

BAYESIAN OBJECT RECOGNITION FOR THE ANALYSIS OF COMPLEX FOREST SCENES IN AIRBORNE LASER SCANNER DATA

Hans-Erik Andersen ^{a, *}, Stephen E. Reutebuch ^b, Gerard F. Schreuder ^a

^a University of Washington, College of Forest Resources, Seattle, WA, 98195 USA - (hanserik, gsch)@u.washington.edu

^b USDA Forest Service, Pacific Northwest Research Station, Seattle, WA, 98195 USA - sreutebuch@fs.fed.us

Commission III, WG III/3

KEY WORDS: Forestry, laser scanning, LIDAR, object, recognition, remote sensing, statistics

ABSTRACT:

Bayesian object recognition is applied to the analysis of complex forest object configurations measured in high-density airborne laser scanning (LIDAR) data. With the emergence of high-resolution active remote sensing technologies, highly detailed, spatially explicit forest measurement information can be extracted through the application of statistical object recognition algorithms. A Bayesian approach to object recognition incorporates a probabilistic model of the active sensing process and places a prior probability model on object configurations. LIDAR sensing geometry is explicitly modelled in the domain of *scan space*, a three-dimensional analogue to two-dimensional *image space*. Prior models for object configurations take the form of Markov marked point processes, where pair-wise object interactions depend upon object attributes. Inferences are based upon the posterior distribution of the object configuration *given* the observed LIDAR. Given the complexity of the posterior distribution, inferences are based upon dependent samples generated via Markov chain Monte Carlo simulation. This algorithm was applied to a 0.21 ha area within Capitol State Forest, WA, USA. Algorithm-based estimates are compared to photogrammetric crown measurements and field inventory data.

1. INTRODUCTION

1.1 Automated forest inventory

While national and local inventories often utilize remotely sensed data for stratified sampling and classification of general forest type, most of these programs remain heavily reliant upon expensive field data for individual tree-level information. At the national level in the United States, individual tree inventory information is collected at considerable expense. It is significant that senior researchers within the USDA Forest Service Forest Inventory and Analysis (FIA) program have recognized the need for the development of automated forest interpretation and measurement algorithms to reduce human intervention and labor costs (Gulden, 2000).

1.2 The LIDAR technology

LIDAR (LIght Detection And Ranging) is an operationally mature remote sensing technology that can provide highly accurate measurements of both forest canopy and ground surface. While specifications vary among systems, LIDAR systems emit from 5,000 - 100,000 pulses per second. In forested areas, individual LIDAR pulses can penetrate the forest canopy through gaps, and can therefore acquire information relating to three-dimensional forest structure as well the underlying terrain surface.

1.3 LIDAR analysis for forest measurement applications

In recent years, there has been increasing interest in the use of LIDAR for automated detection and measurement of forest features. Research efforts in the last fifteen years were focused

on the use of small footprint (0-1 m pulse diameter) LIDAR systems to estimate forest stand level parameters (Nelson *et al.*, 1988; Means *et al.*, 2000). Researchers in Canada used a model-based approach to recover tree heights from LIDAR canopy height measurements (Magnussen *et al.*, 1999). Three-dimensional mathematical morphology has been applied to a high-resolution LIDAR-based canopy surface model to extract individual tree measurements (Andersen *et al.*, 2001).

1.4 Automated individual tree crown recognition through template matching

With the recognition that high-resolution remotely sensed spatial data can support more intensive forest management practices, there has been increasing interest in recent years in the development of algorithms for automated identification and measurement of individual trees using high-resolution, two dimensional digital imagery. Several studies have used a model-based approach to locate individual trees using tree crown template models (Pollock, 1996; Larsen, 1998; Sheng *et al.*, 2001).

Researchers in Scandinavia have attempted to model the relationship between the spatial distribution of individual trees and the position of spectral maxima in a digital image (Dralle and Rudemo, 1997). Another Scandinavian study has used deterministic parameter search methods for maximum likelihood estimation on a spatial point process model to infer the parameters of a disturbance model that relates the true position of tree-tops to those observed on an aerial photograph (Lund and Rudemo, 2000).

* Corresponding author.

2. STUDY AREA AND DATA

2.1 LIDAR data

LIDAR data were acquired with a Saab TopEye system over a 5 km² area within Capitol State Forest, WA in the spring of 1999. The sensor settings and flight parameters are shown in Table 1. Data were provided in the form of an ASCII text file, with GPS time, aircraft position, and coordinate position for the first laser reflection included.

Flying height	200 m
Flying speed	25 m/s
Swath width	70 m
Forward tilt	8 degrees
Laser pulse density	3.5 pulses/m ²
Laser pulse rate	7000 pulses/sec

Table 1. Flight parameters and LIDAR system settings.

The LIDAR vendor also provided a LIDAR-derived digital terrain model (DTM) for the study area with a 4.57-meter (15-ft) resolution.

2.2 Aerial photography

Large-scale (1:7000) normal-color aerial photography was acquired over the study area in 1999. This photography was oriented in an analytical stereoplotter.

3. METHODS

3.1 Bayesian image analysis

In general, Bayesian image analysis provides a means to incorporate prior knowledge or beliefs into the analysis of remotely sensed data (Besag, 1993). These *a priori* beliefs are represented in the form of a prior distribution, or prior model, that is placed over the image and is updated upon observation of the data. Formally, if this prior description of the image is denoted as $p(x)$, then the conditional spatial distribution of this description, given the observed image y , is given by:

$$p(x | y) \propto l(y | x)p(x) \quad (1)$$

In Bayesian parlance, this conditional distribution $p(x|y)$ is referred to as the *posterior* distribution, on which all inferences are based. In Bayesian inference this posterior distribution is always represented as the product of the *likelihood* $l(y|x)$ and the prior $p(x)$. Typically, the goal in Bayesian inference is to calculate expectations or credible intervals (explicit probability statements made regarding the range of a parameter given the observed data).

Bayesian image analysis has traditionally been carried out using digital images consisting of a discrete grid of picture elements (or *pixels*). Often the objective is to reconstruct an "underlying" image that has been distorted through a noise process.

3.2 Bayesian object recognition

More recently, the methods of Bayesian image analysis have been applied to the problem of object recognition (Baddeley and van Lieshout, 1993; van Lieshout, 1995; Rue and Syversveen, 1998; Rue and Hurn, 1999). The objective of this type of analysis is typically to locate and characterize various objects of interest in space, incorporating prior knowledge of the spatial distribution of these objects. Therefore, prior models based upon discrete grid-based neighborhood structures tend to be less appropriate. The description of Bayesian object recognition presented here generally follows van Lieshout (1995).

In Bayesian object recognition, the observed data consist of an image, $y = (y_t; t \in T)$, where T (the image space) is an arbitrary finite set. The class of possible objects U , is an arbitrary set, termed *object space*. Objects can be seen as points u in U , and each determine a subset $R(u) \subset T$ of image space that is occupied by the object. Any particular configuration is a finite set of distinct objects $x = \{x_1, x_2, \dots, x_n\}$. The objective in object recognition is to estimate the (unobserved) true underlying pattern x given the observed image y .

This true configuration x is related to the observed image y through the likelihood function $l(y|x)$. As van Lieshout (1995) describes, the likelihood $l(y|x)$ represents both the deterministic influence of the true configuration x , and the stochastic effects within the remote sensing process that produces the image, y .

In a Bayesian analysis, the prior models will represent our prior beliefs regarding the spatial distribution of objects, and can be formulated to assign low probability to configurations that we do not expect to occur frequently, such as a large number of overlapping objects. The *maximum a posteriori* (MAP) estimator of x is the configuration \hat{x} that maximizes the function $l(y|x)p(x)$, and the prior essentially is a penalty assigned to this maximization. Therefore MAP estimation is also called *penalized maximum likelihood estimation*.

3.3 Bayesian object recognition for the analysis of three-dimensional LIDAR data in forested areas

While Bayesian object recognition has previously been applied to the analysis of two-dimensional images, this approach can also be applied to analyze structure within three-dimensional LIDAR data. In this case, the observed data, y_t , are not defined in terms of a raster image space, T . Instead, the *scan space* becomes a collection of vectors, T , determined by the LIDAR scanning process. Therefore, an individual pulse vector, t , represents the three-dimensional direction of each LIDAR pulse, from the aircraft to the terrain surface. The observed data, y_t , then represent range measurements along these vectors at which point the returning signal intensity exceeded a predetermined threshold (see Figure 1).

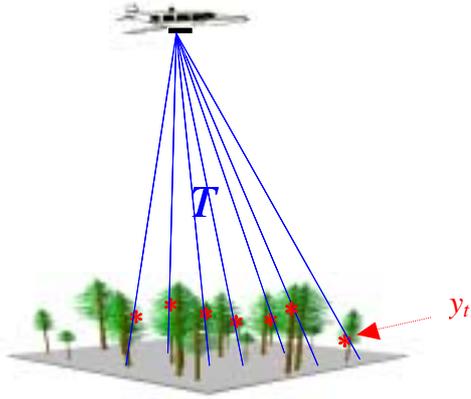


Figure 1. LIDAR sensing geometry (red stars represent LIDAR measurements; blue lines represent pulse vectors composing the scan space T ; y_t represents a single range measurement along pulse vector t).

The distribution of tree crowns over the entire scene is then modeled as an object configuration, x . If individual plants were actually solid objects (e.g. ellipsoids, spheres, etc.) in object space U , these LIDAR measurements, y_t , would represent the location where each vector t intercepted the surface of the object. In the terminology introduced above, these measurements would represent the *signal*, or the deterministic influence of the actual configuration of objects x on the series of LIDAR range measurements that are observed. A more realistic approach, however, would need to account for the fact that plants are not solid geometric objects, and LIDAR pulses actually penetrate a certain distance into the canopy through foliage gaps.

This would incorporate a stochastic element to the LIDAR measurements, y_t , due to the irregular spatial distribution of foliage elements (leaves, branches, etc.) in the pathway of a laser pulse as it intersects a tree crown. Therefore, again using the notation introduced above, the conditional distribution of individual LIDAR measurements, given the signal, is given by a family of densities $g(y_t | x)$. In the context of LIDAR measurement of tree crowns, these probability densities will be related to the laser attenuation function, which in the case of discrete LIDAR systems is directly related to the probability of reflection. If the values of individual LIDAR measurements along a pulse vector can be considered conditionally independent, given the true configuration of objects, x , the likelihood function, representing the joint probability of the data, is given by:

$$l(y | x) = \prod_{t \in T} g(y_t | x) \quad (2)$$

3.3.1 Modelling the distribution of foliage in complex forest scenes: Previous studies of laser transmission through the forest canopy have utilized three-dimensional grid models populated with generalized geometric forms that represent individual plants (Sun and Ranson, 2000).

In this study, a three-dimensional array, with 0.91-m cell size, was used to model the distribution of foliage density throughout

a forested area. The spatial distribution of foliage is a function of individual tree locations, sizes, crown forms, and an average leaf area density (*LAD*). Crown forms were represented as generalized ellipsoids following Sheng *et al.* (2001), where the space occupied by the foliage within an individual tree crown is determined by four parameters: crown width (*cw*), crown height (*ch*), crown curvature (*cc*), tree height (*ht*), and the 2-D coordinate of the crown top (X_{top}, Y_{top}) (see Figure 2).

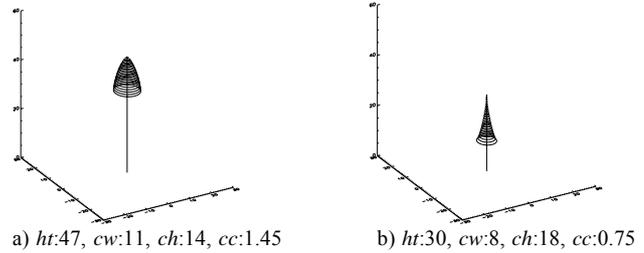


Figure 2. Generalized ellipsoid crown models.

The surface of a tree crown is then given by the following mathematical expression:

$$\frac{(Z + ch + Z_{top})^{cc}}{ch^{cc}} + \frac{\left((X - X_{top})^2 + (Y - Y_{top})^2 \right)^{cc/2}}{cr^{cc}} = 1 \quad (3)$$

where the elevation of the crown top, Z_{top} , is determined by adding the tree height to the elevation of the base as determined from the DTM. Values for *LAD* were obtained from previous research findings (Webb and Ungs, 1993).

3.3.2 Modelling laser-canopy foliage interaction: The analysis of data acquired from active remote sensing technologies requires an understanding of the interactions between the emitted radiation and the physical properties of the target.

In our model, where the laser footprint (0.4 m) is significantly less than the cell size (0.91 m), the probability that a direct light beam that enters a cell exits from the cell without being intercepted is calculated as a function of the leaf area density and the off-nadir angle (θ) of the laser pulse. The model then calculates the probability of penetration to the center of any cell by direct laser energy originating from outside the cell.

Specifically, the probability of a laser beam that enters a canopy cell, z_i , at a specific off-nadir angle, θ_i , reflects from this cell with foliage density LAD_i is given by the following function (Vanderbilt, 1990):

$$p(\text{reflection}) = \left[1 - \exp(-1/\cos \theta_i \cdot LAD_i \cdot G(\theta_i, z_i) \cdot dz_i) \right] \quad (4)$$

where LAD_i is the leaf area density (m^2/m^3) within cell z_i , $G(\theta, z_i)$ is the projection of the vegetation within cell z_i in the direction of θ , and dz_i is the depth of cell z_i . Often, a spherical leaf angle distribution can be assumed, in which case $G(\theta, z_i)$ is

0.5 (Goudriaan, 1988). Therefore, the probability of an *individual* LIDAR pulse reflecting from a specific cell z_i in the grid (and *not* reflecting from the cells $i-1, i-2, \dots, 0$ that it has already passed through) will be given by:

$$g(y_t | x) = \left(\prod_{k=0}^{i-1} [\exp(-1/\cos \theta_t \cdot LAD_k \cdot G(\theta_t, z_k) \cdot dz_k)] \right) \times [1 - \exp(-1/\cos \theta_t \cdot LAD_i \cdot G(\theta_t, z_i) \cdot dz_i)] \quad (5)$$

This function defines a probability density for LIDAR reflection y_t , anywhere along a three-dimensional pulse vector t . In addition, in our model it is assumed that the probability of a laser pulse reflecting if it penetrates to within 6 meters of the terrain (DTM) elevation is 1. In addition, foliage reflectance is assumed to be constant.

The likelihood function is then given by the following expression, which represents the joint probability of the LIDAR data:

$$l(y | x) = \prod_{t \in T} \left[\left(\prod_{k=0}^{i-1} [\exp(-1/\cos \theta \cdot LAD_k \cdot G(\theta, z_k) \cdot dz_k)] \right) \times [1 - \exp(-1/\cos \theta \cdot LAD_i \cdot G(\theta, z_i) \cdot dz_i)] \right] \quad (6)$$

The object configuration \hat{x} that maximizes this function will represent the maximum likelihood estimate (MLE) of the true object configuration x . However, given that MLE does not penalize large numbers of overlapping objects, it is likely that the MLE will be overly sensitive to the data and therefore will not represent a realistic forest object configuration. Through a Bayesian approach, prior knowledge relating to tree distributions and interactions can be incorporated into the model through the specification of the prior model, leading to more accurate estimates of the true object configuration.

3.3.3 Forest object processes: In the Bayesian object recognition approach, the underlying prior distribution, and the resulting posterior probability distribution of the true object configuration *given* the observed image data, usually takes the form of a *spatial point process*, a stochastic geometric model for an irregular, random pattern of points. These models allow for inference to be carried out relating to the *spatial position* of individual objects as well as the *attributes* of these individuals. These models also allow inter-object interaction, as well as possible global properties of a distribution of objects to be incorporated into the spatial model (Ripley, 1991).

If we define the environment $E(A)$ of a set A to be the set of neighbors of points in A , a point process is a *Markov* process if the conditional distribution on A given the rest of the process depends only on the process in $E(A)$. One of the most common Markov point process models is the pair-wise interaction model, which has the form:

$$p(x) = \alpha \beta^{n(x)} \prod_{i < j} g(x_i, x_j) \quad (7)$$

where $n(x)$ is the number of points in x and $g(x_i, x_j)$ is an interaction function (Ripley, 1981). This model therefore places a constant multiplicative penalty on each pair of interacting points. This type of model can be used to represent varying radii of inhibition surrounding biological phenomenon, and therefore can provide a useful model for forest object processes where trees exhibit pair-wise interactions.

A *marked point process* is a point process with a characteristic (*mark*) attached to each point in the process. Therefore a marked point process on R^d is a random sequence $x = [s_n; m_n]$ where the points s constitute an (unmarked) point process in R^d and the m are the marks corresponding to each location s .

In this model, s denotes the location of a tree, while m represents a vector of object attributes including height, crown width, crown height, and crown curvature.

Given a probability distribution for the marks, $\nu(m)$, the prior model, representing the *Markov object process*, takes the following form, where forest object interactions depend upon the individual tree attributes (marks):

$$p(x) = \alpha \beta^{n(x)} \prod_i \nu(m_i) \prod_{i < j} g(x_i, x_j) \quad (8)$$

In our model, two crowns were considered to be overlapping if the ratio of the distance between the center of the crowns and the sum of the crown radii was less than 0.75. The mark distribution was a multivariate normal distribution, with parameters determined from stand observations.

3.3.4 Simulation-based posterior inference for the MAP forest object configuration: In Bayesian analysis, all inferences are based upon the posterior distribution: $p(x | y) \propto l(y | x)p(x)$. The typical objective of Bayesian object recognition is to estimate the true configuration of objects x , given the observed data y . In particular, the *maximum a posteriori* (MAP) estimate, representing the mode of the posterior distribution, is of primary interest in the context of object recognition.

Within our model formulation, the posterior distribution is also a Markov object process. Due to the complex nature of the posterior distribution in this case, posterior inference was conducted via Markov chain Monte Carlo (MCMC) simulation. In MCMC, one constructs a Markov chain with an equilibrium distribution converging to the target distribution (the posterior distribution in the case of Bayesian inference). Ideally, this Markov chain should be constructed so as to efficiently move throughout the set of possible configurations, while maintaining the correct equilibrium distribution.

In this case, proposed moves for the Markov chain include 1) change of object parameters, 2) birth of an object, 3) death of an object, 4) merging of two objects, and 5) splitting a single object into two objects. Due to the change of dimension when proposing to add, delete, merge, or split objects in the configuration, the Metropolis-Hastings-Green reversible jump MCMC algorithm was used to maintain the correct equilibrium distribution (Green, 1995; Hurn and Syversveen, 1998). The change of parameters is achieved by drawing a sample from a multivariate normal distribution centered on the parameter vector for a selected object. The birth of an object is carried out by drawing a sample from the multivariate normal mark distribution for tree objects.

3.3.5 MAP estimation via simulated annealing: Finding the object configuration that maximizes the posterior probability (i.e. MAP estimate) is essentially a combinatorial optimization problem. Simulated annealing, an optimization technique with its origins in statistical mechanics, provides a means of arriving at the global optimum for a given system through a process of first “melting” the system at a high temperature, then gradually lowering the temperature until the system freezes at the optimal configuration (Kirkpatrick *et al.*, 1983). In the context of Bayesian object recognition, it has been shown that samples obtained, via MCMC, from the tempered posterior distribution $[p(x|y)]^{1/H}$ will converge to the MAP solution as $H \rightarrow 0$ (van Lieshout, 1994). In our algorithm, an annealing schedule of $H_{n+1} = \lambda \cdot H_n$ was used, where H is the temperature at iteration n , and λ is the cooling rate.

4. EXAMPLE

The algorithm was run on a 0.21 ha area within a lightly thinned unit of the Capitol State Forest (see Figure 3). The interaction parameter used in the prior distribution was set to e^{-750} , which places a moderately heavy penalty on severely overlapping tree crowns. The intensity parameter for the object process, β , was also set to e^{-100} . For this example, 133900 iterations of the MCMC algorithm were run with the cooling rate, λ , set to 0.999975, and initial temperature (H) of 20. The MCMC algorithm started with zero objects.



Figure 3. Location of example area (delineated in white) within area of Capitol State Forest, WA.

The locations, heights, and tree crown dimensions corresponding to the MAP estimate of the true object configuration within this area are shown in Figure 4, superimposed on the three-dimensional scatter plot of the LIDAR data and terrain model. A two-dimensional representation of the MAP estimate and LIDAR data is shown in Figure 5.

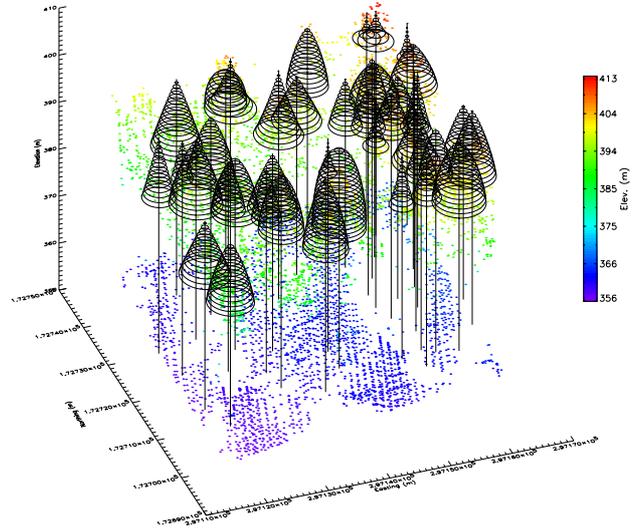


Figure 4. Three-dimensional perspective view of MAP estimate of tree locations and crown dimensions superimposed on LIDAR data. LIDAR pulse footprints are drawn to scale and are color-coded by elevation.

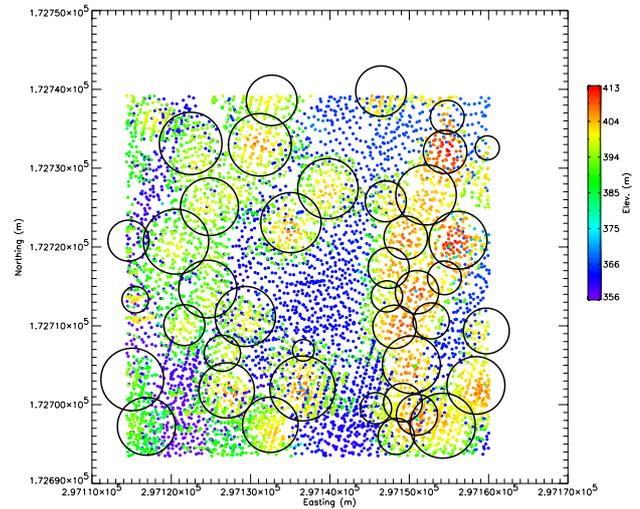


Figure 5. Planimetric view of MAP estimate of crown configuration (black circles) superimposed on LIDAR data. LIDAR pulse footprints are drawn to scale and are color-coded by elevation.

The estimates obtained from the object recognition algorithm were compared to photogrammetric measurements of crown locations made from large-scale (1:7000) aerial photography within a 0.21 ha area (see Figure 6).

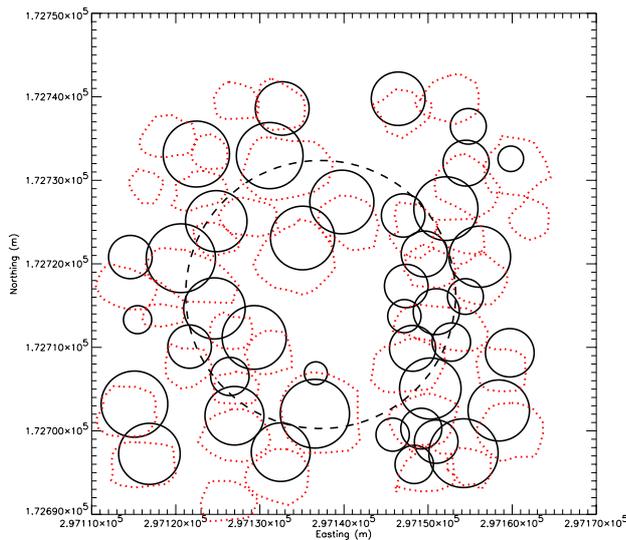


Figure 6. Planimetric view of MAP estimate of crown configuration (black circles), photogrammetric crown measurements (short dashes) and 0.081 ha circular inventory plot boundary (long dashes).

5. DISCUSSION

Results indicate that the algorithm is generally successful in identifying structures associated with individual tree crowns within this forest area. The MAP estimate of the crown configuration generated by the algorithm closely matches the spatial patterns evident in the LIDAR data (Figure 5). The algorithm appears to be very sensitive to the data, and in some areas added spurious small crowns to increase the likelihood of the data.

In general, the MAP estimate of crown locations corresponds to the photogrammetric crown measurements (see Figure 6). It should be noted that accurate recognition and delineation of overlapping tree crowns is difficult even in high-resolution aerial imagery. In this case, there is a systematic discrepancy of 1-4 meters in the north-south direction between algorithm-based crown locations and photo-based crown locations. This offset is probably due to the effect of crown layover and/or misregistration of the aerial photography.

Field data was available for a 0.081 ha circular inventory plot located within the study area (see Figure 6). Interestingly, the number of codominant (overstory) trees found within the plot in the field (14) matches the number found by the algorithm and measured in the photographs.

6. CONCLUSIONS

Bayesian object recognition provides a promising framework for the analysis of complex forest scenes using high-density, three-dimensional LIDAR data. It is clear that modelling assumptions will have a strong influence on the results; for example, it is apparent that crowns with an asymmetrical, irregular shape will be difficult to detect given the constraints of the generalized ellipsoidal crown model used here. The use of more complex models may improve recognition of irregularly shaped crowns.

Future research will focus on comparing algorithm results to field-based measurements and assessing the influence of automated measurement error on stand-level parameter estimates. In addition, Bayesian object recognition offers a flexible modelling approach that allows for fusing the information content from multiple sources of data. Such multiple data sources are becoming more available as vendors offer simultaneous acquisition of georeferenced imagery and LIDAR data. As the data enter the model only through the likelihood function in Bayesian object recognition, other types of remotely sensed data (including aerial photography and high resolution satellite imagery) can be easily incorporated into the model through adjustment of the likelihood function.

7. REFERENCES

- Andersen, H., S. Reutebuch, and G. Schreuder, 2001. Automated individual tree measurement through morphological analysis of a LIDAR-based canopy surface model. In: *Proceedings of the First International Precision Forestry Symposium*, Seattle, WA, USA.
- Baddeley, A., and M. van Lieshout, 1993. Stochastic geometry models in high-level vision. In: Mardia, K.V. and G.K. Kanjii, eds. *Advances in Applied Statistics I*. Carfax, Abingdon, Oxfordshire, pp. 231-256.
- Besag, J., 1993. Towards Bayesian image analysis. In: Mardia, K.V. and G.K. Kanjii, eds. *Advances in Applied Statistics I*. Carfax, Abingdon, Oxfordshire, pp. 107-119.
- Dralle, K., and M. Rudemo, 1997. Automatic estimation of individual tree positions from aerial photos. *Canadian Journal of Forest Research*, 27, pp. 1728-1736.
- Goudriaan, J., 1988. The bare bones of leaf-angle distribution in radiation models for canopy photosynthesis and energy exchange. *Agricultural and Forest Meteorology* 43, pp. 155-169.
- Green, P., 1995. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, 82(4), pp. 711-732.
- Gulden, R., 2000. Remote sensing at the dawn of a new millennium: A Washington DC perspective. In: *Proceedings of the Eighth Biennial Remote Sensing Applications Conference*, Albuquerque, NM, USA.
- Kirkpatrick, S., C.D. Gelatt, Jr., and M.P. Vecchi, 1983. Optimization by simulated annealing. *Science*, 220(4598), pp. 671-680.
- Larsen, M., 1998. Finding an optimal match window for spruce top detection based upon an optical tree model. In: *Proceedings of the International Forum on Automated Interpretation of High Spatial Resolution Digital Imagery for Forestry*, Victoria, BC, Canada.
- Lund, J. and M. Rudemo, 2000. Models for point processes observed with noise. *Biometrika*, 87(2), pp. 235-249.
- Magnussen, S., P. Eggermont, and V.N. LaRiccia, 1999. Recovering tree heights from airborne laser scanner data. *Forest Science*, 45(3), pp. 407-422.

Means, J., S. Acker, B. Fitt, M. Renslow, L. Emerson, and C. Hendrix, 2000. Predicting forest stand characteristics with airborne scanning LIDAR. *Photogrammetric Engineering and Remote Sensing*, 66, pp. 1367-1371.

Nelson, R., W. Krabill and J. Tonelli, 1988. Estimating forest biomass and volume using airborne laser data. *Remote Sensing of the Environment*, 24, pp. 247-267.

Pollock, R., 1996. *The automatic recognition of individual trees in aerial images of forests based upon a synthetic tree crown image model*. PhD thesis, Department of Computer Science, University of British Columbia, Vancouver, BC, Canada.

Ripley, B., 1981. *Spatial Statistics*. Wiley, New York.

Ripley, B., 1991. The use of spatial models as image priors. In: Antonio Possolo, ed., *Spatial Statistics and Imaging, vol. 20 of Lecture Notes – Monograph Series*, Institute of Mathematical Statistics, Hayward, CA, USA.

Rue, H., and M. Hurn, 1999. Bayesian object recognition. *Biometrika*, 86(3), pp. 649-660.

Rue, H., and A. Syversveen, 1998. Bayesian object recognition with Baddeley's delta loss. *Advances in Applied Probability*, 30, pp. 64-84.

Sheng, Y., P. Gong, and G. Biging, 2001. Model-based conifer-crown surface reconstruction from high-resolution aerial images. *Photogrammetric Engineering and Remote Sensing*, 67(8), pp. 957-965.

Sun, G., and K. Ranson, 2000. Modeling lidar returns from forest canopies. *IEEE Transactions on Geoscience and Remote Sensing* 38(6), pp. 2617-2626.

Vanderbilt, V.C., L.F. Silva, and M.E. Bauer, 1990. Canopy architecture measured with a laser. *Applied Optics*, 29(1), pp. 99-106.

van Lieshout, M., 1994. Stochastic annealing for nearest-neighbour point processes with application to object recognition. *Advances in Applied Probability*, 26, 281-300.

van Lieshout, M., 1995. *Stochastic geometry models in image analysis and spatial statistics*. CWI, Amsterdam, the Netherlands.

Webb, W. and M. Unga, 1993. Three-dimensional distribution of needle and stem surface area in a Douglas-fir. *Tree Physiology* 13, pp. 203-212.

Acknowledgements

The Precision Forestry Cooperative at the University of Washington College of Forest Resources and the USDA Forest Service Pacific Northwest Research Station provided financial support and data for this research. The authors would also like to thank Robert McGaughey (USDA Forest Service PNW Research Station) for assistance in carrying out this research.