

INVESTIGATING FULL-WAVEFORM LIDAR DATA FOR DETECTION AND RECOGNITION OF VERTICAL OBJECTS

Christopher E. Parrish, Physical Scientist
National Geodetic Survey
Remote Sensing Division
1315 East-West Hwy
Silver Spring, MD 20910
chris.parrish@noaa.gov

Frank L. Scarpace, Professor Emeritus
University of Wisconsin-Madison
2205 Engineering Hall
Madison, WI 53706
scarpace@wisc.edu

ABSTRACT

A recent innovation in commercially-available topographic lidar systems is the ability to record return waveforms at high sampling frequencies. These “full-waveform” systems provide up to two orders of magnitude more data than “discrete-return” systems. However, due to the relatively limited capabilities of current processing and analysis software, more data does not always translate into more or better information for object extraction applications. In this paper, we describe a new approach for exploiting full waveform data to improve detection and recognition of vertical objects, such as trees, poles, buildings, towers, and antennas. Each waveform is first deconvolved using an expectation-maximization (EM) algorithm to obtain a train of spikes in time, where each spike corresponds to an individual laser reflection. The output is then georeferenced to create extremely dense, detailed X, Y, Z, I point clouds, where I denotes intensity. A tunable parameter is used to control the number of spikes in the deconvolved waveform, and, hence, the point density of the output point cloud. Preliminary results indicate that the average number of points on vertical objects using this method is several times higher than using discrete-return lidar data. The next steps in this ongoing research will involve voxelizing the lidar point cloud to obtain a high-resolution volume of intensity values and computing a 3D wavelet representation. The final step will entail performing vertical object detection/recognition in the wavelet domain using a multiresolution template matching approach.

INTRODUCTION

Airborne lidar systems can be divided into two broad categories, based on the method of recording the return signal: 1) discrete-return systems, and 2) full-waveform systems (Lefsky et al., 2002). The first type records only a few (e.g., one to four) returns per transmitted pulse, while the second type digitizes the entire return waveform at very high sampling frequencies (~ 1 GHz). Until quite recently, full-waveform systems tended to utilize a large spot size (e.g., 10-70 m) and were generally one-of-a-kind experimental research systems, often built and operated by government agencies. Examples included NASA’s Laser Vegetation Imaging Sensor (LVIS) (Blair et al., 1999) and Scanning Lidar Imager of Canopies by Echo Recovery (SLICER) (Lefsky et al., 1999; Blair et al., 1994). In contrast, the discrete-return category generally comprised commercially-available production topographic mapping systems utilizing a small spot size (e.g., 0.15 to 1.0 m).

A very recent and exciting development in the field of topographic lidar is the release by the three largest commercial manufacturers (Optech, Inc., Leica Geosystems, LLC, and TopoSys, GmbH) of systems employing a small spot size and also providing waveform digitization capability. Figure 1 shows a plot of one individual waveform (amplitude vs. time for a single laser shot) captured with an Optech ALTM 3100 and waveform digitizer. The amplitude values shown on the ordinate are the output of the following process: 1) the received optical signal (reflected laser radiation) is collected by the receiver optics and converted to an electrical signal by an avalanche photodiode; 2) the output photocurrent is amplified and then digitized using an 8-bit quantizer and a sampling period of 1 ns (Optech, 2006). In Figure 1, the waveform is divided into three sections separated by vertical lines: the

transmitted pulse (which is saturated or “clipped” for this particular laser shot), the first 220 samples (corresponding to the first 220 ns of the return), and the last 220 samples of the return.




Figure 1. Example of a digitized lidar waveform.

These new systems appear ideally suited for applications such as vertical object detection/recognition, since they combine the advantages of a small spot size (the ability to discriminate between individual vertical objects) and waveform digitization (more information about vertical structure). One particular application area which stands to benefit from these systems is airport obstruction surveying, in which the primary goal is to accurately geolocate and attribute vertical objects, such as trees, towers, buildings, poles, and antennas, on and around airfields and in the runway approach paths. The data collected in these surveys are used by the Federal Aviation Administration (FAA) in designing runway approach procedures, which allow pilots to navigate safely into airports, even in reduced visibility weather condition.

While previous studies have shown that airborne lidar is a potentially viable technology for conducting airport obstruction surveys (Parrish et al., 2005; Parrish et al., 2004), a few problems remain unsolved. Most significantly, reliable, automated methods of extracting and attributing airport obstructions from lidar point clouds do not exist. Current software encounters difficulty with the following tasks: 1) distinguishing between lidar points from vertical objects of interest (e.g., a light pole) and those due to clutter (e.g., birds or other “flyers”) and noise; 2) determining which points in the lidar point cloud correspond to laser returns from a single object (e.g., a tree) vs. separate objects; and 3) classifying extracted objects (e.g., assigning attributes, such as “tree,” “building,” or “antenna”). Because of the difficulties in performing these tasks in an automated manner, lidar airport obstruction surveys currently require extensive manual labor and analysis, which negates the potential cost and time savings.

This research will investigate solutions to these problems by exploiting the additional information contained in the digitized waveforms and novel digital signal processing techniques. Our method involves first deconvolving the waveform data using an expectation-maximization (EM algorithm). The deconvolved waveforms are then georeferenced to produce extremely dense point clouds which greatly improve the ability to extract and attribute vertical objects. Preliminary results obtained using lidar waveform data collected over Madison, Wisconsin indicate that the average number of points on vertical objects obtained using this method is several times higher than using discrete-return lidar data. Without the additional points, some types of vertical objects might be undetectable and/or unrecognizable. Ongoing work involves voxelizing the lidar point clouds to obtain high-resolution 3D grids of intensity values and computing a 3D wavelet decomposition. Vertical object detection/recognition can then be performed in the wavelet domain using a multiresolution template matching approach.

DATA COLLECTION AND PREPROCESSING

The lidar waveform data were collected on June 23-24, 2006 using an Optech ALTM 3100 and waveform digitizer installed on a Dynamic Aviation Beechcraft King Air. Data were acquired for two study areas in Madison, Wisconsin: one centered on the University of Wisconsin (UW) campus and the other to the southwest of campus (Figure 2). These study areas were selected because they contain a variety of different types of vertical objects, including towers, poles, deciduous and coniferous trees, and buildings. The data for the campus area are being used for testing, while the data for the southwest area are used for training. In addition to the lidar waveform data, high resolution digital color (RGB) imagery was acquired for the two project areas on September 14-15, 2006 using an Applanix Digital Sensor System (DSS) medium-format, directly-georeferenced digital camera. The DSS imagery was used to generate 0.3 m resolution orthophotos for the two project areas.

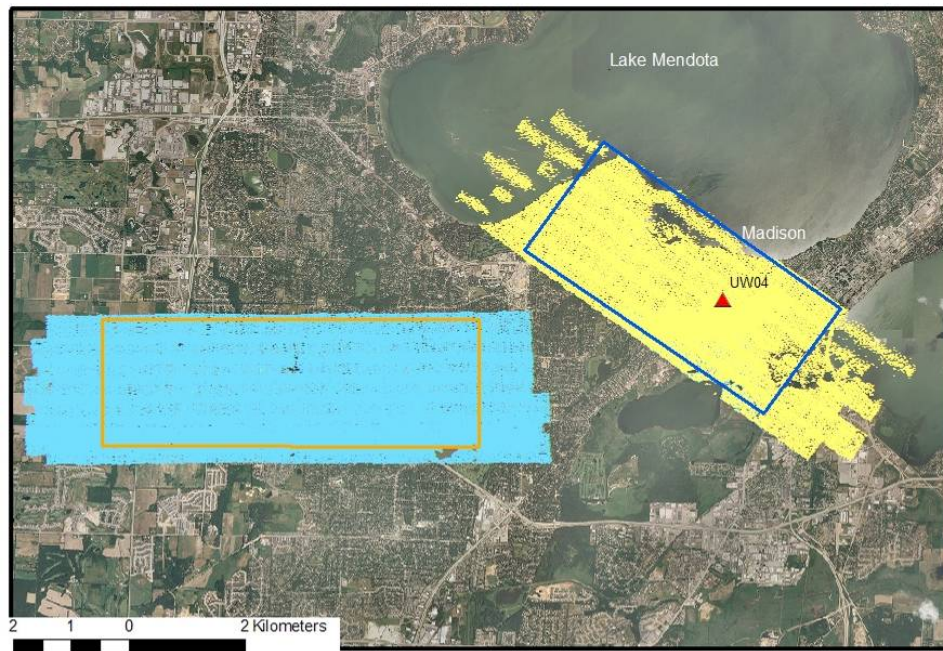


Figure 2. Study areas. Full-waveform lidar data coverage areas are shown in yellow and light blue. The location of the GPS basestation, UW04, is also shown.

The initial data processing was performed entirely with commercial software. First, POSPac v. 4.2 (Applanix Corp.) was used to post-process the raw airborne kinematic GPS observables and inertial data from the IMU to produce the position and orientation solution used in georeferencing the lidar data, referred to as a Smoothed Best Estimate of Trajectory (SBET). The “regular” (discrete-return) lidar data were processed using REALM v.3.5.3 (Optech, Inc.). REALM reads the SBET and raw laser range and scanner angle files and uses these data to georeference and output laser points. The processing parameters were set based on the unique requirements of vertical object detection (Parrish, 2003). First and last return data (including intensity data for each return) were output in delimited text format. All output data were referenced to NAD83 and projected in UTM (Zone 16N). Heights were referenced to the GRS80 ellipsoid.

To enable visual analysis of the waveform data and gain insight into the nature of the returns from different types of targets, software was written to geographically link the waveform data to a digital orthophoto. This software locates the closest first-return point in the discrete-return lidar data to a user-selected location on the orthophoto (or vice versa). The corresponding waveform is then retrieved and plotted alongside the orthophoto and terrestrial hand-held camera photo. Figure 3 shows examples of the output, providing some indication of how the shape of the return waveforms is affected by differing vertical structure.

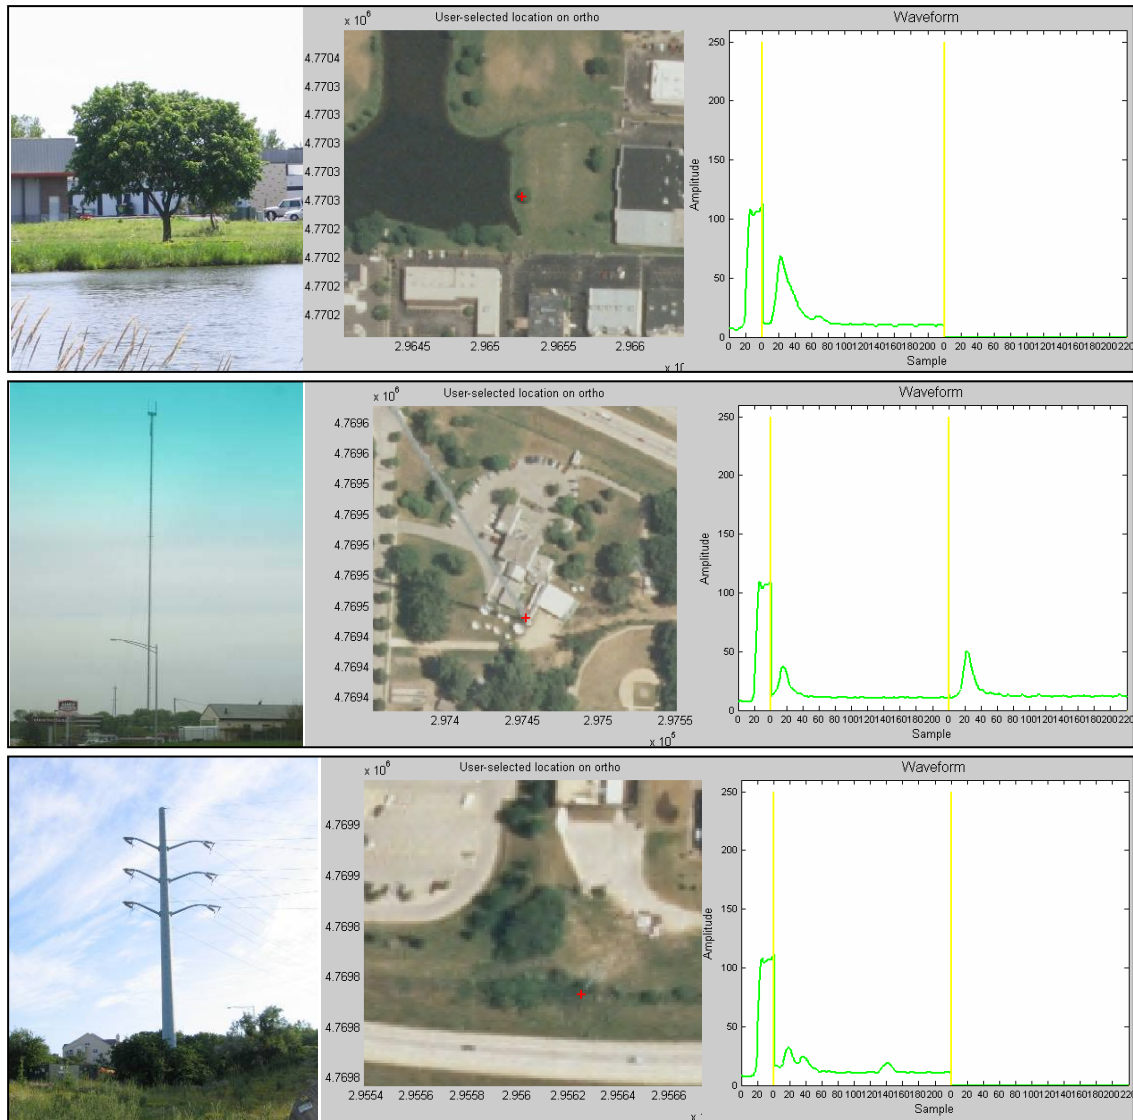


Figure 3. Examples of output of software for geographically linking waveform data with a digital orthophoto. The vertical objects shown here are a tree (top), tower (middle), and transmission tower (bottom). The left-most image in each row is a conventional hand-held photo. The middle image in each row shows the corresponding location (small red plus symbol) on the orthophoto. The right-most image in each row is the corresponding lidar waveform.

Additionally, field work was performed to obtain reference data for a wide variety of vertical objects located within the two project areas. In total, 30 vertical objects were surveyed in the campus project area and 35 in the southwest project area. Most (67%) of the objects were surveyed using post-processed differential GPS. Coordinates for the remainder of the vertical objects were obtained by heads-up digitizing from the 0.3 m resolution orthophotos. The field crew also obtained terrestrial photos of each vertical object using a hand-held digital camera.

METHODS

The first step in our approach to vertical object detection/recognition entails deconvolving and then georeferencing the waveform data. Considered alone, the georeferencing process is fairly straightforward. It is achieved through implementation of the laser geolocation equation (Parrish, 2003; Lindenberger, 1989; Filin, 2001), which relates the (X,Y,Z) position of a laser target in the mapping frame to the laser range, instantaneous scanner

angle, and the sensor position and orientation at the time of the laser shot. When georeferencing waveform data, the issue, however, is which point(s) on the received waveform to use in the time-of-flight calculation used to obtain the laser range. One possibility is simply to use every sample in the digitized waveform that is above a specified threshold in a separate time-of-flight calculation and georeference each corresponding point. This will lead to a point cloud which is very dense, but also blurry and noisy. Therefore, our approach is to deconvolve the received waveform and to georeference the deconvolution output.

Figure 4 further illustrates the motivation for deconvolving the waveforms. The leftmost waveform in the figure is idealized, in that it assumes the transmitted pulse is an impulse. Furthermore, it ignores the complex interactions between the laser pulse and target(s) and any pulse shaping that may occur, so that the return (i.e., the ideal impulse response) is a train of spikes in time. The beauty of this idealized waveform is that each spike in the return corresponds directly to an individual laser reflection (for example, to each branch or leaf in a tree encountered on the downward path of the laser pulse) and the amplitude of the spike corresponds to the strength of the reflection. Thus, the laser geolocation equation could be applied directly to this idealized waveform to create a detailed, accurate, noise-free point cloud. In contrast, the middle waveform in Figure 4 is a more realistic representation of an actual waveform. The actual transmitted pulse is approximately Gaussian with a pulse width of ~10 ns, and the actual return waveform can be considered a blurry, noisy version of the ideal waveform. Simply stated, the goal of the deconvolution is to attempt to “restore” the ideal waveform (i.e., the spike train where each spike corresponds to an individual laser reflection) from the actual received waveform prior to the georeferencing.

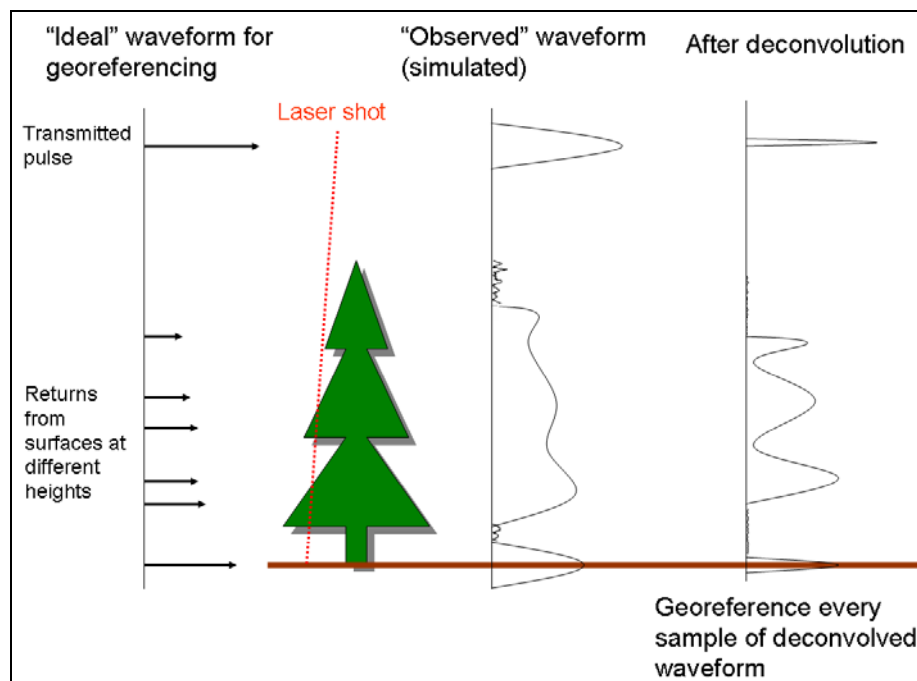


Figure 4. Illustration of the motivation for deconvolving the received waveforms prior to the georeferencing step.

Wiener filtering is one standard, well-known signal restoration approach that can be applied to our problem of recovering the ideal waveform of Figure 4 from the actual received waveform. The Wiener filter produces a minimum mean square error (MMSE) estimate of the desired signal. Unfortunately, the Wiener filter suffers some serious drawbacks in our application. For example, the use of the Wiener filter leads to a tradeoff between deblurring and denoising. Furthermore, the Wiener filter produces an undesirable artifact: ringing (artificial oscillations) around sharp edges.

The primary source of the limitations of the Wiener filter in this application is that the signal is not well characterized in the frequency domain. In modeling the signal as a Gaussian random vector with a flat power spectrum, the Wiener filter encounters difficulty in distinguishing between the signal and white Gaussian noise (WGN). A better method of deconvolving the lidar waveform data is through the expectation-maximization (EM) algorithm developed by Figueiredo and Nowak (2003), which numerically computes the maximum *a posteriori* (MAP) estimate of the signal by alternating between Fourier-based estimation (the E-step) and denoising in the

wavelet domain (or, in our implementation, directly in the time domain) (the M-step). Whereas the Wiener filter forces the signal to fit the Gaussian model and does not utilize any additional information we may have about the signal, the EM algorithm allows us to model the signal as a random set of spikes in time that we know (or believe) to be sparse. This use of additional information (specifically, the information that the ideal signal is sparse) enables better results than can be achieved with the Wiener filter. Another nice feature of the EM algorithm is that it includes a tunable parameter, τ , which controls the number of spikes in the output. Decreasing τ increases the number of spikes in the output, but could increase the number of false alarms. Increasing τ has the opposite effect. For applications, such as airport obstruction surveying, in which the consequences of a miss are far more severe than the consequences of a false alarm, τ can be decreased to increase the number of output points on vertical objects.

Figure 5 shows examples of the deconvolution output from both the Wiener filter and the EM algorithm. As illustrated in this figure, the EM algorithm achieves much better results than the Wiener filter; the deconvolved signals are much sharper, the noise has been suppressed to a much greater extent, and the artifacts produced by the Wiener filter (e.g., ringing) do not occur.




Figure 5. Comparison of Wiener filter vs. EM deconvolution algorithm for two different lidar waveforms. In each plot, the top (green) curve is the original waveform, the middle (blue) curve is the Wiener filtered waveform, and the bottom (red) curve is the output of the EM deconvolution algorithm.

Figure 6 shows a comparison of output point clouds generated using different methods. Subfigure (A) is a photomosaic of the tree line on the west side of Whitney Way in Madison, Wisconsin generated from photos taken with an ordinary hand-held digital camera. Subfigure (B) shows a vertical slice through the point cloud generated from the discrete-return lidar data output from Optech's commercial software, REALM. Subfigure (C) shows a vertical slice through the point cloud generated by georeferencing every sample of the return waveform (no deconvolution). Here, the point cloud is much denser (e.g., the tree canopy is completely filled in), but it is also blurry and noisy, as we would expect. Subfigure (D) shows the results of applying a Wiener filter to the waveforms and then georeferencing the output. Here, we are beginning to get the best of both worlds in that the resulting point cloud contains a lot of information, yet is less blurry and noisy than the point cloud created without the deconvolution step. However, the limitations of the Wiener filter are evident. For example, due to the tradeoff between deblurring and denoising with the Wiener filter, some blur and noise are seen in the output. Another artifact seen in the output is the "false ground" approximately 2 m above and below the actual ground, which is the result of the ringing effect of the Wiener filter. Subfigure (E) shows a vertical slice through the point cloud generated by applying the Figueiredo-Nowak EM deconvolution algorithm to the waveform data, followed by georeferencing. This is clearly the best of the four point clouds for vertical object detection applications. It contains more information than the point cloud generated from the discrete-return lidar data, yet little blur or noise.

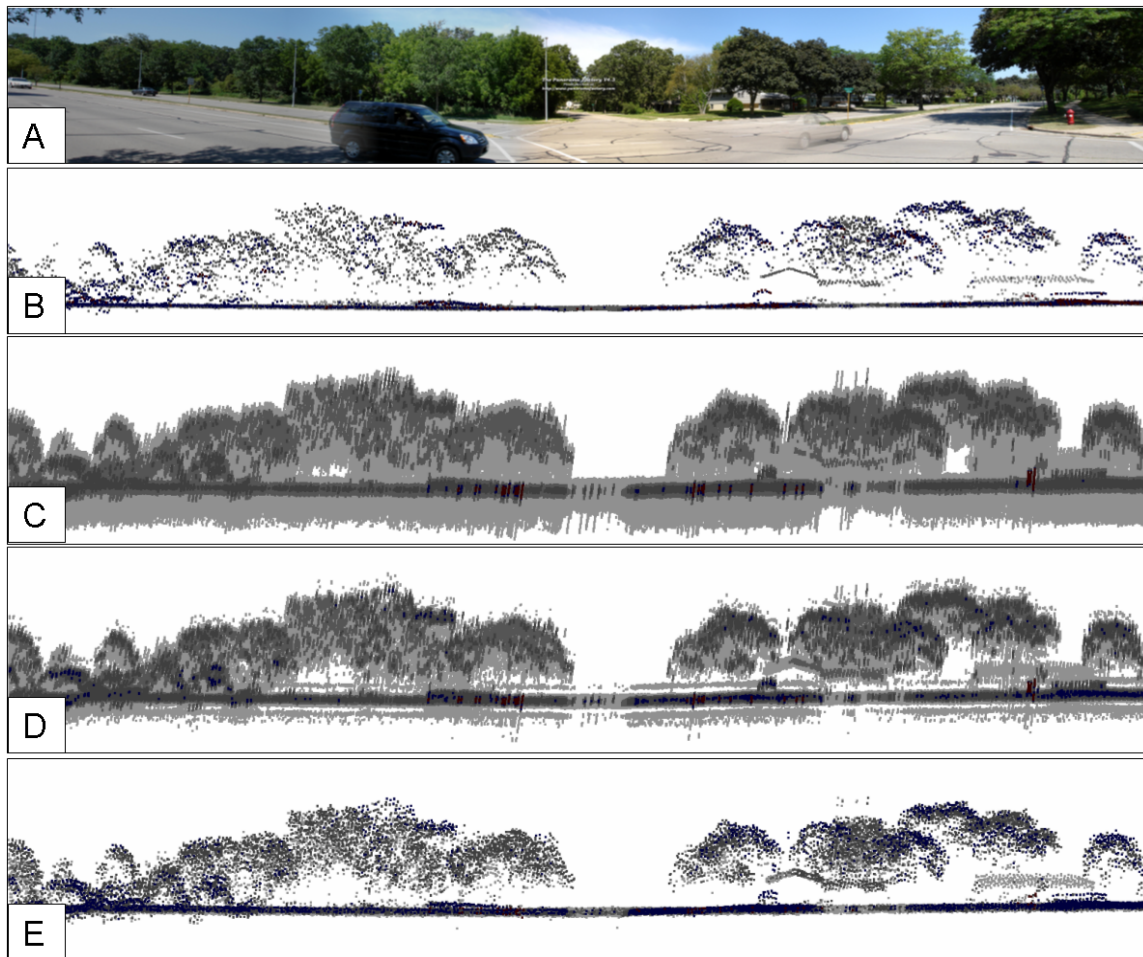


Figure 6. Comparison of point clouds generated using different strategies. A: Photos of the tree line on the west side of Whitney Way in Madison, WI. B: Point cloud generated from discrete-return lidar data using Optech's commercial software, REALM. C: Point cloud generated by georeferencing every sample of the return waveforms (no deconvolution performed). D: Point cloud generated by first Wiener filtering and then georeferencing the waveform data. E: Point cloud generated by applying the EM deconvolution algorithm to the waveform data and then georeferencing the output.

Figure 7 illustrates the remaining steps in our proposed approach to vertical object detection and recognition. A

volume representation of the lidar data (i.e., a 3D grid of lidar intensity values) is computed and a 3D wavelet transform is applied. Finally, vertical object detection/recognition is performed in the wavelet domain using a multiresolution template matching approach. Software to complete the final steps in this proposed method is currently in the final stages of development and testing.

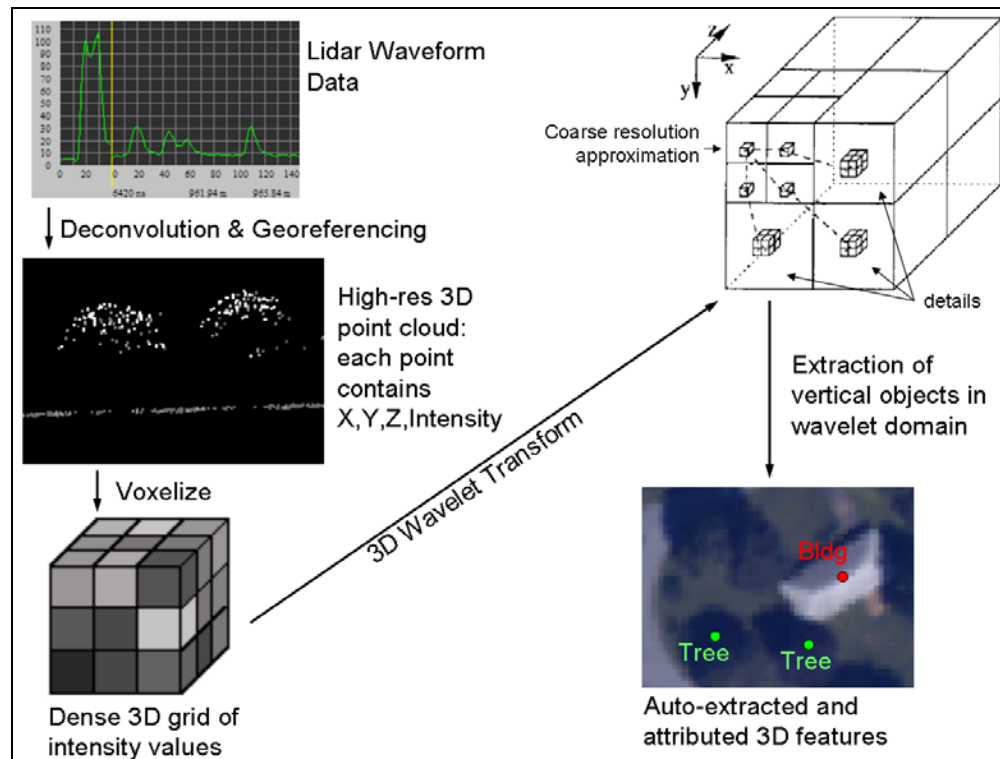


Figure 7. Primary steps in our methodology for extracting and attributing vertical objects in full-waveform lidar data.

PRELIMINARY RESULTS

The first step in our results analysis involves comparing the point clouds generated from the “regular” (discrete return) lidar data using the manufacturer’s commercial software with point clouds generated using the waveform deconvolution and georeferencing algorithms described in this paper. This comparison is performed through a combination of automated and manual methods. A human operator first visually compares subsets of the two point clouds centered on the object of interest using TerraScan v.5.005 (Terrasolid Ltd.). An example is shown in Figure 8. Next, the Measure Point Density tool in TerraScan is used to count the number of points on the vertical object in each point cloud. To date, this analysis has only been performed for 8 of the 35 objects in the towers box, but the increase in the average number of points on each vertical object is remarkable: approximately 300%. These results are even more dramatic considering that the signal-to-noise ratio (SNR) is actually lower for the full-waveform data than for the discrete return data, due to noise added by the digitizer (Optech, 2006). For a small-diameter, low reflectance object, such as a pole or antenna, this could mean an increase from 2-3 points on the object, which is probably too few for successful detection/recognition, to 6-9 points, which may be enough for successful detection/recognition. Furthermore, the visual analysis indicates that the point clouds generated through our methods contain only slightly more noise and clutter than those generated from the discrete-return data.




Figure 8. Left: photo of a deciduous tree. Middle: points on this object in the point cloud generated from the discrete-return lidar data. Right: points on this object in the point cloud generated by deconvolving the lidar waveform data using the EM algorithm and then georeferencing output.

The remaining steps in our results assessment will involve comparing the results of vertical object extraction and attribution obtained through the procedure outlined in Figure 7 with those obtained using a commercial software package for vertical object detection, which uses a lidar-derived DSM as the input. At present, this analysis has not been completed, as the wavelet-based detection/recognition algorithms are still in the final stages of development.

DISCUSSION AND CONCLUSIONS

In this paper, we have described a new approach to extraction and attribution of vertical objects, such as airport obstructions, from lidar data. A critical step in our approach is to obtain dense, detailed, accurate point clouds, which facilitate obstruction detection and recognition. To achieve this goal, we first deconvolve the waveforms using an EM algorithm and then georeference the output. Preliminary results show that our methodology significantly improves the quality of the lidar point clouds for vertical object detection. For the objects tested to date, the average number of points on each object increased by approximately 300% using our method, without a significant increase in noise or clutter. This increase is particularly important for small-diameter, low-reflectance vertical objects, such as poles, which might be undetectable without these additional points.

Ongoing work is being performed to implement the final steps in our approach to vertical object detection and recognition from lidar waveform data. These steps include voxelizing the lidar point cloud to obtain a high-resolution volume of intensity values and then computing a 3D wavelet decomposition. Vertical object detection/recognition is then performed in the wavelet domain using a multiresolution template matching approach.

ACKNOWLEDGEMENTS

The work presented in this paper was conducted at the University of Wisconsin-Madison as part of the lead author's dissertation research (expected completion: August, 2007). The following UW faculty members are recognized for their significant contributions to this work: Dr. Robert Nowak, whose EM deconvolution algorithm was adapted for lidar waveforms, Dr. Amos Ron, who graciously lent his wavelet expertise to this research, and Dr. Alan Vonderhoe, who provided general guidance. Additionally, we wish to thank the following NOAA, National Geodetic Survey employees for their support and assistance: Remote Sensing Division Chief, Michael Aslaksen, Jason Woolard, Jon Sellars, and Stephen White.

REFERENCES

- Blair, J.B., D.B. Coyle, J.L. Bufton, and D.J. Harding (1994). Optimization of an airborne laser altimeter for remote sensing of vegetation and tree canopies. *Proceedings of the International Geoscience and Remote Sensing Symposium-IGARSS 1994*. ESA Scientific and Technical Publications, Noordwijk, pp. 939-941.
- Blair, J.B., D.L. Rabine, and M.A. Hofton (1999). The Laser Vegetation Imaging Sensor: a medium-altitude, digitisation-only, airborne laser altimeter for mapping vegetation and topography. *ISPRS Journal of Photogrammetry & Remote Sensing*, Vol. 54, pp. 115-122.
- Ching, W.-S. (1995). Normalized cross-correlation: the contrast-dependent problem and its solution. *Journal of Electronic Imaging*, 4(3), pp. 278-282.
- Figueiredo, M.A.T., and R.D. Nowak (2003). An EM Algorithm for Wavelet-Based Image Restoration. *IEEE Transactions on Image Processing*, Vol. 12, No. 8.
- Filin, S. (2001). *Calibration of Airborne and Spaceborne Laser Altimeters Using Natural Surfaces*. Ph.D. dissertation, The Ohio State University, Columbus, Ohio, 128 p.
- Kandadai, S. and C.D. Creusere (2003). An experimental study of object detection in the wavelet domain. *Conference Record of the Thirty-Seventh Asilomar Conference on Signals, Systems and Computers*, November 9-12, Pacific Grove, CA, Vol. 2, pp.1620 – 1623.
- Lefsky, M.A., W.B. Cohen, S.A. Acker, G.G. Parker, T.A. Spies, and D. Harding (1999). Lidar Remote Sensing of the Canopy Structure and Biophysical Properties of Douglas-Fir Western Hemlock Forests. *Remote Sensing of Environment*, Vol. 70, pp. 339-361.
- Lefsky, M.A., W.B. Cohen, G.G. Parker and D.J. Harding (2002). Lidar remote sensing for ecosystem studies. *Bioscience* Vol. 52, No. 1 19-30.
- Lewis, J.P. (1995). Fast template matching. *Vision Interface*, pp. 120-123.
- Lindenberger, J. (1989). Test results of laser profiling for topographic terrain survey. *Proceedings 42nd Photogrammetric Week*, Stuttgart, Germany.
- Optech (2006). *Airborne Laser Terrain Mapper (ALTM) Waveform Digitizer* (Doc. No. 0028443/Rev B). Optech Incorporated, Toronto, Ontario, Canada.
- Parrish, C.E. (2003). *Analysis of Airborne Laser-Scanning System Configurations for Detecting Airport Obstructions*. Masters Thesis, University of Florida, Gainesville, Florida, 108 p.
- Parrish, C.E., J. Woolard, B. Kearse, and N. Case (2004). Airborne LIDAR Technology for Airspace Obstruction Mapping. *Earth Observation Magazine (EOM)*, Vol. 13, No. 4.
- Parrish, C.E., G.H. Tuell, W.E. Carter, and R.L. Shrestha (2005). Configuring an airborne laser scanner for detecting airport obstructions. *Photogrammetric Engineering & Remote Sensing*, Vol. 71, No. 1.
- Stoker, J. (2004). Voxels as a Representation of Multiple-Return Lidar Data. *Proceedings, American Society of Photogrammetry and Remote Sensing (ASPRS) Annual Conference*, May 23-28, Denver Colorado.
- Thornton, J., P. Hennings, J. Kovacevic, and B.V.K. Vijaya Kumar (2005). Wavelet packet correlation methods in biometrics. *Proceedings IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP '05)*, March 18-23, Philadelphia, PA.