Generate Digital Elevation Models Using Laser Altimetry (LIDAR) Data

Christopher Weed

Final Report
EE 381K
Multidimensional Digital Signal Processing
December 11, 2000

Abstract

A Laser Altimetry (LIDAR) system aboard an aircraft can yield highly accurate data about the ground surface and vegetation below. Raw LIDAR points must be processed to generate a digital elevation model (DEM), i.e. a digital map of the terrain surface. A number of techniques can be used to generate the DEM, and in this report I investigate several of these including a new algorithm based on recovering the lower envelope of a signal using amplitude demodulation. The default and most widely used technique consists of a series of thresholds based on local mean and variance measures. The amplitude demodulation technique runs much faster than the default technique and should be robust for flat and varying ground surfaces. Visual inspection of the entire image suggests that the amplitude demodulation technique achieves good results.

Introduction

Capable of achieving decimeter-level accuracy, laser altimeters are quickly becoming an indispensable tool for generating high-resolution topographic maps. Laser altimetry (LIDAR) data is collected by flying a laser over the area to be mapped. The laser is shot at the earth and the elapsed time to the first and last return pulse is recorded. The LIDAR system uses the elapsed time, GPS location, and other positioning information to create a set of 3-dimensional points of the surface below. In addition to achieving a high level of accuracy, the ability of the laser to penetrate the vegetation canopy opens the possibility of accurately mapping the ground surface underneath. A map of the ground surface, also called a digital elevation model (DEM), is an invaluable tool for hydrological purposes, such as studying water flow, erosion, and flooding. Unfortunately, the laser does not always penetrate the canopy, and it is necessary to process the data to produce a DEM. Another difficulty necessitated by the high accuracy of the LIDAR system is the volume of data produced by the system. Regions less than 100 km^2 in area can produce greater than 25 gigabytes of data points. Therefore, data reduction and algorithm order are very important factors in any algorithm used for processing the data.

Objective

The objective of this project is to produce an accurate DEM for non-flat ground surfaces. The DEM should remove all vegetation while maintaining natural changes in elevation. Important factors for comparing algorithms include level of user interaction, accuracy of result, and speed of computation. The level of user interaction pertains to how much human input is necessary, including assumptions about the data, adjusting parameters, and post-processing results. Accuracy can be tested by comparing the results to ground truth where it is known, and visual inspection where ground truth is unavailable. The speed of computation pertains to the amount of time needed to produce a 1 m by 1 m DEM starting from the raw LIDAR data points. A reasonable test of computation time is whether a new algorithm is faster than the widely accepted default algorithm.

Background

The first step in many techniques for generating a DEM is to create a digital image by gridding the data points. This will drastically reduce the amount of data that needs to be processed from multiple data points in a grid cell to a single statistic per grid cell. The data points in a grid cell can be modeled as observations from a uniform random variable with the ground as the lower bound and the vegetation top as the upper bound. The first order statistic, i.e. the minimum, is an unbiased estimate of the lower bound [5]. Therefore, the data is gridded by assigning the elevation value of the minimum data point in a grid cell to that cell. In each grid cell, if the laser actually reflected off of the ground, the image value will be very close to the actual ground height. Otherwise, the image pixel will have some value above the ground surface.

The default technique for generating a DEM is generally the commercial implementation included with the LIDAR processing system. The software from different companies vary, but many are based on the same set of operations [1]. These programs start with the minimum-gridded image and threshold values based on the minimum and variance within a moving window. The thresholding is repeated several times with different size windows and different thresholds. Pixels above the thresholds correspond to vegetation pixels, and pixels below the thresholds are considered ground pixels. The elevations of the ground pixels are used to interpolate values to replace the values of the vegetation pixels.

Another useful technique for creating a DEM from LIDAR data is to grid the data and then segment different regions for processing. Anisotropic Diffusion Pyramids (ADP) can be useful for segmenting regions with steep edges, often found around man-made features. Anisotropic Diffusion was developed to maintain edge transitions while smoothing homogenous regions [2]. This is achieved by changing the amount of smoothing according to the magnitude of the local gradient. The pyramid structure allows multiscale segmentation for varying sized regions [3]. Once the regions are segmented by ADP or any number of other possible algoriths, statistics about the regions and boundaries can be used to apply specialized techniques.

Waveform Decomposition uses the actual LIDAR data values in a small region to select ground points [4]. A waveform is generated by binning the height value of points within a circular area (diameter of ~25 m). Generally, the waveform will have multiple peaks related to the ground and reflecting surfaces. Each peak is identified by fitting a Gaussian function to it. The gaussian component with the lowest mean represents the lowest reflecting surface and is assumed to be the ground surface.

Proposed Method

First, the data is gridded into a 1 m by 1 m grid based on the minimum data point in each grid cell. If the laser is able to penetrate to the ground, the minimum data value will accurately reflect the elevation of the ground surface. If all of the laser points in a grid cell reflect off of some part of the vegetation, the minimum data value will be higher than the elevation of the ground surface at that location. The grid cell value will not accurately reflect the elevation of the ground surface, and is always assumed to be higher than the actual ground elevation. The gridded image, therefore, contains pixels with values on the ground surface and above the ground in the vegetation.

The ground surface forms the lower envelope of the vegetated regions, similar to a typical amplitude modulated signal. However, the vegetated image signal is non-stationary and is not modulated with a sine wave. Next, the 2-D signal is detrended to produce a stationary signal that maintains the lower envelope below the origin. The simplest means of doing this is by subtracting the result of running an average window over the image. A simple amplitude demodulator is then applied to the stationary signal. The demodulator follows the lower envelope of the signal and removes any pixels that are above it. This result contains pixels on the ground surface and holes for vegetated pixels.

The amplitude demodulator is analogous to a circuit containing a parallel diode, resistor, and capacitor. This simple amplitude demodulator circuit follows the voltage as it decreases, and exponentially decays toward zero for large voltage increases. For the discrete image the amplitude demodulator follows the detrended signal as it decreases and thresholds values with an exponentially decaying threshold as the signal increases. This algorithm at each pixel only depends on the neighboring pixels and can easily be applied in the vertical, horizontal, and diagonal directions.

The following description explains the amplitude demodulator for a 1-D signal. The demodulator starts by storing 0, as the current ground surface, $y[n-1]=0$. The neighboring pixel, $x[n]$, is then thresholded with the value of the current ground surface, $y[n-1]$, multiplied by the decay rate, a, between 0 and 1. If the value of the neighboring pixel is above the threshold, $x[n] > ay[n-1]$, then that pixel is removed. The threshold value becomes the current ground surface at the neighboring pixel, $y[n] = ay[n-1]$, and the demodulator restarts at $n = n+1$. If the value of a surrounding pixel is below or equal to the threshold, $x[n] < ay[n-1]$, that pixel is left unchanged. The pixel value becomes the current ground surface, $y[n] = x[n]$, and the demodulator restarts at $n = n+1$. This algorithm can be applied to all eight neighboring pixels surrounding each pixel by making two passes over the data. The first pass following the path of a raster scan applies the demodulator to the three pixels below and the one directly to the right of each pixel. The second pass following the reverse path of a raster scan applies the demodulator to the remaining pixels above and to the left of the current pixel. This implementation requires storing two rows of data in memory and is very quick.

In addition to choosing any detrending parameters, the amplitude demodulator requires choosing the decay rate. The decay rate trades off removing vegetation versus following a varying ground surface. A decay rate close to one will remove more points, while a decay rate close to zero will be more likely to following a changing ground surface.

The demodulator works well over vegetated regions, however, the demodulator takes out many of the ground pixels over non-vegetated regions. This is due to the demodulator finding the lower envelope of these regions, whereas all of the points are of interest for the non-vegetated regions. Improperly removed pixel values can be added back to the final image. The pixels are added by including any values that are within a certain height distance from the pixels left in from the modulation step.

Comparison

Applying the amplitude demodulation technique required very little time to process once the minimum image was created. The minimum image required one pass over all of the LIDAR data points and has an order of $O(m)$ for m data points. The inverse distance gridded image and minimum gridded image for Bentsen State Park are shown in Figure 1. The amplitude demodulator and post-processing is on the order of $O(n^2)$ for an n by n image. Visual inspection of the result in Figure 2 show the amplitude demodulator removing more of the vegetation, however some of the natural topology along the riverbed is also removed.

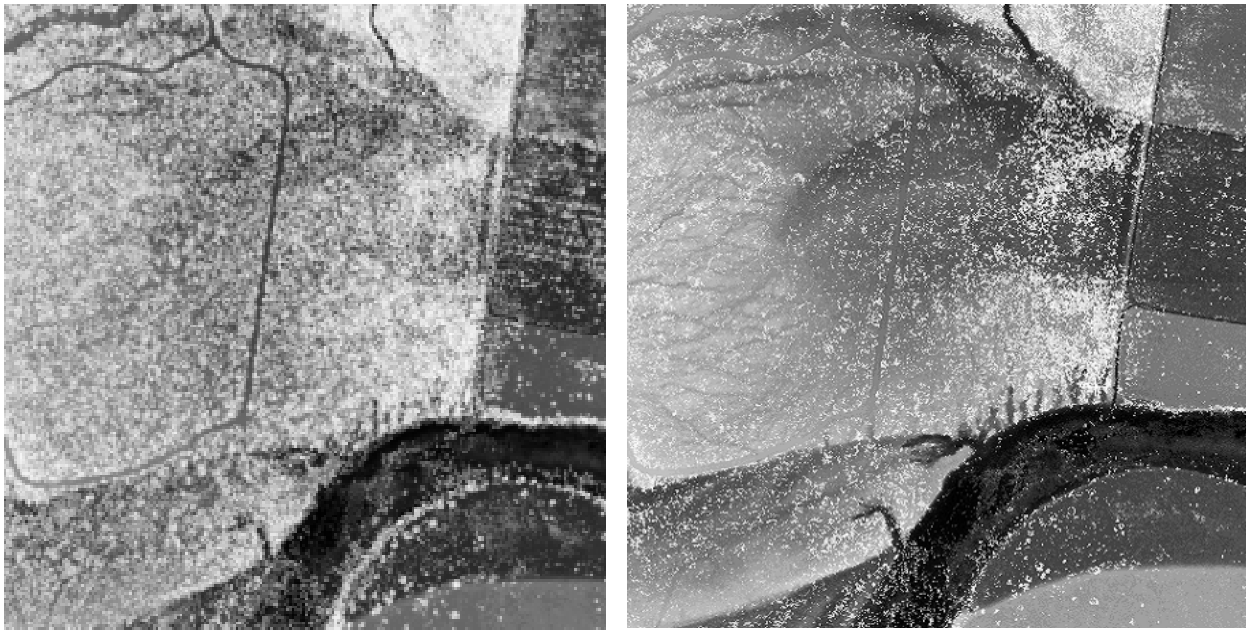

Figure 1. a) Inverse distance gridded image b) Minimum gridded image

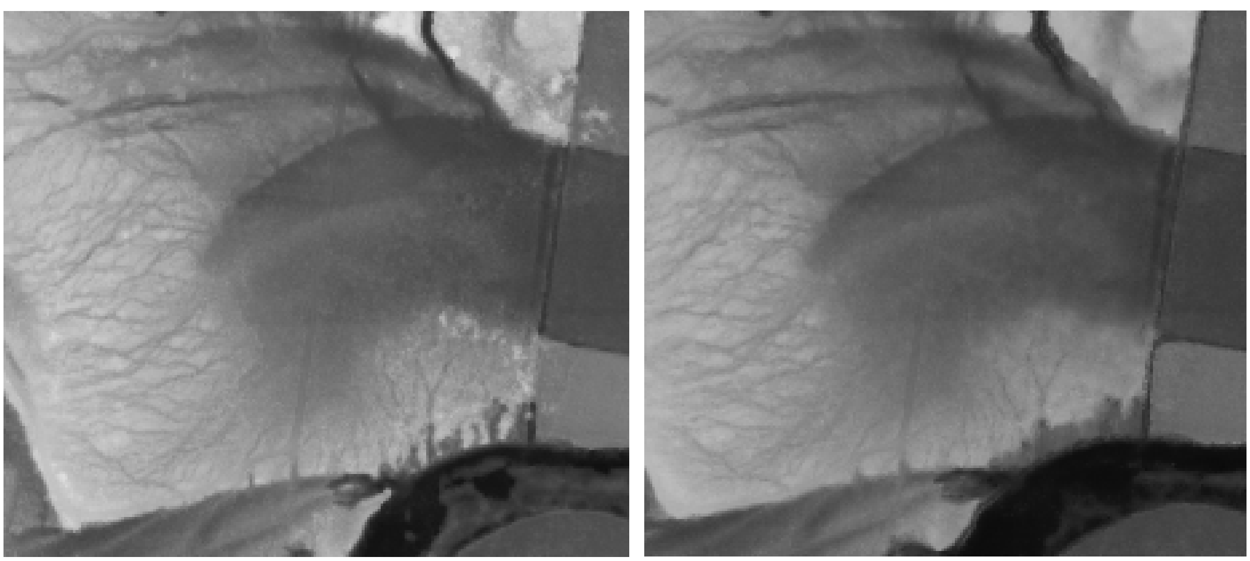

Figure 2. a) DEM from commercial package b) Amplitude demodulated DEM

The amplitude demodulation maintains the underlying ground structure useful for water flow analysis. The commercial package required 8 hours while the amplitude demodulated code took 2 hours.

Conclusion

The amplitude demodulation technique's intuitive relation to the lower envelope and ground surface and good initial results suggest it could achieve superior results to other techniques. The speed of processing is also much faster than the default windowing technique. Possible improvements include improving the detrending for dense canopy regions and segmenting the original image and applying the amplitude demodulation only to vegetated regions.

References

[1] B. Petzold, P. Reiss, and W. Stossel. "Laser scanning - surveying and mapping agencies are using a new technique for the derivation of digital terrain models," *ISPRS J. Photogramm. & Remote Sensin*g, vol. 54, pp. 95-104, 1999.

[2] P. Perona and J. Malik. "Scale-Space and Edge Detection Using Anisotropic Diffusion," *IEEE Transactions on Pattern Analysis and Machine Intelligenc*e, vol. 12, pp. 629-639, 1990.

[3] S. Acton, A. Bovik, and M. Crawford. "Anisotropic Diffusion Pyramids for Image Segmentation," *Proc. IEEE Int. Conf. Image Processin*g, Austin, TX, Nov. 1994.

[4] M. Hofton, J. Minster, and J. Blair. "Decomposition of Laser Altimeter Waveforms," *IEEE Transactions on Geoscience and Remote Sensin*g, vol. 38, pp. 1989-1996, July 2000.

[5] G. Casella, R. Berger. *Statistical Inference*, Duxbury Press, 1990.