

Towards Real-Time Forest Inventory using Handheld LiDAR

Alexander Proudman^a, Milad Ramezani^{a,b,*}, Sundara Tejaswi Digumarti^a, Nived Chebrolu^{a,**} and Maurice Fallon^a

^a*Oxford Robotics Institute, The George Building, 23 Banbury Road, Oxford, UK - OX2 6NN.*

^b*CSIRO Data61, Queensland Centre for Advanced Technologies, 1 Technology Court, Pullenvale QLD 4069, Australia*

ARTICLE INFO

Keywords:

Forestry
SLAM
Real-time Operation
Automated forest inventory

ABSTRACT

While mobile LiDAR sensors are increasingly used to scan in ecology and forestry applications, reconstruction and characterisation are typically carried out offline. Motivated by this, we present an online LiDAR system which can run on a handheld device. Our system is capable of creating 3D point cloud reconstructions of large forest areas, segment and track individual trees, and create an inventory for the detected trees. Segments relating to each tree are accumulated over time, and tree models are completed as more scans are captured from different perspectives. All the processing is performed on-board in an online fashion, and immediate feedback can be provided to the operator via a screen mounted on the device. This is a critical feature as it allows the operator to ensure the desired area is mapped satisfactorily without any gaps or missing sections. We employ a pose-graph based SLAM system with loop closures to correct for drift errors allowing us map large areas accurately. Our mapping system also provides multi-session capability where data captured during different runs times can be automatically merged in a post-processing step. In this work, we estimate the Diameter at Breast Height (DBH) of individual trees an example parameter for the forest inventory. The DBH parameter is estimated online by fitting a cylinder to each tree trunk through a least-squares optimization within an RANSAC loop. We demonstrate our mapping approach operating online in two different forests (both ecological and commercial) with total travel distance spanning several kilometers. Further, we also provide experimental results comparing our DBH estimation to ground-truth measurements recorded manually in an ecological forest (Wytham Woods, Oxford). We demonstrate that our DBH estimates are within ~ 7 cm accuracy for 90% of individual trees detected in the dataset.

1. Introduction

Researchers in ecology and forestry monitor the size and growth of individual trees within a forest to infer arboreal health. It is common to retrieve quantitative metrics such as Diameter at Breast Height (DBH) using manual methods, such as measuring tapes, to ensure consistency with previous surveys. Such metrics go towards forming the forest inventory and describes the state of the forest.

Over the last decade, 3D LiDAR scanners are being used to map forests and extract various metrics from the point cloud data. Automated methods have been developed for tree segmentation and the construction of Quantitative Structure Models (QSMs) [14]. These models can be used to estimate the above-ground biomass and carbon stock by estimating the volume of individual trees [2, 16]. These kind of analysis has traditionally being done on high-quality point cloud data captured with Terrestrial Laser Scanners (TLS), which is typically expensive to obtain. Furthermore, computation of these tree-scale metrics requires data processing of a large-scale point cloud which is often conducted offline as a post-processing step.

*This presented work was done when the author was with Oxford Robotics Institute.

**Corresponding author

✉ alexander.proudman@wadhams.ox.ac.uk (A. Proudman);

milad.ramezani@data61.csiro.au (M. Ramezani); tejaswid@robots.ox.ac.uk (S.T. Digumarti); nived@robots.ox.ac.uk (N. Chebrolu);

mfallon@robots.ox.ac.uk (M. Fallon)

ORCID(s): 0000-0002-7529-9903 (M. Ramezani); 0000-0002-0004-0868 (S.T. Digumarti); 0000-0001-6408-4459 (N. Chebrolu); 0000-0003-2940-0879 (M. Fallon)

In this work, we present an online forestry mapping system capable of providing real-time feedback to the operator on the ground. Our mapping system uses a hand-held LiDAR setup with on-board compute capability (See Fig. 2). This provides an easier and more flexible surveying experience for the operator compared to traditional TLS systems.

The design of the mapping system takes into account the challenges of typical forestry surveys. For example, we use an elastic pose-graph structure as the underlying representation, which allows us to scale the mapping to large areas. We also support multi-session mapping allowing the operator to survey the forest in multiple runs, which may be necessary due to constrained battery powers or operator fatigue. Furthermore, we present an online point-cloud processing tool chain which can segment and track individual trees. The ground surface can also be inferred, allowing us to estimate the DBH of each tree (at 1.4m height) as the scanning process is carried out.

The main contributions of this work are as follows:

- An online mapping system for forest environments using a handheld LiDAR capable of real-time feedback for the operator to aid in the survey mission.
- Segmentation, tracking and estimation of DBH parameter for individual trees in the forest.
- A pose-graph based SLAM system with loop closure integration to correct for odometry drift, and to maintain an accurate tree map over large areas.
- A multi-session mapping capability where data cap-




Figure 1: Visualisation of the system output. **Top:** a high level view showing a 150 m transect. The path taken by the person carrying the sensor in Fig. 8 is shown in black. **Centre and Bottom:** Detail of a section is shown with two different representations: one showing each tree segmented and individually coloured and the other shows a cylinder (primary axis and diameter) fitted at breast height.

tured during multiple runs can be merged in a post-processing step.

- Demonstration of the system on challenging large scale datasets from ecological as well as commercial forests spanning several kilometers.

We also performed a quantitative evaluation of our DBH estimates against measurements taken manually at our test site for comparison against a commercial tool (using the final accumulated point cloud). We further perform some studies to characterise the behaviour of different modules in our system.

Video: <https://tinyurl.com/ecmr-forest-mapping>

2. Related Work

Advances in LiDAR technology have led to it being used in non-traditional fields such as ecology, forestry and remote sensing. Here, we review some of the state-of-the-art approaches for tree segmentation from LiDAR point clouds. We first review techniques developed for high end Terrestrial LiDAR scanners followed by more recent approaches developed for mobile LiDAR scanners.

2.1. Tree Segmentation Using Terrestrial LiDAR

Many existing tree segmentation techniques operate on a point cloud accumulated from multiple static scanning locations using a Terrestrial Laser Scanner (TLS), in post-processing. Raunonen *et al.* [13] and Trochta *et al.* [19] clustered these scans into smaller point clouds to work around large size of these scans and to reduce memory consumption. They then segmented tree-level clouds within each cluster by growing segments using assumptions of fixed inter-cluster distance and orientation to infer connectivity. Methods such as [23] utilize concepts from graph theory to find the connectivity between adjacent points and segment each tree. These approaches, however, rely on multiple assumptions about a tree's architecture as well as assuming minimal interconnection between their crowns.

In [1] Burt *et al.* present a software package named *treeseq* that uses region-growing technique to segment individual trees. Key features in *treeseq*'s design are its independence of forest type, scanning instrument and no assumptions about the tree structure.

2.2. Tree Segmentation Using Mobile LiDAR

Most of the above approaches are intended primarily for very high quality TLS point clouds. Their performance drops when applied to noisier point clouds captured from Aerial Laser Scanners (ALS) and ground based mobile laser scanners (MLS). However, the lower cost and increased portability of mobile laser scanners in comparison to TLS outweighs makes them a desirable pick despite the measurement data being noisier. As a result several works have looked at tree segmentation and automatic inventory generation systems for these scanners.

Heo *et al.* [6] used mobile LiDAR to collect data in an urban area, including parks and streets, to estimate the height of trees and their DBH by calculating the height-above-ground and using a least-squares circle fit approach [11]. They emphasised the advantage of using mobile LiDAR to reduce shadow and occlusion effects, which are more prominent with terrestrial LiDAR systems, especially in an urban forest environment. They utilised a Stencil LiDAR system produced by Kaarta¹ which includes a Velodyne sensor.

Similarly, Zhou *et al.* [24] collected LiDAR data with a Velodyne VLP-16 LiDAR sensor. In offline processing the authors estimated the DBH of the trees. They removed the points residing on the ground and estimated the DBH using Random Sample Consensus (RANSAC) algorithm on

¹<https://www.kaarta.com/>



Figure 2: The handheld device used in our system. The labelled components work together to provide online feedback on the forestry data collection

the segments produced by an Euclidean-based clustering algorithm [18].

Westling *et al.* [20] scanned individual avocado trees with a 5 m spacing using a GeoSLAM Zebedee 1 handheld device. The authors first voxelised point clouds and conducted a graph-based search over the voxels to find all paths connecting to a root voxel, which was considered to be the tree node. The tree node is segmented from the ground by comparing the height of points locally within a search radius.

More recently data driven and learning based approaches have also been used for tree segmentation. Digumarti *et al.* digumarti2018automatic train a random forest classifier to segment individual trees into their component structures. Convolutional Neural Networks (CNNs), trained on simulated tree data are used in [17] to segment trees from colour and depth images. Windrim and Bryson [21] detect trees in point clouds collected using ALS by encoding the point clouds as a 2D raster image and detecting points of high density. This is followed by applying a segmentation network, adapted from *PointNet* [12], to segment trees. Krisanski *et al.* [9] further extend the idea of learning based tree segmentation and present a sensor agnostic approach that works point clouds of different densities.

The above presented approaches typically work as post-processing techniques applied after data collection. This does not allow the operator to perceive the reconstruction during scanning. Our motivation is to develop a LiDAR-driven technique which can reconstruct point clouds, extract individual trees and estimate their structural parameters in dense forests with rough terrain in real-time. A real-time mapping system can ensure full coverage by providing feedback to the operator, to help identify gaps in the scanned region, enable rescanning and ensure that the environment is fully scanned and processed at runtime.




Figure 3: An overview of the system pipeline, the frequencies of communication between blocks are given.

3. Approach

Our device consists of an Ouster OS0-128 LiDAR and an Intel Realsense D435i, each with a built-in IMU (Fig. 2). However, we only use the data from the LiDAR sensor and its IMU measurements for our LiDAR odometry and tree tracking systems. The LiDAR has 128 beams and a 90° vertical field of view (FOV). While it cannot completely cover the environment vertically, the 90° vertical FOV is sufficient to identify individual trees within about 20 m of the sensor.² The sensors used in systems such as Zeb-Revo or Zebedee² have a narrow field of view and require the sensor to be either rotated or oscillated to cover the environment entirely.

Our proposed system architecture is shown in Fig. 3. Due to the system's online nature, the pipeline is different from many existing methods, and alterations have been made to parallelise and speed up the overall pipeline.

In the following sections, we describe the algorithms and methods used by each functional block within the system architecture. We start with a LiDAR-inertial odometry module to estimate the pose of the sensor which is accurate over short distances (Sec. 4.1). The estimated odometry is fed to the SLAM module (Sec. 4) which builds a pose-graph structure as the sensor moves in the environment. Local maps created by accumulating laser scans are attached to the nodes of the pose graph. When a loop closure is detected on revisiting a place, the pose estimates are corrected for drift. The pose-graph SLAM module allows us to scale the mapping to larger areas in the forest, and extend the mapping capability to multiple-sessions.

Parallel to the pose-graph SLAM module, we process the incoming LiDAR scans to segment and track trees. The *Tree Tracking* block (Sec. 5), not used in the typical LiDAR forest inventory systems, segments, tracks and fuses each tree detection. Finally, we also compute the DBH parameter for each of the detected trees (Sec. 6). Advantages of using this online fused data approach are that:

- Prior inventory data can be used to automatically identify trees.

²<https://geoslam.com/>




Figure 4: Pose-graph representation of the SLAM system. Each node represents the pose of the sensor whereas the edges represent the constraints coming from multiple sources.

- Success of the data collection can be evaluated online, including the sensor's coverage. This allows the operator to have a real time feedback of the reconstruction and ensure that desired area is being mapped completely without any gaps/missing sections in the surroundings.
- Data collection and analysis only takes the amount of time needed to walk along a selected transect (a narrow section of land along which measurements are taken).

4. Localization and Mapping

4.1. LiDAR Odometry

Our LiDAR odometry system [22] is a factor-graph based windowed smoother which fuses Inertial Measurement Unit (IMU) readings with measurements from a multi-beam LiDAR scanner mounted on a handheld device (Fig. 2). The particular configuration used in this work, uses IMU pre-integration to remove motion distortion from scans and to initialise Iterative Closest Point (ICP) registration, specifically the implementation of Pomerleau [10], to determine the relative transform for every 2 metres of distance travelled.

The IMU measurements and the relative transforms then form constraints in a factor graph to estimate the poses in a sliding window optimisation utilising iSAM2 [7] (as part of the GTSAM library). Additionally, since IMU provides gravity information, we can align the clouds with the gravity direction. This alignment helps determine the base of trees using the z direction. Evaluation of our LiDAR odometry system is discussed in Sec. 7.1.

4.2. Pose-Graph SLAM with Loop Closures

The LiDAR odometry module provides accurate local tracking over short distances (~ 100 m). However, as we travel further, the pose estimated by the odometry system accumulates drift. This drift will in turn effect other parts of the pipeline, and can potentially result in incorrect association of the previously generated maps and the trees segmented in them.

We overcome this problem by implementing a pose-graph SLAM system which corrects for odometry drift using loop closure constraints. Fig. 4 shows the pose-graph structure where nodes represent the pose of the sensor, and the edges represent constraints from multiple sources. New nodes are




Figure 5: Mapping over multiple sessions as post-processing step. Loop closures between the two sessions are computed (in purple) and jointly optimized. A global map is then computed using the jointly optimized poses.

added to the pose-graph as the device moves through the environment. We initialize a new node after traveling a distance of 5 m. With each node x_i in the pose-graph, we associate a local sub-map $\mathcal{M}_{x_i}$ obtained by accumulating the LiDAR scans using the odometry estimate computed in Sec. 4.1. Here, we assume that the drift accumulated within a local map is small, and does not effect the tree segmentation and tracking procedures. The point cloud of each local map is transformed into a co-ordinate system whose origin corresponds to a pose from the SLAM node.

The consecutive nodes in the graph are constrained with LiDAR odometry (shown in orange), whereas loop closure constraints (shown in pink) are obtained on re-visiting a place within the existing map. The entire pose-graph is optimized at regular intervals, which ensures that the current pose estimates are computed with all the constraints in the graph. During optimization, the loop closure constraints are responsible for correcting the drift error.

In our online implementation, we detect the loop closures between non-consecutive nodes in the graph based on their distances in the global coordinate frame. We consider nodes within a distance of 3 m as loop closure candidates. We then retrieve the local sub-maps associated with the corresponding nodes, and perform a scan matching step using ICP. If a successful registration is obtained between the two sub-map point clouds, we add a loop closure constraint in the pose-graph given by the transformation estimated by ICP procedure. Through the regular addition of loop closure constraints and reoptimization, the drift error is corrected.

Finally, we obtain the combined map of the environment $\mathcal{M}$ by transforming the local sub-maps $\mathcal{M}_{x_i}$ in to the global frame using the pose of the corresponding node in the pose-graph.

The pose-graph SLAM system with loop closure integration allows us to extend our tree mapping approach to larger areas than was possible using LiDAR odometry only.

4.3. Multi-Session Mapping

While the pose-graph SLAM system in Sec. 4.2 allows us to map large areas, forestry survey missions can require multiple scanning sessions. This may be due to the limited battery capacity of the hand-held device, or to provide a break for the operator to limit fatigue. In certain cases, multiple operators would scan the forests separately to cover larger areas. In all these scenarios, it becomes critical for an effective mapping system to be able to merge maps gener-

ated from multiple sessions.

To achieve multi-session mapping capability, we propose an approach where we exploit the pose-graph structures built during individual missions, and merge them in a post-processing step. At the end of each session, we save the pose-graph as well as the corresponding local maps to the disk. Once all the sessions are completed, we use a place recognition system by Giseop *et al.* named *ScanContext* [8] to propose loop closure constraints between the different sessions. *ScanContext* computes a compact descriptor for each local map in our pose-graph. The descriptor captures the distribution of the points in the 3D space and creates a discriminative signature for the point cloud. We compute these descriptors for each local map in all of the sessions, and compares them to retrieve a set of loop-closure candidates. The loop-closure candidates correspond to the places where the local maps created during different sessions overlap with each other. Note that our approach assumes that the individual sessions have an overlap with at least one other session in the entire mission.

We retrieve these loop closure candidates for each session pair, and finally create a combined pose graph consisting of individual pose-graphs along with the inter-session loop closures. An illustrative figure with two pose-graphs $\mathcal{PG}_1$ and $\mathcal{PG}_2$ along with the loop closures between them is shown in Fig. 5. We then optimize the combined pose-graph to obtain the final poses of the missions in a common coordinate frame. Finally, we obtain the combined map over all sessions by transforming each local map into a common coordinate frame. This is similar to the final map generation step performed for a single pose-graph.

5. Tree Segmentation and Tracking

In parallel to the pose-graph SLAM process, we also segment and track trees from the LiDAR scans. The tree segmentation and tracking uses the point clouds accumulated with LiDAR odometry. The segmentation results are then integrated into pose-graph SLAM system as illustrated in Fig. 3.

To roughly extract the tree trunks, branches, canopy and shrubbery in the original point cloud, Euclidean segmentation is carried out. Using the PCL library³, the cloud is first downsampled using a voxel filter to lessen the radial variation of the point density and to speed up the segmentation process as also done by [1].

After voxel filtering the point cloud, it is reformulated into a k -d tree to quickly find the nearest neighbours of every point. These points are clustered together into groups of points which are within some distance threshold of each other.

Note that the parameters of the voxel filtering and Euclidean segmentation methods are tuned based on the LiDAR sensor's characteristics to maximise the chances that the point clusters contain individual trees.

³<https://pointclouds.org/>




Figure 6: Filtering process on the elevation map: Original elevation map (top) is filtered to remove spikes (middle). Finally, holes are filled using a morphological closing filter on the grid map (bottom).

5.1. Elevation Mapping

To perform *Base Segmentation* of the trees being tracked, we employ an open-source sensor-centric elevation mapping framework⁴ which is built upon a universal Grid Map library [5]. The algorithm, presented in [4], generates consistent elevation maps of the environment surrounding the sensor, at the same frequency of the LiDAR output (10 Hz). In this work, we use a resolution of 16 cm for the grid of $32\text{m} \times 32\text{m}$.

Nevertheless, due to the presence of foliage and branches, the terrain created by the elevation mapping software can contain phantom spikes, making it difficult to detect the precise base of certain trees. To solve this problem, we compute the slope of each grid cell. Cells with slope greater than a threshold are removed. This removes the majority of spikes, resulting in smooth terrain with some remaining holes. Finally, we utilise a morphological closing filter on the grid map to fill the holes.

Applying the chain of filters is relatively time consuming so we update the elevation map every 8 m of travel distance based on the state estimate from our LiDAR odometry. This travel distance allows online operation. We analyse the computation time of each component including elevation mapping in the next section. Fig. 6 shows this process.

5.2. Tree Tracking

To facilitate a formal definition of a tree the following minimal, ‘tree descriptor’ t was defined as follows:

- A unique id,
- Major axis of the tree incline ($\mathbf{I}$ a vector),
- Diameter at Breast Height (DBH) (D),
- 3D position of the tree base ($\mathbf{b}$),

⁴https://github.com/ANYbotics/elevation_mapping

- The minimum and maximum height of the point cloud defining the tree (*min* & *max*).

Every tree descriptor is derived from a corresponding point cloud of the tree points P ($t = f(P)$). These points are accumulated over time, gathering information from different angles as well as reducing the effect of occlusion and improving the estimate of the tree descriptor as Heo *et al.* [6] indicated.

Most of the descriptors characteristics are easily evaluated with the exception of the major axis of the tree. We used fitting of a line, to the set of tree points (P), to find this axis. The algorithm is designed to ignore branches and any potentially segmented crown or ground points and is described in Sec. 5.2.1.

Internally, the *Tree Tracker* holds an inventory of n distinct trees T defined by a descriptor and set of tree points $T_i \sim t_i, P_i$ for $0 \leq i \leq n$.

Referring to the system diagram, Fig. 3, incoming clusters C from the *Segmentation* block are either discarded or assigned to a new/existing tree descriptor. If assigned to an existing descriptor, the cluster of points is merged with the corresponding set of tree points and the descriptor is re-evaluated $P_{new} = P_{old} \cup C, t_{new} = f(P_{new})$.

An assign/discard decision is made by evaluating a tree descriptor for each input cluster. A cluster is confirmed and subsequently assigned only if it has characteristics within the tolerances of given parameters. The parameter tolerances that resulted in the most consistent tree assignment are:

- Requiring the major axis of the tree to be close to vertical, i.e. $|\mathbf{I} \cdot [0, 0, 1]| < \theta_{\text{threshold}}$
- Requiring a minimum height, i.e. $\max - \min < h_{\text{threshold}}$

Should a cluster satisfy the above properties we deem that it contains points belonging to a tree. It is then compared to the existing set of confirmed trees T and merged if a match is found. In order for a match between two clusters to be found, their major axis must converge to within some threshold distance at the base of the highest cluster or at a plane segmenting them. Fig. 7 (left) demonstrates an example of this.

Alg. 1 summarises the *Tree Tracking* pipeline. The evaluation of a match and the subsequent merge process is assessed until no more matches are found for an input cluster. This is important to ensure all potential clusters belonging to a tree are merged.

5.2.1. Robust estimation of the tree's major axis

As mentioned earlier, the calculation of the tree's major axis, or a clusters major axis, is done using regression on a subset of the tree's points, to find a line of best fit. Beginning with regression utilising every point in the cluster gives an initial estimate of the incline ($\mathbf{I}$) and its intercept.

However, since the objective is to find the axis of the tree's trunk, least squares regression can be affected when




Figure 7: An example of *Tree Tracking*. **Left:** The two clusters C with their major axis shown and the plane of between them. **Right:** Line of best fit on all tree points (green) and on the points excluding branch and ground points (red).

the branches, ground or crown points are included in the segmentation. Therefore, points that are a multiple of the standard deviation away from the calculated line or outside the central 90% of the tree's height are ignored. The regression is re-calculated with the subset of points. This is iterated multiple times with successively smaller tolerances on the inclusion of points. This result is shown in Fig. 7 (right).

5.3. Base Segmentation

As the ground has not been removed prior to segmentation, points that are not necessarily portions of the tree have the potential to be segmented. As a result of this the lowest point in P (minimum z value) may not be an accurate

Algorithm 1: Tree Tracking Algorithm.

```

input: Set of point clouds  $C$ 
output: Set of updated trees  $T$ 
begin
  for Cluster  $C_i \in C$  do
    Calculate the tree descriptor  $t_i = f(C_i)$ 
    if  $t_i$  is classified as being a tree then
      if a first match for  $t_i$  is found in  $T$  ( $T_j$ ) then
        merge cluster into the set of tree points
         $P_j = P_j \cup C_i$  and re-calculate the tree
        descriptor  $t_j = f(P_j)$ 
        while a subsequent match for the
        updated  $t_j$  is found in  $T$  ( $T_{imp}$ ) do
          merge the two sets of tree points
           $P_j = P_j \cup P_{imp}$  and re-calculate
          the tree descriptor  $t_j = f(P_j)$ 
          delete  $T_{imp}$  from  $T$ 
        end
      end
      else
        Push back the  $t_i, C_i$  into the set  $T$ 
      end
    end
  end
end

```

estimate of the tree's base. We segment the base as a result.

The two most recent overlapping elevation maps, from Sec. 5.1, are converted into a point cloud so that a search over these 'ground points' can be achieved. Points within a 2 metre radius of the lowest point in P are identified, using a nearest neighbour method from the PCL library. From this set, the closest point to the tree's axis is found and set as the tree's base $\mathbf{b}$.

6. Estimation of Diameter at Breast Height (DBH)

To estimate the DBH of each individual tree, we employ a procedure similar to that used in the literature for general cylinder fitting [11]. This procedure is broken down into three parts: firstly, segmenting the points at breast height, secondly, projection of these points onto a plane, and lastly, fitting a circle to the projected points.

Utilising the method described by Zhou *et al.* [24] and Heo *et al.* [6], our system segments the LiDAR points, for recently updated trees in T . Points (S), from the accumulated point cloud (P), that are located within a 10 cm height range centered 1.4 m above the tree's base are segmented. These points are projected onto a plane with normal ($\hat{\mathbf{n}}$) equal to the tree's incline ($\mathbf{I}$) to ensure the set of points can be modelled closely to a circle. The alternative is to use elliptical fitting for trees angled from the vertical. In practice the sensor only captures data from a single side of the tree, which may cause overfitting when using an elliptical distribution, especially for trees which have irregularly curved surfaces.

The potential problem with this method is that these points are taken from the downsampled point cloud so are likely to be inaccurate. However, an advantage of this method is that points should cover a greater proportion of the tree's surface if the LiDAR sensor measured it from a set of different angles.

RANSAC circle fitting is then used on the set of resultant 2D points, which is robust to potential outliers from errors in segmentation [15].

7. Experiments and Evaluation

The described system was implemented in C++ and the Robot Operating System (ROS) was used for communication between the hand-held device and the software modules.

We evaluated our system at two locations. The first one is an ecological forest called Wytham Woods in Oxford, UK. The woods are spread over 400 ha in area and has been used for ecology research for over 80 years. It is a Smithsonian ForestGEO site; which is a collection of forestry plots located world-wide used for collective ecology research. Within this scheme, it has taken part in 3 mass censuses and about 16200 trees have been manually measured. The second location for our evaluation is a commercial pine-tree forest in one hour north of Helsinki, Finland spread over several kilometers. These two locations differ in the type and densities of trees as well as different profiles of the underlying terrain.




Figure 8: Visualisation of the trees selected for measurement within Wytham Woods, Oxford, as well as the path taken using the handheld device. The location of the chosen transect is $51^{\circ}46'60.59''\text{N}$, $001^{\circ}20'21.47''\text{W}$.

We also perform a quantitative analysis of our DBH estimations by collecting ground truth data for a particular transect at Wytham Woods. We selected 55 trees along a ~ 150 m transect in the forest. The selected trees had various diameters, ranging from ~ 10 cm to ~ 80 cm, which could be reasonably measured manually at breast height and had clearly visible bases.

The transect of the forest used for the evaluation is shown in Fig. 8. The DBH of every tree was measured using a tape measure at a height of 1.4 m from the ground prior to the LiDAR data collection. Two types of LiDAR sensors were used: The mobile LiDAR device specified earlier and a Leica BLK terrestrial LiDAR⁵ (for comparison). The path taken by the mobile LiDAR can be seen in Fig. 8 and Fig. 1, while a terrestrial LiDAR map was made along the central path to generate the ground truth for evaluation of our LiDAR odometry system and to provide a performance baseline.

7.1. LiDAR Odometry

To evaluate the odometry component of the system, we used the reconstruction created by the Leica BLK as a prior map. We then registered the individual point clouds from the hand-held LiDAR against the prior map to generate an accurate 6 degree-of-freedom ground truth trajectory at 10 Hz. The estimated trajectory from the lidar odometry system (without access to the Leica map) was then compared against the ground truth trajectory using Relative Translation Error (RTE) at travel distances of between 10 metres and 50 metres to produce Fig. 9.

We aim to accumulate and process LiDAR scans for individual trees as we walk up and pass them at ranges up to 20 m, thus having a rate of just a few centimetres on this scale is satisfactory.

7.2. Tree Tracking

The *Tree Tracking* system performed satisfactorily within the constraints of online performance and the density of foliage and trees on the test transect. Occasionally, trees were

⁵<https://leica-geosystems.com/en-gb/products/laser-scanners/scanners/blk360>




Figure 9: Relative translation error of the LiDAR odometry at different subtrajectory presented as a series of boxplots. Each box indicates the 25 and 75 percentiles of the estimation errors, the horizontal line through the box the median, and the whiskers are the maximum and minimum excluding outliers.

clustered together when they should not have been, this is most often due to low hanging interconnecting branches or foliage close to a tree's base. Out of the 55 measured trees 3 of them contained points which belonged to other objects. It was observed that this was also partially due to a transient effect in *Tree Tracking* in which assignment between nearby clusters is uncertain when the tree is first detected. This problem is shown in Fig. 10.

On the other hand, from the 55 measured trees, one failed to be automatically segmented due to its small size, having a DBH ~ 0.11 m. This indicates that our current hardware setup and software configuration limits the performance of the system detecting small trees with DBH below about 10 cm.

7.3. DBH Estimation with LiDAR360

To provide a baseline, we used the commercial package LiDAR360, in particular its TLS forestry module, which is designed for tree segmentation and DBH estimation, on both our handheld LiDAR and Leica BLK datasets.

To evaluate the difference between the estimated values of DBH and the manual ground truth measurements, we use the Root Mean Square Error (RMSE) metric.

LiDAR360 successfully identified 53 of the 55 trees in the Leica BLK data. The remaining two, unidentified trees, were manually selected, then DBH estimation was done, across all 55 trees, resulting in an overall RMSE of 0.046 m. The mobile LiDAR point clouds were accumulated, af-




Figure 10: An example of tree tracking failing for trees in close proximity, the ground points clearly interconnect the two trees. **Left:** The two trees shown entirely. **Center and Right:** Segmented tree with incorrectly assigned points.




Figure 11: Estimation of DBH using the LiDAR360 Forestry module (using the mobile LiDAR data). **Left:** Plot of the 55 measured trees vs estimated values as extrapolated from the package. **Right:** Boxplot of the errors between the estimated and measured values, RMSE = 0.056 m.




Figure 12: Estimation of DBH using the designed system. **Left:** Plot of 54 measured trees vs estimated values, one tree was not segmented. **Right:** Boxplot of the errors between the estimated and measured values, RMSE = 0.07 m excluding the outliers and 0.14 m including them.

ter odometry, and also tested using the the same software. Every tree was successfully segmented in the accumulated point cloud and the resulting RMSE was 0.056 m.

The results for LiDAR360's forestry module are shown in Fig. 11. The graphical results for the terrestrial LiDAR was similar to that of the mobile LiDAR so the graph is omitted.

7.4. DBH Comparison with Proposed System

The results for our system are shown in Fig. 12, these results contain data points for every segmented tree (54/55). Three of the trees were poorly segmented which resulted in outlier measurements in the figure. The RMSE of these estimates, not including the outliers is 0.07 m. The RMSE is slightly higher than that of the commercial software, indicating reasonable performance of the system but some room for improvement. Table 1 summarises the DBH estimates and segmentation results for the proposed system and LiDAR360 using Leica BLK and accumulated mobile LiDAR data.

7.5. Timing Analysis

To demonstrate the key feature of our presented system, that it can run online, we carried out an evaluation of the computation time for each component. Fig. 13 demonstrates the time required for the two major components of the system

Algorithm	Detections	RMSE (m)	Mode	Data used
LiDAR360	53/55	0.046	Offline	Leica BLK
LiDAR360	55/55	0.056	Offline	Ouster (accumulated)
Proposed	54/55	0.07/0.14	Online	Ouster (live)

Table 1

Summary of the DBH results for the systems evaluated.




Figure 13: Evolution of computation time, taken over the 622 sec to walk along the Wytham Woods transect, for the two major components. **Top:** The total time for tree segmentation and fitting. **Bottom:** Time for the elevation map filtering process (computed every 8 m of travel distance). The mean and standard deviation are shown by the red solid and dashed lines respectively.

Process	Tree Tracking	Circle Fitting	Euclidean Clustering	Base Segmentation	Total Time
Time (ms)	88	33	66	20	232
std (ms)	±30	±8	±33	±6	±44

Table 2

Computation times for the *Tree Tracking* modules.

— (individual) scan processing in *Tree Tracking* and (accumulated) elevation map filtering. A breakdown of individual modules of the *Tree Tracking* pipeline is given in Table 2.

The *Tree Tracking* system takes about 0.23 sec (about 4 Hz) to process an individual scan, which is almost two times faster than the output of our 2 Hz LiDAR odometry.

The elevation map filtering process takes 3.4 sec on average. The elevation map is processed every ~8 m and walking speeds of ~1 m/s are expected. Therefore, on average, this process takes less than half the time available. Future work aims to optimise this block to run closer to sensor frame rate.

7.6. Mapping with Loop Closures

In this experiment, we demonstrate that our mapping system is capable of scaling to larger areas of forest. We demonstrate our approach on two additional large-scale datasets recorded at Wytham woods, Oxford and a commercial forest one hour north of Helsinki, Finland respectively.

Fig. 14 illustrates the mapping results from the Wytham woods dataset. This dataset covers an area of about 0.5 ha traversed in a lawn-mower pattern with a spacing of about 10 m. Several loop-closures have been detected towards the end of each row allowing the SLAM system to correct for drift error. Fig. 14 shows the point cloud reconstruction as well as the extracted trees. We also see the terrain profile of




Figure 14: Mapping results on Wytham Woods dataset surveyed in a lawn-mower pattern. **Top:** Accumulated point cloud reconstruction and the extracted trees. **Middle:** Cross-sectional views of the clouds shown above. **Bottom:** Zoomed in view of the extracted trees, showing the path of the camera in black.

the forest in the cross sectional view.

Fig. 15, we highlight the effect of the loop closures by plotting the estimated trajectory using LiDAR odometry only (in black), and the SLAM trajectory which uses loop closures (in green). This dataset was recorded the commercial forest in Finland with the length of the total trajectory being 1.3 Km. The trajectory begins and ends at the same position (indicated by Start/End) with several loops made inside the forest area. The odometry estimate drifted over time, and accumulated an error of 25 m on returning to the start position. In contrast, we see that SLAM trajectory returns to the same location on finish indicating that there is no drift error. This has been possible thanks to the corrections made due to the loop closures.

7.7. Mapping over Multiple Sessions

In this experiment, we demonstrate the multi-session mapping capability of our system. We record the data in a commercial pine tree forest in Finland over three sessions. In




Figure 15: Effect of loop closures on the estimated trajectory of the device. Odometry estimate (in black) and SLAM pose estimate with loop closures (in green) for a 1.3 Km long trajectory. The odometry estimate drift can be clearly seen at the start/end position. The drift accumulated for this trajectory is around 25 m.

each session, the human operator starts and ends at a operator vehicle and walks down an access road to a plot of trees. The first session inspected a relatively sparse forest with clear separation between trees. The second was somewhat denser and overlapped with the first. Finally the third area was very dense with un-thinned pine saplings. Fig. 16 illustrates our mapping results of our approach where the area covered in the three sessions is shown. The total travel distance of the combined missions is over 4 kms with each session being roughly equal in length. Note that the three sessions have overlapping areas which is recognized by the loop closure module. An example of such a loop closure is shown in Fig. 19.

This dataset consists of several challenging elements, including long corridor like passages, dense forest patches with limited LiDAR range due to occlusions, transitions in between areas with different densities. Further, as the dataset is captured with a human operator walking on rough surface, the sensor module is often subject to shaking and fast motions. Despite these challenges, our mapping system was successfully able to map each of sessions as well as to merge them together in a post-processing step.

The performance of the *ScanContext* loop proposal mechanism was impressive. Despite being originally developed for automotive localisation, it was still successful in our forestry application. We noticed that loops could be found in the forests with managed tree separation. Performance degraded in unmanaged forests where trees were typically smaller samplings and with little clear separation. A further study would be necessary to properly evaluate place recognition algorithms in this type of environment.

7.8. Mapping with Different Sparsity Levels

In this section, we show that our mapping and segmentation pipeline can deal with different sparsity of LiDAR without additional tuning. To test this capability, we recorded datasets using Ouster LiDAR OS0-128 and Ouster LiDAR OS1-64, which are 128 beam and 64 beam LiDARs respectively. As seen in Fig. 20, OS0-128 provides a much denser point cloud in comparison to OS1-64. The two sensors also have different vertical field of views (FOVs) and maximum ranges. The experimental results suggest that a narrow FOV and range are sufficient for our application of estimating DBH and the motion from odometry. Our algorithms have a voxel filtering step at the beginning which ensures that the subsequent steps are robust to sensor changes.

8. Conclusion and Future Work

This paper presented a mobile LiDAR scanning system for mapping large forest areas, automatic segmentation of trees and the online estimation of their DBH. The architecture of a tree tracking functional block was introduced to facilitate the online operation of this system. We tested the system's performance using data from a natural forest, Wytham Woods, as well as a commercial forest one hour north of Helsinki, Finland to prove that acceptable results can be achieved online which compares favourably to the performance achieved by a commercial software package conducted in post-processing.

In this work, we focused on extracting tree trunks and computing the DBH parameter. In the future, we would like to extend the breadth of metrics used to characterise trees. In particular, we would like to quantify information on the branching of the tree, as well as properties related to the tree crown. We plan to integrate information from the cameras as well as extract semantics to estimate a richer set of quantitative and qualitative metrics for the forest inventory.

One of the limitations identified in the experimental analysis was the downgraded performance on small trees. We like to address this issue by considering a adaptive segmentation algorithm such as in [3] to improve performance on the smaller trees.

Finally, by more fully using the screen we intend to provide feedback to the operator to better direct their actions during operation.

9. Acknowledgements

This research is supported by the UKRI/EP SRC ORCA Robotics Hub (EP/R026173/1). M. Fallon is supported by a Royal Society University Research Fellowship. The authors would like to Taneli Vuornos (Silvere) for the collection of the data used in Sec. 7.7.

Credit authorship contribution statement

Alexander Proudman: Conceptualization, Methodology, Software, Writing. **Milad Ramezani:** Conceptualization, Methodology, Software, Writing. **Sundara Tejaswi**




Figure 16: Mapping results on a large-scale dataset from a commercial pine tree forest in Finland. Figure shows the point cloud reconstruction of the forest obtained by merging all the local maps after joint pose-graph optimization.




Figure 17: Pose-graphs from individual sessions visualized in different colors. Loop closures both with-in the same session, and in-between different sessions in shown in red. The total travel distance of all the 4 km.




Figure 18: Zoomed-in view of a portion of the forest shown in Figure 16 showing dense clusters corresponding to individual trees. The trajectory traversed (green) and the detected loop closures (red) are also shown.

Digumarti: Conceptualization, Methodology, Software, Writing. **Nived Chebrolu:** Conceptualization, Methodology, Software, Writing. **Maurice Fallon:** Conceptualization, Methodology, Software, Writing.




Figure 19: Example of a loop closure detected between Session 1 and Session 2. Top: The corresponding point clouds overlaid on each other after ICP registration. Point cloud from Session 1 is shown in red whereas white points are from Session 2. Bottom: Corresponding camera images of the loop closure candidates are shown. Note that the loop closures are detected only from the LiDAR data, and the images are shown for illustrative purposes only.

References

- [1] Burt, A., Disney, M., Calders, K., 2019. Extracting individual trees from lidar point clouds using treeseg. *Methods in Ecology and Evolution* 10, 438–445.
- [2] Calders, K., Newnham, G., Burt, A., Murphy, S., Raemonen, P., Herold, M., Culvenor, D., Avitabile, V., Disney, M., Armston, J., et al., 2015. Nondestructive estimates of above-ground biomass using terrestrial laser scanning. *Methods in Ecology and Evolution* 6, 198–208.
- [3] Digumarti, S.T., Nieto, J., Cadena, C., Siegwart, R., Beardsley, P.,





Figure 20: Reconstruction and tree extraction results using LiDAR scanners with different sparsity levels. Top: 128 beam LiDAR resulting in a denser point cloud reconstruction. Bottom: 64 beam LiDAR resulting in a sparser point cloud reconstruction. The proposed tree extraction approach is robust to different sparsity levels.

2018. Automatic Segmentation of Tree Structure From Point Cloud Data. *IEEE Robotics and Automation Letters* 3, 3043–3050. doi:10.1109/LRA.2018.2849499.
- [4] Fankhauser, P., Bloesch, M., Hutter, M., 2018. Probabilistic Terrain Mapping for Mobile Robots With Uncertain Localization. *IEEE Robotics and Automation Letters* 3, 3019–3026.
- [5] Fankhauser, P., Hutter, M., 2016. A Universal Grid Map Library: Implementation and Use Case for Rough Terrain Navigation, in: *Robot Operating System (ROS)*. Springer, pp. 99–120.
- [6] Heo, H.K., Lee, D.K., Park, J.H., Thorne, J.H., 2019. Estimating the heights and diameters at breast height of trees in an urban park and along a street using mobile LiDAR. *Landscape and Ecological Engineering* 15, 253–263.
- [7] Kaess, M., Johannsson, H., Roberts, R., Ila, V., Leonard, J.J., Dellaert, F., 2012. iSAM2: Incremental smoothing and mapping using the Bayes tree. *Intl. Journal of Robotics Research* 31, 216–235.
- [8] Kim, G., Kim, A., 2018. Scan context: Egocentric spatial descriptor for place recognition within 3d point cloud map, in: *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pp. 4802–4809. doi:10.1109/IROS.2018.8593953.
- [9] Krisanski, S., Taskhiri, M.S., Gonzalez Aracil, S., Herries, D., Turner, P., 2021. Sensor agnostic semantic segmentation of structurally diverse and complex forest point clouds using deep learning. *Remote Sensing* 13, 1413.
- [10] Pomerleau, F., 2013. Applied registration for robotics: Methodology and tools for ICP-like algorithms. Ph.D. thesis. ETH Zurich.
- [11] Pratt, V., 1987. Direct least-squares fitting of algebraic surfaces. *SIG-GRAPH* 21, 145–152.
- [12] Qi, C.R., Su, H., Mo, K., Guibas, L.J., 2017. Pointnet: Deep learning on point sets for 3d classification and segmentation, in: *Proc. of the IEEE Int. Conf. Computer Vision and Pattern Recognition*, pp. 652–660.
- [13] Raunonen, P., Åkerblom, M., Kaasalainen, M., Casella, E., Calders, K., Murphy, S., 2015. MASSIVE-SCALE TREE MODELLING FROM TLS DATA. *ISPRS Annals of the Phot., Remote Sensing and Spatial Inf. Sciences* 2.
- [14] Raunonen, P., Kaasalainen, M., Åkerblom, M., Kaasalainen, S., Kaartinen, H., Vastaranta, M., Holopainen, M., Disney, M., Lewis, P., 2013. Fast Automatic Precision Tree Models from Terrestrial Laser Scanner Data. *Remote Sensing* 5, 491–520.

- [15] Schnabel, R., Wahl, R., Klein, R., 2007. Efficient RANSAC for Point-Cloud Shape Detection. *Computer Graphics Forum* 26, 214–226. doi:https://doi.org/10.1111/j.1467-8659.2007.01016.x.
- [16] Gonzalez de Tanago, J., Lau, A., Bartholomeus, H., Herold, M., Avitabile, V., Raunonen, P., Martius, C., Goodman, R.C., Disney, M., Manuri, S., et al., 2018. Estimation of above-ground biomass of large tropical trees with terrestrial LiDAR. *Methods in Ecology and Evolution* 9, 223–234.
- [17] Tejaswi Digumarti, S., Schmid, L.M., Rizzi, G.M., Nieto, J., Siegwart, R., Beardsley, P., Cadena, C., 2019. An Approach for Semantic Segmentation of Tree-like Vegetation, in: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pp. 1801–1807. doi:10.1109/ICRA.2019.8793576.
- [18] Trevor, A.J., Gedikli, S., Rusu, R.B., Christensen, H.I., 2013. Efficient Organized Point Cloud Segmentation with Connected Components. *Semantic Perception Mapping and Exploration (SPME)*.
- [19] Trochta, J., Krůček, M., Vrška, T., Král, K., 2017. 3D Forest: An application for descriptions of three-dimensional forest structures using terrestrial LiDAR. *PLOS ONE* 12, e0176871. doi:10.1371/journal.pone.0176871.
- [20] Westling, F., Underwood, D.J., Bryson, D.M., 2020. Graph-based methods for analyzing orchard tree structure using noisy point cloud data. *arXiv preprint arXiv:2009.13727*.
- [21] Windrim, L., Bryson, M., 2020. Detection, segmentation, and model fitting of individual tree stems from airborne laser scanning of forests using deep learning. *Remote Sensing* 12, 1469.
- [22] Wisth, D., Camurri, M., Fallon, M., 2019. Robust Legged Robot State Estimation Using Factor Graph Optimization. *IEEE Robotics and Automation Letters* 4, 4507–4514.
- [23] Zhong, L., Cheng, L., Xu, H., Wu, Y., Chen, Y., Li, M., 2016. Segmentation of individual trees from TLS and MLS data. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 10, 774–787.
- [24] Zhou, S., Kang, F., Li, W., Kan, J., Zheng, Y., He, G., 2019. Extracting Diameter at Breast Height with a Handheld Mobile LiDAR System in an Outdoor Environment. *Sensors* 19, 3212.



Alexander Proudman completed his M.Eng from the University of Oxford. He is also associated with Wadham College, Oxford. He worked with the Dynamic Robot Systems Group at the Oxford Robotics Institute for his masters thesis. His research interests are in developing robust SLAM systems and real-time semantic interpretation from sensor data.



Milad is a Postdoctoral Fellow in the Embodied AI research team of the Robotics and Autonomous Systems Group at CSIRO Data61. He studied computer vision and photogrammetry at K.N. Toosi University of Technology, Iran. He carried out his PhD research in localisation via GPS/IMU integration and visual-inertial odometry at the University of Melbourne. He joined the DRS Group in Oxford Robotics Institute as a postdoc under Assoc. Prof. Maurice Fallon (2018–2021). His research focused on robust and accurate localisation and mapping systems deployable on various robotics platforms.



Tejaswi is a Postdoc at the Dynamic Robot Systems Group at the Oxford Robotics Institute. He received his B.Tech in Electrical Engineering from IIT Jodhpur, India and his MSc. and PhD. in Robotics from ETH Zurich, Switzerland. From 2019–2021 he was a postdoctoral research associate at ACFR, University of Sydney, Australia, where he worked on semantic segmentation of large scale point clouds of forests and learning-based approaches to interpret novel camera sys-

tems. His research focusses on utilizing semantics and geometry to improve navigation, scene understanding and 3D reconstruction, with applications in natural environments.



Nived Chebrolu is a post-doctoral research associate at the Oxford Robotics Institute, University of Oxford. His research interests are in developing robust localization and mapping techniques for field robotics applications. He obtained his Ph.D. from the University of Bonn in 2021, where he developed long-term registration techniques for SLAM systems which were deployed on agricultural fields within the Flourish H2020 EU project.



Maurice Fallon received his B.Eng (electrical engineering) from University College Dublin, Ireland and his PhD (acoustic source tracking) from University of Cambridge, UK. From 2008-2012 he was a postdoc and research scientist in MIT Marine Robotics Group working on SLAM. Later, he was the perception lead of MIT's team in the DARPA Robotics Challenge. From 2017 he has been a Royal Society University Research Fellow and Associate Professor at University of Oxford, UK. He leads the Dynamic Robot Systems Group in Oxford Robotics Institute. His research focuses on probabilistic methods for localization, mapping, dynamic motion planning, and multi-sensor fusion.