Generate Digital Elevation Models Using Laser Altimetry (LIDAR) Data

Literature Survey

Christopher Weed

October 2000

Abstract

Laser altimetry (LIDAR) data must be processed to generate a digital elevation model (DEM), a digital map of the terrain surface.

The raw LIDAR points can be processed and then gridded to form a map of the ground surface. For example, a waveform generated by binning a set of points over a limited region can be decomposed into a set of Gaussian components. The component with the lowest mean represents the ground elevation. The data contributing to the lowest component is then gridded to form the final DEM.

The LIDAR data can be gridded to form different digital images and then the images are processed to create a map of the ground surface. A common and simple technique uses varying window sizes and thresholds at each pixel of the image. This can produce good results when properly tweaked but with very little statistical justification. Anisotropic Diffusion Pyramids' (ADPs) ability to preserve edge information while smoothing homogenous regions can be useful for segmenting buildings or areas of vegetation which need to be removed.

1 Introduction

Laser altimetry (LIDAR) systems have the ability to generate high-resolution maps of the ground surface, man-made features, and/or vegetated areas. Depending on the application, information about one or all of these features can be important. However, the data produced by the instrument is only a set of 3-D X,Y,Z points. (Some systems also produce intensity information) Each one of these data points could have been produced from the laser reflecting off of a building, the ground, or some part of a plant. Generally, the LIDAR system will produce a first and last return pulse for each laser pulse that can correspond to the laser penetrating through some of the vegetation. If the laser pulse hits the ground or a building directly, the first and last return pulse will be nearly identical. For generating a digital elevation model (DEM), the last return points are the ones of greatest importance.

The most important issue for generating a DEM is height accuracy. Many simple algorithms work well for simple land features and vegetation, but as the terrain, buildings, and plant types vary and intersperse, results will suffer.

As with any computer program speed and memory is also important. The most important differentiating factor about speed and memory is whether the algorithm works directly with the raw LIDAR data, which can run into the hundreds of millions of 3-D points, or only needs gridded images generally less than fifty million pixels. Computer resources are also dependent upon the necessary image resolution, which is generally based on 1 m by 1 m pixels. The code does not need to run in real-time since processing the data prior to generating the DEM is not done in real-time.

2 Background

2.1 Default technique

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]. They grid the data by assigning the elevation value of the minimum data point in a grid cell to that cell. This creates an image of elevation heights which estimates the ground surface. In each grid cell, if a data point actually reached 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. Each pixel is then thresholded by some amount above the minimum value of a large moving window. The threshold is generally based on some type of variance information in the window. The thresholding is repeated several times with smaller windows and smaller thresholds. The user will often have to tweak the window size and threshold parameters to achieve the result they want. A small window size will fail to remove points on large buildings, and a large window size will smooth the terrain and remove small discontinuities on the ground. A low threshold will also remove natural variation in the ground surface, while a high threshold will fail to adequately remove all non-ground points. This technique works best for flat terrain, and would require different parameters for different terrain types. A variation on this technique uses the variance information of data points in a grid cell instead of a larger window.

After this process is finished, the image will contain some data values and many pixels that were removed by thresholding. The thresholded pixels can be filled in by a number of interpolation techniques, but we found the natural neighbor gridding algorithm to achieve the best results. This program creates a mesh of triangles from the data and projects it down onto the new grid. Whichever triangle the center of the grid cell falls in, that grid cell is determined by an inversed distance weighting of the data at the triangle vertices.

Once the gridding is finished, the image may need to be manually inspected alongside the first gridded image for possible misclassification problems as shown in the following figure. A shaded relief image is often useful for finding problem areas. The areas generally most likely to be incorrect are dense forests with low penetration, or small ground discontinuities.

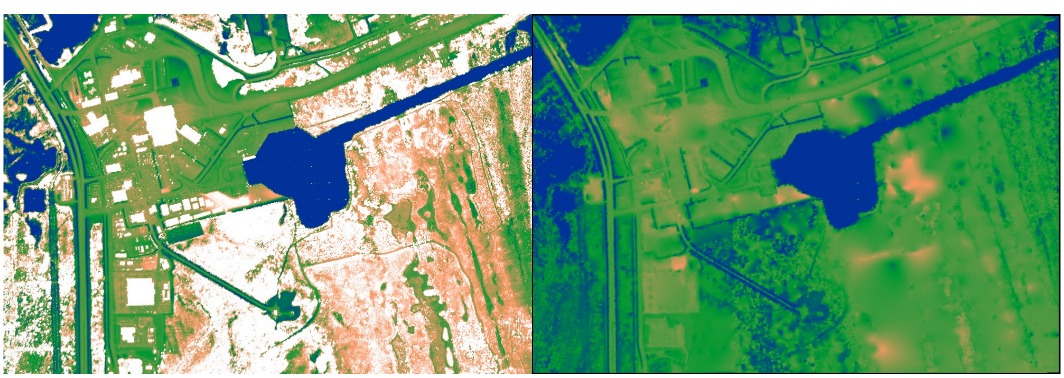

Figure1. Map before and after windowing technique

2.2 Anisotropic Diffusion Pyramids

Anisotropic Diffusion Pyramids are just one of the many segmentation techniques that can be used to distinguish the ground from man-made features. The most distinguishing factor about buildings and most man-made features is steep edges. Anisotropic Diffusion was developed to maintain edge transitions while smoothing homogenous regions [2].

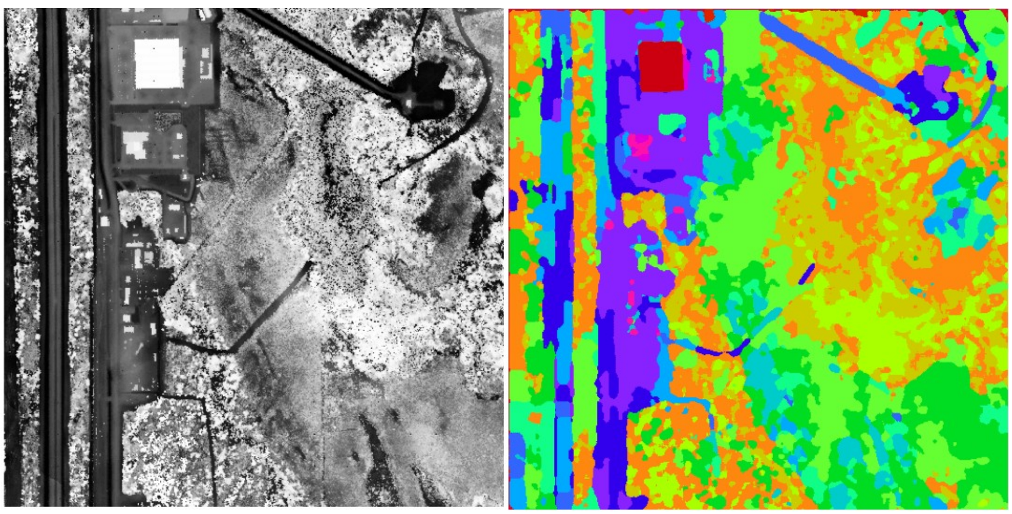

Figure 2. Image before and after Anisotropic Diffusion

This is achieved by varying the smoothing constant according to the magnitude of the local gradient. The pyramid structure allows multiscale segmentation for varying sized regions [3]. Using edge information to determine building regions can remove some of the problems with window size determination for different features in the image. The segmented buildings are thresholded before applying the previous window method. This is just one example of a family of techniques where the DEM is developed by segmenting regions and applying different techniques to different regions.

2.3 Waveform Decomposition

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 meters). Generally, the waveform will have multiple peaks related to the ground and reflecting surfaces. Each of the peaks will be differentiated by fitting a Gaussian function to them. The Gaussian component with the lowest mean represents the ground surface for that circular area. The other Gaussian components will represent significant reflecting regions in the vegetation or on a man-made object.

This approach is made up of a series of estimation problems, beginning with estimating the number of Gaussian components. This is based on the number of inflection points in a smoothed version of the waveform. The positions and half-widths of the Gaussian components are estimated by the location of the inflection points. Next, a non-negative Least Squares Method is used to determine the amplitude of each of the Gaussian components. To avoid modeling the noise the Gaussian components are ranked and only a limited number are used to maintain a certain level of error between the model and the data. The estimates serve as the initial values to the Levenburg-Marquadt optimization algorithm which minimizes the model error. The Levenburg-Marquadt method is a very popular nonlinear LSM which trades off the advantages of both the steepest descent and Newton's method. If the error is not smaller than a set amount, more Gaussian components are used to estimate the data.

This algorithm can be adapted to use non-Gaussian components to get a better fit to the spikes in the waveform. This method seems to find good initial values for achieving a fairly quick and accurate minimization. However, the minimization must be performed many times for a large data set. This technique will only work if enough ground shots are achieved to produce a spike for each region.

A region with low ground cover and changing ground surface would likely cause difficulty in differentiating the ground from the vegetation. This technique has the added benefit of producing Gaussians at regions that have reflecting structure that would likely be useful for analysis of vegetation properties.

3 Conclusion

The current set of approaches for developing DEMs from LIDAR data leaves a lot to be desired. For trivial and simple data sets, variations on the default windowing technique can yield good results. However, the level of user interaction is very high and it contains little statistical justification. Adding the Anisotropic Diffusion or other stages to this approach may improve the results, but ultimately a more rigorous model will be necessary to continue improving. The model of these computations are generally on the order of O(n^2) for an n by n image. Waveform decomposition has a better statistical foundation, but calculating such a high number of minimizations seems excessive. Possibly applying this technique to the gridded data values would cut computation while achieving a similar result. Also waveform decomposition does not make use of the information between neighboring circular regions. The computational complexity for decomposition is based on the number of data points, for example each iteration of the minimization would calculate the error, for a overall complexity of O(m^2) for m data points. None of these techniques is well suited for changing terrain. They all assume locally or regionally flat areas.

I plan to concentrate on modeling the ground with a nonstationary process, maintain a low level of user interaction, and base the algorithm on the gridded data to keep computation low. One idea I am considering is using a Kalman filter to find the mean of a row of the data signal. This mean is subtracted from the data to get a stationary signal. The bottom of the waveform of the stationary signal represents the ground surface. To select only values on the bottom surface I want to implement something similar to an amplitude demodulator that would find a signal from the negative amplitude. Another idea I am considering is using estimates of the top of the vegetation canopy along with estimates of the ground values. Using the correlation of these two signals I want to improve the estimation of the ground values. Generally, I think the best model of most of these images is a stochastically varying signal with a variance that goes through step changes, and finally with buildings interspersed.

Bibliography

[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 Sensing*, 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 Intelligence*, 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 Processing*, Austin, TX, Nov. 1994.

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