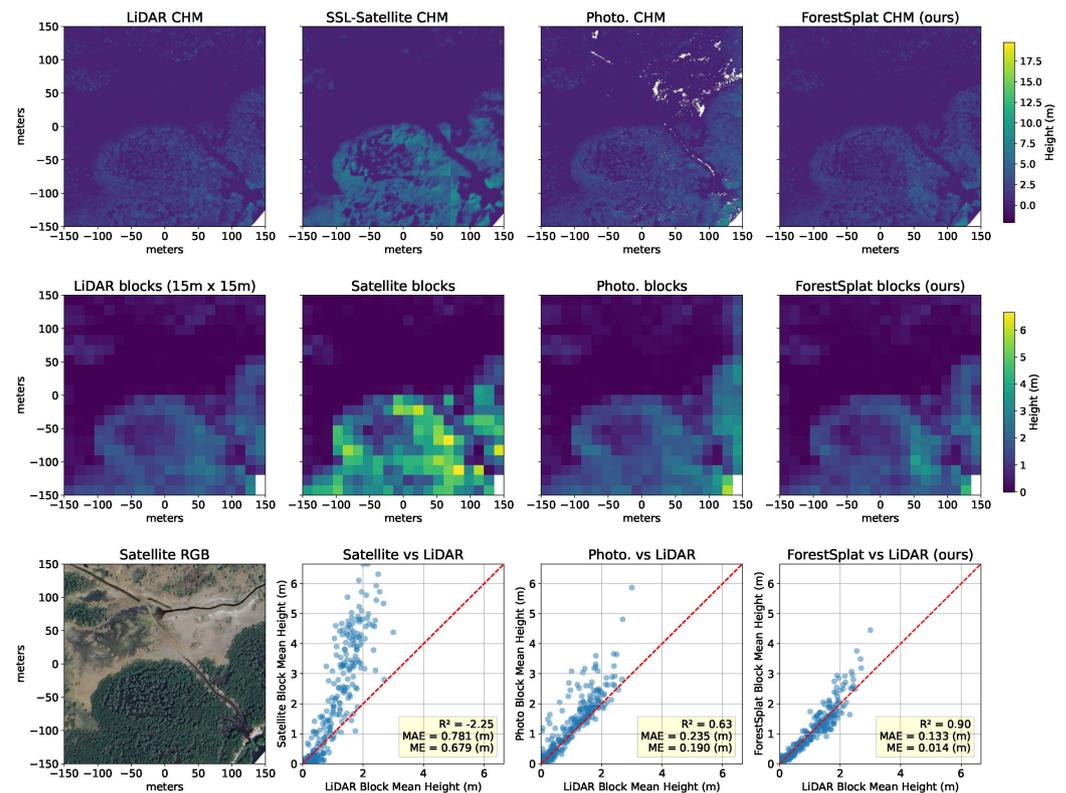


Figure 4. Visual comparison of CHMs generated by different methods for partition #19, compared against LiDAR. Top row: The CHMs at high resolution of GSD at 0.1 m are shown for Airborne LiDAR, SSL-Satellite [17], photogrammetry-based Pix4D [34], and ForestSplat (proposed pipeline). Second row: CHMs downsampled to 15 m × 15 m blocks to facilitate direct comparison with LiDAR-based block averages. Last row: Satellite RGB image of the study area, providing context for the canopy structure and scatter plots evaluating the block-wise correlation of each method with the LiDAR ground truth.

4.2.2. Merging CHMs Result

Here, we present a visual comparison of all methods for generating merged CHMs, including ForestSplat, crop-based CHMs from SSL-Satellite [17], and four sub-sites from photogrammetry [34], with LiDAR serving as the ground truth. The visual results, shown in Figure 6, highlight that the proposed ForestSplat method achieves a consistent match with the LiDAR CHM. In contrast, the photogrammetry-based method demonstrates the poorest performance in the merging process, failing to cover many areas and leaving significant gaps. Meanwhile, the SSL-Satellite method struggles to ensure smooth transitions between adjacent crop regions, resulting in noticeable discontinuities and sharp edges in the merged CHMs.

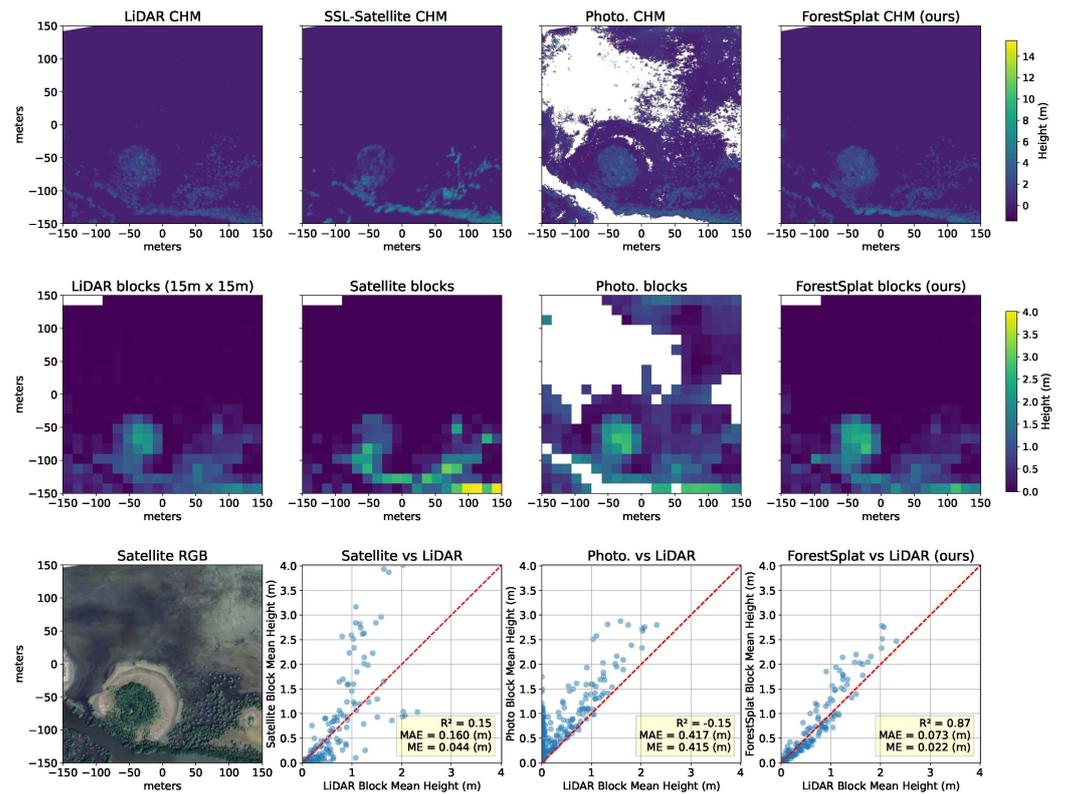


Figure 5. Visual comparison of CHMs generated by different methods for partition #24, compared against LiDAR [17,34].

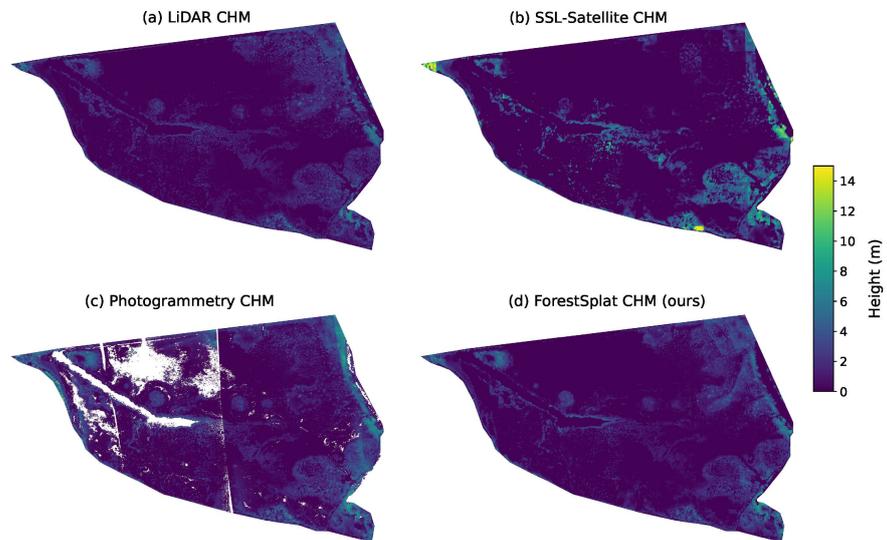


Figure 6. Visual comparison results of merged CHMs [17,34].

4.2.3. Changing Blocks Size

To evaluate the performance of each method under different spatial resolutions, we analyzed the results of canopy height models (CHMs) using varying block sizes: 15 m × 15 m, 10 m × 10 m, and 5 m × 5 m. The scatter plots in Figure 7 illustrate the correlation between the predicted block mean heights and the LiDAR ground truth for SSL-Satellite, photogrammetry-based, and the proposed ForestSplat method at these resolutions.

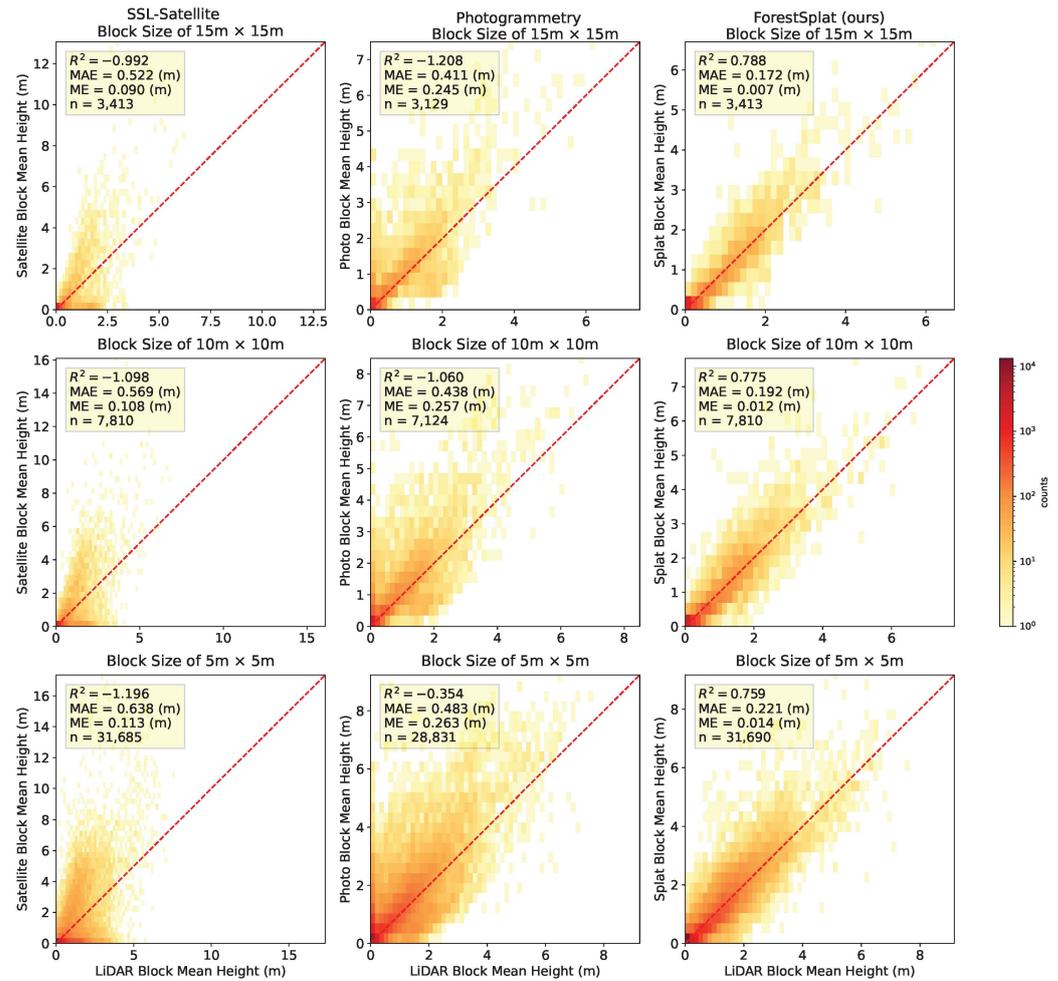


Figure 7. Scatter plots of estimated CHM block mean heights versus LiDAR ground truth at varying block sizes (15 m × 15 m, 10 m × 10 m, 5 m × 5 m), comparing SSL-Satellite [17], photogrammetry-based [34], and ForestSplat (ours).

Across all block sizes, ForestSplat consistently demonstrates superior performance, achieving the highest R^2 values (e.g., 0.788 at 15 m × 15 m, 0.775 at 10 m × 10 m, and 0.759 at 5 m × 5 m) and the lowest mean absolute error (MAE) and mean error (ME). These results indicate that ForestSplat maintains high accuracy and robustness in capturing canopy height distributions at finer resolutions, even when block sizes decrease.

In contrast, the photogrammetry-based method shows moderate performance at larger block sizes (e.g., $R^2 = -1.208$ at 15 m × 15 m), but its reliability deteriorates significantly as the block size decreases, with lower R^2 scores and increasing error. Additionally, large gaps and deviations from the LiDAR ground truth are observed in smaller blocks.

The SSL-Satellite method performs poorly across all block sizes, with consistently negative R^2 values (e.g., $R^2 = -0.992$ at 15 m × 15 m and $R^2 = -1.196$ at 5 m × 5 m). This method struggles to capture fine-grained canopy height variations and exhibits higher MAE and ME compared to the other approaches.

These results demonstrate that ForestSplat is robust to changes in block sizes, offering reliable estimations at varying spatial resolutions, while the other methods exhibit significant limitations, particularly at smaller block sizes or very high resolutions. This is because ForestSplat leverages 3D Gaussian Splatting, which preserves fine-scale details by continuously optimizing scene representation at a sub-pixel level. On the other hand, the SSL-Satellite method relies on coarse-resolution satellite imagery, leading to blurred height estimations at smaller block scales. Photogrammetry-based CHMs often require high image overlap

($\geq 70\%$) for accurate reconstruction; however, since our dataset had only 30% side overlap, its reliability at finer spatial resolutions was significantly reduced. Additionally, the point clouds produced by the photogrammetry-based method remain relatively coarse, whereas ForestSplat refines them further—first through MVS (Multi-View Stereo) to improve density and accuracy, and then through Gaussian Splatting, which continuously optimizes and smooths the representation for a more precise canopy height model (CHM).

4.2.4. Adjustment for Tree Growth

To account for an assumed tree growth of 5 cm per month and the two-month gap between the LiDAR data collection and image acquisition, all CHM partitions were adjusted by -10 cm with $\text{CHM}(u, v) > 1$ m. As shown in Table 3, this adjustment led to a general improvement in performance metrics, highlighted in green, with notable enhancements in R^2 (e.g., $+0.015$ on average) and a reduction in MAE (-0.005 m on average). While some partitions experienced slight reductions in performance, marked in red, the overall metrics demonstrate better alignment with the LiDAR ground truth after this correction, achieving an average accuracy of 90.91% within 0.5 m. This adjustment reflects the significance of accounting for tree growth when comparing CHMs to ground truth data.

Table 3. Performance metrics for CHM partitions of ForestSplat before and after a -10 cm adjustment to account for tree growth, with improvements in green, reductions in red, and unchanges in gray.

ID	Partition	#Blocks	RMSE	MAE	R^2	ME	Acc. at 0.5 m
#10	-01_02 -10 cm	12 12 (+0.000)	0.146 0.146 (+0.000)	0.114 0.116 (+0.002)	0.766 0.766 (+0.000)	-0.053 -0.074 (-0.021)	100.000% 100.000% (+0.000)
#11	000_02 -10 cm	100 100 (+0.000)	0.214 0.202 (-0.012)	0.132 0.128 (-0.004)	0.927 0.935 (+0.008)	0.021 -0.010 (-0.031)	94.000% 96.000% (+2.000)
#12	001_02 -10 cm	177 177 (+0.000)	0.534 0.503 (-0.031)	0.321 0.297 (-0.024)	0.792 0.815 (+0.023)	0.198 0.148 (-0.050)	76.840% 78.530% (+1.690)
#13	002_02 -10 cm	24 24 (+0.000)	0.767 0.715 (-0.052)	0.623 0.578 (-0.045)	0.140 0.253 (+0.113)	0.528 0.458 (-0.070)	54.170% 50.000% (-4.170)
#16	-02_01 -10 cm	134 134 (+0.000)	0.164 0.152 (-0.012)	0.095 0.092 (-0.003)	0.767 0.799 (+0.032)	-0.013 -0.024 (-0.011)	97.760% 98.510% (+0.750)
#17	-01_01 -10 cm	371 371 (+0.000)	0.181 0.178 (-0.003)	0.110 0.115 (+0.005)	0.774 0.779 (+0.005)	-0.054 -0.068 (-0.014)	98.650% 98.650% (+0.000)
#18	000_01 -10 cm	400 400 (+0.000)	0.149 0.167 (+0.018)	0.103 0.115 (+0.012)	0.874 0.842 (-0.032)	-0.099 -0.113 (-0.014)	99.500% 99.250% (-0.250)
#19	001_01 -10 cm	398 398 (+0.000)	0.245 0.230 (-0.015)	0.137 0.133 (-0.004)	0.889 0.902 (+0.013)	0.044 0.014 (-0.030)	93.720% 95.230% (+1.510)
#20	002_01 -10 cm	83 83 (+0.000)	0.416 0.393 (-0.023)	0.291 0.272 (-0.019)	0.911 0.921 (+0.010)	0.133 0.073 (-0.060)	80.720% 81.930% (+1.210)
#22	-03_000 -10 cm	22 22 (+0.000)	0.864 0.809 (-0.055)	0.625 0.587 (-0.038)	0.441 0.510 (+0.069)	0.593 0.537 (-0.056)	54.550% 54.550% (+0.000)
#23	-02_000 -10 cm	330 330 (+0.000)	0.341 0.315 (-0.026)	0.175 0.164 (-0.011)	0.779 0.812 (+0.033)	0.102 0.080 (-0.022)	90.000% 90.610% (+0.610)
#24	-01_000 -10 cm	396 396 (+0.000)	0.176 0.160 (-0.016)	0.077 0.073 (-0.004)	0.839 0.868 (+0.029)	0.033 0.022 (-0.011)	96.210% 96.970% (+0.760)
#25	000_000 -10 cm	400 400 (+0.000)	0.148 0.147 (-0.001)	0.073 0.072 (-0.001)	0.803 0.807 (+0.004)	-0.003 -0.010 (-0.007)	97.750% 97.750% (+0.000)
#26	001_000 -10 cm	381 381 (+0.000)	0.506 0.489 (-0.017)	0.367 0.357 (-0.010)	0.603 0.630 (+0.027)	0.036 -0.021 (-0.057)	74.280% 73.750% (-0.530)
#27	002_000 -10 cm	22 22 (+0.000)	1.217 1.167 (-0.050)	0.846 0.805 (-0.041)	-0.834 -0.686 (+0.148)	0.603 0.535 (-0.068)	50.000% 50.000% (+0.000)

Table 3. Cont.

ID	Partition	#Blocks	RMSE	MAE	R ²	ME	Acc. at 0.5 m
#31	-01_001	14	0.056	0.026	0.641	-0.005	100.000%
	-10 cm	14 (+0.000)	0.057 (+0.001)	0.027 (+0.001)	0.628 (-0.013)	-0.006 (-0.001)	100.000% (+0.000)
#32	000_001	61	0.367	0.255	0.715	-0.053	80.330%
	-10 cm	61 (+0.000)	0.363 (-0.004)	0.247 (-0.008)	0.722 (+0.007)	-0.093 (-0.040)	83.610% (+3.280)
#33	001_001	88	0.551	0.428	0.441	0.195	65.910%
	-10 cm	88 (+0.000)	0.519 (-0.032)	0.404 (-0.024)	0.505 (+0.064)	0.112 (-0.083)	67.050% (+1.140)
$E(X)$		3413.0	0.284	0.177	0.773	0.034	90.419%
$E(X)_{-10\text{cm}}$		3413.0 (+0.000)	0.272 (-0.011)	0.172 (-0.005)	0.788 (+0.015)	0.007 (-0.027)	90.918% (+0.499)

4.2.5. Additional Results

In this section, we present additional example results that demonstrate the high fidelity of the proposed rendered images compared to satellite images. We also provide qualitative vegetation prediction results obtained using the proposed method on the rendered images.

In Figure 8, we present two examples of rendered image tiles, each covering an area of 50 m × 50 m, compared to satellite images. These examples demonstrate that the proposed method can achieve the ultra-high resolution required for MRV. The rendered images have a ground sampling distance (GSD) of 1 cm per pixel, with each tile containing 5000 × 5000 pixels. The last column displays predicted vegetation masks generated from the corresponding rendered images.

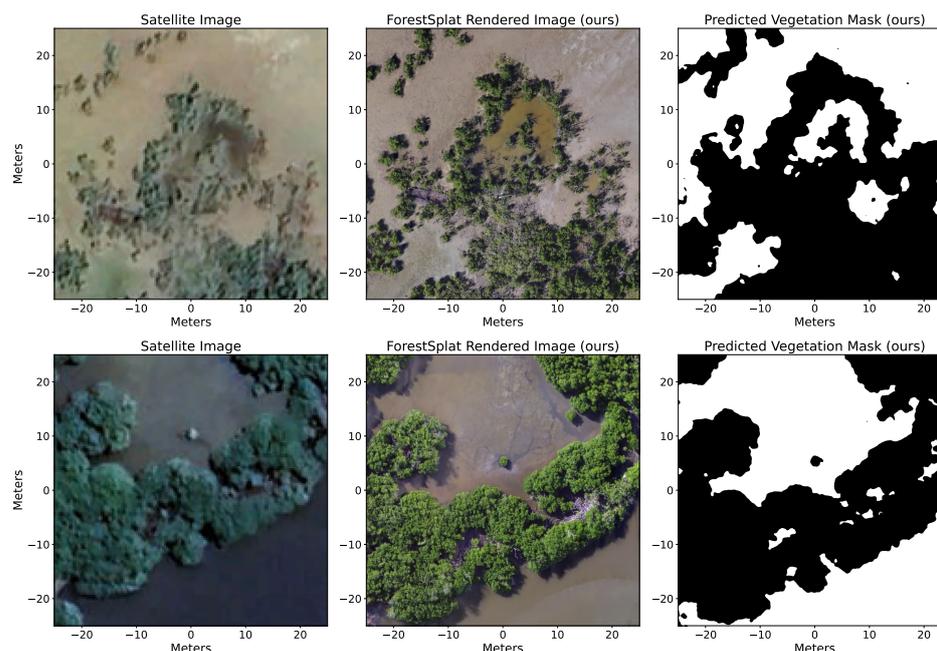


Figure 8. Example of high-fidelity image tiles of 50 m × 50 m (middle) rendered by the proposed method, compared to satellite images (left). On the right, we present an example prediction result of the proposed vegetation filter applied to the rendered images (middle), with black areas indicating predicted trees. The first row is located at 17.935166°N, -66.165946°W, and the second row is at 17.938328°N, -66.170666°W.

Since no hand-labeled ground truth is available for these images, we cannot provide quantitative metrics for these vegetation predictions. However, we found that the current filter is sufficient for removing incorrect height estimates from non-vegetation regions. Nonetheless, we believe that further improving the vegetation filter could potentially enhance CHM estimation accuracy. Therefore, we identify this as a potential direction for future work.

5. Discussion

Discussion. This work demonstrates the promise of using computer vision-based 3D reconstruction (ForestSplat) as a highly accurate yet more scalable alternative to airborne LiDAR for reforestation MRV. ForestSplat relies solely on a low-cost camera to achieve competitive results, with the potential to reduce operational costs by up to 100× compared to airborne LiDAR scans.

ForestSplat proves most effective in forestry settings where ground visibility is preserved within localized areas. As reforested canopies close over time, accurate geographical understanding of the ground can still be maintained by incorporating temporal dimensions into the modeling process. However, in forest conservation scenarios with enclosed, dense canopies or challenging terrains, such as primary tropical rainforests, the performance of Gaussian Splatting may be limited due to its inability to perceive the ground effectively.

Limitation. While ForestSplat demonstrates high accuracy and cost-effectiveness for canopy height estimation, there are several limitations to consider. First, the method relies on high-quality aerial imagery, meaning that variations in lighting conditions, camera calibration, and flight stability can impact reconstruction quality. Second, dense forest canopies, such as those found in tropical rainforests, may obstruct ground visibility, leading to errors in height estimation. Additionally, complex terrains, such as steep slopes or mountainous regions, may introduce distortions in 3D reconstruction if not properly corrected with terrain-aware adjustments or external Digital Terrain Models (DTMs). Another limitation is the dependence on GNSS accuracy for aligning SfM reconstructions to real-world coordinates; in GNSS-denied environments, alternative localization methods, such as visual-inertial odometry (VIO) or SLAM, may be required.

Furthermore, while LiDAR serves as the primary comparison ground truth in this study, it is not without its own errors, particularly in sparse vegetation regions or dense canopies where penetration is limited, potentially affecting the accuracy of reference canopy height maps. Additionally, variations in LiDAR point density and ground filtering algorithms can introduce biases in height estimation, especially in regions with mixed vegetation structures. Differences in georeferencing accuracy and processing pipelines between LiDAR and photogrammetry-based CHMs may also contribute to local misalignments, affecting direct comparisons.

Finally, scalability remains a challenge—while ForestSplat has been validated on a 200-acre site, further optimizations in data processing and computational efficiency will be needed for large-scale deployments spanning thousands of acres. Addressing these challenges is crucial to ensuring ForestSplat's broader applicability and effectiveness in more diverse and demanding environments.

Future work. To overcome these limitations, future work will explore the applicability of ForestSplat to diverse forest types, including tropical rainforests, temperate forests, and coniferous woodlands, where differences in canopy density and structure may require adaptations in Gaussian Splatting parameters and multi-temporal data collection. Additionally, extending the method to complex terrain conditions, such as mountainous regions, will involve integrating Digital Terrain Models (DTMs) and exploring GNSS-free camera pose optimization for areas with poor GNSS signals. To enhance scalability, future research will investigate distributed processing techniques and cloud-based CHM computation to extend ForestSplat's usability to large-scale forest monitoring.

Another important direction is improving tree-level segmentation of 3D Gaussian Splat models to enable more precise biomass estimation and carbon stock analysis. By refining individual tree extraction, ForestSplat could provide valuable insights into forest structure, growth dynamics, and carbon sequestration potential, with applications in forest conservation and climate impact studies. Validation will involve comparisons against

comprehensive ground truth data, including hand-measured transect data and additional scans from airborne and terrestrial LiDAR systems. Developing a cost-effective and scalable MRV system with LiDAR-comparable accuracy is vital for advancing forestry management and nature-based climate solutions. Such a system would empower foresters to transition from simple drone scans to detailed, tree-specific intelligence, enhancing forestry practices such as planting strategies, maintenance, and carbon credit generation. We hope this work contributes to accelerating the growth of the emerging forestry-based carbon removal sector.

6. Conclusions

This study introduces the ForestSplat pipeline for high-fidelity mapping of reforestation projects. Through both quantitative and qualitative evaluations on a 200-acre reforestation dataset, we demonstrate the clear advantages of ForestSplat over existing methods in high-resolution canopy height map estimation. The proposed pipeline excels in capturing tree structure variability, even at fine spatial scales, significantly improving both accuracy and resolution. Beyond its cost-effectiveness, ForestSplat provides a scalable alternative to traditional airborne LiDAR, enabling accurate large-scale forest monitoring with consumer-grade drones. Unlike satellite-based CHM estimation, which suffers from resolution constraints and occlusions, ForestSplat leverages 3D Gaussian Splatting to optimize scene representation at a sub-pixel level, preserving fine structural details. We believe that the ultra-high-resolution canopy height maps generated by ForestSplat can substantially enhance the monitoring of forest degradation, restoration efforts, and forest carbon dynamics while maintaining exceptionally low cost requirements.

Author Contributions: Conceptualization, B.S. (Belal Shaheen), M.D.Z. and M.W.; Methodology, B.S. (Belal Shaheen), M.D.Z., B.-T.B., S. and M.W.; Validation, T.H.; Investigation, M.W.; Data curation, M.D.Z.; Writing—original draft, B.-T.B.; Writing—review & editing, B.S. (Belal Shaheen), M.D.Z. and M.W.; Visualization, B.S. (Belal Shaheen) and B.-T.B.; Supervision, S.C. and M.W.; Funding acquisition, M.M. and B.S. (Ben Scheelk). All authors have read and agreed to the published version of the manuscript.

Funding: This research was supported by the Ocean Foundation’s Blue Resilience Initiative through both technical support and funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries and data availability requests can be directed to the corresponding author or accessed at <https://coolant.earth>.

Acknowledgments: We thank Hanspeter Pfister for his guidance. This research was conducted in the Jobos Bay National Estuarine Research Reserve in Puerto Rico. We extend our gratitude to the JBNERR for providing us access to their protected mangrove sites and enabling us to conduct our experiments.

Conflicts of Interest: Author Belal Shaheen, Matthew David Zane, Bach-Thuan Bui, Shubham, and Michael Wu are employee at Coolant Climate, Inc. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

1. Bastin, J.F.; Finegold, Y.; Garcia, C.; Mollicone, D.; Rezende, M.; Routh, D.; Zohner, C.M.; Crowther, T.W. The global tree restoration potential. *Science* **2019**, *365*, 76–79. [[CrossRef](#)] [[PubMed](#)]
2. Husch, B.; Beers, T.W.; Kershaw, J.A., Jr. *Forest Mensuration*; John Wiley & Sons: Hoboken, NJ, USA, 2002.
3. Köhl, M.; Magnussen, S.; Marchetti, M. *Sampling Methods, Remote Sensing and GIS Multiresource Forest Inventory*; Springer: Berlin/Heidelberg, Germany, 2006; Volume 2.

4. Kauffman, J.B.; Arifanti, V.B.; Basuki, I.; Kurnianto, S.; Novita, N.; Murdiyarso, D.; Donato, D.C.; Warren, M.W. *Protocols for the Measurement, Monitoring, and Reporting of Structure, Biomass, Carbon Stocks and Greenhouse Gas Emissions in Tropical Peat Swamp Forests*; Center for International Forestry Research: Bogor, Indonesia, 2016.
5. Tomppo, E.; Gschwantner, T.; Lawrence, M.; Mcroberts, R.; Godinho-Ferreira, P. National Forest Inventories: Pathways for Common Reporting. In *National Forest Inventories: Pathways for Common Reporting*; Springer: Berlin/Heidelberg, Germany, 2010; pp. 3–18.
6. Luoma, V.; Saarinen, N.; Wulder, M.A.; White, J.C.; Vastaranta, M.; Holopainen, M.; Hyyppä, J. Assessing precision in conventional field measurements of individual tree attributes. *Forests* **2017**, *8*, 38. [[CrossRef](#)]
7. Goetz, S.J.; Hansen, M.; Houghton, R.A.; Walker, W.; Laporte, N.; Busch, J. Measurement and monitoring needs, capabilities and potential for addressing reduced emissions from deforestation and forest degradation under REDD+. *Environ. Res. Lett.* **2015**, *10*, 123001. [[CrossRef](#)]
8. Dubayah, R.; Luthcke, S.; Sabaka, T.; Nicholas, J.; Preaux, S.; Hofton, M. *GEDI L3 Gridded Land Surface Metrics, Version 2*; ORNL DAAC: Oak Ridge, TN, USA, 2021.
9. Markus, T.; Neumann, T.; Martino, A.; Abdalati, W.; Brunt, K.; Csatho, B.; Farrell, S.; Fricker, H.; Gardner, A.; Harding, D.; et al. The Ice, Cloud, and land Elevation Satellite-2 (ICESat-2): Science requirements, concept, and implementation. *Remote Sens. Environ.* **2017**, *190*, 260–273. [[CrossRef](#)]
10. Silva, C.A.; Duncanson, L.; Hancock, S.; Neuenschwander, A.; Thomas, N.; Hofton, M.; Fatoyinbo, L.; Simard, M.; Marshak, C.Z.; Armston, J.; et al. Fusing simulated GEDI, ICESat-2 and NISAR data for regional aboveground biomass mapping. *Remote Sens. Environ.* **2021**, *253*, 112234. [[CrossRef](#)]
11. Schwartz, M.; Ciais, P.; Ottlé, C.; De Truchis, A.; Vega, C.; Fayad, I.; Brandt, M.; Fensholt, R.; Baghdadi, N.; Morneau, F.; et al. High-resolution canopy height map in the Landes forest (France) based on GEDI, Sentinel-1, and Sentinel-2 data with a deep learning approach. *arXiv* **2022**, arXiv:2212.10265. [[CrossRef](#)]
12. Lang, N.; Jetz, W.; Schindler, K.; Wegner, J.D. A high-resolution canopy height model of the Earth. *Nat. Ecol. Evol.* **2023**, *7*, 1778–1789. [[CrossRef](#)]
13. Potapov, P.; Li, X.; Hernandez-Serna, A.; Tyukavina, A.; Hansen, M.C.; Kommareddy, A.; Pickens, A.; Turubanova, S.; Tang, H.; Silva, C.E.; et al. Mapping global forest canopy height through integration of GEDI and Landsat data. *Remote Sens. Environ.* **2021**, *253*, 112165. [[CrossRef](#)]
14. Mugabowindekwe, M.; Brandt, M.; Chave, J.; Reiner, F.; Skole, D.L.; Kariryaa, A.; Igel, C.; Nsengiyumva, O.; Uwizeye, F.K.; Byiringiro, J.C.; et al. Nation-wide mapping of tree-level aboveground carbon stocks in Rwanda. *Nat. Clim. Change* **2023**, *13*, 91–97. [[CrossRef](#)]
15. Goetz, S.J.; Baccini, A.; Laporte, N.T.; Johns, T.; Walker, W.; Kellndorfer, J.; Houghton, R.A.; Sun, M. Mapping and monitoring carbon stocks with satellite observations: A comparison of methods. *Carbon Balance Manag.* **2009**, *4*, 2. [[CrossRef](#)]
16. Nascetti, A.; Yadav, R.; Brodt, K.; Qu, Q.; Fan, H.; Shendryk, Y.; Shah, I.; Chung, C. BioMassters: A Benchmark Dataset for Forest Biomass Estimation using Multi-modal Satellite Time-series. *Adv. Neural Inf. Process. Syst.* **2024**, *36*.
17. Tolan, J.; Yang, H.I.; Nosarzewski, B.; Couairon, G.; Vo, H.V.; Brandt, J.; Spore, J.; Majumdar, S.; Haziza, D.; Vamaraju, J.; et al. Very high resolution canopy height maps from RGB imagery using self-supervised vision transformer and convolutional decoder trained on aerial lidar. *Remote Sens. Environ.* **2024**, *300*, 113888. [[CrossRef](#)]
18. Dubovik, O.; Schuster, G.L.; Xu, F.; Hu, Y.; Bösch, H.; Landgraf, J.; Li, Z. Grand challenges in satellite remote sensing. *Front. Remote Sens.* **2021**, *2*, 619818. [[CrossRef](#)]
19. Fassnacht, F.E.; White, J.C.; Wulder, M.A.; Næsset, E. Remote sensing in forestry: Current challenges, considerations and directions. *For. Int. J. For. Res.* **2024**, *97*, 11–37. [[CrossRef](#)]
20. Guan, H.; Skutsch, M.; Paneque-Gálvez, J.; Ghilardi, A. Remote sensing of forest degradation: A review. *Environ. Res. Lett.* **2020**, *15*, 103001.
21. Chen, X.; Wang, R.; Shi, W.; Li, X.; Zhu, X.; Wang, X. An Individual Tree Segmentation Method That Combines LiDAR Data and Spectral Imagery. *Forests* **2023**, *14*, 1009. [[CrossRef](#)]
22. Kaartinen, H.; Hyyppä, J.; Yu, X.; Vastaranta, M.; Hyyppä, H.; Kukko, A.; Holopainen, M.; Heipke, C.; Hirschmugl, M.; Morsdorf, F.; et al. An international comparison of individual tree detection and extraction using airborne laser scanning. *Remote Sens.* **2012**, *4*, 950–974. [[CrossRef](#)]
23. Zhang, Y.; Liao, Q.; Ding, L.; Zhang, J. Bridging 2D and 3D segmentation networks for computation-efficient volumetric medical image segmentation: An empirical study of 2.5 D solutions. *Comput. Med. Imaging Graph.* **2022**, *99*, 102088. [[CrossRef](#)]
24. Csillik, O.; Kumar, P.; Mascaro, J.; O’Shea, T.; Asner, G.P. Monitoring tropical forest carbon stocks and emissions using Planet satellite data. *Sci. Rep.* **2019**, *9*, 17831. [[CrossRef](#)]
25. Kerbl, B.; Kopanas, G.; Leimkuehler, T.; Drettakis, G. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Trans. Graph. (TOG)* **2023**, *42*, 1–4. [[CrossRef](#)]

26. Yu, Z.; Chen, A.; Huang, B.; Sattler, T.; Geiger, A. Mip-splatting: Alias-free 3d gaussian splatting. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 16–22 June 2024; pp. 19447–19456.
27. Yan, C.; Qu, D.; Xu, D.; Zhao, B.; Wang, Z.; Wang, D.; Li, X. Gs-slam: Dense visual slam with 3d gaussian splatting. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 16–22 June 2024; pp. 19595–19604.
28. Lin, J.; Li, Z.; Tang, X.; Liu, J.; Liu, S.; Liu, J.; Lu, Y.; Wu, X.; Xu, S.; Yan, Y.; et al. VastGaussian: Vast 3D Gaussians for Large Scene Reconstruction. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA, 16–22 June 2024; pp. 5166–5175.
29. Zhou, X.; Lin, Z.; Shan, X.; Wang, Y.; Sun, D.; Yang, M.H. DrivingGaussian: Composite Gaussian Splatting for Surrounding Dynamic Autonomous Driving Scenes. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA, 16–22 June 2024; pp. 21634–21643.
30. Giang, K.T.; Song, S.; Jo, S. TopicFM: Robust and interpretable topic-assisted feature matching. In Proceedings of the AAAI Conference on Artificial Intelligence, Washington, DC, USA, 7–14 February 2023; Volume 37, pp. 2447–2455.
31. Schönberger, J.L.; Frahm, J.M. Structure-from-Motion Revisited. In Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 27–30 June 2016.
32. Pan, L.; Baráth, D.; Pollefeys, M.; Schönberger, J.L. Global structure-from-motion revisited. In Proceedings of the European Conference on Computer Vision, Milan, Italy, 29 September–4 October 2024; Springer: Cham, Switzerland, 2025; pp. 58–77.
33. Fischler, M.A.; Bolles, R.C. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM* **1981**, *24*, 381–395. [[CrossRef](#)]
34. Pix4D S.A. *Pix4Dmatic*, Versions 1.70; Pix4D S.A.: Prilly, Switzerland, 2024.
35. Liu, T.; Ye, X.; Zhao, W.; Pan, Z.; Shi, M.; Cao, Z. When Epipolar Constraint Meets Non-Local Operators in Multi-View Stereo. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), Paris, France, 1–6 October 2023; pp. 18088–18097.
36. Ye, V.; Turkulainen, M.; Li, R.; Kerr, J.; Yi, B.; Pan, Z.; Ye, J.; Tancik, M.; Kanazawa, A. gsplat: A Gaussian Splatting Implementation for 3D Reconstruction. Version 1.0. 2023. Available online: <https://github.com/nerfstudio-project/gsplat> (accessed on 2 August 2024).
37. Shuai, Q.; Guo, H.; Xu, Z.; Lin, H.; Peng, S.; Bao, H.; Zhou, X. Real-Time View Synthesis for Large Scenes with Millions of Square Meters. *arXiv* **2024**, arXiv:2404.01133v1.
38. Lu, T.; Yu, M.; Xu, L.; Xiangli, Y.; Wang, L.; Lin, D.; Dai, B. Scaffold-gs: Structured 3d gaussians for view-adaptive rendering. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 16–22 June 2024; pp. 20654–20664.
39. Liu, Y.; Luo, C.; Fan, L.; Wang, N.; Peng, J.; Zhang, Z. Citygaussian: Real-time high-quality large-scale scene rendering with gaussians. In Proceedings of the European Conference on Computer Vision, Milan, Italy, 29 September–4 October 2024; Springer: Cham, Switzerland, 2025; pp. 265–282.
40. Zhang, W.; Qi, J.; Wan, P.; Wang, H.; Xie, D.; Wang, X.; Yan, G. An easy-to-use airborne LiDAR data filtering method based on cloth simulation. *Remote Sens.* **2016**, *8*, 501. [[CrossRef](#)]
41. Ronneberger, O.; Fischer, P.; Brox, T. U-net: Convolutional networks for biomedical image segmentation. In Proceedings of the Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, 5–9 October 2015; Proceedings, Part III 18; Springer: Cham, Switzerland, 2015; pp. 234–241.
42. Esri. *ArcGIS Desktop: Release 10.8*; Environmental Systems Research Institute: Redlands, CA, USA, 2020.

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.