Empirical Comparison of Full-Waveform Lidar Algorithms: Range Extraction and Discrimination Performance

Christopher E. Parrish, Inseong Jeong, Robert D. Nowak, and R. Brent Smith

Abstract

As third-party lidar software manufacturers are increasingly adding support for full-waveform data, a common question is which algorithm(s) to implement. To this end, a new approach is needed to compare and contrast various lidar waveform processing strategies from a practical, operational perspective. Quality and type of information output, processing speed, suitability for particular applications, robustness against poor parameter selection, and more subjective measures related to user experience are of interest. This paper describes a new empirical method of comparing range extraction and discrimination performance of different algorithms, based on a ranging-lab setup with multiple, adjustable screen targets, with precisely-measured separations. We present the results of comparing three different algorithms described in the scientific literature. The results show distinct differences and also indicate that there is no "one-size-fits-all" approach: the choice of a specific algorithm and adjustable parameter settings are highly application-dependent.

Introduction

Prior to 2004, the commercial, topographic lidar market was dedicated to discrete-return systems, which record only a few ($n \leq 5$) individual ranges per transmitted pulse. In principle, each recorded return in a discrete-return system corresponds to an individual laser reflection (i.e., an echo from one particular reflecting surface, which could be ground, or some elevated feature, such as a tree, pole, building, etc.). Ranging in discrete-return systems is done through hardware subsystems, typically comprising a constant fraction discriminator and time-interval meter (Baltsavias, 1999; Parrish *et al.*, 2005). By recording just a few individual ranges, discrete-return systems obviate the need for heavy data storage and processing requirements. An alternative to the discrete-return concept is full-waveform (FW) lidar. In FW systems, each backscattered laser pulse received by the system is digitized at a high sampling rate (e.g., 500 MHz to 1.5 GHz). This process generates digitized waveforms (amplitude versus time) that are stored for subsequent processing and analysis. In contrast to discrete-return lidar, FW systems enable software-based ranging through processing of the recorded waveforms.

The FW concept is not new; in fact, it has been utilized since the 1970s in bathymetric lidar systems, such as the National Aeronautics and Space Administration (NASA) Airborne Oceanographic Lidar (AOL) (Hoge *et al.*, 1980), and in large-footprint, experimental, vegetation-mapping systems, including NASA's Scanning Lidar Imagery of Canopies by Echo Recovery (SLICER) (Means *et al.*, 1999), and Laser Vegetation Imaging Sensor (LVIS) (Blair *et al.*, 1999). However, commercially-available, small-footprint topographic lidar systems have only been available since around 2004, starting with the Riegl LMS-Q560 scanner and turnkey systems built upon it, and followed shortly by systems of several other lidar manufacturers (e.g., Hug *et al.*, 2004; Lemmens, 2007). Mallet and Bretar (2009) provide a review of full-waveform lidar and description of current systems.

Since the introduction of small-footprint FW lidar systems into the commercial market, many advantages have been described in the published literature. First, FW systems greatly improve the target resolution (also referred to as "vertical discrimination distance"), defined as the minimum separation of targets in the range direction, such that each can be individually resolved. Discrete-return systems tend to have poor target resolution, typically due to a sizable (up to ~3.5 m) dead zone after each recorded return, resulting from the inherent limitations of the receiver electronics (Nayegandhi *et al.*, 2006; Wagner *et al.*, 2008). Using high-sampling-rate digitizers and software-based ranging, FW systems enable great improvement in target resolution. Furthermore, it may be possible to perform ranging more accurately in software, and certainly more alternatives are available in the ranging technique (Wagner *et al.*, 2004; Harding, 2009). As a related benefit, FW systems enable generation of much denser, detail-rich point clouds, which provide enhanced information about vertical structure

Christopher E. Parrish is with the NOAA/NGS Remote Sensing Division, University of New Hampshire, Center for Coastal and Ocean Mapping, Jere A. Chase Ocean Engineering Lab, 24 Colovos Road, Durham, NH 03824 (chris.parrish@noaa.gov).

Inseong Jeong is with NOAA/NGS Remote Sensing Research (DST), 1315 East West Highway, Silver Spring, MD 20910.

Robert D. Nowak is with the University of Wisconsin-Madison, Department of Electrical and Computer Engineering, 3627 Engineering Hall, 1415 Engineering Drive, Madison, WI 53706.

R. Brent Smith is with Optech Incorporated, 300 Interchange Way, Vaughan, Ontario, Canada, L4K 5Z8.

Photogrammetric Engineering & Remote Sensing
Vol. 77, No, 8, August 2011, pp. 825–838.

0099-1112/11/7708–0825/$3.00/0

(Persson *et al.*, 2005; Reitberger *et al.*, 2006; Parrish and Nowak, 2009). In addition to (or, in some cases, instead of) range information, radiometric and other types of information related to surface characteristics can also be extracted from waveform data.

The transition from discrete-return to FW in commercial, topographic lidar can also be described as a fundamental change in philosophy, in which fixed, hardware-based subsystems are replaced with customizable software alternatives. Sophisticated digital signal processing overcomes the need for sophisticated hardware, and limitless customization is enabled. All ranging strategies (whether implemented in hardware or software) involve engineering tradeoffs. The major difference is that a hardware-based solution remains fixed, once implemented by the lidar system manufacturer; in software, the user can select and/or tune the algorithm to best suit their application, be it floodplain mapping, forestry, building modeling, coastal marsh vegetation mapping, or airport obstruction surveying, for example.

However, there is also an intrinsic tradeoff involved in the FW approach: in enabling more to be done in software, FW systems necessitate additional processing and algorithms to perform that processing. To date, several waveform processing strategies (and variations of the general approaches) have been described in the scientific literature (see, e.g., Hofton *et al.*, 2000; Persson *et al.*, 2005; Jutzi and Stilla, 2006; Chauve *et al.*, 2007; Mallet and Bretar, 2009, and references therein). Unfortunately, third-party commercial software for processing and analyzing full-waveform data has lagged a bit behind both the new hardware development and the scientific community's algorithm development. Thus, as many third-party software developers are currently at the stage of adding full-waveform processing capabilities, a common dilemma is which waveform processing algorithm(s) to implement. To make this determination, the software developers would like to know how the different algorithms compare, not just from a theoretical perspective, but in terms of real-world experience in implementing and using the different algorithms operationally.

The objective of this paper is to help address these needs by: (a) presenting a new empirical approach to comparing lidar waveform processing algorithms in terms of range extraction and discrimination performance, and (b) using this experimental technique to compare three lidar waveform processing strategies described in the literature. Specifically, we present results of an experiment conducted using an Optech, Inc. topographic lidar system in a controlled ranging lab environment with multiple, adjustable targets. These results enable a robust comparison of the three algorithms, including advantages and disadvantages of each, and conclusions that may assist software providers in implementing FW support.

Methods

The three algorithms selected to test in this study are listed below. There are certainly different (or additional) algorithms that we could have selected, including, for example, the B-splines approach described in Roncat *et al.* (2010), the Average Square Difference Function (ASDF) method described in Wagner *et al.* (2007), the decomposition approach of Chauve *et al.* (2007) using refined peak detection and Lognormal or generalized Gaussian functions, or well-known signal processing techniques that could be adapted for lidar waveform processing, such as matched filtering (Turin, 1960). Ultimately, our selection was guided by three main criteria. First, we sought algorithms that were fundamentally different from one another, or broadly representative of different approaches to waveform processing. Second, we considered only algorithms for which source code was available or which were described in sufficient detail in the published literature that we could faithfully and efficiently reconstruct the authors' steps in implementing them. Third, in order to keep the scope of the study manageable, we limited the total number of algorithms tested to three. The fact that the experimental design for comparing the different algorithms was new and needed to be tested and refined as a critical component of the study was a primary factor in limiting the number of algorithms tested.

To enable fair comparisons of target resolution, the parameters of each algorithm were adjusted to maximize target discrimination while constraining the mean false alarm rate to $<1\%$ (see the Experiment Section for further details of the parameter tuning). As an aside, another form of comparison would be to produce receiver operating characteristic (ROC) curves for the different algorithms, but for our target resolution comparisons and application areas of interest, such as, airport obstruction surveying, we found it most insightful to constrain the false alarm rate to a specified value and examine how target discrimination falls off with decreasing target separation for the three algorithms.

Gaussian Decomposition

Gaussian decomposition is currently the most widely-applied topographic lidar waveform processing strategy, based on numbers of known implementations and published papers (e.g., Hofton *et al.*, 2000; Persson *et al.*, 2005; Wagner *et al.*, 2006; Reitberger *et al.*, 2006). In Gaussian decomposition, each lidar waveform is modeled as a linear combination of Gaussian components (i.e., an N-component Gaussian mixture model):

$$\hat{y}(t) = \sum_{i=1}^{N} \alpha_i \exp\left\{ -\frac{1}{2\sigma_i^2}(t-\mu_i)^2 \right\} \quad (1)$$

where α_i is the amplitude of the i^{th} Gaussian component, σ_i is its width, and μ_i is the location of its mean on the time axis. Thus, the waveform processing entails fitting parameterized Gaussians to the received waveform. Different implementations vary in the optimization algorithm used in fitting the Gaussian components and in how N, the number of components, is selected. The underlying assumption is that each Gaussian component is the result of an individual laser reflection, and the estimated parameters $\{\alpha_i, \sigma_i, \mu_i\}$ are related to the backscattered signal strength, surface characteristics, and range. Wagner *et al.* (2006) provide the theoretical framework for this approach.

Based on the selection criteria listed above, the Gaussian decomposition algorithm described in Hofton *et al.* (2000) was selected for implementation and testing in this study. After applying a smoothing filter to reduce noise, parameters of initial Gaussians (i.e., positions and half widths) are estimated by a nonnegative least squares method. Next, Gaussians are flagged as "important" if: (a) their half width is no less than the half width of T_0 pulse, and (b) their amplitude is larger than three times the noise standard deviation. The other Gaussians are ranked by closeness to the important Gaussians. Parameters of the important Gaussians are refined using the Levenberg-Marquardt (LM) technique. If the residual between the original and fitted waveform does not meet a predefined level, one more Gaussian is added, and the LM method is reapplied. Based on extensive experimentation, there were three slight modifications that we found to be necessary or

advantageous for the type of data collected in this study (see Experiment Section):

- The half width of the smoothing Gaussian filter (see parameter descriptions in Experiment Section) was chosen slightly differently than in Hofton *et al.* (2000). We first tried the half width of the T_0 (transmit) pulse, but determined that this setting smoothed the input too much; through experimentation, we found that setting the half width of the smoothing filter to 2.5 nanoseconds provided the best results with our data.
- A small value of w_I (minimum full width half maximum of "important" Gaussian components) was picked to prioritize the amplitude criterion over the width criterion in selecting important Gaussian components.
- We did not include the epsilon parameter for adding additional Gaussian components, as we found that this did not improve the results with the types of waveforms recorded in our ranging lab experiment.

EM Deconvolution

In contrast to Gaussian decomposition, the deconvolution approach to lidar waveform processing is not a fitting strategy, but rather, an attempt to recover the unknown signal from degraded (blurry, noisy) observations by solving the inverse problem. The approach is based on the following observation model:

$$y[n] = h[n]*x[n] + \eta[n] \tag{2}$$

where $y[n]$ is the observed (discrete-time) waveform, $h[n]$ models the blur (which can be estimated using different approaches; here we use the system impulse response function), $\eta[n]$ is additive noise (usually assumed to be Gaussian, white noise), $x[n]$ is the "true" (unknown, undegraded) signal, and $*$ denotes convolution. The output of the deconvolution algorithm is an estimate of the undegraded signal, $\hat{x}[n]$. Although not as widely implemented as Gaussian decomposition for processing of topographic lidar waveforms, deconvolution approaches have been demonstrated on topographic lidar waveforms (e.g., Walter, 2005; Nordin, 2006; Parrish, 2007; Parrish and Nowak, 2009; Wang *et al.*, 2009; Roncat *et al.*, 2010) and on simulated bathymetric lidar waveforms (e.g., Johnstone *et al.*, 2004). A key feature is convex optimization, and, hence, the ability to avoid suboptimal solutions (Parrish and Nowak, 2009).

The deconvolution algorithm tested in this research was implemented as described in Parrish (2007) and Parrish and Nowak (2009). The derivation of the underlying EM deconvolution algorithm can be found in Figueiredo and Nowak (2003), in which it is shown to produce *maximum a posteriori* (MAP) estimates of a signal in the presence of blur and noise. The MAP estimator is based on a prior assumption that the signal to be recovered consists of relatively few localized "spikes," each corresponding to an individual reflection. This is called a "sparse" signal model, and the MAP estimator is designed for such situations. Because the MAP estimate cannot be computed in closed form, an iterative procedure is used. Our formulation for processing lidar waveforms can be expressed in the time domain as:

$$\begin{aligned} \text{E}-\text{step: } \hat{z}^{(t)}[n] &= \hat{x}^{(t)}[n] + h^*[n]*\left(y[n] - h[n]*\hat{x}^{(t)}[n]\right) \\ \text{M}-\text{step: } \hat{x}^{(t+1)}[n] &= \frac{\max\left\{(\hat{z}^{(t)}[n])^2 - \tau\sigma_\eta^2, 0\right\}}{\hat{z}^{(t)}[n]} \end{aligned} \tag{3}$$

where $\hat{x}^{(t)}[n]$ is the estimate of the undegraded signal $x[n]$ at the t^{th} iteration of the algorithm; the nonnegative parameter τ is inversely related to the prior probability of a spike at a given location, and can be tuned for application-specific performance; $\hat{z}^{(t)}[n]$ is the estimate of the missing data in the EM procedure (see Figueiredo and Nowak, 2003); $*$ denotes the convolution operator; and other variables are as previously defined. The expectation (E) step (partial deconvolution) and maximization (M) step (denoising) of Equation 3 are iterated until convergence. The only difference between the implementation in this work and that described in Parrish and Nowak (2009) was in the stopping criterion used in this iterative algorithm. Based again on extensive testing, we adopted the following stopping criteria for use in this work: (a) the number of detected peaks does not change from one iteration to the next, and (b) the width of each peak is one sample, or (c) 100,000 iterations are reached without satisfying 1 and 2.

Hybrid (Deconvolve – Decompose) Approach

Another general approach to waveform processing entails a two step process: (1) the waveforms are Wiener filtered to reduce noise and remove blurring due to the transmit pulse width, and (2) Gaussians are fitted to the output. Jutzi and Stilla (2006) describe the process, as well as the underlying theory and motivation. Because this approach combines Wiener deconvolution and decomposition, we refer to it in this paper as a hybrid approach.

The implementation in this study follows Jutzi and Stilla (2006), but with one modification. In Jutzi and Stilla's approach, Wiener deconvolution is first applied in the frequency domain to estimate what is termed the surface response. Then, Gaussians are used to fit the estimated surface response in the time domain using the Levenberg-Marquardt algorithm. In our implementation, the first step is the same as Jutzi and Stilla's approach, i.e., Wiener deconvolution of the original waveform. But, in order to alleviate ringing which was pronounced in the Wiener deconvolved waveforms (but not observable in the raw waveforms with the system used in this study), Gaussian lowpass filtering is applied before implementing Gaussian fitting. We found that further smoothing helps greatly improve the correct detection rate and reduce the false alarm rate by reducing the ringing introduced in the Wiener filtering.

Experiment Design

To meet the research objectives described earlier, it is advantageous to use data collected in a controlled lab environment. It is certainly possible to use data collected from an airborne platform in empirical tests of waveform processing, as was done for bathymetric lidar waveforms by Allouis *et al.* (2010). However, in experiments using data collected from an airborne platform, there are typically numerous confounding factors, such as the exact distribution of points on the ground and on vertical objects, the weather (including, for example, the strength and direction of the wind which can greatly affect the returns from tree canopy and other vertical objects which may sway in the wind), atmospheric conditions, the angle of incidence of laser pulses intercepted by targets, and the presence and location within the scene of vehicles: people, birds, etc. Additionally, in obtaining ground-truth data for targets in the area covered by the airborne data, it is typically impossible to precisely characterize every target in the scene, no matter how extensive the ground survey. Due to a combination of the above factors, it can be difficult to distinguish between correct detections and false alarms in data acquired from an aircraft. Thus, the methods developed in this work involve a controlled ranging lab experiment in which the target characteristics and positions and other acquisition variables are precisely known.

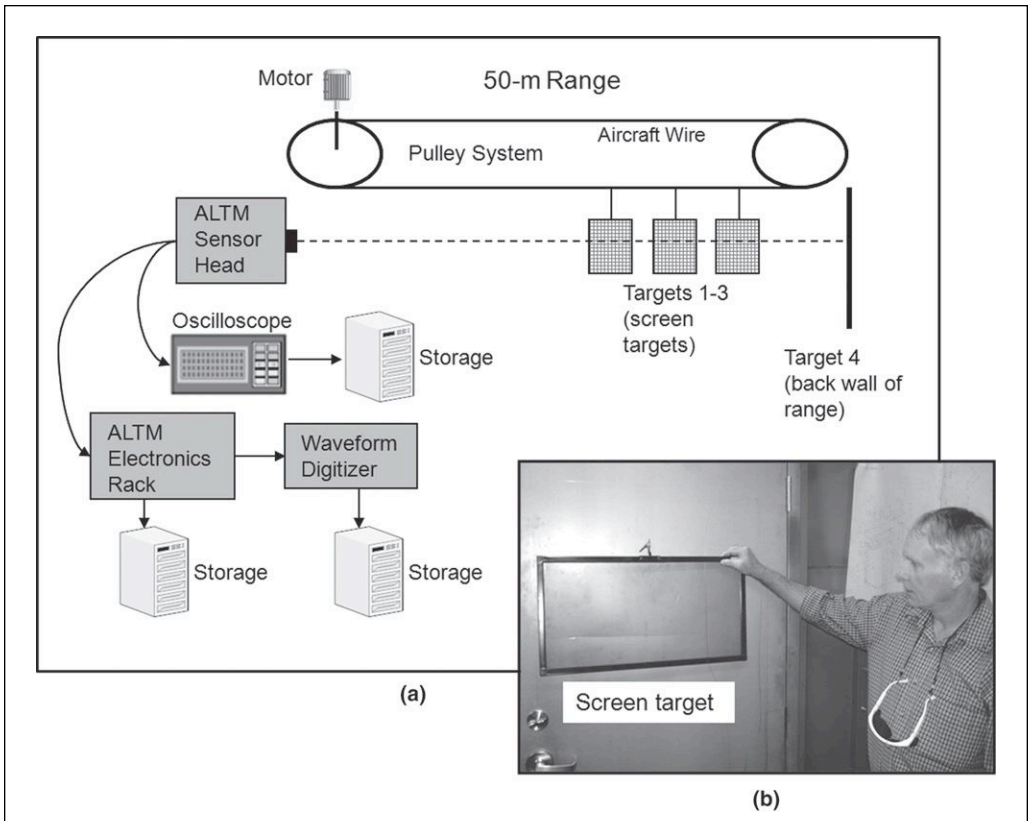

Figure 1. (a) schematic illustration of experiment design, and (b) photo of one of the screen targets. Each screen target was oriented such that its surface normal was parallel to the laser beam propagation direction.

Our empirical test procedure for comparing different lidar waveform processing techniques is based on the ranging lab setup depicted in Figure 1. The three adjustable screen targets comprise a key feature of the method. These are constructed of ordinary wire mesh, similar to that found on a window screen. Laser pulses incident on each screen target produce a detectable reflection (in fact, the reflectance [and, hence, the received optical power] can be varied in a reliable, low-tech manner by applying different shades of gray paint to the front side), while most of the incident photons pass through the screen mesh. The pulse shape is, thus, relatively unaltered at each interaction with a screen target, while the amplitude is decreased slightly. The amplitude of the return pulses can be kept nearly constant by using paint of higher reflectance (at the laser wavelength) on each successive screen target. Equally importantly, the separations of the screen targets can be accurately and directly measured using an engineer's scale, survey tape, or other distance measurement device. By starting the screen targets at relatively large separations and iteratively decrementing the separations, the achievable target resolution with each waveform processing strategy can be precisely determined. Furthermore, by acquiring discrete-return lidar simultaneously, the improvement over discrete-return lidar can be quantified. Additionally, other system settings, such as the pulse repetition frequency (PRF) and beam divergence, can also be varied in the experiment, as desired.

A similar empirical approach was successfully demonstrated by Tolt and Larsson (2007). The authors experimentally investigated the capability to detect a mine behind a tuft of grass in full-waveform lidar data and compared the results of two processing methods. Jutzi and Stilla (2006) also used an experimental setup with two targets to test their waveform processing algorithm. Our methods extend this approach though the introduction of the multiple screen targets with adjustable separations, which can be decremented iteratively, as described above.

Data Collection

The experiment was conducted in Optech's ranging lab on 21–22 October 2008 following the methodology described in the preceding section. The equipment used included an Optech ALTM Gemini (1,064 nm wavelength, 167 kHz max PRF, 0.25 mrad (1/e) beam divergence), an Optech waveform digitizer (8-bit amplitude resolution, 1 ns sample rate, 70 kHz max acquisition rate), and a Tektronix TDS 3052 two channel color digital phosphor oscilloscope (500 MHz bandwidth, 8 GS/s sample rate, 9-bit vertical resolution). As depicted in Figure 1, four targets were used in the 50 m range (Figure 1a). Targets 1, 2, and 3 were screen targets (Figure 1b), while Target 4 was the back wall of the range. Based on the general experiment design, the positions of Targets 1 and 2 were adjustable which enabled the Target 1-2 and Target 2-3 separations to be decremented at each iteration of the experiment. Targets 3 and 4 were immovable, and, therefore, the Target 3-4 separation remained fixed at 4 m throughout.

Each time the target separation was adjusted, discrete-return ALTM data, digitized waveforms, and oscilloscope traces were recorded with the laser firing first at a pulse repetition frequency (PRF) of 50 kHz, and then at 70 kHz. For each setup (defined by a specific set of target separations and PRF), 16,383 ($\sim 2^{14}$) waveforms were recorded, each

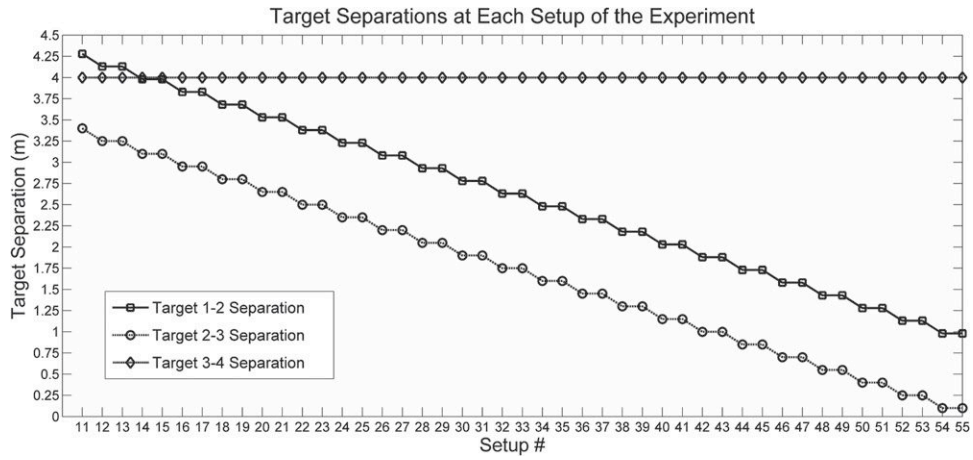

Figure 2. Target separations for each setup. The PRF was 50 kHz for even-numbered setups and 70 kHz for odd-numbered setups. The initial target separations at Setup 11 were 4.28 m, 3.40 m, and 4.00 m. For the last two setups (No. 54 and 55), the target separations were 0.98 m, 0.10 m, and 4.00 m.

corresponding to a separate laser shot. The reason for testing two different PRFs is that the size and shape of the transmitted pulse (and, hence, the system impulse response function) are functions of the PRF, and it was of interest to investigate how varying these parameters affected the results. After acquiring data at the two different PRFs, the Target 1-2 separation and Target 2-3 separation were each decreased by 15 cm. The above steps were repeated until the final setup, at which the smallest target separation was just 10 cm (see Figure 2). Each time Targets 1 and 2 were moved, their displacements were measured with a meter stick along the aircraft wire from which the targets were hung. Although it would have been preferable to have had a more accurate distance measurement device for obtaining the reference target separations, the measurements were made carefully and had an estimated uncertainty of 0.5 cm or better.

In the first ten iterations of the experiment, large target separations were used to test the equipment and overall experiment design and make slight modifications, as needed. These modifications included swapping out one of the screen targets for a one of higher reflectivity (bright white paint on the front side) and adjusting the attenuation factor using different flip-in optical filters to avoid saturating the digitizer. Because of the adjustments made in these first ten setups, and because the data from these setups provided little interesting information (at the large separations, all four targets were easily detected in both the discrete-return and full-waveform data), only the data from Setups 11 through 55 were used in the remainder of the study. The initial target separations at Setup 11 were 4.28 m, 3.40 m, and 4.00 m. Figure 2 shows how the Target 1-2 and 2-3 separations were decreased in 15 cm increments through Setup 55, the last iteration of the experiment.

One additional aspect of the experiment is worth noting here. In order to minimize the amount of valuable ranging lab time needed to carry out the experiment, once the targets were put into position for each setup, we immediately powered on and began firing the laser at 50 kHz and logging data. However, as a likely consequence of not allowing sufficient time for the system (including the receiver and digitizer) to stabilize, the 50 kHz data were noisier than the 70 kHz data (σ_η = 0.71 versus 0.46 digitizer counts; see Table 2 in the Data Analysis Section), and greater variability in the shape, amplitude, and shot-to-shot noise level of returns was observed. By the time we began acquiring at 70 kHz, the system had had greater time to stabilize, and returns were less noisy and noticeably more consistent from shot to shot. The greater shot-to-shot variability and higher noise level in the 50 kHz returns turned out to enable some interesting comparisons, as will be discussed later.

Data Analysis

The first step in the data analysis, prior to extracting ranges from the return waveforms, was to quantitatively analyze the digitized T_0 (transmit) pulses recorded during the experiment. This was done for four reasons: (1) to quantify differences between transmitted pulses at the two different PRFs used in this study, (2) to investigate shot-to-shot variation in transmitted pulses, (3) to evaluate how well the T_0 pulse can be modeled as a Gaussian, and (4) to estimate the lidar system impulse response function (alternatively referred to herein as the reference signal) needed for the waveform processing algorithms. This analysis consisted of visually inspecting several hundred digitized T_0 pulses for each PRF, as well as computing pulse widths (full width at half maximum (FWHM)) and peak amplitudes, and fitting Gaussians to the digitized T_0 pulse for all 884,682 waveforms collected over the two-day experiment. The coefficient of determination (R^2) of the Gaussian fit was used as an indication of how well each digitized T_0 pulse is modeled by a Gaussian. Statistics for each computed parameter for the T_0 pulses are shown in Table 1, broken down by PRF. Figure 3 shows the envelope of T_0 pulses within 1 standard deviation of the mean and Gaussian fits at both PRFs.

The pulse widths were found to be ~2 ns larger than anticipated; Optech's rule-of-thumb FWHM values are 9 ns at 50 kHz and 11 ns at 70 kHz, whereas the mean pulse widths shown in Table 1 are approximately 11 and 13 ns at 50 and 70 kHz, respectively. Increasing the PRF from 50 kHz to 70 kHz increased the pulse width by 16 percent and decreased the peak amplitude by 40 percent, on average. The pulse broadening with PRF did follow the expected

TABLE 1. T_0 (TRANSMITTED) PULSE STATISTICS

	PRF	Mean	StDev	Min	Max
Pulse width (FWHM) (ns)	50kHz	11.121	0.215	10.206	12.200
	70kHz	12.893	0.371	11.000	14.813
Peak amplitude (digitizer counts)	50kHz	55.321	1.185	50.000	61.000
	70kHz	32.996	1.115	28.000	39.000
Coefficient of determination (R^2) of Gaussian fit	50kHz	0.989	0.002	0.973	0.996
	70kHz	0.992	0.003	0.974	0.999

relationship. Briefly, when the PRF is increased, this drains the energy from the laser active medium faster. As a result, more round trips inside the laser cavity are needed to reach the lasing threshold, and the more round trips to reach the threshold, the longer the duration of emitted pulses.

The coefficients of determination of the Gaussian fit indicate that the transmitted pulse is well modeled by a Gaussian at both PRFs. Also, at both PRFs the pulses exhibit the same characteristic shape, rising just slightly faster than the fitted Gaussian, reaching a slightly greater peak amplitude, and falling off slightly more slowly on the trailing edge (see Figure 3). The standard deviations of the pulse widths and amplitudes, as well as the results of the visual inspection, indicate that there was relatively little pulse-to-pulse variability in the digitized T_0 pulses. All of the waveform processing algorithms investigated in this study require an estimate of the reference signal. Based on the quantitative and visual analysis of the digitized T_0 pulses described above, we decided to use a Gaussian fit to the T_0 pulse at each of the two PRFs as our estimate of the reference function. Importantly, the same reference function estimates were used consistently in all algorithms being compared.

From the T_0 pulse widths in Table 1, it is also possible to compute the theoretically-expected target resolution, as given in Wehr (2009):

$$\Delta R_{tar} = \frac{c \cdot w_{T0}}{2} \tag{4}$$

where c is the speed of light, and w_{T0} is the transmit pulse width (referred to as pulse length in Wehr, 2009). This gives theoretically-expected target resolutions of 1.67 m at 50 kHz and 1.93 m at 70 kHz. However, these should be considered conservative estimates, as studies have shown the capability to improve the target resolution well beyond that predicted by Equation 4 (e.g., Jutzi and Stilla, 2006).

The next step was to tune and run the three algorithms for extracting ranges from the recorded waveforms. An important point is that there is some subjectivity in tuning each algorithm. Thus, great effort in this study was devoted to the goal of ensuring that each algorithm was performing at an optimal level (i.e., maximum detection rate), subject to the false alarm rate constraint listed above. Table 2 shows the parameters for each algorithm and the final settings. The primary parameters adjusted during the tuning stage were: *Hw*, the half width half maximum of a Gaussian filter for smoothing original waveform, in the Gaussian decomposition; τ, inversely related to the prior probability of observing a reflection in the waveform, in the EM deconvolution; and *K*, the ratio of the power spectral density (PSD) of noise to PSD of signal in the Wiener filter (inversely related to SNR), in the hybrid approach. The noise standard deviation, σ_η, was estimated at each PRF as the sample standard deviation of over 330 manually-selected, noise-only samples, and was kept consistent throughout the three algorithms. The half width of the smoothing filter used in the Gaussian decomposition and hybrid algorithms was set experimentally, as described earlier, and thereafter kept constant.

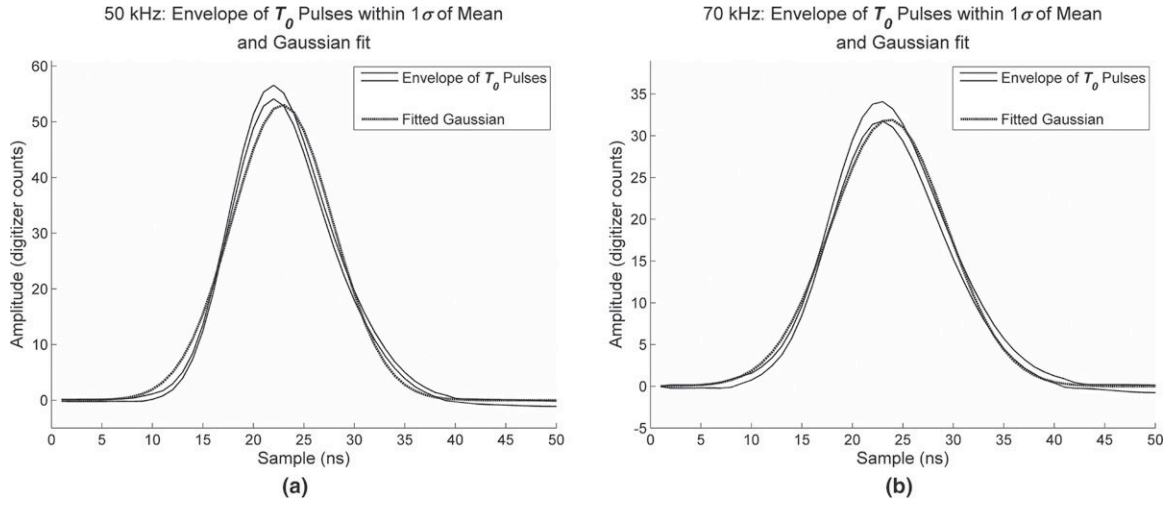

Figure 3. Envelope of T_0 (transmit) pulses within 1σ of mean and fitted Gaussians at (a) 50 kHz, and (b) 70 kHz.

TABLE 2. PARAMETERS IN THE ALGORITHMS TESTED IN THIS RESEARCH AND SETTINGS USED IN MAXIMIZING DETECTION RATE, WHILE CONSTRAINING AVERAGE FALSE ALARM RATE IN SETUPS 11 TO 20 TO <1% FOR BOTH PRFS (70 KHZ AND 50 KHZ)

Processing Algorithm	Parameters	Final setting
Gaussian Decomposition	*Hw*: half width half maximum of a Gaussian filter for smoothing original waveform σ_η: noise standard deviation w_l: minimum full width half maximum of important Gaussian	*Hw* = 2.5 σ_η = 0.714 (50 KHz), 0.462 (70 KHz) w_l = 0.118
Deconvolution	τ: inversely related to the prior probability of observing a reflection in the waveform σ_η: noise standard deviation	τ = 0.1 σ_η = 0.714 (50 KHz), 0.462 (70 KHz)
Hybrid (Deconvolve-Decompose)	*K*: ratio of PSD of noise to PSD of signal in Wiener filter (inversely related to SNR) *Hw*: half width half maximum of a Gaussian filter for smoothing original waveform σ_η: noise standard deviation w_l: minimum full width half maximum of important Gaussian	*K* = 1.5 *Hw* = 2.5 σ_η = 0.714 (50 KHz), 0.462 (70 KHz) w_l = 0.118

The primary points of comparison in this study were: (a) target resolution, (b) target separation accuracy, (c) run time, and (d) the user's experience in running the different algorithms. Due to the stochastic nature of the outcome of any one particular laser shot, for the target resolution comparisons in this work, we define the target resolution, ΔR_{tar}^{P}, as:

$$\Delta R_{tar}^{P} \equiv \min \{s{:}Pr(D_s) > P\} \quad (5)$$

where s is the target separation (distance between adjacent targets), D_s is the event that the targets are correctly resolved (i.e., individually detected) at separation s, and P is a user-defined probability that can be adjusted, based on the considerations of the particular application (forestry, coastal wetlands mapping, or airport obstruction surveying, for example). Thus, in words, $\Delta R_{tar}^{50\%}$ is the minimum target separation such that, on any given laser shot, the probability of correctly resolving the targets is greater than 50 percent.

The second point of comparison for the different waveform processing strategies is target separation accuracy. In this study, direct determination of the ranging accuracy achieved with the different algorithms was not possible, since no independent measurement of the absolute distance from the lidar to the targets was made. However, because the target separations were independently and accurately measured during the experiment, the ranging accuracy could be indirectly assessed through comparison of the target separations determined through processing of the waveforms with the known (reference) separations.

The final two points of comparison were run time and user's experience. To enable meaningful run-time comparisons, all tests were performed on the same computer: a Dell™ Optiplex 960 Desktop, with an Intel™ Core2 Duo CPU E8600 (3.33 GHz processor), 3.25 GB RAM, running Windows® XP. The final category, user's experience, captures all other observations made in running the different algorithms tested in this study. While these user observations are less quantifiable and more subjective than the other comparison categories, we consider them an important outcome of the study, since our primary objective is to compare the different waveform processing algorithms from an operational user's perspective.

Results and Discussion

Comparisons of detection rate as the target separation is decreased are shown in Figure 4. Starting with the Figure 4a (70 KHz), the detection rate falls off fastest with the discrete return data (as expected), then Gaussian decomposition, then the hybrid approach, then the EM deconvolution approach. Until setup 31, greater than 40 percent (correct) detection rate is obtained for all three waveform-based methods. The trend is roughly the same in Figure 4b (50 KHz), except that here the hybrid approach outperforms the EM deconvolution.

One interesting and counterintuitive observation can be made in comparing the curves in Figure 4a and 4b. Namely, we would expect the results at 50 KHz to be consistently better than at 70 KHz, since the transmit pulse width is greater (by ~2 ns) at the higher PRF (see Table 1), leading to greater blurring. However, the trend is actually reversed in the EM deconvolution output: the 70 KHz results are much better than the 50 KHz results. The likely cause of this unexpected result is the greater shot-to-shot variability and higher noise (0.71 versus 0.46 digitizer counts; see Table 2) in the 50 KHz data, due to not allowing the system sufficient time to stabilize before each 50 KHz acquisition (see the Experiment Section). Closer analysis of the results revealed that the EM deconvolution algorithm actually continued reporting the correct number of targets (four) in the 50 KHz data, even as the target separation became quite small, just as it did in the 70 KHz data. However, the greater shot-to-shot variability in returns led to a much greater spread in the reported positions on the time axis of detected targets (in particular, Targets 2 and 3) with the EM deconvolution algorithm (as illustrated in Figure 5), in turn leading to greater variability in reported target separations. And, whenever the reported target separation differed from the nominal separation by more than our established threshold,

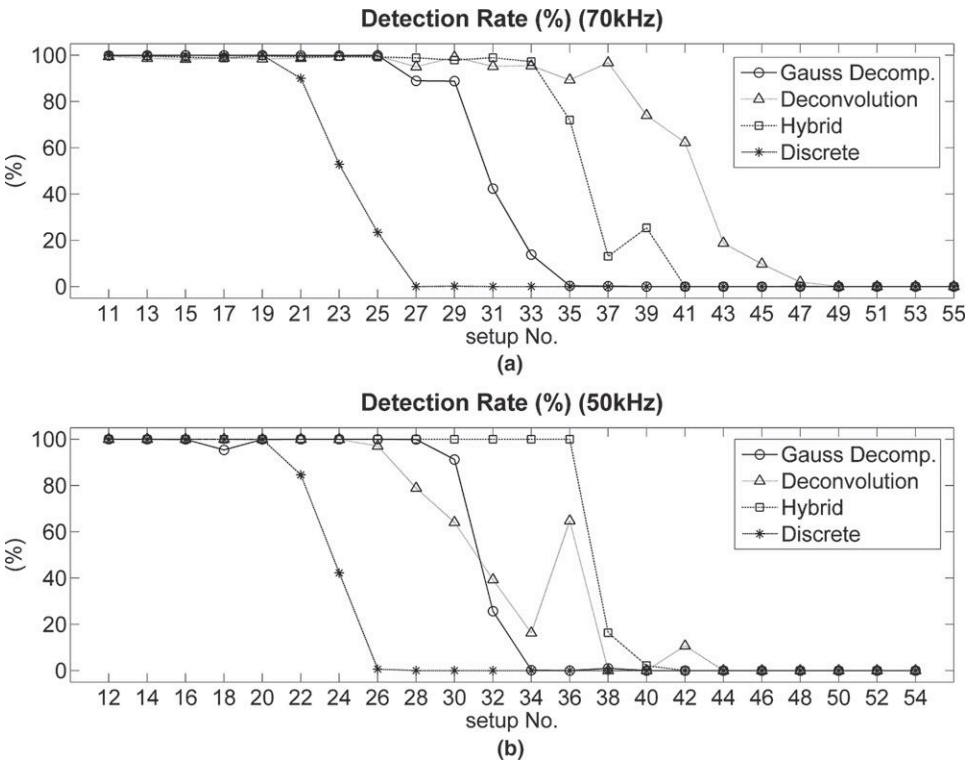

Figure 4. Detection rate plots for the three different algorithms tested at both (a) 70 kHz, and (b) 50 kHz.

our program for comparing the output of the different waveform processing algorithms declared the result a "miss." The peak in the curve that occurs at Setup 36 corresponds to a Setup for which the observed shot-to-shot variability in returns was not as high.

This behavior of the EM deconvolution algorithm can be understood through the fact that the algorithm is based on the assumption that the noise and the system impulse response function (used as our estimate of the blurring function) parameters are accurately known and stationary; in the case of our 50 kHz data, this assumption appears to have been invalidated by the insufficient system stabilization time. The Gaussian decomposition and hybrid approaches are less sensitive to uncertainty in these parameters, as they can adjust the size and shape of fitted Gaussians. In typical airborne lidar acquisition, however, the system should have sufficient time to stabilize before logging. In short, we believe the results shown in Figure 4a are more representative of the true capabilities of the EM deconvolution algorithm; Figure 4b illustrates the degradation in results when the starting assumptions about the data are invalid.

To further understand the results shown in Figure 4, we also plotted and visually examined the output of the three algorithms for a number of waveforms. Figures 6, 7, and 8 illustrate how detection results changes as target separation is decremented, i.e., as setup number increases. In this case, we have focused on the 70 kHz results, in which the returns were more stable. In Figure 6, the target separations are relatively large, and all three methods perform comparably. In Figure 7, the Gaussian decomposition method misses one target, because separation 2-3 is too close to resolve. In contrast, the hybrid method succeeds in detecting four targets; this is due to the impact that Wiener filtering removes waveform ambiguity that occurs in closely spaced targets before implementing Gaussian decomposition. As shown in Figure 8, when the targets are much closer, both the Gaussian decomposition and hybrid method lose target detection capability. But, in this case, all four targets are still detected by EM deconvolution. As an aside, it should be noted that the absolute positions of the targets on the time axis in these figures are inconsequential, as we considered only the target separations (i.e., their relative positions). The absolute positions on the time axis are a function of the range offset, which must be determined as part of the lidar calibration procedure, regardless of the ranging algorithm used.

The next point of comparison for the three algorithms tested in this study was target separation accuracy. Comparisons of the measured target separations against the known (reference) separations are shown in Figures 9 and 10. (Because it would have been difficult to provide meaningful comparisons of target separation accuracy beyond the point at which large numbers of target misses began to be reported, we limited the comparisons of target separation accuracy to Setups 11 to 31 for 70 kHz and 12 to 30 for 50 kHz.)

Table 3 further summarizes these comparisons by averaging the target separation accuracy results across the setups for each algorithm and at each PRF. The standard deviations arguably provide more insight than the means, since any bias in reported target ranges can, at least in theory, be removed through an appropriate calibration technique (although, of course, in practice this may not always be easy). As a very general observation from Figures 9 and 10, for all three algorithms, the accuracy tends to degrade (differences between measured and reference target separations tend to increase) as the targets get closer, as one would expect. In the early setups (relatively large

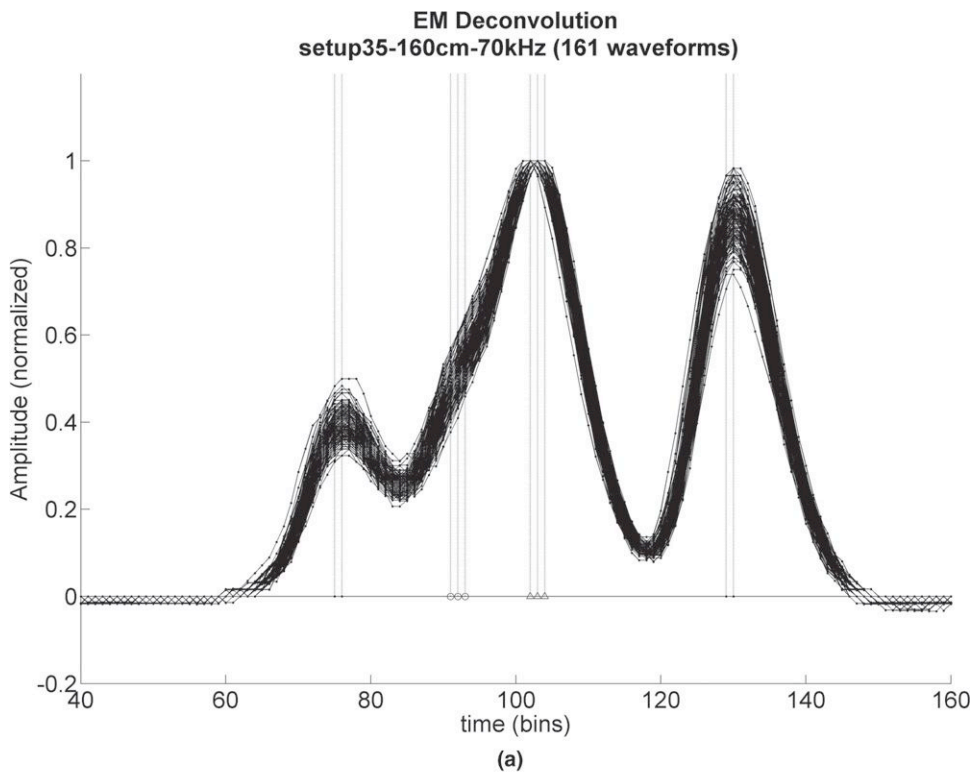

Figure 5. Example of greater shot-to-shot variability in (b) 50 kHz returns versus (a) 70 kHz returns and effect on the spread in detected target positions on the time axis (indicated by spread in dashed, vertical lines) with the EM deconvolution algorithm. In this case, the greater spread in the Target 2-3 separation led the result to sometimes be declared a "miss," even when the algorithm reported the correct number of detected targets (four).

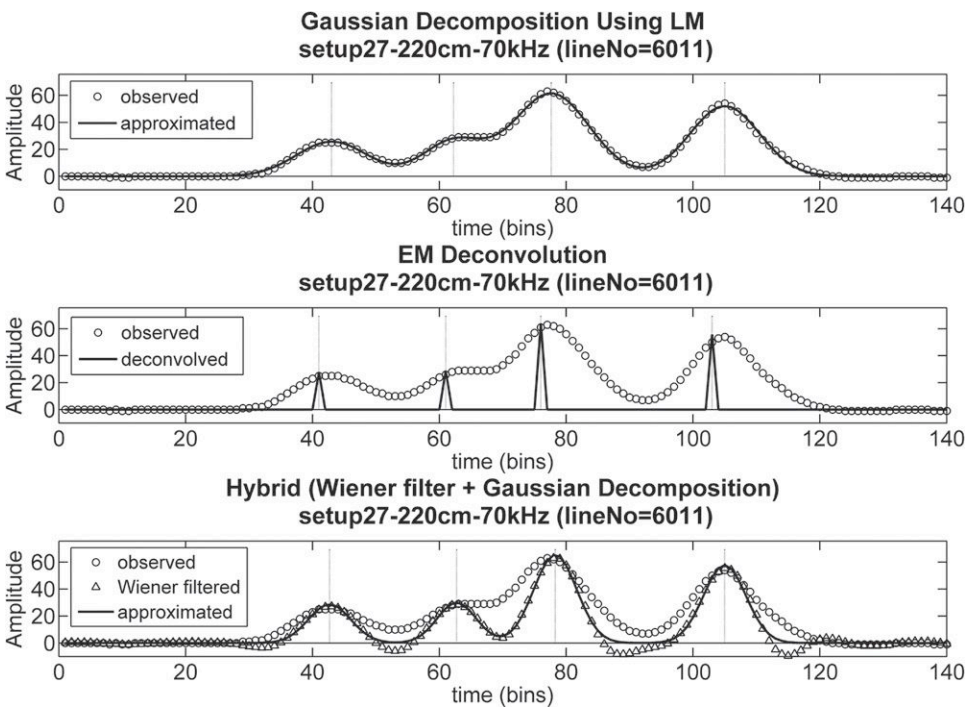

Figure 6. Example of a waveform from Setup 27. The target separations are fairly large, and all four targets are resolved easily with each waveform processing algorithm.

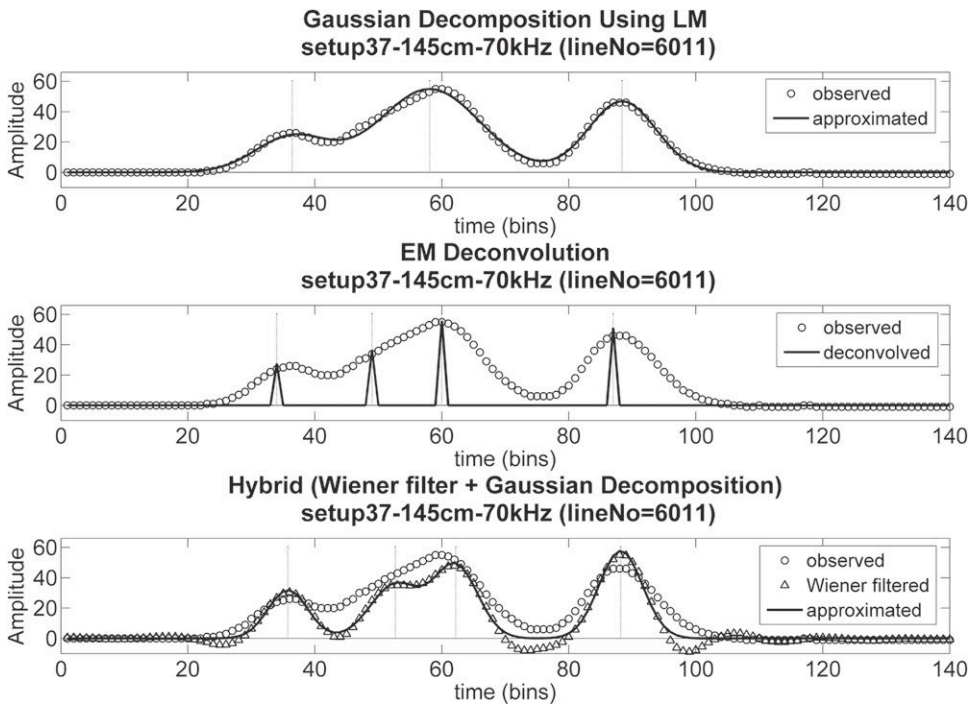

Figure 7. Example of a waveform from Setup 37. All four targets are correctly resolved using both EM deconvolution and the hybrid approach, but not with the Gaussian decomposition approach.

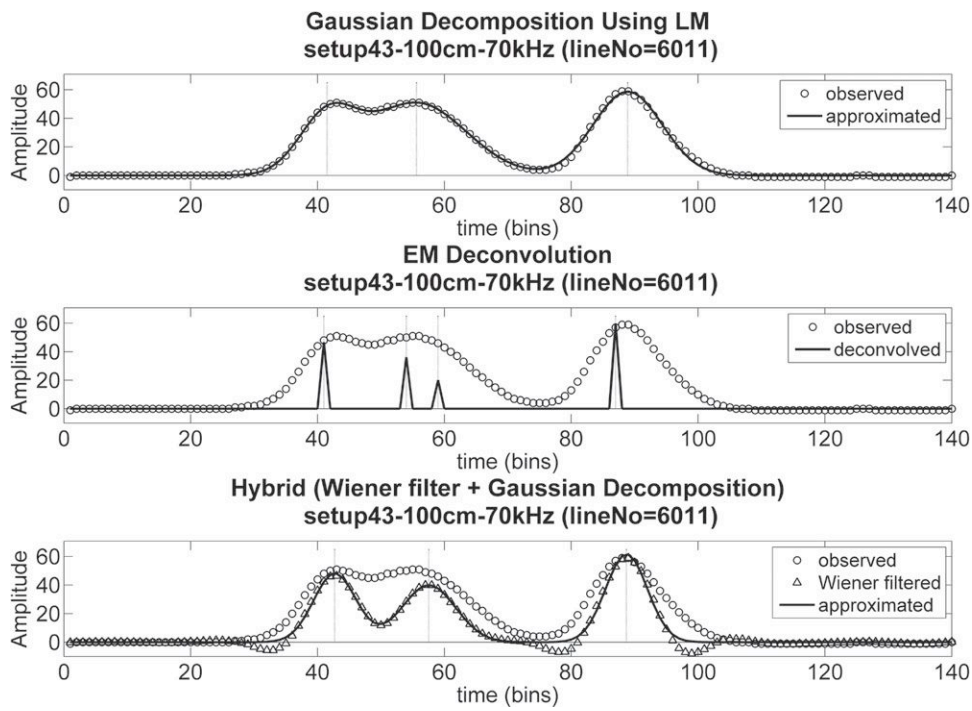

Figure 8. Example of a waveform from Setup 43. Here, only the EM deconvolution resolves all four targets.

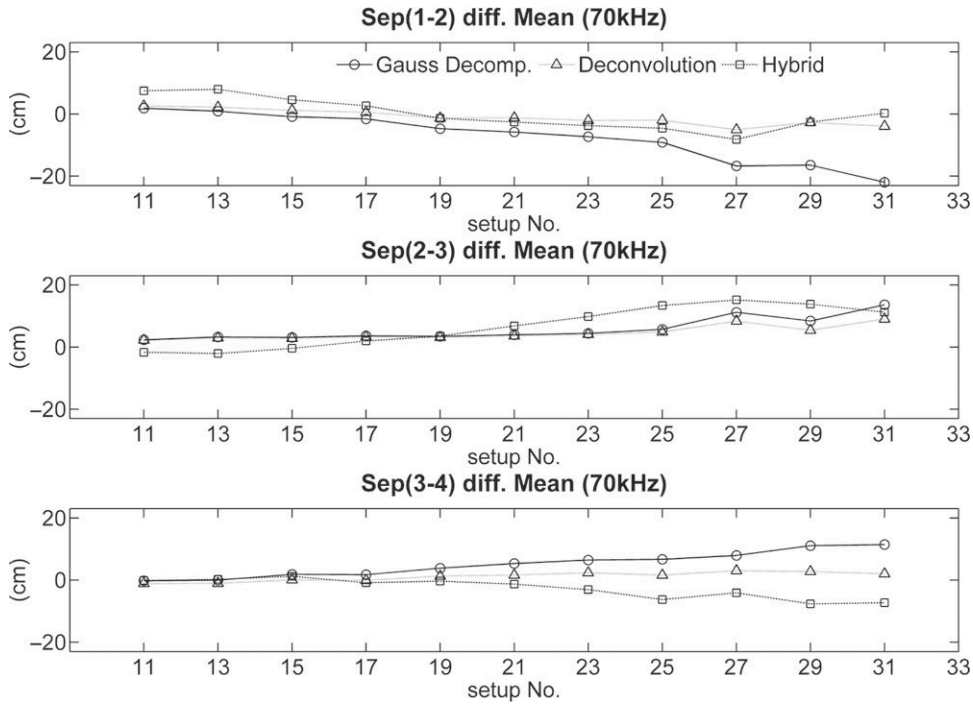

Figure 9. Target separation accuracy at 70 kHz (Setups 11 to 31) shown in terms of average differences between nominal (reference) separation and waveform lidar based measurement.

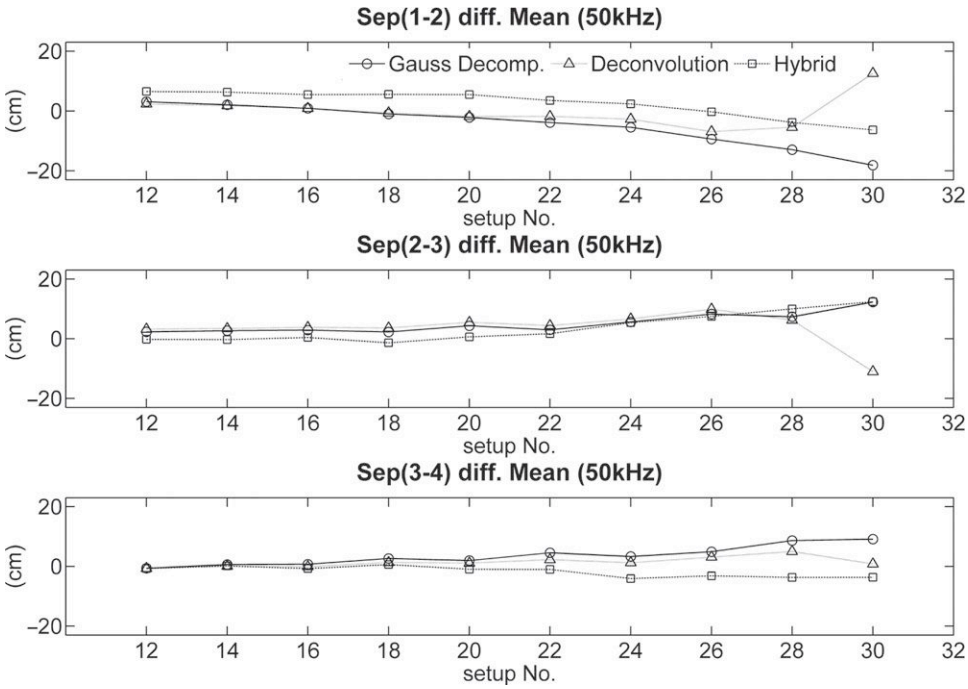

Figure 10. Target separation accuracy at 50 kHz (Setups 12 to 30) shown in terms of average differences between nominal (reference) separation and waveform lidar based measurement.

TABLE 3. TARGET SEPARATION ACCURACY RESULTS AVERAGED ACROSS SETUPS 11 TO 31 FOR EACH ALGORITHM AT EACH PRF; ALL VALUES ARE SHOWN IN METERS

		Gaussian Decomposition			EM Deconvolution			Hybrid (Wiener Deconvolution)		
		Sep. (1-2)	Sep. (2-3)	Sep. (3-4)	Sep. (1-2)	Sep. (2-3)	Sep. (3-4)	Sep. (1-2)	Sep. (2-3)	Sep. (3-4)
70KHZ	mean	−0.0742	0.0574	0.0511	−0.0106	0.0452	0.0114	0.0000	0.0651	−0.0268
	Std.	0.0791	0.0373	0.0407	0.0246	0.0227	0.0146	0.0514	0.0655	0.0319
50KHZ	mean	−0.0468	0.0510	0.0359	−0.0016	0.0357	0.0142	0.0250	0.0362	−0.0172
	Std.	0.0692	0.0331	0.0329	0.0537	0.0551	0.0167	0.0453	0.0487	0.0174

target separation), all three algorithms provide good accuracy (especially if any bias is removed) as indicated by the relatively flat lines near the left hand side of the plots. However, the target separation accuracy tends to deteriorate, as indicated by the diverging curves near the right-hand side of the plots, slightly before the point at which the target resolution begins to fall off (for example, Figure 4). Again, this agrees with intuition: for any given algorithm, as the target separations are iteratively decreased, first the accuracy of the measured target separations begins to degrade, and then, at some point after that, the algorithms start to lose the ability to resolve the targets at all.

The user's experiences in running the three algorithms are summarized below, based on detailed, written notes made while processing the data. Additionally, the run time and other quantifiable results are further summarized in Table 4.

User's Experience with Gaussian Decomposition

- The parameters used to vary the amount of denoising (Gaussian smoothing of observed waveform) must be tuned very carefully; lack of attention to this step limits target separability. The smoothing filter width heavily influences the results. (With actual airborne survey data, these issues may be addressed in the calibration stage.)
- Saturated returns (e.g., due to specular reflections, which are occasionally observed, even in airborne data), can cause the program to abort, unless the code is modified to check for and handle this case.
- The criterion used in defining "important Gaussians" is also very important. We found that using an amplitude of approximately 3× the noise standard deviation provided the best results.
- An interesting observation is that the width of fitted Gaussians in the return is sometimes less than the width of the T_0 pulse. We are unsure why this occurs, but one possibility is that it is simply due to the fact that, in the Optech digitizer, the T_0 receiver is different than the return pulse receiver and has a different rise time.

User's Experience with EM Deconvolution

- It is important to tune the parameter τ to achieve the desired balance between target separability and false alarm rate.
- Selection of the EM iteration stopping criterion greatly affects run time. After experimenting with various stopping criteria, we determined that a criterion based on spike positions remaining constant from one iteration to the next and the width of each spike being one sample gave the best results.

TABLE 4. SUMMARY OF KEY RESULTS. THE LISTED RUN TIMES ARE THE MEASURED TIMES FOR PROCESSING 10 PERCENT OF ALL WAVEFORMS ACQUIRED OVER THE TWO-DAY EXPERIMENT (I.E., ~1,683/ SETUP). THE TARGET RESOLUTION IS THE MINIMUM DISTANCE AT WHICH ADJACENT TARGETS ARE CORRECTLY RESOLVED WITH A PROBABILITY OF >50% ($\Delta R_{tar}^{50\%}$). THE BIAS-REMOVED TARGET SEPARATION ACCURACIES ARE THE VALUES OBTAINED BY AVERAGING THE LISTED STANDARD DEVIATIONS IN TABLE 3 ACROSS THE TARGET 1-2, 2-3, AND 3-4 SEPARATIONS.

Performance metric	Gaussian Decomposition	EM Deconvolution	Hybrid Approach (Wiener Deconvolution)
Run time	212.2 min	436.7 min	419.5 min
Target resolution ($\Delta R_{tar}^{50\%}$)	70 KHZ: 205 cm (Setup 29) 50 KHZ 190 cm (Setup 30)	70 KHZ: 115 cm (Setup 41) 50 KHZ: 190 cm (Setup 30)	70 KHZ: 160 cm (Setup 35) 50 KHZ: 145 cm (Setup 36)
Target separation accuracy (bias removed)	70 KHZ: 0.052 m 50 KHZ: 0.045 m	70 KHZ: 0.021 m 50 KHZ: 0.042 m	70 KHZ: 0.050 m 50 KHZ: 0.037 m

- To avoid limiting the ranging precision by the sampling rate of the digitizer, the dimension of the reconstruction must be increased. This enables the ranging to be done with arbitrary precision, but at the expense of increased runtime.

User's Experience with Hybrid Approach (Wiener Deconvolution)

- Wiener-filtered signal may be a better basis for Gaussian decomposition, since both noise and blurring are removed in significant amount. But ringing is introduced in the Wiener filtering step. This ringing effect can lead to many false alarms as targets are closer, i.e. latter setups.
- Since the second step in this approach involves fitting Gaussians, the general comments made regarding the Gaussian decomposition also apply here.

Some general observations can be made from the results summarized in Table 4 and Figures 4 through 10. First, we found that all three waveform-based ranging strategies easily outperformed the discrete return (hardware-based) ranging, in terms of target resolution. While this was completely expected, it further confirms the great improvement in target resolution that is possible with full-waveform data and software-based ranging. We also were able to improve upon the theoretically-expected target resolution predicted by Equation 4 with the best algorithm(s) at each PRF, even with a fairly stringent restriction on false alarm rate. We further found that both the EM deconvolution and hybrid (deconvolve-decompose) strategies tended to outperform the Gaussian decomposition approach in resolving very close targets. The EM deconvolution algorithm performed especially well on the 70 kHz data, continuing to resolve targets for several setups beyond both the Gaussian decomposition and hybrid algorithms. As noted earlier, the EM deconvolution algorithm was found to be particularly affected by the insufficient system stabilization time that led to higher noise and greater shot-to-shot variability in the 50 kHz returns, as this caused the underlying assumptions of the algorithm to be violated. The algorithm would have to be modified to handle this type of data, if it was expected to be encountered frequently. The Gaussian decomposition approach offered the shortest run times by a factor of ~2. Although outside the scope of our comparisons, which focused on range extraction and discrimination performance, it is important to note that the Gaussian decomposition also provides additional information in the widths of extracted Gaussian components. We do not consider this parameter here, but other studies have addressed its physical interpretation and uses (Wagner *et al.*, 2006 and 2008).

Conclusions and Perspectives

In this work, we have presented a new empirical procedure for testing and comparing different lidar waveform processing algorithms. The experiment is carried out in a ranging lab, and uses multiple, screen targets, the separations of which can be directly measured and incrementally adjusted. The methodology was shown to enable a robust characterization of different waveform processing strategies, as well as both quantitative and subjective comparisons of the differences. Here, we used these methods to test three waveform processing methods described in the scientific literature: Gaussian decomposition, EM deconvolution, and a hybrid (deconvolve-decompose) approach.

The results obtained through this experimental approach enabled us to make some general observations regarding the differences between the three waveform processing strategies tested in this study. The EM deconvolution algorithm provided the best target resolution and target separation accuracy results when the system and noise parameters were well known and stable, outperforming the other two algorithms in the 70 kHz case. However, in its present form, it is less robust to poorly-known or non-stationary parameters. A possible modification to the EM deconvolution algorithm would enable it to better handle the case of unknown or time-varying pulse shapes. For example, the problem could be reformulated as sparse recovery from an over-complete dictionary, consisting of Gaussian pulses of varying widths and positions, rather than just varying positions (essentially a compressed sensing approach). However, the extent to which this would be beneficial for actual airborne survey data is presently unknown, since, in typical survey situations, the system is given sufficient time to stabilize before data acquisition.

The hybrid (deconvolve-decompose) approach provided better target resolution and slightly better target separation accuracy results than the Gaussian decomposition at both PRFs. It also outperformed the EM deconvolution on the 50 kHz data and had a slightly shorter run time than the EM deconvolution. The Gaussian decomposition approach was shown to have the shortest run times (by a factor of ~2), and provided good results at relatively large target separations. Although outside the scope of our investigations, Gaussian decomposition also provides additional information in the widths of extracted Gaussian components.

These results further indicate that there is no single best waveform processing strategy across all applications; different waveform strategies have different advantages and disadvantages. The bottom line is that choice of a specific lidar waveform processing algorithm is very application-dependent and should be left to the end-user. Thus, one recommendation for commercial software developers in implementing full-waveform support is to provide the user with options for different waveform processing algorithms, including user-selectable parameter settings within each. Documentation should also be provided to assist the user in selection of a

particular algorithm and settings. These recommendations are consistent with the general philosophy behind full-waveform data, which is to replace fixed, hardware-based subsystems with user-customizable software alternatives. A fusion-based approach to waveform processing (i.e., combining the results of multiple processing strategies) may ultimately prove most beneficial for a variety of applications and is a recommended topic for continued research.

Other potential extensions of this work would further characterize the performance of different algorithms for complex targets that occur in the natural world. This could include extending the ranging lab experiment using more targets, additional combinations of target separations, and targets with distinct reflecting behaviors. Additionally, it would be interesting to attempt a similar experiment with airborne data. The drawbacks are that the environment could not be controlled to the extent that is possible in the ranging lab, and the reference data would be subject to greater uncertainty. However, the ability to assess a nearly limitless range of natural targets and configurations (including parameters related to target roughness, reflectance, etc.) is a strong motivating factor for the proposed extension. These are additional recommended topics for further research.

References

Allouis, T., J.-S. Bailly, Y. Pastol, and C.L. Le Roux, 2010. Comparison of LiDAR waveform processing methods for very shallow water bathymetry using Raman, near-infrared and green signals, *Earth Surface Processes and Landforms*, Vol. 35, pp. 640–650.

Baltsavias, E.P., 1999. Airborne laser scanning: basic relations and formulas, *ISPRS Journal of Photogrammetry and Remote Sensing*, Vol. 54, pp. 199–214.

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 and Remote Sensing*, Vol. 54, pp. 115–122.

Chauve, A., C. Mallet, F. Bretar, S. Durrieu, M. Pierrot-Deseillingny, and W. Puech, 2007. Processing full-waveform lidar data: Modeling raw signals, *Proceedings of ISPRS Workshop on Laser Scanning 2007 and SilviLaser 2007*, 12–14 September, Espoo, Finland.

Figueiredo, M.A.T., and R.D. Nowak, 2003. An EM algorithm for wavelet-based image Restoration, *IEEE Transactions on Image Processing*, 12(8):906–916.

Harding, D., 2009. Pulsed laser altimeter ranging techniques and implications for terrain Mapping, *Topographic Laser Ranging and Scanning: Principles and Processing* J. Shan and C.K. Toth, editors), CRC Press, Taylor and Francis Group, Boca Raton, Florida.

Hofton, M.A., J.B. Minster, and J.B. Blair, 2000. Decomposition of laser altimeter waveforms, *IEEE Transactions of Geoscience and Remote Sensing*, 38(4):1989–1996.

Hoge, F.E., R.N. Swift, and E.B. Frederick, 1980. Water depth measurement using an airborne pulsed neon laser system, *Applied Optics*, 19(6):871–883.

Hug, C., A. Ullrich, and A. Grimm, 2004. LITEMAPPER-5600 – A waveform digitizing lidar terrain and vegetation mapping system, *International Archives of Photogrammetry and Remote Sensing*, Vol. 36, Part 8/W2, pp. 24–29.

Johnstone, I., G. Kerkyacharian, D. Picard, M. Raimondo, 2004. Wavelet deconvolution in a periodic setting, *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, Vol. 66, pp. 547–573.

Jutzi, B., and U. Stilla, 2006. Range determination with waveform recording laser systems using a Wiener Filter, *ISPRS Journal of Photogrammetry and Remote Sensing*, 61(2):95–107.

Lemmens, M., 2007. Airborne lidar sensors: 2007 Product Survey, *GIM International* , Vol. 21, No. 2.

Mallet, C., and F. Bretar, 2009. Full-waveform topographic lidar: State-of-the-art, *ISPRS Journal of Photogrammetry and Remote Sensing*, Vol. 64, pp. 1–16.

Means, J.E., S.A. Acker, D.J. Harding, J.B. Blair, M.A. Lefsky, W.B. Cohen, M.E. Harmon, and W.A. McKee, 1999. Use of large-footprint scanning airborne lidar to estimate forest stand characteristics in the western Cascades of Oregon, *Remote Sensing of Environment*, 67(3):298–308.

Nayegandhi, A., J.C. Brock, C.W. Wright, and M.J. O'Connell, 2006. Evaluating a small footprint, waveform-resolving lidar over coastal vegetation communities, *Photogrammetric Engineering & Remote Sensing*, 72(12):1407–1417.

Nordin, L., 2006. *Analysis of Waveform Data from Airborne Laser Scanner Systems*, MS thesis, Luleå University of Technology, Luleå, Sweden.

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*, 71(1):37–46.

Parrish, C.E., 2007. *Vertical Object Extraction from Full-Waveform Lidar Data Using a 3D Wavelet-based Approach*, PhD dissertation, University of Wisconsin-Madison, Madison, Wisconsin.

Parrish, C.E., and R.D. Nowak, 2009. An improved approach to lidar airport obstruction surveying using full-waveform data, *Journal of Surveying Engineering*, 135(2):72–82.

Persson, Å., U. Söderman, J. Töpel, and S. Ahlberg, 2005. Visualization and analysis of full-waveform airborne laser scanner data, *Proceedings of ISPRS WG III/3, V/3 Workshop "Laser scanning 2005,"* 12–14 September, Enschede, The Netherlands.

Reitberger, J., P. Krzystek, and U. Stilla, 2006. Analysis of full waveform LIDAR data for tree species classification, Symposium of ISPRS Commission III: Photogrammetric Computer Vision (PCV06), *International Archives of Photogrammetry, Remote Sensing, and Spatial Information Sciences*, 36(3):228–233.

Roncat, A., G. Bergauer, and N. Pfeifer, 2010. Retrieval of the backscatter cross-section in full-waveform lidar data using b-splines, *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 38(Part 3B):137–142.

Tolt, G., and H. Larsson, 2007. Waveform analysis of lidar data for targets in cluttered Environments, *Proceedings of Electro-Optical Remote Sensing, Detection, and Photonic Technologies and Their Applications*, 18 September, Florence, Italy, SPIE, Vol. 6739, 67390A.

Turin, G.L., 1960. An introduction to matched filters, *IRE Transactions on Information Theory*, 6(3):311–329.

Wagner, W., A. Ullrich, T. Melzer, C. Briese, and K. Kraus, 2004. From single-pulse to full-waveform airborne laser scanners: Potential and practical applications, *International Archives of Photogrammetry and Remote Sensing*, 35(Part B3):201–206.

Wagner, W., A. Ullrich, V. Ducic, T. Melzer, and N. Studnicka, 2006. Gaussian decomposition and calibration of a novel small-footprint full-waveform digitising airborne laser scanner, *ISPRS Journal of Photogrammetry & Remote Sensing*, 60(2):100–112.

Wagner, W., A. Roncat, T. Melzer, and A. Ullrich, 2007. Waveform analysis techniques in airborne laser scanning, *International Archives of Photogrammetry and Remote Sensing*, XXXVI(Part 3/W52):413–418.

Wagner, W., M. Hollaus, C. Briese, and V. Ducuc, 2008. 3D vegetation mapping using small-footprint full-waveform airborne laser scanners, *International Journal of Remote Sensing*, 29(5)1433–1452.

Walter, M.D., 2005. *Deconvolution Analysis of Laser Pulse Profiles from 3-D LADAR Temporal Returns*, MS thesis, Air Force Institute of Technology, Wright-Patterson Air Force Base, Ohio.

Wang, Y., J. Zhang, A. Roncat, C. Künzer, and W. Wagner, 2009. Regularizing method for the determination of the backscatter cross section in lidar data, *Journal of the Optical Society of America*, 26(5):1071–1079.

Wehr, A., 2009. LiDAR systems and calibration, *Topographic Laser Ranging and Scanning: Principles and Processing* (J. Shan and C.K. Toth, editors), CRC Press, Taylor and Francis Group, Boca Raton, Florida.

(Received 22 October 2010; accepted 25 February 2011; final version 30 March 2011)