

University of Nevada, Reno

**Discovertree - An Automated Tool To Generate Stem Maps From Terrestrial Laser
Scanner Point Clouds**

A thesis submitted in in partial fulfilment of the requirements for the degree of Master of Science
in Natural Resources & Environmental Science

by

Theodore Hartsook

Dr. Jonathan Greenberg/Thesis Advisor

August 2021

Copyright © 2021 by Theodore Hartsook

All rights reserved.



THE GRADUATE SCHOOL

We recommend that the thesis
prepared under our supervision by

entitled

be accepted in partial fulfillment of the
requirements for the degree of

Advisor

Committee Member

Graduate School Representative

David W. Zeh, Ph.D., Dean
Graduate School

Abstract

Terrestrial laser scanning (TLS) is increasingly used in forestry to quickly and nondestructively capture a variety of tree attributes such as diameter, height, and volume. However, in order for these attributes to be measured, the individual trees must first be segmented from the point cloud. Tools for manual or semi-automatic tree segmentation are widely available, but a fully automated and generalizable tool does not yet exist. The first step in tree segmentation is the creation of a stem map consisting of the position and size of all trees in the point cloud. We developed a novel stem mapper, "Discovetree" that uses Hough transforms combined with a machine learning algorithm calibrated with field data. Our algorithm outperformed a similar existing tool, TreeLS, at both the tree and stand level. Our study examines the consequences of tree shape on common representations (i.e. circles). The results suggest that stem mappers benefit from being tuned on real-world data and that analytical approaches that can represent trees by shapes other than circles may be needed to achieve the same levels of accuracy that in-field or manual mensuration can achieve.

Dedication

To Peter Eayre, for fostering my love of science and technology. I follow in your footsteps every day.

Acknowledgements

I would like to thank my committee for their guidance, wisdom, and patience. I would also like to thank all my family and friends for their support in all my endeavors.

Table Of Contents

Thesis Abstract	i
Dedication	ii
Acknowledgements	iii
Preface	viii
Chapter 1. Introduction	1
Chapter 2. Discovetree	2
Abstract	2
Introduction	3
Methods	7
Overview	7
Study sites and data collection	7
Calibration/validation data	9
Circle detection using Hough transforms	11
Preprocessing	14
Identifying clusters and Hough transformer parameters	14
Attributing circles	16
Filtering detections via gradient boosting	20
Comparison vs. TreeLS	20
Accuracy assessment	21
Understanding errors	21
Field data by circle ratio	22
Results	23
Tree level accuracies	23
Tree presence/absence	23
DBH	25
Plot level accuracies	26
Number of trees	26
Basal area	27
Accuracy by species	27
Accuracy by DBH	29
Accuracy by circle ratio	31
Gradient boosting performance	32
Discussion/Conclusion	34
Failures of circular assumption	35
Importance of parameter tuning	38
References	39
Chapter 3. Conclusion	46

List of Tables

Table 1: Confusion matrix for initial Hough detections	13
Table 2: Discovertree model attributes	19
Table 3: Confusion matrix for loose accuracies	24
Table 4: Confusion matrix for strict accuracies	24
Table 5: Species sample size and accuracy	28

List of Figures

Figure 1: Map of study sites	8
Figure 2: Diagram of scanner positions	9
Figure 3: A rasterized point cloud slice	10
Figure 4: Ground truth digitization	11
Figure 5: Overprediction of circles	13
Figure 6: Discovertree process	13
Figure 7: Identifying clusters	16
Figure 8: Loose and strict accuracy	17
Figure 9: Diagram of pixel metrics	18
Figure 10: Calculating circle ratio	22
Figure 11: Circle ratios for field data	23
Figure 12: DBH regression for loose accuracies	25
Figure 13: DBH regression for strict accuracies	26
Figure 14: Tree count regression	26
Figure 15: Basal area regression	27
Figure 16: Species accuracy comparison	29
Figure 17: Accuracy and DBH correlation	30
Figure 18: Accuracy and DBH comparison	30
Figure 19: Accuracy and circle ratio correlation	31
Figure 20: Accuracy and circle ratio comparison	32
Figure 21: Gradient boosting weights	33
Figure 22: Gradient boosting ROC	33
Figure 23: Gradient boosting PRC	34
Figure 24: Circular assumption failure cases	37

List of Equations

Equation 1: Hough transform for circles	12
Equation 2: Determining minimum radius	14
Equation 3: Determining maximum radius	15

Preface

This chapter is in preparation for submission to two journals: *IEEE Transactions on Geoscience and Remote Sensing* and *Remote Sensing of the Environment*.

Introduction

The purpose of this thesis was the creation of a novel approach to mapping the position and sizes of trees from terrestrial laser scanning (TLS) data. The approach, "Discovertree", combined computer vision techniques with machine learning, calibrated with a large, diverse set of training data collected across the northern Sierra Nevada Mountains in California. I compared my approach to an existing algorithm "TreeLS". Algorithms to produce stem maps from TLS data are a critical first step in the larger goal of automatic segmentation of trees from point clouds. While there are many existing segmentation tools, many rely on manually tuned parameters and assume that trees are a series of cylinders.

Abstract

Terrestrial laser scanning (TLS) is increasingly used in forestry to quickly and nondestructively capture a variety of tree attributes such as diameter, height, and volume. However, in order for these attributes to be measured, the individual trees must first be segmented from the point cloud. Tools for manual or semi-automatic tree segmentation are widely available, but a fully automated and generalizable tool does not yet exist. The first step in tree segmentation is to create a stem map consisting of the position and size of all trees in the point cloud. I developed a novel stem mapper, "Discovetree," that uses Hough transforms combined with a machine learning algorithm, calibrated with field data. This algorithm outperformed a similar existing tool, TreeLS, at both the tree and stand level. This study examines the consequences of tree shape on common representations (i.e. circles). The results suggest that stem mappers benefit from tuning themselves on real-world data and that tree cross sections are not represented best with a circle.

Introduction

Spatially explicit individual tree measurements, or “stem maps” and their associated taxonomic and structural status (e.g. diameter-at-breast-height (DBH), tree height, crown spread, height-to-live-crown, as well as the proportion of a tree that is trunk, branch, or leaf), are core measurements that are used to parameterize fire models, calculate the economic value, assess carbon sequestration, and monitor tree growth, health, mortality, and harvesting (Biernet et al., 2006; Luoma et al., 2017; Price et al., 2017). Stem maps allow for repeated tracking of individual tree growth over time (Vepakomma, 2011; Yan et al., 2021) and, scaled up to the level of forest stands, they can be used to derive plot level metrics such as basal area and stem density. When scaling to landscapes, the spatially explicit nature of stem maps are an appropriate basis of calibration and validation data for high-resolution remote sensing data such as airborne LiDAR and aerial and spaceborne hyperspatial optical datasets (Fricker et al., 2015; Popescu et al., 2003).

The current approach to generating stem maps and associated structural and taxonomic data on trees is dependent on the organization collecting the data and its protocols. Trees qualify for capture depending on their size relative to the plot and the region. Location is measured by the tree’s position relative to the plot center or another reference point (US Forest Service [USFS], 2021). This is typically captured with a rangefinder or a compass and a measuring tape, while the reference point is captured with a high-accuracy GPS. However, field measurements of trees are associated with relatively high spatial errors dependent on the quality of the GPS (Danskin et al., 2009; Fadili et al., 2019; Valbuena et al., 2010). DBH is measured primarily using calipers and diameter tapes (USFS, 2021). Diameter tapes are fairly consistent, even when the tree trunk is not circular. Calipers have more room for variation in measurements, particularly when the tree trunk is not circular, which is more common in trees with large DBHs. (Luoma et al., 2017; Liu et al., 2011; Williamson, 1975). Williamson (1975) found that out of roundness (noncircular trunks) is more frequent in larger trees and when

measured closer to the ground. Small measurement errors of DBH at the tree level propagate to the plot, the stand, and ultimately to forest level metrics such as basal area (Luoma et al., 2017).

Increasingly, a novel set of technologies has been developed to model a 3D representation of a forest stand with sufficient fidelity that the position and structure of trees can be determined afterwards in a digital environment. These technologies include terrestrial laser scanning (TLS) and ground-based photogrammetry via structure from motion techniques (SfM). TLS is the use of a LIDAR (light detection and ranging) unit on the Earth's surface (as opposed to air or spaceborne). LIDAR emits many pulses of light and measures the return time and intensity, measuring the distance between the scanner and its surroundings. These distances are placed into a local coordinate system using the IMU (inertial measurement unit) and then translated into a global coordinate system through an onboard GPS. The results of these scans are saved as point clouds, where each return is represented by a point in 3D space.

These point clouds contain much of the data that is commonly collected for forest inventories, without requiring additional fieldwork. The typical fieldwork for a plot includes measuring DBH and tree height, location, and species for each tree, often requiring at least two people (Hyypä et al., 2020). DBH measurements using under canopy LIDAR have lower root mean square errors (RMSEs) than repeated field measurements, which have been found to deviate consistently (Hyypä et al., 2020; Liu et al., 2011). Scanning the plot also allows for additional analysis without making another visit to the field.

TLS technologies afford a new set of approaches to quantify structural information such as volume (Burt et al., 2019; Hackenberg et al., 2015; Raumonon et al., 2013), tree height (Bienert, 2006; Burt et al., 2019; De Conto et al., 2017; Hackenberg et al., 2015; Raumonon et al. 2013), DBH (Bienert et al., 2006; Burt et al., 2019; De Conto et al., 2017; Hackenberg et al., 2015; Raumonon et al. 2013), height-diameter profiles (Burt et al., 2019; De Conto et al., 2017; Hackenberg et al., 2015; Raumonon et al. 2013), and crown area/volume (Burt et al., 2019;

Hackenberg et al., 2015). However, a prerequisite of developing these more complex structural attributes is to create a stem map containing the locations and DBH of potential trees (Burt et al., 2019; De Conto et al., 2017; Hackenberg et al., 2015).

Stem maps from TLS data are typically generated based on a similar set of assumptions, namely that a tree trunk, in a cross section, will appear as roughly circular (2D) (Beinert et al., 2006; De Conto et al., 2017; Gorte & Winterhalder, 2004) or cylindrical (3D) (Burt et al., 2019). The minimum requirement for a stem map is position and size, meaning that computationally simpler 2D techniques are all that would theoretically be needed. However, 3D analysis leverages the full dimensionality of the point cloud, so the stem mapping process can consider the full vertical extent of the point cloud instead of just a cross section. Trees don't grow perfectly perpendicular to the ground so the 3D cylindrical assumption is more robust than a 2D circular cross section.

Regardless of the dimensionality of the analysis, the resulting stem maps are 2D because they only include the position and size of the tree. Because of this, the limitations of automated stem maps are similar to those created through fieldwork. Trees that split or have branches beginning below breast height require special consideration when calculating DBH. When trunks have irregular shapes, a single measure of DBH will be unable to fully describe the tree, requiring either a certain tolerance of error or the averaging of multiple DBH measurements. Many automated stem mapping methods require the user to have prior knowledge of the point cloud (e.g. point density) and/or trees (e.g. maximum/minimum size) (De Conto et al., 2017; Hackenberg et al., 2015). Even near-automatic techniques assume trees can be accurately approximated as a series of cylinders (Burt et al., 2019).

Common challenges to stem mapping with TLS data arise from violations of the underlying circular/cylindrical assumptions of trees. These assumptions can be violated due to data quality issues stemming from occlusion, and also from the characteristics of the trees under investigation, namely that trees may have non-circular cross sections. Occlusion occurs

when one object interferes with the perception of a second object, causing it to appear only partially, as a merged object, or not appear at all (Cheong & Chew, 2017). For instance, in TLS scans this occurs when a large object like a tree or rock obscures additional trees. Large shrubs and branches also introduce occlusion into the point cloud by altering the outline of trunks. Even partial occlusion can lead to challenges because less of the tree can be detected (Trochta et al., 2013). Increased sampling density can decrease the occlusion problem, but typically does not completely eliminate occlusion in more dense forest stands, so any stem mapping approach needs to account for partial occlusion (Trochta et al., 2013).

Non-circularity has been observed in a variety of trees around the world and is more likely in larger and older trees (Koizumi & Hirai, 2006; Nogueira et al., 2006; Pulkinen, 2012; Williamson, 1975). The failure of the assumption of circularity leads to an overestimation of DBH and any further measurements that rely on it, such as basal area or volume from an allometric equation (Nogueira et al., 2006; Pulkinen, 2012). Existing tools require the user to identify these problems in advance, if they accommodate them at all (Beinert et al., 2006; Burt et al., 2019; De Conto et al., 2017; Gorte & Winterhalder, 2004; Hackenberg et al., 2015).

The objective of this thesis was to develop a fully automated, generalizable approach to produce a stem map (the position and size) of all trees in a forest plot combining Hough transforms at multiple height slices, that were then linked together and augmented with pixel metrics. These detections were combined with field data to train a gradient booster for stem mapping and to evaluate the accuracy in a diverse natural system. Furthermore, I evaluated my errors in the context of the circularity assumption to determine the degree to which occlusion and non-circularity impact the detection of trees and the accurate estimate of their DBH. With all of these analyses, I compared my approach against TreeLS, a tree segmentation package that also uses a Hough transform for initial stem mapping (De Conto et al., 2017).

Methods

Overview

My goal was to demonstrate a fully automated stem mapping algorithm using terrestrial laser scanning data of forest plots collected across the Sierra Nevada Mountains, in California and Nevada, U.S.A. I used a combination of Hough transforms and pixel metrics as inputs for a gradient boosting-based stem mapper, calibrated and validated from a large dataset of stem maps derived through a combination of fieldwork and manual measurement of the TLS data. I compared this approach to a similar existing method TreeLS (De Conto et al., 2017). I examined the effects of species, size, and shape on stem mapping accuracy.

Study sites and data collection

The study sites cover the Lake Tahoe Basin and the Plumas National Forest. Both forests are mixed conifer forests in a Mediterranean climate, with hot dry summers and cool wet winters. The elevation of the plots ranged from 215 m to 2818 m above sea level, with a mean of 1707 m and a median of 1797 m.

Plot locations were selected with a two-stage stratified random sampling design to ensure a representative sample of forest types across Plumas National Forest and the Lake Tahoe Basin (Figure 1). The first stage of sampling placed 10 km grid cells within the boundaries of the sample sites. The second stage placed a 30 x 30 m plot centered within each grid cell randomly selected to be on publicly accessible land, between 60-120 m from the nearest road, and have more than 10% tree cover based on a LANDSAT analysis. Plot locations were also stratified by elevation and avoided any areas with burn scars or signs of recent logging. A total of 108 plots were scanned and used in this analysis during leaf-on periods in 2017 and 2018.

Each plot represents a 30 x 30 m square (plus a 5 m buffer) centered on the plot center using 9-13 separate scanner positions, with 2-4 scans per position (dependent on stand density)

to maximize the vertical distribution of data (Figure 2). Additional scan positions were added to maintain a visual line of sight between reflectors. Scans were taken with a Riegl VZ-400i Terrestrial Laser Scanner. All scans were coregistered within-plot using RiSCAN PRO (Riegl Laser Measurement Systems GmbH, 2020), and then subsequently to an airborne laser scanning (ALS) base map using LAsTools (RapidLasso GmbH, 2019), and CloudCompare (Girardeau-Montaut, 2020). The average number of points per plot was 125,830,459 and the average density was 7.86 points/cm².

Stem mapping was conducted during the initial setup of plots. First, the most common species present in each plot was identified. Any remaining trees of a different species with a DBH greater than or equal to 15 cm were chosen for stem mapping. The distance and azimuth of each tree within 30 m of the plot center was recorded.

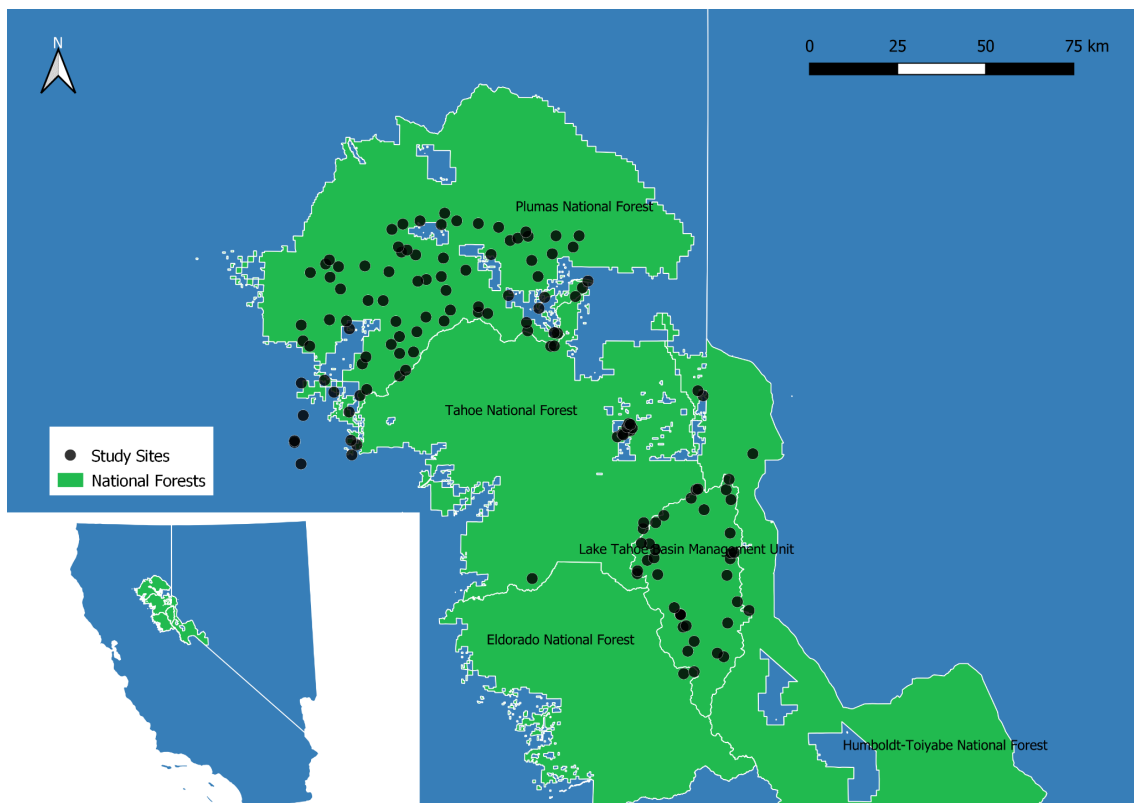


Figure 1. The study sites include 108 plots across the Sierras that have been scanned over a 2 year period.

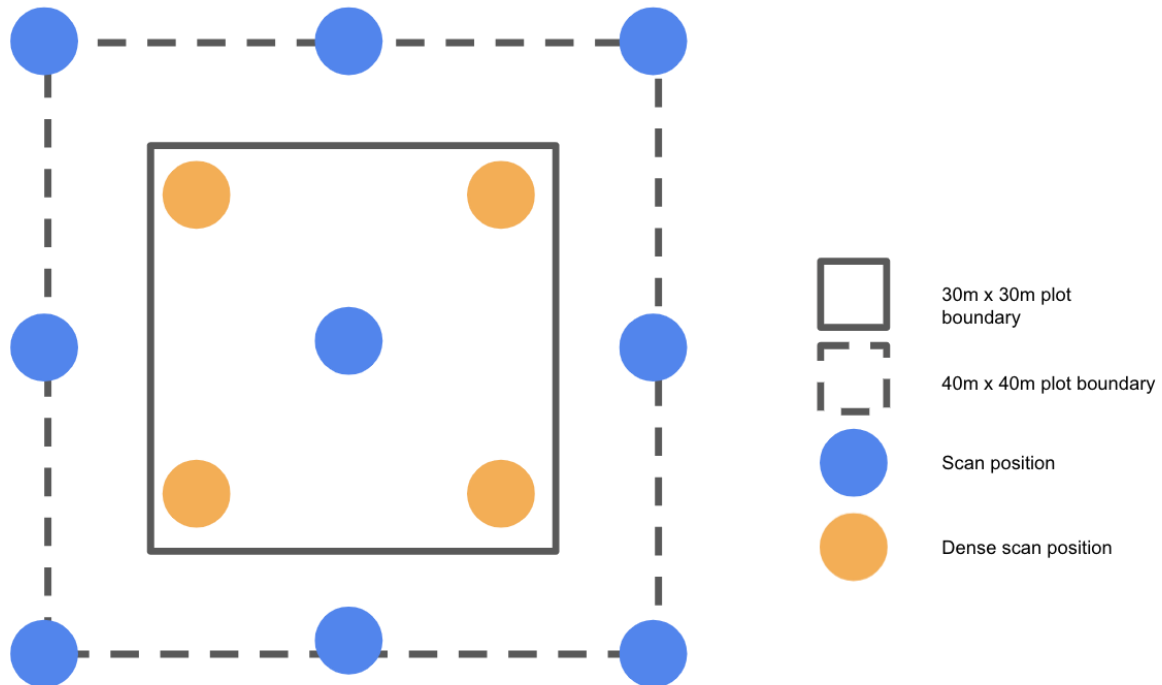


Figure 2. Scanning set up with scan positions (blue circles). The interior scan positions (orange circles) were done on a plot-by-plot basis to maintain line of sight between positions. The final data was trimmed to the 30m x 30 m boundary, but the 40m x 40 m boundary was used for processing to prevent edge effects.

Calibration/validation data

The calibration/validation dataset was a set of circles representing the size, position and species of trees located within the plots. Each point cloud of a TLS plot was normalized to height above ground using the `lasground` (-fine, step size of 1 m) and `lasheight` (-replace_z) functions from `LAStools` (RapidLasso GmbH, 2019). From this point cloud, I subsetted all points within 5 cm of diameter at breast height (1.37 m above ground). These points were rasterized at 1 cm resolution with the raster pixel value determined by the number of points falling in each cell grid, with 0 being used for pixels with no returns.



Figure 3. A rasterized cross section of a point cloud 1.37 m above ground.

The stem maps collected in the field and the DBH-sliced rasters were used as the starting point for a virtual measurement of trees performed within photointerpretation in a GIS. Each tree touching or within 40 m of the plot center with a DBH equal to or greater than 15 cm was digitized as a circle. Edge trees were digitized with the best circle that fit the visible arc. In cases where a circle was not able to fully fit the tree, the largest circle that best contained the trunk pixels was drawn (Figure 4). The species information collected in the field was linked with these circles, resulting in a stem map containing the positions, DBH and species of all trees in the field plots.

I performed a stratified random sampling of the field plots, divided 70-30 into a training and testing set, stratified by basal area to accommodate for differences in stem density. The training set contained 83 plots and 8092 trees and the testing set contained 25 plots and 1262 trees. Pseudo-absence data was created by selecting 1225 random points within each plot that were more than 4 cm from the nearest ground truth tree. The 4 cm buffer was selected to match

the accuracy tolerance for detections (any detection within 4 cm of the ground truth's center was considered a true positive). 7 of the 14 species present in the ground truth dataset had a sample size of at least 50: *Abies magnifica* (California red fir; 89), *Pseudotsuga menziesii* (Douglas fir; 135), *Pinus jeffreyi* (Jeffrey pine; 176), *Pinus contorta* (lodgepole pine; 145), *Tsuga mertensiana* (mountain hemlock; 85), *Pinus ponderosa* (ponderosa pine; 76), and *Abies concolor* (white fir; 443).




Figure 4. Ground truth trees, digitized so that the circle contains the full perimeter of the tree.

Circle detection using Hough transforms

A cross section of a tree from a TLS scan appears as a circular ring in the raster, so the initial identification step utilized Hough transforms. A Hough transform is a parametric search function that can efficiently search an image for circles using a 3D accumulator array (Eq. 1) (Yuen, 1985). Hough transforms are thought to be robust against occlusion and small deviations from circularity, making them a reasonable choice for stem mapping (Beinert et al., 2006; De Conto 2017; Van Leeuwen, 2010). I used the Hough circle transform available in Scikit-Image (van der Walt S et al., 2014).

$$(x - a)^2 + (y - b)^2 = r^2$$

$$x = a + r \cdot \cos(\theta)$$

$$y = b + r \cdot \sin(\theta)$$

Equation 1. A parametric definition of a circle. (a,b) is the center of the circle and (x,y) is a given point on the edge of the circle. By changing θ and r , any given parameter can be searched.

Each detection has an x,y position, a radius, and a weight indicating how well the pixels in the image match a circle of that description. The weights range from 0 (no pixels match the circle) to 1 (all pixels match the circle). A high weight (i.e. greater than 0.5) is strong evidence for a circle. Because the Hough transform returns all possible circles, it is inherently overpredictive and some way to limit the sensitivity of detections is required. This can be done by setting limits on the weight (i.e. a circle must have a weight greater than or equal to w) or by limiting the number of returns (i.e. accept n circles with the highest weight). Filtering by weight produces a subset of well defined circles and greatly reduces the number of false positives. However, occlusion can reduce the weight of legitimate circles. Since an upper limit of the number of circles for each size bin was identified earlier, filtering by number of returns was chosen instead.

The image segmentation method described in the next section removes the need for tuning of most Hough parameters (weights, maximum size) and can translate the remaining parameters based on a forestry context (minimum tree size for inclusion, upper limit for number of trees per unit area). The general approach for tuning is to reduce the search space as cautiously as possible, to minimize the loss of any true trees. There are a variety of reasons a tree may not be detected with a high weight. Occlusion may create incomplete circles, limiting the maximum weight. Not all trees are best approximated with a circle, due to bark fissures and/or having a more elliptical shape. Branches can also introduce irregularities in the shapes of

trees. The result of this is that while weight is a useful filter, it is unable to capture many legitimate trees. Therefore additional filtering mechanisms are needed (Table 1).

Table 1. Confusion matrix for initial Hough detections

		Observed		
		Presence	Absence	Sum
Predicted	Presence	1730	256952	258682
	Absence	1	351	352
	Sum	1731	257303	259034




Figure 5. Left: The input raster. Center: the ground truth data. Right: The Hough detections.



Figure 6. The seven steps of the Discovertree approach. The first step is to classify the point cloud so points can be subset by height above ground. Slices of the point cloud are rasterized at multiple height levels, then divided into clusters. Hough transforms are parameterized for each region, then a series of pixel metrics are calculated for each Hough detection. The detections are connected between different height layers and finally fed into a gradient boosting classifier.

Preprocessing

The first step was to classify ground points in each of the N=108 point clouds using the LAStools package (RapidLasso GmbH, 2019). Next, I normalized the heights of each point to the ground (instead of elevation). Then, horizontal cross sections were extracted at 1.37 meters, 2 m, 3 m, 4 m, and 5 m above ground level. These cross sections were then converted into a raster format for 2D image analysis with a cell size of 1 cm per pixel (Figure 3).

Identifying clusters and Hough transform parameters

A Hough transform has several parameters that can be tuned for effective identification. The most important are the minimum and maximum radius of potential circles, minimum distance between circles, and an upper limit on the number of circles.

Minimum radius is a constant equal to half the DBH cutoff for trees, rounded down to the nearest pixel. The DBH cutoff was 0.15 m DBH, and the pixel size of the rasters is 0.01 m. Therefore any detections with a radius lower than 7 cm are excluded from the search space. The minimum distance between circle detections is also related to minimum radius. Because trees typically do not grow inside another, I make the assumption that, at most, stems can be touching but not overlapping. While there are exceptions to this assumption, they are relatively rare and beyond the scope of this research, which focuses on a generalized approach. Therefore the minimum distance between trees is set to the minimum radius.

$$r_{\min} = \text{DBH}_{\min}/2$$

Equation 2. The minimum radius was set to half the DBH cut off, rounded down (7 cm).

Maximum radius requires additional information. Because the dataset does not have an upper limit on DBH, the maximum radius is unknown for any given plot until stem mapping is completed. Therefore, the raster needs some preprocessing to identify a reasonable upper limit.

This can be done through a series of transformations. First, a distance transformation calculates the distance of each pixel to its nearest neighbor. This creates a raster with large, continuous patches of empty space separated by peaks of connected pixels. Applying a ridge filter to this distance transform creates boundaries between these patches. Connected components labelling segments portions of the image and isolates them for further analysis. All analysis was done using the Scikit-Image library (van der Walt S et al., 2014).

The maximum radius of any potential is now limited by the size of each component. The largest radius can be no more than half the fetch of the polygon. To simplify the calculation, the bounding box of the component is used to identify the shortest side (breadth).

$$r_{\max} = \text{breadth}/2$$

Equation. 3 The maximum radius was set to half the length of the shortest side of the component.

The difference between r_{\min} and r_{\max} is often fairly large due to the size of the search space. Since minimum distance is identical to r_{\min} , searching the entire range of possible circles is inefficient and generates excessive false positives. Dividing the search space into bins alleviates this problem and allows the minimum distance to scale with the size of circles being searched. For instance, a component with a breadth of 4 m would need to identify all circles with a radius between 0.075 m and 2 m. If this were done with one pass, the minimum distance between circles would have to be set to 0.075 m. While this is desirable to identify small trees, it will generate an excessive number of overlapping false positives on larger circles. By dividing the search space into 0.5 m bins, the minimum distance scales to reduce the amount of unhelpful overlap.

With r_{\min} and r_{\max} determined, the remaining parameter is the number of circles expected. This parameter is also not known until stem mapping is complete, but an upper limit can be estimated based on the area of the component. Because the search spaces are divided into 0.5 m bins, each search is tuned based on expected stem density. As the size of the potential trees

increases, the number of circles expected is lowered. When searching for smaller trees, the number of circles expected is set quite high to ensure that even weak detections are included. However when searching the same component for larger trees, the number of potential circles is set much lower, often as low as one.




Figure 7. From left to right: input raster, distance transform, cluster, and the cluster area overlaid on the input raster.

Attributing circles

The Hough algorithm included some rudimentary filtering, but the goal of the initial filtering was to minimize omissions. The resulting dataset has a large number of false positives and multiple detections per tree (average 11 detections per tree) that require subsequent filtering. To achieve this, I first attributed each Hough circle detected in the previous step at the various heights with a set of metrics that I used with a machine learning algorithm (gradient boosting, described in the next section) to produce the final map. The circles which correctly matched the calibration data were identified based on proximity and size. If a detected tree's area contained the center of a ground truth tree, it was classified as containing a tree. If this detection's center was within 4 cm of the ground truth and the DBH was within 4 cm as well, it was also classified as a tree match. The attributes are found in table 2.




Figure 8. Left: A loose Discovertree detection (red) and the ground truth (blue). Right: A strict Discovertree detection (green) and the ground truth (blue).

Calculating three additional metrics for each detection helps with this problem. Each metric is a pixel count of certain portions of the detected circles. The first is a core pixel count. Because the TLS cannot penetrate the bark, a legitimate detection will not have any returns within the interior. The second and third metrics are pixel counts of an inner and outer ring around the circumference of the circle. Hough detections are strongest where pixels follow along a circle, rather than the center of a tree. By counting the pixels inside and outside the circumference I can derive additional information about whether the Hough detection may be over or underfitting. When these three metrics are combined with the detection's weight, the resulting filters are much more effective (Figure 8).

To further improve detection, additional information is needed. The Hough detection process was repeated at multiple height slices of the point cloud. As the height increases, fewer shrubs and other occluding elements are included in the raster. This continues until the canopy starts to appear, which increases occlusion significantly. Detections at multiple height slices can be connected through taper functions (Li et al., 2021). A legitimate tree detection should have a similar counterpart at higher slices, with a slightly smaller r and an xy position that is within a

certain threshold. To recreate this, each detection is matched to its most similar detection at additional height slices of 2, 3, 4, and 5 meters above ground.

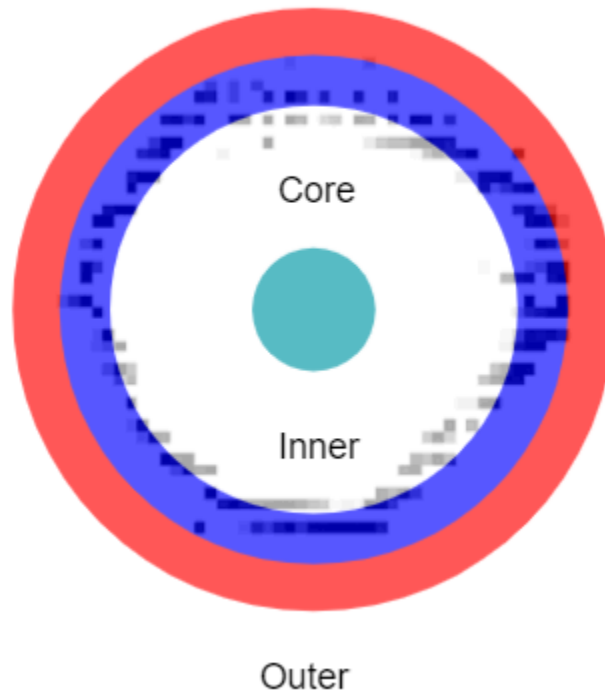


Figure 9. The outer buffer for pixel counts is shown in red, the inner buffer for pixel counts is shown in blue, and the core pixel count buffer is shown in green.

Table 2. The attributes used to train Discovetree.

Attribute	Description
Plot	Link detections between plots
Detection ID	Unique ID for each detection
X,Y position	Physical location of center of detection in real world coordinates
Radius	Size of detection in real world coordinates. Hough transforms search for radius instead of DBH
Hough weight	Primary attribute to filter detections
Core count	A buffer around the center of the detection with a size of $0.2 \times \text{Radius}$. Chosen because laser scans cannot penetrate the trunk, so a valid detection should have no pixels in the interior.
Inner count	A buffer around the interior of the detection's perimeter. Chosen for similar reasoning as core count. Pixels in the inner ring suggest the detection is not centered on the tree and/or that the tree is not circular.
Outer count	A buffer around the exterior of the detection's perimeter. Pixels in the outer ring suggest the same possibilities as inner count, as well as the possibility of shrubs, branches, etc. surrounding the shape.
XY distance to nearest detection at next height slice	Consecutive detections at different height slices with a small distance between them suggest a distinct structure exists (hopefully a tree, but could be something else like a rock), while large distances between detections suggest the opposite
Radius difference with nearest detection at next height slice	Tree trunks taper as height increases, so it is assumed for the size of a detection to lower or stay the same as height increases. Any large swings in radius, particularly an increase in size, suggest a different geometric structure (such as the sphere formed by some shrubs) exists.
Closest ground truth	Each detection was linked to the closest ground truth tree, allowing for multiple accuracy definitions to be evaluated.
XY distance to closest ground truth	The distance between the center of the detection and the closest ground truth tree. Used for accuracy evaluation.
R difference to closest ground truth	The difference between the radius of the detection and of the closest ground truth tree. Used for accuracy evaluation.
Species	Species of the ground truth tree. Used for accuracy evaluation.
Contains tree	This is the looser definition of accuracy. Any detection whose area included the center of a ground truth tree was classified as containing a tree.
Tree match	This is the stricter definition of accuracy. Any detection with a location and DBH within 4cm of the ground truth was classified as a tree match.

Filtering detections via gradient boosting

The goal of the final filtering step was to preserve all circles that were identified as "correct" via the calibration dataset, and to remove all circles that were not. After some exploratory data analysis on a subset of the training data, I selected a gradient boosting classifier implemented in the Python package scikit-learn (Pedregoa et al., 2011). Gradient boosting is a type of ensemble machine learner that has shown significant success in a variety of complex classification and regression tasks that focuses on combining multiple models to reduce overall error (Natekin & Knoll, 2013). Using the calibration data described in the Calibration/Validation Data section, I did a parameter sweep for the best gradient boosting classifier parameters. The learning rate was set to 0.3, max depth to 10, and the number of estimators to 500. Once gradient boosting was complete, the detections predicted as true were checked for overlaps. If two or more predictions overlapped, the one with the highest weight was selected and the remainder were set to false.

Comparison vs. TreeLS

I compared Discovetree to an approach with a similar methodology and goals, TreeLS (De Conto et al., 2017). Both approaches use 2D Hough transforms to identify initial seed points for a stem map. However, Discovetree's self tuning and the use of a machine learner are unique. As an additional preprocessing step for TreeLS, the input point clouds were tiled into 14x14m subsections (10 m with a 4 m buffer). The height step was set to 0.5 m, the maximum height was set to 15 m, pixel size to 3 cm, 20 RANSAC (random sample consensus) iterations, with a p-value of 0.9 and inliers set to 0.99. These parameters were chosen to follow the methods found in De Conto (2017).

Accuracy assessment

The primary accuracy statistics used in this project were precision, recall, F-score, confusion matrices, accuracy, and Cohen's kappa coefficient. For the gradient boosting classifier I included the area under curve and precision recall curve. For trees correctly detected, I also calculated the residuals between predicted and observed DBH. At the plot level, I compared the predicted tree count and basal area versus observed tree count and basal area. I calculated accuracy metrics using two definitions of true positive. The first, "looser" definition of a true positive was any detection whose area included the center of a ground truth tree. The second, "stricter" definition of a true positive was any detection whose location and DBH were within 4cm of a ground truth tree. All accuracy statistics were compared between Discovetree and TreeLS.

Understanding errors

In order to understand errors, I looked at three primary factors: tree shape, species, and size. Size data was used for regressions between the predicted stem maps and the observed ground truth. Accuracy data was subset to identify any correlation between errors and species or size. To evaluate tree shape, I created a ratio to measure how well trees met the circular assumption. The first step was to find the tree and extract its surrounding area (Figure 9 A1 and A2). The second step is to overlay the circle from the field data (Figure 9 B1 and B2). The third step was a morphological close to improve the shape of the trunk (Figure 9 C1 and C2). The fourth step was an edge detection (Figure 9D1 and D). The fifth step is a second morphological close to enhance the trunk once more (Figure 9 E1 and E2). The sixth step was to use the morphological close E as a mask on the identified circle B (Figure 9 F1 and F2). This allows for a ratio to be calculated between the pixels present and the pixels I would have expected to see if the trunk was a full circle. Example 1 has a circle ratio of 0.90 while example 2 has a circle ratio of 0.67.

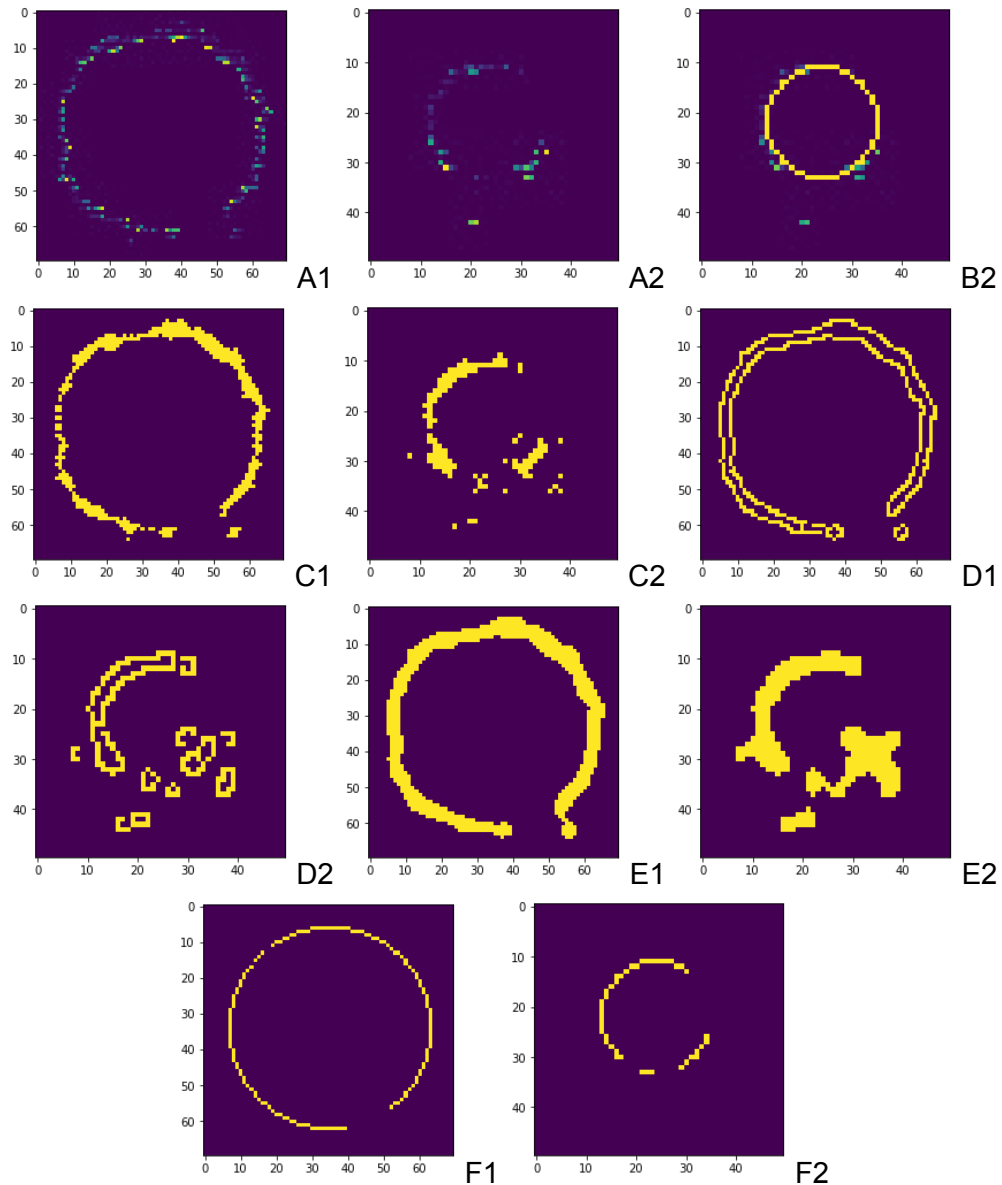


Figure 10. Two examples of calculating circle ratio. The left hand column (A1) is for a tree with minimal occlusion, while the right hand column (A2) is for a tree with significant occlusion. B1 and B2 show their respective circles overlaid. C, D, and E show the close, edge detection, and second close to enhance the trunk. F shows the final results used to calculate circle ratio.

Field data by circle ratio

677 trees in the field data had a circle ratio between 0.9 and 1, and 1017 had a circle ratio of at least 0.7.

Field data by circle ratio




Figure 11. Most trees were well represented by a circle, with 677 trees having a circle ratio between 0.9 and 0.1. However, a significant portion of the trees were not, with ~20% having a circle ratio lower than 0.7.

Results

Tree Level Accuracies

Tree Presence/Absence

Using the loose definition of tree accuracy, Discovetree had a precision of 0.96 and a recall of 0.68 with an F-score of 0.79. The mean difference in location was 0.02 m with a standard deviation of 0.02 m. The mean difference in DBH was 0.03 m with a standard deviation of 0.03 m. The kappa was 0.65. By comparison, TreeLS had a precision of 0.42 and a recall of 0.13. The F-score was 0.19 (Table 3). The mean difference in location was 0.01 m with a standard deviation of 0.01 m. The mean difference in DBH was 0.02 m with a standard deviation of 0.01 m. The kappa was 0.09.

Table 3. Confusion matrix for Discovertree and TreeLS with loose accuracy

Discovertree TreeLS (Loose)		Observed		
		Presence	Absence	Sum
Predicted	Presence	858 158	39 217	897 375
	Absence	404 1104	1225 1224	1629 2328
	Sum	1262 1262	1264 1441	2526 2703

Using the stricter definition of accuracy, where both the position and DBH are needed to be within 4 cm of ground truth, Discovertree had a precision of 0.70 and a recall of 0.50. The F-score was 0.58. The mean difference in location was 0.01 m with a standard deviation of 0.01 m. The mean difference in DBH was 0.02 m with a standard deviation of 0.01 m. The kappa was 0.33. By comparison, TreeLS had a precision of 0.42 and a recall of 0.13. The F-score was 0.19 (Table 4). The mean difference in location was 0.01 m with a standard deviation of 0.01 m. The mean difference in DBH was 0.02 m with a standard deviation of 0.01 m. The kappa was -0.03.

Table 4. Confusion matrix for Discovertree and TreeLS with stricter accuracy

Discovertree TreeLS (Strict)		Observed		
		Presence	Absence	Sum
Predicted	Presence	629 158	268 217	897 375
	Absence	633 1104	1225 1224	1859 2328
	Sum	1262 1262	1493 1441	2756 2703

DBH

Using the looser definition of accuracy, for true positives, Discovertree underestimated DBH slightly with an RMSE of 0.04 m and a bias of -8%, with a Pearson's coefficient of 0.98. TreeLS overestimated DBH for smaller trees, but overall underestimated DBH with an RMSE of 0.08 m and a bias of -7%, with a Pearson's coefficient of 0.95.

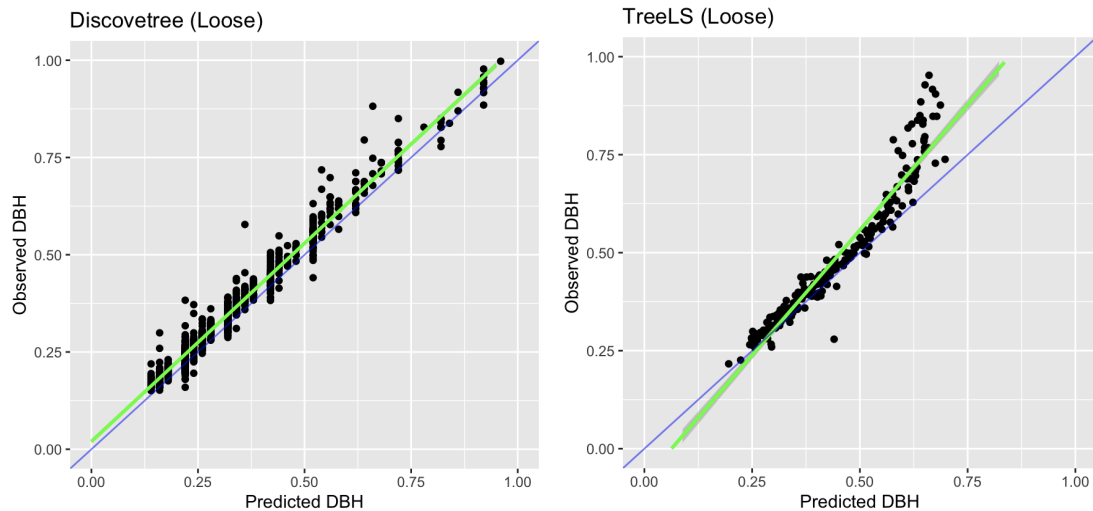


Figure 12. Regression between predicted and observed DBH using loose accuracy (true positives only). The green line represents the linear model while the blue line represents a 1:1 line.

Discovertree also underestimated DBH when using the stricter accuracy definition, with an RMSE of 0.02 m, a bias of -6% and a Pearson's coefficient of 0.99. TreeLS consistently underestimated DBH with an RMSE of 0.02 m, a bias of -4%, and Pearson's coefficient of 0.99.




Figure 13. Regression between predicted and observed DBH using strict accuracy(true positives only). The green line represents the linear model while the blue line represents a 1:1 line.

Plot level accuracies

Number of trees

Discovertree underestimated the number of trees on the plot level, with an RMSE of 17.23 trees and a bias of -32%, with a Pearson's coefficient of 0.96. TreeLS underestimated tree count with an RMSE of 43.99 trees and a bias of -6%, with a Pearson's coefficient of 0.58.

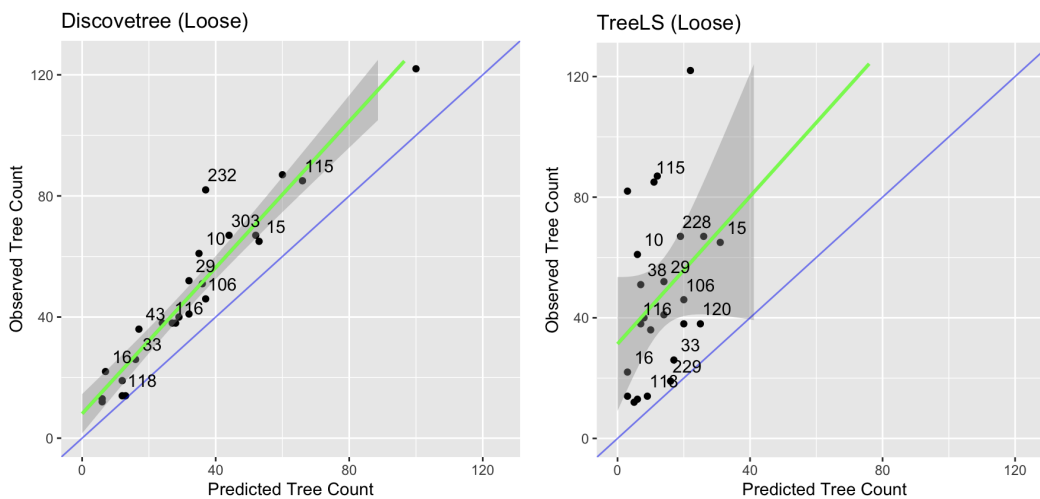


Figure 14. Regression between predicted and observed tree count. The green line represents the linear model while the blue line represents a 1:1 line. 95% confidence interval shown in gray.

Basal area

Discovertree significantly underestimated basal area on the plot level, with an RMSE of 3.5 m² and a bias of -44%, with a Pearson's coefficient of 0.79. TreeLS underestimated basal area with an RMSE of 5.1 m² and a bias of -63%, with a Pearson's coefficient of 0.43.




Figure 15. Regression between predicted and observed basal area. The green line represents the linear model while the blue line represents a 1:1 line. 95% confidence interval shown in gray.

Accuracy by species

Discovertree was able to identify at least some of all 14 species in the dataset. With the looser definition, the lowest accuracy was black oaks at 46%, while the highest accuracy was tanoak at 82% (Table 5 and Figure 15). With the stricter definition, the lowest accuracy was black cottonwood at 25%, while the highest accuracy was lodgepole pine at 67%. TreeLS was able to identify 11 of the 14 species in the dataset, as it was unable to detect any bigleaf maples, black cottonwoods, or black oaks. With the looser definition, the lowest accuracy was incense cedar at 7%, while the highest accuracy was tanoak at 36%. With the stricter definition, the lowest accuracy was again incense cedar at 3%, while the highest accuracy was again tanoak at 27%.

Table 5. The sample size and detection rates for all species present, using both definitions of accuracy.

Species name	Sample size	Discovertree Accuracy (Loose)	Discovertree Accuracy (Precise)	TreeLS Accuracy (Loose)	TreeLS Accuracy (Precise)
Bigleaf maple	8	62.5%	62.5%	0.0%	0.0%
Black cottonwood	8	50.0%	25.0%	0.0%	0.0%
Black oak	13	46.2%	38.5%	0.0%	0.0%
California red fir	89	57.3%	41.6%	14.6%	7.9%
Douglas fir	135	64.4%	49.6%	22.2%	14.8%
Incense cedar	29	55.2%	37.9%	6.9%	3.4%
Jeffrey pine	176	80.7%	60.8%	32.4%	22.7%
Lodgepole pine	145	75.2%	66.9%	20.0%	20.0%
Mountain hemlock	85	74.1%	42.4%	10.6%	3.5%
Ponderosa pine	76	65.8%	48.7%	13.2%	7.9%
Singleleaf pinyon	14	50.0%	35.7%	21.4%	14.3%
Sugar pine	20	55.0%	20.0%	15.0%	15.0%
Tanoak	11	81.8%	63.6%	36.4%	27.3%
White fir	443	65.7%	46.0%	17.2%	9.9%

Species Accuracy Comparison

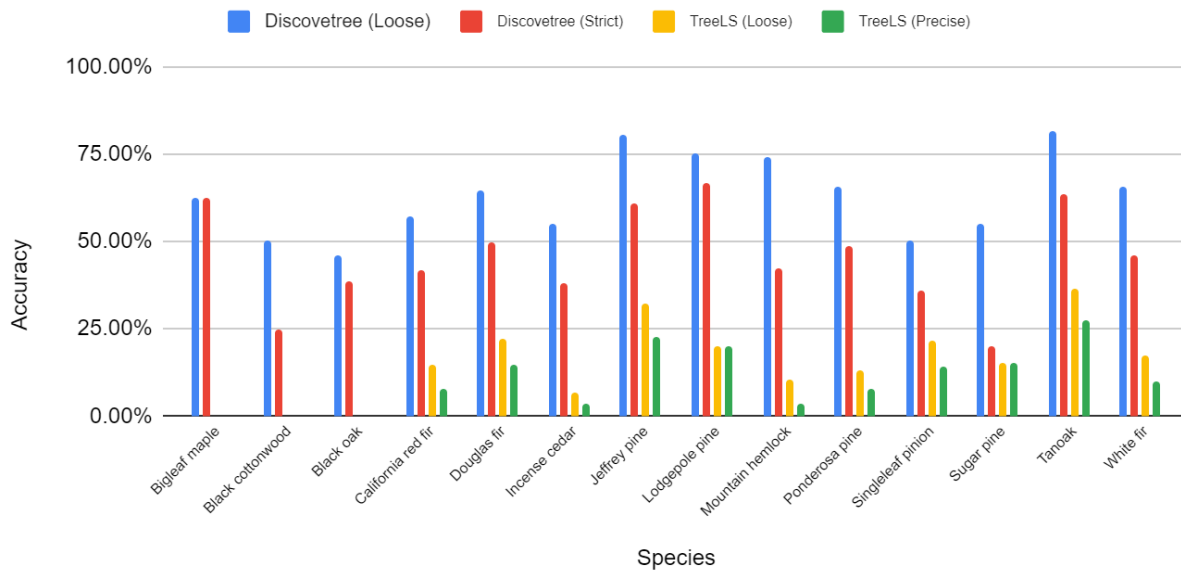


Figure 16. Discovertree outperformed TreeLS for every species.

Both Discovertree and TreeLS were able to identify at least some of all the common species which had at least 50 observations. For Discovertree with the looser definition, the lowest accuracy was California red fir at 57%, while the highest accuracy was Jeffrey Pine at 81%. With the stricter definition, the lowest accuracy was again California red fir at 42%, while the highest accuracy was lodgepole pine at 67%. For TreeLS with the looser definition, the lowest accuracy was mountain hemlock at 11%, while the highest accuracy was Jeffrey pine at 32%. With the stricter definition, the lowest accuracy was mountain hemlock at 4%, while the highest accuracy was Jeffrey pine at 23%.

Accuracy by DBH

Discovertree was most effective at identifying trees with a DBH less than 1 m. TreeLS was most effective at identifying trees with a DBH between 0.35 and 0.75 m. Discovertree's accuracy was negatively correlated to DBH with an RMSE of 0.45 m and a Pearson's coefficient

of -0.73. TreeLS's accuracy was negatively correlated to DBH with an RMSE of 0.35 m and a Pearson's coefficient of -0.47

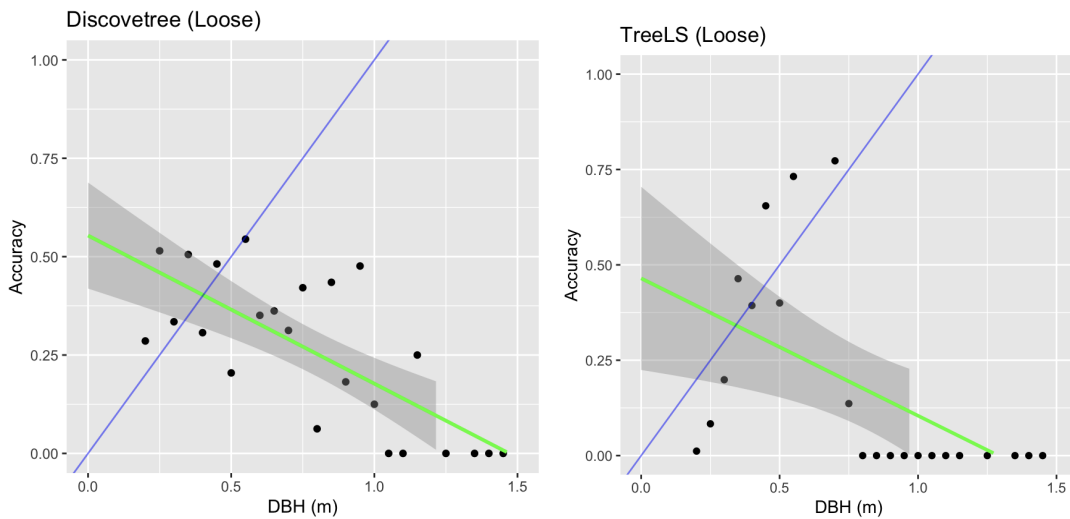


Figure 17. The accuracy of Discovetree was negatively correlated with DBH. TreeLS was positively correlated with DBH for the trees it was able to detect. The green line represents the linear model while the blue line represents a 1:1 line. 95% confidence interval shown in gray.

DBH Accuracy Comparison




Figure 18. Discovetree performed better over a larger size range than TreeLS.

Accuracy by circle ratio

Discovetree's accuracy was positively correlated with circle ratio (Figure 19). Discovetree's accuracy for trees with a circle ratio greater than 0.9 was 83% (loose) and 67% (precise). Discovetree was unable to identify any trees with a circle ratio less than 0.4. TreeLS accuracy had a much weaker correlation with circle ratio and was unable to identify trees with a circle ratio less than 0.1 (Figure 19). In Discovetree, circle ratio was strongly correlated to accuracy with an RMSE of 0.23 and a Pearson's coefficient of 0.98. In TreeLS, circle ratio was correlated to accuracy with an RMSE of 0.50 and a Pearson's coefficient of 0.77 (Figure 19).




Figure 19. Circle ratio was correlated with accuracy for both stem mappers, but more so for Discovetree. The green line represents the linear model. 95% confidence interval shown in gray.

Circle Ratio Accuracy Comparison

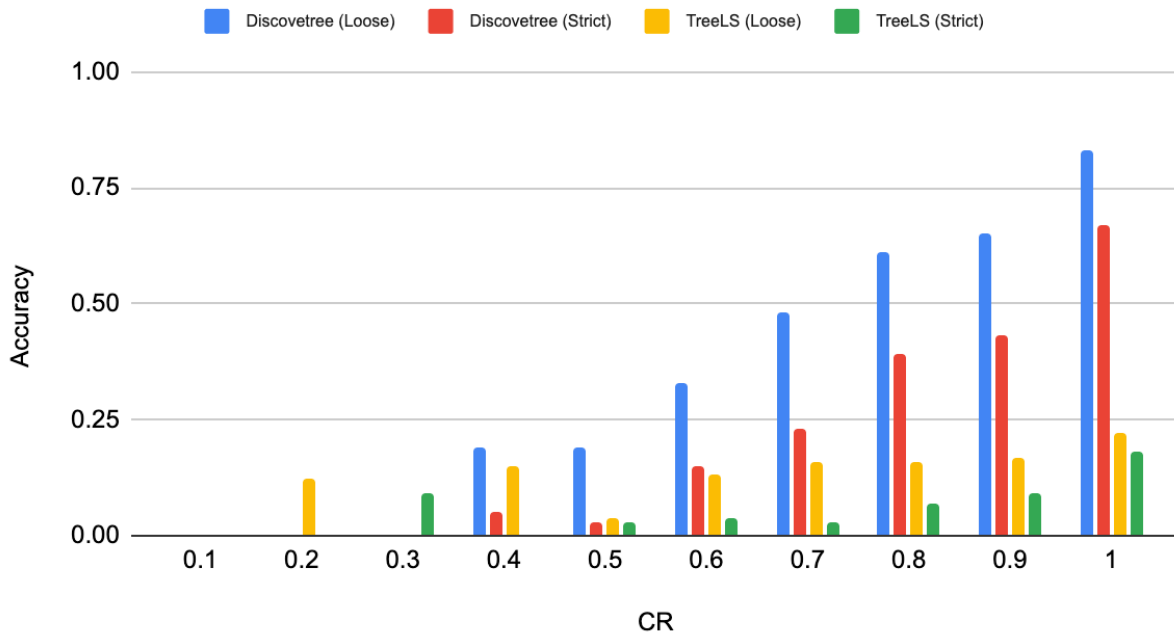


Figure 20. Discovertree's accuracy was positively associated with circle ratio. The correlation was stronger for the loose accuracy than the strict accuracy.

Gradient boosting performance

The resulting model found that the Hough weight, distance to nearest circle at 2 meters, and the pixel metrics of the DBH detection were the most important attributes. The receiver-operator characteristic (ROC) and precision-recall curve (PRC) areas were calculated, with areas of 0.94 and 0.65 respectively. This indicates that Discovertree performed reasonably well despite the unbalanced classification problem.

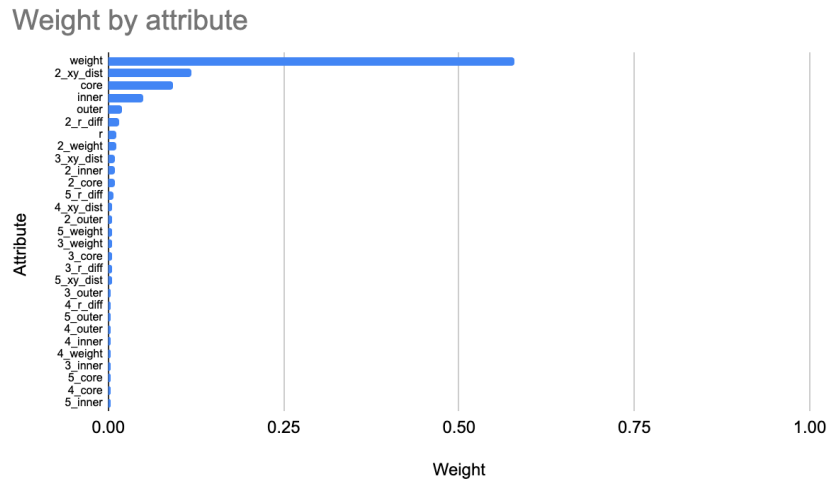


Figure 21. The pixel metrics at 1.37 m above ground were most important for the gradient booster.

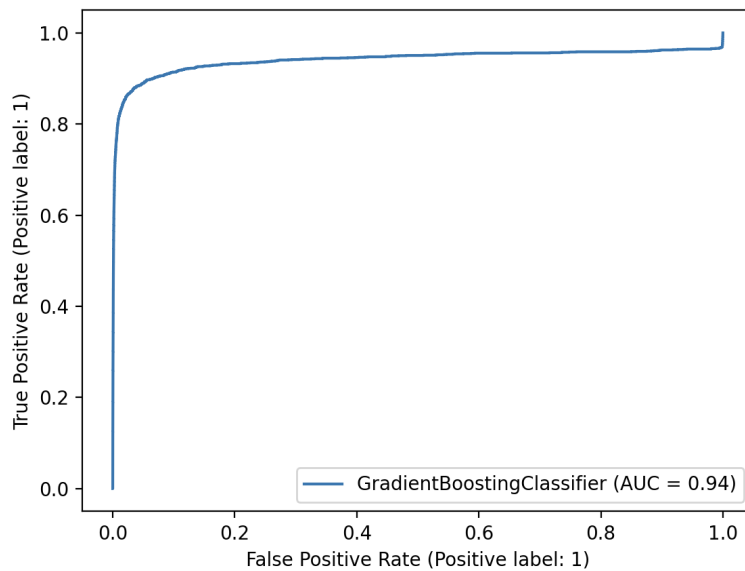


Figure 22. The area under the ROC curve is almost 1, indicating an effective classifier.

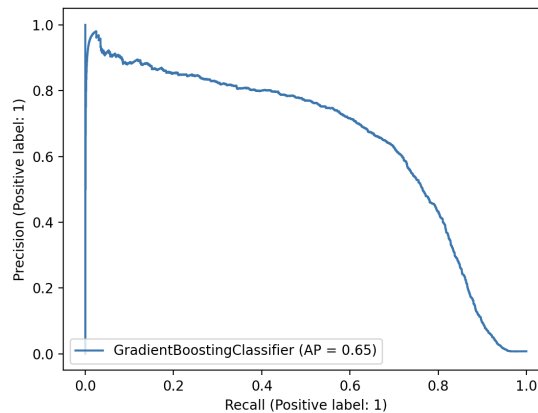


Figure 23. The area under the PRC is lower than the ROC, but still above 0.5, indicating good performance on imbalance classification problems.

Discussion/Conclusion

Discovetree outperformed TreeLS in every accuracy metric I examined. Discovetree's accuracy was positively associated with circle ratio, while TreeLS had a weaker positive association. Discovetree was able to identify trees across a larger size range than TreeLS. Discovetree and TreeLS both underpredicted basal area and tree count across plots.

Discovetree's higher recall compared to TreeLS is a promising step towards the goal of automatically identifying every tree in the plot, but there remains work to be done. While predictions of DBH had only small errors, missed trees led to a large underestimation of tree count and basal area at the plot level. Discovetree's accuracy seemed to be fairly independent of size and species given a sufficient sample size. Neither Discovetree nor TreeLS was capable of finding all trees in a plot or substituting for a field inventory. However, Discovetree could greatly reduce the work required to produce stem maps in a GIS, since users would only have to identify remaining stems. At this stage, these approaches cannot replace field inventories without additional manual edits, but Discovetree is closer to that goal. The use of real world field data to tune parameters contributed to a higher performance. This is an attractive option to having users tweak parameters until they find acceptable results, as it places the focus on the

end goal (good data) rather than the intermediate step of good parameters. Calculating and linking pixel metrics are an expensive and complex operation, respectively, so it is important to ensure that they are contributing something to the classification process. The particularly high importance of pixel metrics at 1.37 m and 2 m above ground shows a promising area to focus and expand on. Perhaps adding more slices at lower heights above ground instead of or in addition to the 3-5 m slices would improve performance even more.

There are several procedural differences between Discovetree and TreeLS that may have contributed to the performance difference. Discovetree automates its own parameters as much as possible, while TreeLS requires a nontrivial amount of manual decision making. Most importantly, the use of the machine learning classifier teaches Discovetree some of the natural variation in the field data.

Failures of circular assumption

An important discovery from this research is on the assumption of circularity. 46% of trees had a circle ratio less than 0.9, while 19% of trees had a circle ratio of less than 0.7. High resolution stem maps, basal area estimates, etc. may not want to use circles (or even cylinders). Further research on species, size, and shape introduce a possible new error component when utilizing DBH. If circles do consistently underestimate the observed basal area, this could potentially be corrected for. This correction factor could reduce errors propagating from the individual tree to the landscape level, complementing the use of TLS as a sampling instrument for larger scale RS data. More complex models of trunk shape that are infeasible to acquire in the field can be automatically generated via TLS scans. For instance, DBH mensuration in the field usually consists of one measurement or two measurements (longest and shortest cross sections), allowing for at best an ellipse. However, TLS scans capture the full perimeter, including any trunk hollows or other variations.

There are three failure cases for the circular assumption that also extend to cylindrical models in 3D space (Figure 24). The first failure case occurs when an ellipse is a better representation of the cross section than a circle (Figure 24, 1a and 1b). In these cases, the true center of the tree and/or its DBH cannot be accurately captured. In the case of 1a, the predicted DBH is close to the observed value (an underestimation of the DBH provided by a measuring tape), but the center of the tree does not match the center of the detection. In the second case of 1b, the center of the tree and detection are accurate, but the DBH is underestimated more significantly. In both cases, the predicted basal area is smaller than the observed value.

The second failure case for the circular assumption is when occlusion occurs. When occlusion is high, the tree is difficult or even impossible to detect. At lower levels, Hough transforms can still detect the trees, but with a larger margin of error. Case 2 is a relatively rare example of Discovetree overestimating, rather than underestimating, the DBH (Figure 24, 2). The detection was fit to the strong arc on the upper left side, rather than fitting to the visible perimeter. The third case is when trees have an irregular shape that cannot be easily modelled with a geometric primitive. In case 3, no single circle or even an ellipse can represent the trunk (Figure 24, 3). These failure cases can co-occur, and often do.

The first failure case, of ellipses, is relatively simple to solve. Hough transforms can be modified to search for ellipses instead of circles by searching for two additional parameters (length/breadth instead of radius, and orientation), although this increases the search space. Occlusion is easiest to solve by increasing the number of initial scan positions, and a limited amount of occlusion can be removed using downsampling and/or morphological closing. The third case is more challenging to solve. A potential solution is to perform an initial search using circles/ellipses and then perform a more detailed search as a second pass.




Figure 24. 1a (Upper left). Ellipse 1b (Upper right) Ellipse and irregular shape. 2 (Lower left) Occlusion and probable ellipse. 3 (Lower right) Irregular shape.

The results of this study suggests that in many cases an ellipse would represent the trunk's cross section better than a circle. Future studies should examine how ellipse detection performs in comparison, not just in the case of identifying trees but also how it treats non-tree structures such as bushes or fallen logs. How circles and ellipses perform at different resolutions is also a potential area of study. One possible link between circularity and accuracy could be pixel/voxel resolution. While Discovetree was using a pixel resolution of 0.01 m^2 , TreeLS used the recommended voxel size of 0.03 m^3 .

Importance of parameter tuning

One of the limitations with current research into stem mapping and tree segmentation is that parameters tend to not generalize very well between ecosystems. As mentioned above, Discovetree's ability to tune and teach itself was a major difference from TreeLS. The performance increase is promising, while leaving room for improvement. Given the distribution of circle ratios in the field data, a method for a stem mapper to independently pick between circles, ellipses, and potentially other shapes is one possible goal. Another benefit of using the field data for training is that it somewhat accounts for occlusion. Because the Hough weights and other pixel metrics of a detection are lower when data is occluded, Discovetree learned to tolerate them through the training process.

This study compared the use of a novel stem mapper Discovetree and an existing tool TreeLS on 1262 trees across 108 plots in the Sierra Nevadas. Discovetree had a higher precision and recall than TreeLS. Looser definitions of accuracy (presence) are a promising way to increase recall without having a large negative impact on size and position accuracy. The circular assumption of tree trunks made in most stem maps leads to an underestimation of DBH, basal area, and other derived values. The use of a machine learning to tolerate deviations from ideal conditions is a promising way to increase performance.

References

- Aicardi, I., Dabove, P., Lingua, A. M., & Piras, M. (2016). Integration between tIs and uav photogrammetry techniques for forestry applications. *iForestBiogeosciences and Forestry*, 10 (1), 41.
- Akay, A. E., Öguz, H., Karas, I. R., & Aruga, K. (2009). Using lidar technology in forestry activities. *Environmental Monitoring and Assessment*, 151, 117– 125.
doi:10.1007/s10661-008-0254-1
- Anderson, K., Hancock, S., Disney, M., & Gaston, K. J. (2016). Is waveform worth it? a comparison of lidar approaches for vegetation and landscape characterization. *Remote Sens Ecol Conserv*, 2 (1), 5–15. doi:10.1002/rse2. 8
- Aschoff, T., & Spiecker, H. (2004). Algorithms for the automatic detection of trees in laser scanner data. *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 36 (8/W2), 71–75.
- Asner, G. P., Mascaro, J., Muller-Landau, H. C., Vieilledent, G., Vaudry, R., Rasamoelina, M., . . . van Breugel, M. (2012). A universal airborne lidar approach for tropical forest carbon mapping. *Oecologia*, 168 (4), 1147– 1160. doi:10.1007/s00442-011-2165-z
- Bienert, A., Maas, H., & Scheller, S. (2006). Analysis of the information content of terrestrial laserscanner point clouds for the automatic determination of forest inventory parameters. in workshop on 3d remote sensing in forestry. In *Workshop on 3d remote sensing in forestry*, International Society for Photogrammetry and Remote Sensing.
- Burt, A., Disney, M., & Calders, K. (2019). Extracting individual trees from lidar point clouds using treeSeg. *Methods Ecol Evol*, 10, 438–445
- Cheong, Y. Z., & Chew, W. J. (2018). The application of image processing to solve occlusion issue in object tracking. *MATEC Web of Conferences*, 152.
- Danskin, S. D., Bettinger, P., Jordan, T. R., & Cieszewski, C. (2009). A Comparison of GPS Performance in a Southern Hardwood Forest: Exploring Low-Cost Solutions for Forestry

- Applications. *Southern Journal of Applied Forestry*, 33 (1), 9–16. doi:10.1093/sjaf/33.1.9.
eprint: <https://academic.oup.com/sjaf/article-pdf/33/1/9/23387455/sjaf0009.pdf>
- De Conto, T., Olofsson, K., Görgens, E. B., Rodriguez, L. C., & Almeida, G. (2017).
Performance of stem denoising and stem modelling algorithms on single tree point
clouds from terrestrial laser scanning. *Computers and Electronics in Agriculture*, 143,
165–176.
- Fadili, M., Renaud, J.-P., Bock, J., & Vega, C. (2019). Registree: A registration algorithm to
enhance forest inventory plot georeferencing. *Annals of Forest Science*, 76.
- Fricker, G. A., Wolf, J. A., Saatchi, S. S., & Gillespie, T. W. (2015). Predicting spatial variations
of tree species richness in tropical forests from high-resolution remote sensing.
Ecological Applications, 25 (7), 1776–1789. doi:<https://doi.org/10.1890/14-1593.1>.
- Girardeau-Montaut, D. (2020). Cloudcompare
- Gorte, D., B. Winterhalder. (2004). Reconstruction of laser-scanned trees using filter operations
in the 3d raster domain. *International Archives of Photogrammetry, Remote Sensing and
Spatial Information Sciences*, 36 (8/W2).
- Hackenberg, J., Morhart, C., Sheppard, J., Spiecker, H., & Disney, M. (2014). Highly accurate
tree models derived from terrestrial laser scan data: A method description. *Forests*, 5 (5),
1069–1105.
- Hackenberg, J., Spiecker, H., Calders, K., Disney, M., & Raunonen, P. (2015). Simpletree —an
efficient open source tool to build tree models from tls clouds. *Forests*, 6 (11),
4245–4294.
- Hudak, A. T., Kato, A., Bright, B. C., Loudermilk, E. L., Hawley, C., Restaino, J. C., . . . Weise, D.
R. (2020). Towards Spatially Explicit Quantification of Pre- and Postfire Fuels and Fuel
Consumption from Traditional and Point Cloud Measurements. *Forest Science*, 66 (4),

- 428–442. doi:10.1093/ forsci/ fxz085. eprint: <https://academic.oup.com/forests/article-pdf/66/4/428/33635281/fxz085.pdf>
- Hyypä, E., Yu, X., Kaartinen, H., Hakala, T., Kukko, A., Vastaranta, M., & Hyypä, J. (2020). Comparison of backpack, handheld, under-canopy uav, and above-canopy uav laser scanning for field reference data collection in boreal forests. *Remote Sensing*, *12*.
- Király, G., & Broly, G. (2007). Tree height estimation methods for terrestrial laser scanning in a forest reserve. *Int. Arch. Photogramm. Remote Sens. Spat. Inf. Sci.*, *36*.
- Koizumi, A., & Hirai, T. (2006). Evaluation of the section modulus for tree-stem cross sections of irregular shape. *Journal of Wood Science*, *52*, 213–219..
- Koizumi, A., Hirai, T., Ikeda, K., & Sawata, K. (2011). Nondestructive measurement of cross-sectional shape of a tree trunk. *Journal of Wood Science*, *57*, 276–281.
- Leeuwen, M. V., Coops, N. C., & Wulder, M. A. (2010). Canopy surface reconstruction from a lidar point cloud using hough transform. *Remote Sensing Letters*, *1*, 125–132.
doi:10.1080/01431161003649339
- Li, D., Guo, H., Jia, W., & Wang, F. (2021). Analysis of taper functions for *larix olgensis* using mixed models and tls. *Forests*, *12* (2). doi:10.3390/f12020196
- Liang, X., & Hyypä, J. (2013). Automatic stem mapping by merging several terrestrial laser scans at the feature and decision levels. *Sensors*, *13* (2), 1614–1634.
- Lin, Y., & Herold, M. (2016). Tree species classification based on explicit tree structure feature parameters derived from static terrestrial laser scanning data. *Agricultural and Forest Meteorology*, *216*, 105–114.
- Lindberg, E., Holmgren, J., Olofsson, K., & Olsson, H. (2012). Estimation of stem attributes using a combination of terrestrial and airborne laser scanning. *European Journal of Forest Research*, *131*, 1917–1931. doi:10.1007/ s10342-012-0642-5

- Liu, S., Bitterlich, W., Cieszewski, C., & Zasada, M. (2011). Comparing the use of three dendrometers for measuring diameters at breast height. *Southern Journal of Applied Forestry*, 35, 136–141. doi:10.1093/sjaf/35.3.136
- Luoma, V., Saarinen, N., Wulder, M. A., White, J. C., Vastaranta, M., Holopainen, M., & Hyypä, J. (2017). Assessing precision in conventional field measurements of individual tree attributes. *Forests*, 8 (2). doi:10.3390/f8020038
- McDaniel, M. W., Nishihata, T., Brooks, C. A., Salesses, P., & Iagnemma, K. (2012). Terrain classification and identification of tree stems using groundbased lidar. *Journal of Field Robotics*, 29 (6), 891–910. doi:https://doi.org/10.1002/rob.21422. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/rob.21422>
- McRoberts, R. E. (2008). Using satellite imagery and the k-nearest neighbors technique as a bridge between strategic and management forest inventories. *Remote Sensing of Environment*, 112 (5), 2212–2221. Earth Observations for Terrestrial Biodiversity and Ecosystems Special Issue. doi:https://doi.org/10.1016/j.rse.2007.07.025
- Natekin, A., & Knoll, A. (2013). Gradient boosting machines, a tutorial. *Frontiers in Neurorobotics*, 7.
- Nogueira, E. M., Nelson, B. W., & Fearnside, P. M. (2006). Volume and biomass of trees in central amazonia: Influence of irregularly shaped and hollow trunks. *Forest Ecology and Management*, 227 (1), 14–21. Perspectives on Site Productivity of Loblolly Pine Plantations in the Southern United States. doi:https://doi.org/10.1016/j.foreco.2006.02.004
- Olofsson, K., Holmgren, J., & Olsson, H. (2014). Tree stem and height measurements using terrestrial laser scanning and the ransac algorithm. *Remote Sensing*, 6 (5), 4323–4344. doi:10.3390/rs6054323
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., . . . Edouard Duchesnay. (2011). Scikit-learn: Machine learning in python. *Journal of Machine*

- Learning Research*, 12 (85), 2825–2830. Retrieved from <http://jmlr.org/papers/v12/pedregosa11a.html>
- Popescu, S., Wynne, R., & Nelson, R. (2003). Measuring individual tree crown diameter with lidar and assessing its influence on estimating forest volume and biomass. *Canadian Journal of Remote Sensing*, 29, 564–577. doi:10. 5589/m03-027 .
- Powers, D. M. (2014, May). *What the f-measure doesn't measure*. Flinders University.
- Price, A., Hapka, A., Gardiner, B., Macdonald, E., & Mclean, P. (2017). *Assessing the stem straightness of trees*. Forestry Commission.
- Pulkkinen, M. (2012). On non-circularity of tree stem cross-sections: Effect of diameter selection on cross-section area estimation, bitterlich sampling and stem volume estimation in scots pine. *Silva Fennica*, 46. doi:10.14214/ sf.924
- RapidLasso GmbH. (2019). Lastools version 191111.
- Raunonen, P., Kaasalainen, M., Akerblom, M., Kaasalainen, S., Kaartinen, H., Vastaranta, M., . . . Lewis, P. (2013). Fast automatic precision tree models from terrestrial laser scanner data. *Remote Sensing*, 5 (2), 491–520. doi:10. 3390/rs5020491
- Reese, H., Nilsson, M., Pahén, T., Hagner, O., Joyce, S., Tingelöf, U., . . . Olsson, H. (2004). Countrywide estimates of forest variables using satellite data and field data from the national forest inventory. *Ambio*, 32, 542–8. doi:10. 1639/0044-7447(2003)032[0542:CEOFVU]2.0.CO;2
- Riegl Laser Measurement Systems GmbH. (2020). Riscan pro 2.0.
- Roedig, E., Cuntz, M., Heinke, J., Rammig, A., & Huth, A. (2017). Spatial heterogeneity of biomass and forest structure of the amazon rain forest: Linking remote sensing, forest modelling and field inventory. *Global Ecology and Biogeography*, 26. doi:10.1111/geb.12639

- Saito, T., & Rehmsmeier, M. (2015). The precision-recall plot is more informative than the roc plot when evaluating binary classifiers on imbalanced datasets. *PLOS ONE*, *10* (3), 1–21. doi:10.1371/journal.pone.0118432
- Trochta, Jan & Král, Kamil & Janík, David & Adam, Dušan. (2013). Arrangement of terrestrial laser scanner positions for area-wide stem mapping of natural forests. *Canadian Journal of Forest Research*. *43*. 10.1139/cjfr-2012-0347.
- Tuominen, S., Pitkänen, J., Balazs, A., Korhonen, K., Hyvönen, P., & Muinonen, E. (2014). Nfi plots as complementary reference data in forest inventory based on airborne laser scanning and aerial photography in finland. *Silva Fennica*, *48*. doi:10.14214/sf.983
- US Forest Service. (2021). *Forest inventory and analysis national core field guide volume i: Field data collection procedures for phase 2 plots*.
- Valbuena, R., Mauro, F., Rodríguez-Solano, R., & Manzanera, J. (2010). Accuracy and precision of gps receivers under forest canopies in a mountainous environment. *Spanish Journal Of Agricultural Research*, *8*, 1047–1057.
- Vepakomma, U., St-Onge, B., & Kneeshaw, D. (2010). Response of a boreal forest to canopy opening: Assessing vertical and lateral tree growth with multi-temporal lidar data. *Ecological Applications - ECOL APPL*, *21*. doi:10.1890/09-0896.1
- van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., Boulogne, F., Warner, J. D., Yager, N., . . . the scikit-image contributors. (2014). Scikit-image: Image processing in python. *PeerJ*, *2*, e453. doi:10.7717/peerj.453
- Wilkes, P., Lau, A., Disney, M., Calders, K., Burt, A., de Tanago, J. G., . . . Herold, M. (2017). Data acquisition considerations for terrestrial laser scanning of forest plots. *Remote Sensing of Environment*, *196*, 140–153.
- Williamson, R. L. (1975). Out-of-Roundness in Douglas-fir Stems. *Forest Science*, *21* (4), 365–370. doi:10.1093/forestscience/21.4.365. eprint: <https://academic.oup.com/forestscience/article-pdf/21/4/365/23060661/forestscience0365.pdf>

- Xia, S., Wang, C., Pan, F., Xi, X., Zeng, H., & Liu, H. (2015). Detecting stems in dense and homogeneous forest using single-scan tls. *Forests*, 6 (11), 3923– 3945.
doi:10.3390/f6113923
- Yan, S., Jing, L., & Wang, H. (2021). A new individual tree species recognition method based on a convolutional neural network and high-spatial resolution remote sensing imagery. *Remote Sensing*, 13, 479. doi:10.3390/rs13030479
- Yang, Q., Su, Y., Jin, S., Kelly, M., Hu, T., Ma, Q., . . . Guo, Q. (2019). The influence of vegetation characteristics on individual tree segmentation methods with airborne lidar data. *Remote Sensing*, 11 (23), 2880. doi:10.3390/rs11232880
- Yuen, H. K., Princen, J., Dlingworth, J., & Kittler, J. (1989). A comparative study of hough transform methods for circle finding. In *Alvey vision conference*.

Conclusion

This thesis demonstrated the importance of both field data and parameter tuning in the creation of automated stem maps from terrestrial laser scanning (TLS) data. Discovertree's workflow improves on existing techniques while removing the need for manual parameter tuning. This approach appeared to work better across species and size classes compared to a similar algorithm. Beyond the specific workflow presented herein, I believe the framework of using machine learning calibrated against field data can be used as a preprocessing step for the more complex task of tree segmentation, making it compatible with other workflows. The analysis on trunk circularity in natural systems has implications for management as well. I discovered that a noticeable portion of trees did not meet this assumption, resulting in a systematic underestimation of DBH and any metrics derived from it, such as basal area.