

PROCESSING TREE POINT CLOUDS USING GAUSSIAN MIXTURE MODELS

David Belton, Simon Moncrieff and Jane Chapman

Cooperative Research Centre for Spatial Information (CRCSI)
Department of Spatial Sciences, Curtin University of Technology
Perth WA, Australia

d.belton@curtin.edu.au, s.moncrieff@curtin.edu.au, jane.chapman@student.curtin.edu.au

Commission V, WG 3

KEY WORDS: Laser Scanning, Principal Component Analysis, Classification, Gaussian Mixture Models

ABSTRACT

While traditionally used for surveying and photogrammetric fields, laser scanning is increasingly being used for a wider range of more general applications. In addition to the issues typically associated with processing point data, such applications raise a number of new complications, such as the complexity of the scenes scanned, along with the sheer volume of data. Consequently, automated procedures are required for processing, and analysing such data. This paper introduces a method for modelling multi-modal, geometrically complex objects in terrestrial laser scanning point data; specifically, the modelling of trees. The model method comprises a number of geometric features in conjunction with a multi-modal machine learning technique. The model can then be used for contextually dependent region growing through separating the tree into its component part at the point level. Subsequently object analysis can be performed, for example, performing volumetric analysis of a tree by removing points associated with leaves. The workflow for this process is as follows: isolate individual trees within the scanned scene, train a Gaussian mixture model (GMM), separate clusters within the mixture model according to exemplar points determined by the GMM, grow the structure of the tree, and then perform volumetric analysis on the structure.

1 INTRODUCTION

Laser scanning has been adopted for many traditional surveying and photogrammetric applications. Examples of these applications include: modelling for industrial and engineering applications, deformation monitoring, volumetric analysis, topographic surveys. However, the use of laser scanning is moving beyond such traditional applications to be adopted for an increasing number, and variety of applications, including cultural and heritage recording, archaeology, asset management, and modelling vegetation. Subsequently, in addition to the pre-processing issues typically related to laser scanning, such as registration and calibration, the complexity of the scenes being scanned, the volume of data, and the potentially complex and heterogeneous nature of the objects in the scene all need to be accounted for. This has led to the need for semi-automated, and automated tools to aid in analysis undertaken in these new applications. To achieve a greater level of automation, a number of techniques used in fields such as signal processing, machine learning and computer vision can be adopted to point cloud data.

This paper details an example of one such technique, combining feature extraction and machine learning for the modelling of trees from point clouds captured in a forestry landscape. To determine the model for a tree, a number of feature sets based on Principal Component Analysis (PCA) performed on local neighbourhoods are considered, combined, and used to train a multi-modal modelling technique, a Gaussian Mixture Model (GMM). The use of a mixture model is appropriate for two reasons: first, a multi-modal approach is required due to the differing properties of the components of a tree (leaf, branches, and trunk),

and second, the model can be used to cluster the tree into its component parts. This second property can be leveraged to cluster the tree data at a point level, which can then be applied to context based region growing, thus enabling the structural analysis of the tree by removing points corresponding to leaves. From this clustering, the points belonging to the main tree structure can be identified. These points can then be further processed to model the structure of a tree. Such a case is presented based on examining horizontal cross sections of the tree and extracting a graphical representation in the form of a tree graph. This can then be used as context for further analysis. An example provided is the examination of the volume of the main tree structure, where the graphical representation is used to segment the tree into parts to allow for simple fitting of cylinders through RANSAC to approximate the carbon content of a tree.

2 BACKGROUND

2.1 Tree modelling

There is much interest in modelling the properties of trees and forest through the use of laser scanning (van Leeuwen et al., 2011). A large portion of previous work has concentrated on airborne laser scanning (ALS). This is because large areas of forest and vegetation can be captured quickly. Information extracted includes the canopy layers, type and number of trees, and heights (Hyypä et al., 2006). As resolution has increased and waveform modelling became more common, methods were developed for extracting individual trees. However, this still lacks in point density compared to terrestrial systems. While it could

identify and separate out different trees, it was difficult to adequately model various attributes such as structure, surface area and volume.

Terrestrial laser scanning, and mobile laser scanning offered an increase resolution in point density. It has been used in conjunction with ALS to model the properties of the trees and match them with those observed from the ALS point data to infer more precise attributes of an entire forest (Lindberg et al., 2012). The modelling of trees from TLS data can be categorized as either extracting the skeleton of the tree, or partitioning the tree into simple shapes. These shapes normally consist of cylinders and conic sections (Pfeifer et al., 2004; Chaperon and Goulette, 2001). Skeletonisation methods have been predominately based on extracting the tree graph representation from oct-trees (Bucksch et al., 2009), from regularly gridded voxels and rasters (Gorte, 2006), and from extracting the medial axis representation (Su et al., 2011). Meshing techniques can also be applied to represent the tree and for calculation of the volume and surface area, but benefit from the use of context such as the skeleton of the tree (Xu et al., 2007).

A majority of existing techniques are designed to work on single trees. In forestry applications the data is comprised of multiple trees. Often each tree has to be isolated before applying the various modelling techniques. In airborne applications trees can be isolated through examination distribution of the points in terms of density and height. This delineates the crowns of the trees, and allows individual trees to be isolated (Hyyppä et al., 2006). While the density of the TLS is not consistent, a similar technique can be applied. This is performed by extracting the points with the highest distance from the extracted digital terrain model, and performing the crown segmentation on these points if there is sufficient density (Brolly et al., 2009). Alternatively, a slice can be taken through the data for a given interval above the ground and below the canopy. This is taken high enough above the ground so as not to include shrubs and low level vegetation, and low enough so as to not include the upper level of the canopy. The majority of the remaining points are sampled from the tree trunks. A circle or cylinder detection algorithm can be applied to find the tree trunks, and isolated the component trees (Brolly et al., 2009).

2.2 Feature selection

In most cases, automated processing methods used features that can be categorised as being based on either geometric or spectral information. Geometric information is classed as information derived from the point position and measurement in the 3D coordinate space. Spectral information is primarily derived from the intensity return of the laser, or from co-registered imagery for either an external or integrated camera. In the case of this paper, the features used for classification will be derived from the geometric information.

With the intensity return, it is based on factors such as signal strength, wavelength, surface reflectance and the range

and incident angle. It has been shown that it can be combined with RGB channels to perform spectral classification, similarly to vegetation index used in remote sensing (Lichti, 2005), especially for near-infrared lasers. However, because the return value is dependent on the range and surface incident angle, it is not consistent for surfaces over a scene or when data is collected from multiple setups unless it can be corrected (Kukko et al., 2008).

For colour information, there is often a temporal difference between the capture of the points and the imagery. If the structure being sampled is static, then this often has little effect. For trees however, factors such as wind often change the position of the upper canopy, making it much more dynamic in nature. As such, there is often discrepancies in the matching of the pixels to the points for the upper layers of the trees, especially for leaves. In addition, for forest scenes, the high presence of shadows (and their movement) can also effect the colour information, making it inconsistent between multiple setups. Finally, the small size of the leaves and details in the upper canopy, and the difference in resolution of the laser beam and the camera pixels can lead to over saturation of the colour information for such points. Consequently, this paper focuses on examining classification using simple derived geometric information on the point neighbourhood.

Geometric information is usually extracted from a local neighbourhood of points to approximate the true geometric properties for a point. This can then be used to derive higher levels of information. This is usually done in three ways; local surface fitting, geometric primitive or prototype fitting, or principal component analysis (PCA). Each method produces similar or complimentary results.

Local surface fitting involves fitting a surface to the point neighbourhood, normally a polynomial and generally based on the local surface orientation at the neighbourhood. The surface is restricted to 2^{nd} order surface on the assumption that the neighbourhood is small enough to be modelled by a simple surface, thus preventing over fitting due to errors and noise. The surface parameters are then examined to extract measures of curvature, directions of change, and surface types such as inflections, and local maxima and minima (Crosilla et al., 2008).

Fitting local geometric primitives, using methods such as Random Sample Consensus (RANSAC), is most applicable to the analysis of structured information, for instance industrial scenes. The process involves fitting geometric primitive surface types, including planes, cylinders, cones, tori and spheres, to seed neighbourhoods to determine the best descriptor class for the surface. This works best for man-made, regular scenes compared to naturally occurring free-form surfaces.

The third method for extracting geometric information is to use PCA on the local neighbourhood of points. This involves examining the eigenvalue decomposition of the the covariance matrix. The eigenvalues and eigenvectors define the linear combinations of elements that describe

the maximum variation of the points distributed throughout a neighbourhood. The underlying surface and structure highly effects the sampling and distribution of points, and this information is represented through the eigenvalues and eigenvectors. Consequently, PCA is suitable for describing and classifying points based on the surrounding local neighbourhood. Information that can be approximated comprises: the surface normal, curvature, texture, and shape, if the point is close to a boundary, edge or corner point, and the local surface orientation and coordinate system. PCA has also been applied to the surface normals to examine underlying shape in the form of tensor voting, and the determination of the principal curvature directions.

3 CLASSIFICATION

The scanner used to capture the information was a Leica C10 scan station. This is a time of flight instrument, with a quoted point precision of 6mm and a range of 135m (Leica Geosystems HDS, 2008). The region of capture was in a national park close to Walpole in Western Australia during spring, containing old growth forests, primarily consisting in this case of giant Red Tingle trees. These trees can range in height to approximately 40m, as is the case in the example presented. Multiple setups were used around the area to capture the point coordinate data. This was to ensure adequate coverage and density. The data was down-sampled to a maximum point spacing of 0.02m to reduce the amount of redundant data to process. The goal of classification in this paper is to class the points as either leaves, trunk and branches, or unknown. The type of trees in this region maintain their leaves all year round. As such, to get an accurate representation of the main tree structure, data captured from leaves need to be classified and removed. The next section outlines the geometric features examined, and how they were applied to the classification process. A Gaussian mixture model (GMM) was used to cluster the points into different models based on features. Tree model were then examined and combined to generate the different classes.

3.1 Features from Principal Component Analysis

A variety of different geometric features and their combinations were explored. These geometric features were derived from performing PCA over a local neighbourhood of points within a set radius. The first step is to calculate the covariance matrix of the local neighbourhood of coordinate points, described as:

$$C = \frac{1}{k} \sum_{i=1}^k (p_i - \bar{p})(p_i - \bar{p})^T \quad (1)$$

p_i is the i^{th} point in the neighbourhood of k points, with \bar{p} denoting the centroid of the neighbourhood, or the mean coordinate value. The eigenvalues are found by decomposition into the form given by:

$$C = \sum_{i=0}^2 \lambda_i e_i e_i^T \quad (2)$$

The simplest information that will be used are the eigenvalues themselves. As mentioned, they represent the distribution of the points in the local neighbourhood, and the variance in the direction of the associated eigenvectors. The smallest eigenvalue, λ_0 , represents the variance of the points to a planar surface, with the surface normal direction represented by e_0 . For branches, trunks, ground points and other smooth surfaces, it will be small. For regions of high variations, such as neighbourhoods containing leaves, shrubs, it will be high. For non-planar surfaces, this value is affected by the neighbourhood size. As the neighbourhood size increases, so will the variation in the normal direction, and hence the value for λ_0 will also increase. This can be compensated for by dividing the value by the total population variation, to get the curvature approximation κ (Pauly et al., 2002) as denoted by:

$$\kappa \approx \frac{\lambda_0}{\lambda_0 + \lambda_1 + \lambda_2} \quad (3)$$

The other eigenvalues λ_1 and λ_2 denote the variance in the tangential direction of the best fit planar surface. If both are large, then the points are distributed across the neighbourhood in both directions. This is common for surfaces such as the ground and trunk, where the surface is large enough. When the structure is narrower, it causes the distribution to be elongated, resulting in a smaller λ_1 value. Such occupancies include branches and narrow trunks where the neighbourhood radius is less than the radius of the trunk or branch. Twigs and sticks will have small λ_0 and λ_1 values, and most of the neighbourhood will be distributed in the direction of e_2 , denoted by a large λ_2 value.

In addition, ratios and differences between the eigenvalues can also be used (Gumhold et al., 2001). These describe the neighbourhood distribution in terms of the relation between the eigenvalues. Locally planar smooth surfaces such as the tree trunk will have a low difference $\lambda_2 - \lambda_1$ and a high difference in $\lambda_1 - \lambda_0$. A highly linearly distributed surface/structure such as a branch will have a high difference $\lambda_2 - \lambda_1$ and a low difference in $\lambda_1 - \lambda_0$. A highly curved or randomly distributed surface such as a neighbourhood containing leaves will have a low difference $\lambda_2 - \lambda_1$ and a low difference between $\lambda_1 - \lambda_0$.

PCA can also be performed on the normal directions \hat{n} of the points in a local neighbourhood (Jiang et al., 2005). The covariance matrix is formed as:

$$C^{(n)} = \frac{1}{k} \sum_{i=1}^k (\hat{n}_i - \bar{\hat{n}})(\hat{n}_i - \bar{\hat{n}})^T \quad (4)$$

with the eigenvalue decomposition performed as in equation 2. The eigenvalues are examined in a similar fashion, except instead of encapsulating the distribution of points, it denotes the change in surface. A high $\lambda_1^{(n)}$ and $\lambda_2^{(n)}$ denotes change in curvature in two directions (such as a branch intersection or regions containing leaves), a high $\lambda_2^{(n)}$ and a low $\lambda_1^{(n)}$ denotes a change in only one direction (such as a branch or trunk), and a low $\lambda_1^{(n)}$ and $\lambda_2^{(n)}$ denotes no change in surface curvature locally (such as for a

Feature Set	Neighbourhood Radius	Tree Clusters	Mixed Clusters	Leaf Clusters	Tree points	Leaf points
eigenvalues	0.1m	1	2	3	8%	92%
	0.2m	2	0	4	15%	85%
	0.3m	2	0	4	13%	87%
	0.4m	2	0	4	11%	89%
	all	2	1	3	11%	89%
	all (PCA)	3	0	3	16%	84%
all features	0.1m	1	2	3	16%	84%
	0.2m	1	1	4	6%	94%
	0.3m	2	1	3	19%	81%
	0.4m	2	1	3	12%	88%
	all	2	0	4	15%	85%
	all (PCA)	3	0	3	14%	86%

Table 1: Different combinations of features and neighbourhood sizes tested for with GMM.

flat surface or wide trunk).

3.2 Neighbourhood size

The size of the neighbourhood will determine the resolution of the features extracted. Smaller neighbourhoods enable smaller structures to be detected, but will suffer from more noise due to the lack of redundancy. For larger neighbourhood sizes, the increase redundancy reduces the effects of noise, but at the cost of losing finer resolution detail. Because of this, the feature values at different radii are used. This allows clustering to take into account different resolutions of the underlying structure. In this case, the neighbourhood size were used for radii of 0.1m, 0.2m, 0.3m, 0.4m and 0.5m. Large surfaces such as trunks exhibit more consistent feature values. Smaller structures such as branches exhibit changes in the feature values, and become somewhat more consistent for higher radius values (assuming a single structure within the neighbourhood). Where the structure undergoes change or is not consistent, there is little consistency.

3.3 Classification and clustering using Gaussian Mixture Modelling

Based on the features described in the previous section, the points were clustered using a Gaussian mixture model (GMM) (Reynolds, 2008), an unsupervised parametric clustering method. A GMM provides a representation of the distribution of points in a given feature space. This is done by assuming that the overall distribution is formed by a mixture of n Gaussian distributions in the feature space. The attributes of these n Gaussian distributions are then determined such that their convolution of the n Gaussian distributions best models the overall distribution of the points in the feature space. Each of these Gaussian models can then represent a local cluster of points in the feature space.

Different feature values were used as an input for the GMM. This included utilising only the eigenvalues at fixed neighbourhood sizes, utilising all the features described in section 3.1 at fixed neighbourhood sizes, and finally combining the features over varying neighbourhood sizes. This is briefly outlined in table 1. Six clusters were chosen to best model the combined distribution of the points and to encapsulate the different properties between the tree trunk, branches, small twigs and sticks, leaves and random noise. The number of clusters that contain points from the main

tree structure, points sampled from leaves, and a mixture is provided in table 1. Example of these clusters using all the features and multiple neighbourhood radii are presented in figure 1

From examination of the resulting clusters, performing GMM on the features over multiple neighbourhood radii appeared to differentiated the different attributes of the tree better than a fixed radius. This is because the fixed radius value heavily influences the feature values, as outlined in section 3.2. However, when examining multiple radius values the feature space comprises of fifteen dimensions, which is a high dimensionality feature space. To reduce this, PCA was applied to the feature space to reduce its dimensionality. Only those combinations that were deemed to contribute significantly were kept, which in this case comprised of six linear combinations of the original feature dimensions. The results from the GMM on the reduced feature space is presented in figure 1.

The clusters were then examined to see which class they contained based on classes for leaves, tree trunk and branches, and unknown. This was done through manual examination, and the examination of the mean value of each cluster in the feature domain and whether they were close together and described the aspects of each class. Where multiple clusters contained the same class, the Gaussian model for each cluster were combined to create a single model for the class. The resulting classes are shown in figure 2.




Figure 2: merged classes representing (a) leaves and (b) the trunk and branches from the Gaussian mixture model

4 TREE STRUCTURE MODELLING

From the classification of the points into leaves and tree structure, the process of modelling the tree and the various attributes such as volume, surface area, and the structure and shape of the tree is simplified. Techniques such as skeletonisation and surface meshing generally perform better on a tree structure when leaves are removed. In this section, the extraction of two attributes from the classified data is examined; the tree structure and the approximate volume of wood.

4.1 Tree Structure generation

The method used to generate the skeleton in this case was based on examining horizontal slices of the tree from ground




Figure 1: Merged clusters from GMM representing (a) main trunk, (b) branches, (c) leaves and (d) mixed points from trunk and branches

level up through the canopy iteratively. At each stage, the points in each horizontal slice were clustered and an ellipse was found for each cluster. The ellipse was based on either best fit ellipse if the cluster was significant, or on the minimally bounding ellipse if the fitting routine failed or the cluster was too small.

A similar approach of fitting ellipses has been applied for identify pipes and their paths in industrial settings (Mapurisa and Sithole, 2012). The ellipse method works well for the main trunk and the initial branching sections. This is because there is few occlusions and the points in the cluster represent a good sampling of the elliptical cross section. For higher up in the canopy, the ellipse fitting exhibits poor performance. This is due to the organic nature of the structure, high incidents of occlusion, limited surface sampling and the radii of the branch approaching the noise level in point sampling. The last makes it difficult to differentiate between the local neighbourhood variance being the result of sampling noise or the radius of the branch. Since the path of the structure is of importance, the smallest ellipse that encapsules the points is fitted, and its centre used to predict the path of the tree structure.

The centre of the ellipse/cluster was used as a node in the graph. An edge joining the nodes between consecutive layers was created if the ellipses for the two nodes overlapped. From this graph representation, a cyclic detection algorithm was applied to identify cycles in the graph, and remove them by merging common nodes in the cycle that were on the same layer. The results are presented in figure 3, with figure 3(a) showing the found ellipses, and figure 3(b) the final tree skeleton representation. This skeleton representation provides context for the volume analysis. The slicing horizontally through the data will not produce the most optimal cross section of the branches. This could be modified that the cross section at a give point is nominally parallel to the surface of the canopy, or it could be defined locally for a cluster based on the direction of the axes of the tree branch for lower levels in the graph.

4.2 Volume measurement for carbon capture

The volume of the tree is an important attribute. It is often used for studies on carbon capture and to predict the effect




Figure 3: (a) tree represented by ellipses at vertical intervals, and (b) the tree represented by the extracted skeleton

of ecosystems on climate change. A simple allometric for calculating the volume of trees is described in Dean and Wardell-Johnson (2010) and approximates the volume of the tree based on the diameter of the tree at breast height and the height.

$$TotalVolume = \frac{\left(\frac{d}{200}\right)^2 + \pi}{3} h \quad (5)$$

Using this method resulted in an approximate volume of 34.3m³. This value can be seen as conservative when compares to the volume calculated for the modelling of the tree using more complex representations such as meshes or cylindrically objects. This can be demonstrated as in figure 4(a), where the tree is represented by manually fitting cylinders and pipe elements to the data, and results in a volume of 67.174m³. The previous allometric method only encapsulated approximately 49.1 percent of this volume. However, this manual procedure is quite labour intensive.

Consequently, a more automated approach was applied using the extracted skeleton. The points were segmented into individual components based on the paths from the tree graph between branching nodes. For each component a cylinder fitting routine was applied to model that section of the tree by cylinders. The results of this approach is illustrated in figure 4(b). The volume of the combined automatically fitted cylinders was determined to be 74.457m³.

Care must be taken when taken into account the cylinders fitted to segments in the upper layers. In these cases the radii of these smaller cylinders are close to the point spacing and the error in their sampling. Because of this, any instance where the cylinder radius was under 5cm were not included.




Figure 4: (a) manual cylinder fitting and (b) cylinder fitting by region growing on the segments

5 CONCLUSION AND FUTURE WORK

This paper contains a methodology for classifying and processing point cloud data from trees. The techniques were applied to data captured from a forest of Red Tingle trees in the Southwestern part of Australia using multiple setups of a Leica C10 scan station. The classification was performed using Gaussian mixture model as an unsupervised clustering method to separate the leaves from the rest of the tree structure. Several features were examined, with the eigenvalues calculated from multi-resolution local neighbourhoods producing the best results. A skeletonisation method based on examining the data at different vertical intervals was applied to the points classified as belonging to the tree structure. This was to provide context for fitting cylinder sections to the point cloud to calculate the volume. It was demonstrated that a significant amount of the volume was not being capture using traditional allometric approaches. Future work will concentrate on extending the modelling of multiple trees in a forest scene consisting of varying types. This will also include the examination of changes in volume and structure over time.

ACKNOWLEDGEMENTS

This work has been supported by Curtin University of Technology, the Cooperative Research Centre for Spatial Information, who activities are funded by the Australian Commonwealth's Cooperative Research Centres Programme. The authors would also like to thank Grant Wardell-Johnson from the Curtin Institute for Biodiversity and Climate and Dr Tom Schut for help in data collection.

References

Brolly, G., Kiraly, G. et al., 2009. Algorithms for stem mapping by means of terrestrial laser scanning. *Acta Silvatica et Lignaria Hungarica* 5, pp. 119–130.

Bucksch, A., Lindenberg, R. C. and Menenti, M., 2009. Skeltre - fast skeletonisation for imperfect point cloud data of botanic trees. In: *Proceedings of the 2nd Eurographics conference on 3D Object Retrieval, EG 3DOR'09*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, pp. 13–20.

Chaperon, T. and Goulette, F., 2001. Extracting cylinders in full 3D data using a random sampling method and the gaussian image. In: *VMV '01: Proceedings of the Vision Modeling and Visualization Conference 2001*, , pp. 35–42.

Crosilla, F., Visintini, D. and Sepic, F., 2008. A statistically proven automatic curvature based classification procedure of laser points. In: *XXI ISPRS Congress*, Beijing, China, Jul. , pp. 3–11.

Dean, C. and Wardell-Johnson, G., 2010. Oldgrowth forests, carbon and climate change: Functions and management for tall openforests in two hotspots of temperate australia. *Plant Biosystems - An International Journal Dealing with all Aspects of Plant Biology* 144(1), pp. 180–193.

Gorte, B., 2006. Skeletonization of laser-scanned trees in the 3d raster domain. In: A. Abdul-Rahman, S. Zlatanova and V. Coors (eds), *Innovations in 3D Geo Information Systems, Lecture Notes in Geoinformation and Cartography*, Springer Berlin Heidelberg, pp. 371–380.

Gumhold, S., Wang, X. and MacLeod, R., 2001. Feature extraction from point clouds. In: *10th International Meshing Roundtable*, Sandia National Laboratories, , pp. 293–305.

Hyypä, J., Yu, X., Hyypä, H. and Maltamo, M., 2006. Methods of airborne laser scanning for forest information extraction. In: *Workshop on 3D Remote Sensing in Forestry*, Vol. 14.

Jiang, J., Zhang, Z. and Ming, Y., 2005. Data segmentation for geometric feature extraction from lidar point clouds. *Geoscience and Remote Sensing Symposium, 2005. IGARSS '05. Proceedings. 2005 IEEE International 5*, pp. 3277–3280.

Kukko, A., Kaasalainen, S. and Litkey, P., 2008. Effect of incidence angle on laser scanner intensity and surface data. *Applied Optics* 47(7), pp. 986–992.

Leica Geosystems HDS, 2008. [Website] Leica C10 Datasheet, <http://www.leica-geosystems.com/hds>. accessed: September 2013.

Lichti, D. D., 2005. Spectral filtering and classification of terrestrial laser scanner point clouds. *The Photogrammetric Record* 20(111), pp. 218–240.

Lindberg, E., Holmgren, J., Olofsson, K. and Olsson, H., 2012. Estimation of stem attributes using a combination of terrestrial and airborne laser scanning. *European Journal of Forest Research* 131(6), pp. 1917–1931.

Mapurisa, W. T. and Sithole, G., 2012. Deformation detection in piping installations using profiling techniques. In: *XXII ISPRS Congress*, Melbourne, Australia, ISPRS, , pp. 141–146.

Pauly, M., Gross, M. and Kobbelt, L. P., 2002. Efficient simplification of point-sampled surfaces. In: *VIS '02: Proceedings of the conference on Visualization '02*, IEEE Computer Society, Boston, Massachusetts, pp. 163–170.

Pfeifer, N., Gorte, B. and Winterhalder, D., 2004. Automatic reconstruction of single trees from terrestrial laser scanner data. In: *Proceedings of 20th ISPRS Congress*, , pp. 114–119.

Reynolds, D. A., 2008. Gaussian mixture models. *Encyclopedia of Biometric Recognition*.

Su, Z., Zhao, Y., Zhao, C., Guo, X. and Li, Z., 2011. Skeleton extraction for tree models. *Mathematical and Computer Modelling* 54(3), pp. 1115–1120.

van Leeuwen, M., Hilker, T., Coops, N. C., Frazer, G., Wulder, M. A., Newnham, G. J. and Culvenor, D. S., 2011. Assessment of standing wood and fiber quality using ground and airborne laser scanning: A review. *Forest Ecology and Management* 261(9), pp. 1467–1478.

Xu, H., Gossett, N. and Chen, B., 2007. Knowledge and heuristic-based modeling of laser-scanned trees. *ACM Transactions on Graphics (TOG)* 26(4), pp. 19.