FILTERING OF LASER ALTIMETRY DATA USING A SLOPE ADAPTIVE FILTER

G. Sithole
Department of Geodesy, Faculty of Civil Engineering and Geosciences
Delft University of Technology
The Netherlands
g.sithole@citg.tudelft.nl

Commission III, Working Group 4

KEY WORDS: Laser, DEM, Filtering

ABSTRACT

A point set obtained by laser altimetry represents points from not only the ground surface but also objects found on it. For civil works applications points representing the surface of non-ground objects have to be removed from the point set in a filtering process. This paper describes modifications made to an existing "slope based" filtering algorithm, and presents some results obtained from the use of the filter. The "slope based" filter operates on the assumption that terrain slopes do not rise above a certain threshold, and that features in the data that have slopes above this threshold do not belong to the natural terrain surface. However, this assumption limits the use of the filter to terrain with gentle slopes. To overcome this limitation, the filter was modified in manner that the threshold varies with respect to the slope of the terrain. The results of tests carried out using the modified filter confirm that the modification reduces the number of Type I errors (ground points in steep terrain are not filtered off). Further numerical comparison of the filter output with a reference data set for the same site (obtained photogrammetrically) show that the filter generates relatively minimal Type II errors. The output of the modified slope filter was also compared with the output from a filtering found in the commercial software package, "Terrascan".

1 INTRODUCTION

Airborne laser altimetry has gradually become a mainstream tool for abstracting high accuracy and high-density digital terrain surfaces. However, the point set obtained by laser altimetry represents points from both the ground surface and objects found on the ground surface. For civil works applications points representing such objects have to be removed from the point set.

In the filtering process points classified as non-ground are discarded. The large number of points in a laser data set necessitates a high degree of automation in the classification of points. Some filtering techniques that have been developed are described in Kraus and Pfeifer 1998, Vosselman 2000, Axelsson, 2000, Elmqvist, 2001. Most of the criteria used in classifying points have focused on simple geometric characteristics of a point relative to its neighborhood. To further improve the accuracy of classification some filters iterate the classification process. Other classifiers work on the premise that ground points and non-ground points in the laser scanner data set are stochastically separable.

The filters do not work under all circumstances, and efforts have been put into improving the filters (e.g., Schickler and Thorpe, 2001). This paper describes the modification of a slope-based filter with a view to improving the performance of the filter in steep sloped terrain.

The slope-based filter developed by Vosselman (2000) uses the slope of the line between any two points in a point set as the criteria for classifying ground points. The technique relies on the premise that the gradient of the natural slope of the terrain is distinctly different from the slopes of non-terrain objects (trees, buildings, etc.). Any feature in the laser data that has slopes with gradients larger than a certain predefined threshold therefore does not belong to the natural terrain surface. However, this assumption limits the use of the filter to terrain with gentle slopes. To overcome this limitation, the filter was modified so that the threshold varies with respect to the slope of the terrain.

In the first part of this paper, the modifications to the slope-based filter are discussed. The modified filter was implemented and tested using the Vaihingen test field. The preliminary results of the tests are presented in the second part of the paper. Finally, the paper concludes by discussing the implications of the results for future filtering strategies.

2 THE MECHANICS OF THE FILTER

2.1 Slope Based Filter

The basic mechanics of the slope-based filter is illustrated in Figure 1. The vertex of an inverted cone sweeps under each point in the point-set to be filtered. Wherever the cone cuts the point set, then the point at the vertex of the cone is filtered off. In Figure 1(a) the point, p_i, at the vertex of the cone is not filtered off because the cone does not cut the surface. In the implementation of the filter, an inverted bowl whose shape is defined by a probabilistic function designed to minimize classification error replaces the cone. For simplicity, a cone is considered here.

Another way to visualize the method is shown in Figure 1(b). The curved surface shown in the Figure represents the slope of the vectors from the point, p_i, to every other point on the surface. From here, onwards this surface will be referred to as the *point-slopes* surface. The plane is the negative of the absolute value of the gradient of the cone's generators. If the plane cuts the *point-slopes* surface point, p_i, is filtered off. In Figure 1 (b) the *cutoff* plane does not cut the *point-slopes* surface so point, p_i, is not filtered off. In the implementation the *cutoff* plane is not planar but rather curves upwards the further it gets from point, p_i. The curvature is determined by a probabilistic function derived from a training data set.

2.2 Modification

The main parameter of the slope-based filter is the gradient of the cone's generators. Adjusting this gradient has the effect of moving the *cutoff* plane up or down. The steeper the gradient the lower the *cutoff* plane and vice-versa. If the gradient of the cone's generators is such that the *cutoff* plane cuts the *point-slopes* surface point, p_i, is filtered off. This is illustrated in Figure 2. The classifier is expressed as:

$$\forall p_j \in A : hp_i - \Delta h(d(p_i, p_j), m) \le hp_j \qquad (1)$$

Where: p_j is a point in the data set ($p_i \neq p_j$).
hp_i and hp_j are the heights of p_i and p_j respectively.
hp_i - $\Delta h(p_i,p_j)$ is the height of a point directly above or below p_j and on the lateral surface of the cone whose vertex is located at p_i.
m is the absolute value of the gradient of the cone's generators. The negative value of *m* is the height of the cutoff plane.
A is the set of laser points to be filtered in order to extract the DEM.

The next parameter of the filter is the radius of the base of the cone. This parameter defines the operating the range of the filter. In the examples shown in Figure 1 and Figure 2, the operating range of the filter is infinite; the *cutoff* plane extends to infinity. In the implementation, the scope of the *cutoff* plane was restricted. This is because the *point-slopes* surface tends to flatten out the further one moves away from point, p_i, thus reducing the effectiveness of the filter.

The classifier in equation 1 works well if the slope of the terrain is gentle. However, in steep sloped terrain discriminating between the ground surface and features such as buildings and vegetation becomes difficult. To overcome the problem the classifier was modified in such a way that the *cutoff* plane shifts up or down with respect to the position of the cone in the terrain. In the original filter, the *cutoff* plane is held fixed for every point in the point set.

The *cutoff* plane is tuned to the slope of the terrain at point, p_i. Phrased differently as the cone sweeps underneath each point in the point set its slope changes in tune with the maximum slope of the terrain at point, p_i.

The classifier given in equation 1 is now be expressed as

$$\forall p_j \in A : hp_i - \Delta h(d(p_i, p_j), m_i) \le hp_j \qquad (2)$$

Where: m_i is the height of the cutoff plane.

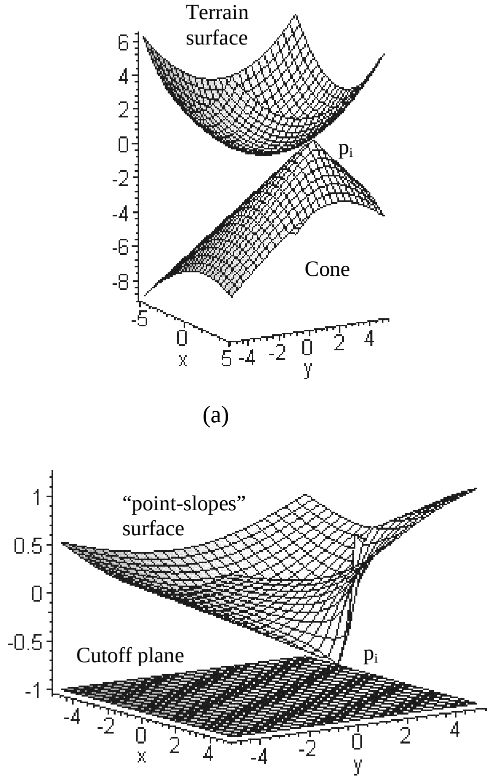

Figure 1 Mechanics of the original filter.

Choosing the value of m_i. Setting the value of m_i equal to the maximum slope of the terrain at the point at which the cone's vertex touches the surface of the terrain is not enough. Figure 3 shows why. Shown in the Figure is a terrain x-section. The x-section of a cone is also shown. The cone sweeps underneath the surface, with its vertex always in contact with the surface. The cone is shown at three points on the surface. At each point $|m_i|$ is set equal to the maximum slope of the terrain at that point. In Figure 3(a), the cone touches the surface at a point where the

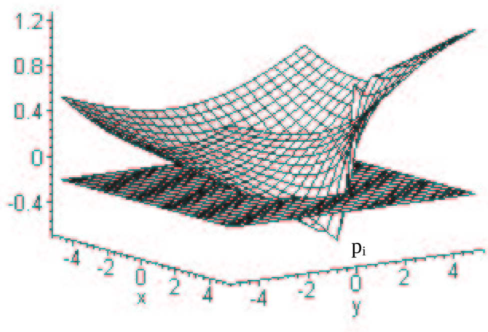

Figure 2 "Cutoff" plane cutting the "point-slopes" surface. Point, p_i is filtered.

surface is concave. The cone does not cut the surface and the point (at the vertex) is correctly accepted as a part of the surface. However, on convex slopes as shown in Figure 3(b), the cone cuts the surface and the classifier fails. In this case the slope is multiplied by a constant factor (>1). The filter can also fail in gently sloped terrain as shown in Figure 3(c). A characteristic of the terrain in this area is the small amplitude and large frequency of the surface (exaggerated in Figure 3(c)). In these areas, where the value of $|m_i|$ is small (flat ground) points will be incorrectly rejected. To overcome this problem a minimum threshold is set for m_i.

To summarize, m_i has to be pre-multiplied by a constant factor (for convex slopes) and then thresholded (for flat terrain). The classifier is now expressed as:

$$\forall p_j \in A : hp_i - \Delta h(d(p_i, p_j), m_i, s_{mi}, m_{\min}) \leq hp_j \tag{3}$$

Where: m_i is the maximum slope of the terrain at the point the cone's vertex touches the surface,
s_{mi} is a predefined factor by which m_i multiplied
m_{min} is the minimum threshold for $s_{mi}*m_i$.

The DEM is expressed as:

$$DEM = \left\{ \begin{array}{l} p_i \in A \mid \forall p_j \in A : \\ hp_i - \Delta h(d(p_i, p_j), m_i, s_{mi}, m) \leq hp_j \end{array} \right\} \tag{4}$$

Slopemap. To tune the cutoff plane to the surface of the terrain, a rough model of the terrain is needed. This model was generated in the form of a minimum height image in which the pixel values are local height minima (assumption: the minimum value in any neighborhood belongs to the terrain). A slopemap image was generated from the minimum height image. The values from the slopemap were then used to tune the *cutoff* plane.

The slopemap image was then dilated. This was done because the minimum height image is a discrete representation of the terrain, and as a result of this, there can be ambiguity in assigning slope values to points located at the edges of pixels. In this way, dilation ensures that the position of test points relative to the current point does not affect the classifier.

3 TEST DATA

The laser data used in the study is from the Vaihingen test field (part of the OEEPE data set). Features found on the site are urban areas, forests, hills, a river and a quarry. The outstanding feature of the Vaihingen data set is the data gaps (the result of a flight planning error) and the presence of large outliers (points with very low or very high heights). Because of the way in which the slopemap was generated the slope values calculated at the edges of gaps are very large, which makes the filter very generous at the edges of gaps. This was corrected by setting all the extremely large values in the slope map to zero.

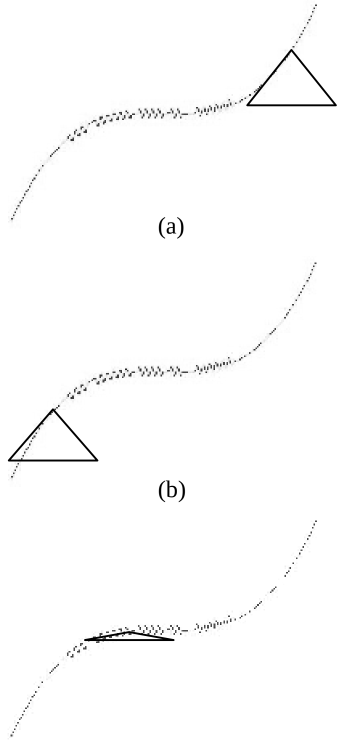

(a) (b) (c)

Figure 3 Mechanics of the modified filter. The slope of the cone's lateral surface adjusts to the slope of the terrain. If the slope of the cone's lateral surface is set equal to the slope of the terrain, the filter fails in cases b and c.

Using a TIN generated from the minimum height image would have avoided some of the problems associated with the slopemap. However, the slopemap was opted for, because it has fewer computational overheads.

Low lying outlying points had to be removed using a maximum height difference function (Vosselman and Maas, 2001). High outlying points were removed during normal filtering.

For ground truth a reference data set composed of 2428 points obtained by photogrammetry, from 1:13000 photography was used. The points are spread out in a regular grid pattern. The grid spacing is approximately 25m. In general the reference data is estimated to have a standard deviation of 0.25m – 0.3m. However, in some areas (with vegetation or bad texture) the standard deviation was estimated to be as much as 0.5m. There are also gaps in the reference data, located in built-up areas and areas of dense vegetation and the results presented in Table 1 should be read with this in mind.

4 RESULTS

4.1 Comparison with reference data

The data set was filtered using different parameter settings, and the filtered data were compared against the reference data. The comparison was achieved by generating a TIN from the filtered data and extracting corresponding heights for the reference points

Table 1 Filtering results (statistics in the last four columns are based on the sample after cutting off outliers)

	Min Slope	Slope Factor	Sample Count after Cutoff	RMS (m)	Mean (m)	Std. Dev. (m)	Median (m)
Slope	0.10		1226	0.2820	0.17	0.2263	0.17
	0.20		1238	0.2665	0.14	0.2241	0.14
	0.30		1241	0.2701	0.14	0.2300	0.14
Slope Adaptive	0.00	1.00	1673	0.3011	0.17	0.2465	0.17
	0.00	1.25	1718	0.2978	0.16	0.2485	0.16
	0.00	1.50	1718	0.2978	0.16	0.2485	0.16
	0.00	2.00	1718	0.2978	0.16	0.2485	0.16
	0.15	1.00	1718	0.2978	0.16	0.2485	0.16
	0.15	1.25	1789	0.2924	0.15	0.2494	0.15
	0.15	1.50	1789	0.2924	0.15	0.2494	0.15
	0.15	2.00	1799	0.2911	0.15	0.2507	0.15
	0.30	1.00	1718	0.2978	0.16	0.2485	0.16
	0.30	1.25	1799	0.2921	0.15	0.2501	0.15
	0.30	1.50	1799	0.2921	0.15	0.2501	0.15
	0.30	2.00	1798	0.2918	0.15	0.2515	0.15

from this TIN. Shown in the charts in Figure 4, are the differences (errors) of the heights in the TIN from their correspondences in the reference data.

The main characteristic of the results of the filtering where differences in a band of ±1m. However, there were a few differences (tightly bunched together) exceeding +1.5m, Figure 4(a). A visual check showed that most of these outliers were from the same area. A positive difference here means that the reference data is higher than the laser data. Considering that the reference data is older than the laser data, an explanation for the large outliers could be that there might have been an excavation in these areas after the aerial photography. Because of this, all differences beyond ±1.5m were discarded (which resulted in the loss of about 80 points). The data provided in Table 1 and the distribution shown in Figure 4(b) are from the set of differences after the ±1.5m outliers were discarded.

Table 1 shows statistics for the differences. What is noteworthy is that although the slope adaptive filter gives a high point count, the standard deviation of the filtered points does not change much. This indicates that the modifications are delivering the desired results in steeper slopes without allowing non-ground points to pass through the filter (Type II errors).

In Figure 5(a) and Figure 5(b) it can be seen that the slope adaptive filter is most effective at a minimum slope of 0.15 and a slope factor of 1.25. Using larger minimum slopes and slope factor values gives lower gains and will result in more Type II errors.

A problem with the reference data is that it generally represents areas in the terrain that are not covered by dense vegetation or human artifacts. In such areas filters have a small chance of failure. Because of this, reference data is not very useful for evaluating filters, unless the data coverage extends to built-up and vegetated areas. Therefore, the results in Table 1, Figure 4 and in

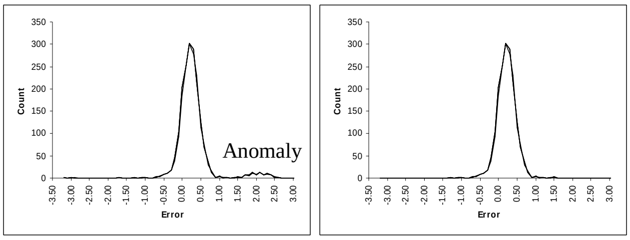

(a) Distribution with outliers (b) Distribution without outliers

Figure 4 Distribution of height differences between reference and filtered data.

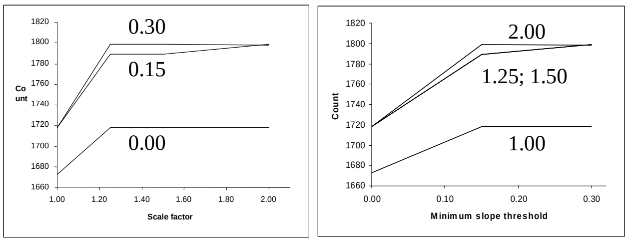

(a) Variation of slope factor (b) Variation of min. slope

Figure 5 Effect of the variation of minimum slope with respect to slope factor.

Figure 5 cannot be extended to areas covered by dense vegetation and human artifacts. For this reason, visual comparisons were done, and the results are described in the next section.

4.2 Visual Comparison

In Figure 6, the slope-based filter and the slope adaptive filter are visually compared. Figure 6(b) through Figure 6(f) are images generated from the filtered data. The areas that have pass through the filter are shown in black or gray (in the case of Figure 6(b)).

Figure 6(b) shows the result of the slope-based filter, used with a minimum slope of 0.3. The areas shaded in light gray represent points filtered by the slope-based filter when used with a minimum slope of 0.1. There are three sites (1, 2 and 3 in Figure 6(a)) were the terrain slopes are steeper than 0.3. In these sites, all points are filtered off. The minimum slope could have been increased but this would have generated many Type II errors. At site 4, there is a quarry, and here too the sides of the terraces are filtered off. The modifications to the slope-based filter are meant to avoid these incorrect rejections of terrain points on steep slopes.

Figure 6(c) and Figure 6(d) shows the same area filtered with the slope adaptive filter using a minimum slope of 0.0. Two slope factors have been used (1.0 and 1.5). Very steep slopes at sites 1 to 4 have not been filtered off. However, this gain has been at the expense of the filter’s performance in urban areas. The reason for the loss of performance in urban areas is the size of buildings. The operating range of the filter is often smaller than the size of a building and consequently the slope map adapts to the roofs of the buildings. Because of this the central part of large buildings are not filtered off.

The anticipated failure of the filter in terrain with gentle slopes (Figure 3(c)), when using a minimum slope of 0.0 is seen in the appearance of furrows in Figure 6(c). Because the pixel size of the

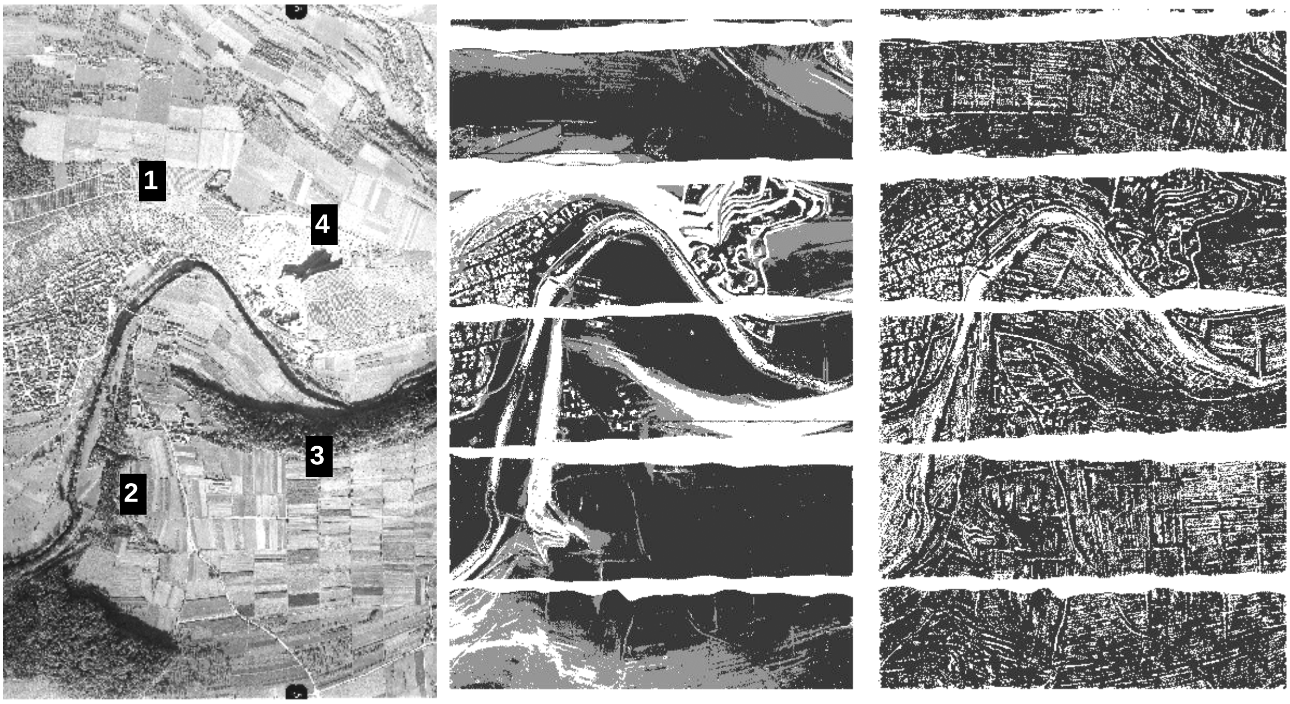

(a) Photograph (b) m =0.3 (c) m_{min} = 0.00, s_{mi} = 1.00

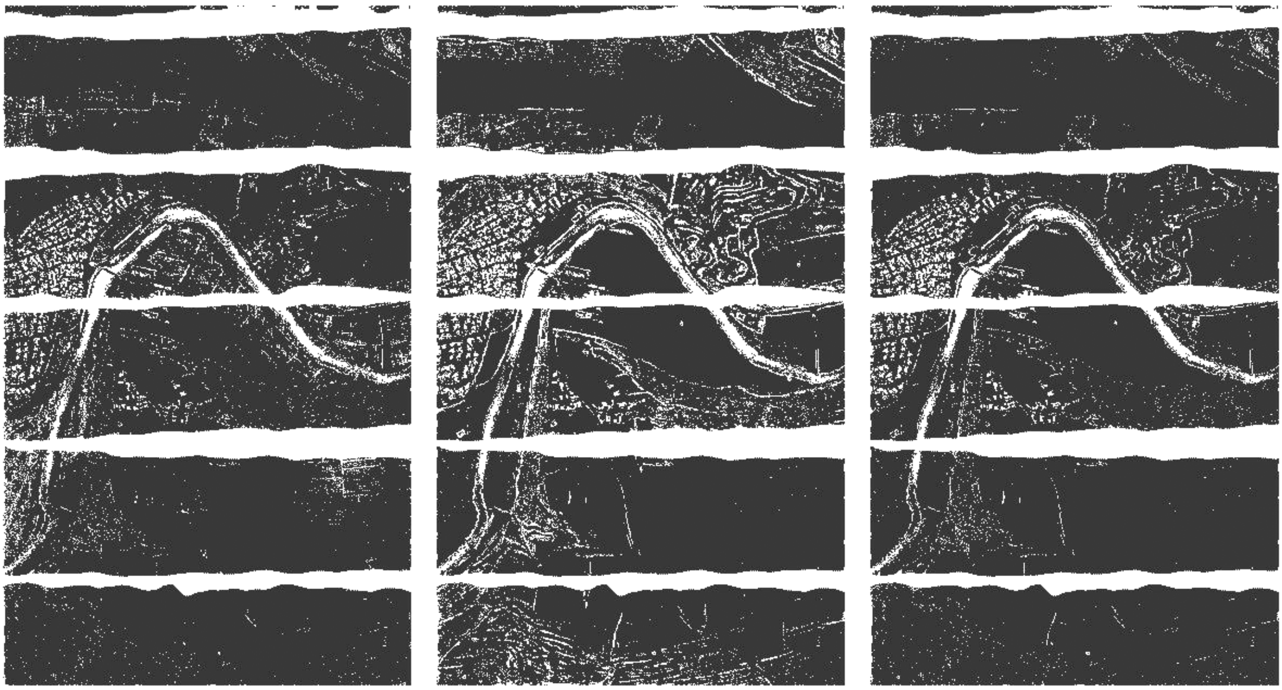

(d) m_{min} = 0.00, s_{mi} = 1.50 (e) m_{min} = 0.15, s_{mi} = 1.00 (f) m_{min} = 0.15, s_{mi} = 1.50

Figure 6 Comparison of filtering results

slopemap is larger than the distance between the furrows, the slope values in the slopemap will also be small (in gently sloped terrain). The problem is corrected by using a minimum slope of 0.15 as in Figure 6(e) and Figure 6(f). The effect of not using a minimum slope can also be seen when comparing figure Figure 6(d) and Figure 6(f). Although a slope factor is used, ground points are incorrectly rejected (compare right side of second strip from bottom).

A drawback of dilating the slope map is increased Type II errors at the foot of steep slopes. Because of the dilation, slope values for some of the gentler slopes will be very high (Figure 7). This effect is further worsened when the slope values are multiplied by a slope factor. It is not easily noticeable, but comparison of Figure 6(b) with Figure 6(e) shows that the modifications to the slope-based filter results in Type II errors on riverbanks and in urban areas. This effect will become more evident when examining the profiles in Figure 9.

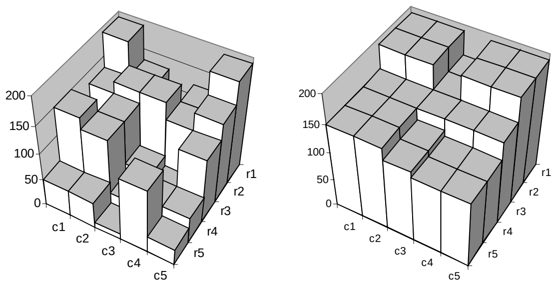

Slopemap　　　Dilated slopemap

Figure 7 Slopemap (exaggerated for effect)

4.3 Comparison with Terrascan Filter

The filtering algorithm used in Terrascan (Axelsson, 2000) starts from a sparse TIN and iteratively refines it to the laser point set. At every iteration, points are added to the TIN if they are below data derived thresholds. The data derived thresholds are distances to TIN facets and angles to the facet nodes.

Figure 7 shows the results of filtering using the Terrascan filter using an iteration angle of 2 degrees and 8 degrees. The Terrascan filter was designed for urban environments and this can be seen in Figure 8. Most buildings have been filtered off. The filter also did well on the wall of the quarry. There is a very big building at the base of the quarry. The Terrascan filter completely removed this. However, it also removed surrounding terrain. The slope adaptive filter on the other hand fails to remove this building (see Figure 6, site 4).

Another interesting aspect of the Terrascan filter is its response to the gaps between the strips. When a small iteration angle is used (in Figure 8(a) it is 2) in some places, the filter erodes the edges of the strips. Increasing the iteration angle solves this problem as shown by the gray areas in Figure 8(b).

The images in Figure 6 and Figure 8 are useful for understanding the response of the filter in relation to the terrain coverage and morphology. However, they do not provide a means to examine the filtered points. For this purpose, profiles were generated, Figure 9 and Figure 10.

4.4 Profiles

Figure 9 shows three terrain cross sections. In the Figure an unfiltered section is followed by its corresponding slope adaptive filtered result (minimum slope = 0.15, slope factor = 1.0). Figure 9(a) shows a steep slope covered with high and low vegetation. The adaptive filter successfully removes high and low lying vegetation without eroding terrain points.

Figure 9(b) shows a situation in which the filter both fails and succeeds. The terrain is covered by vegetation and human artifacts. The vegetation is successfully filtered off, but not all the buildings. A slopemap with a resolution of 10m was used and because of this, buildings (size>10) are captured in the slopemap.

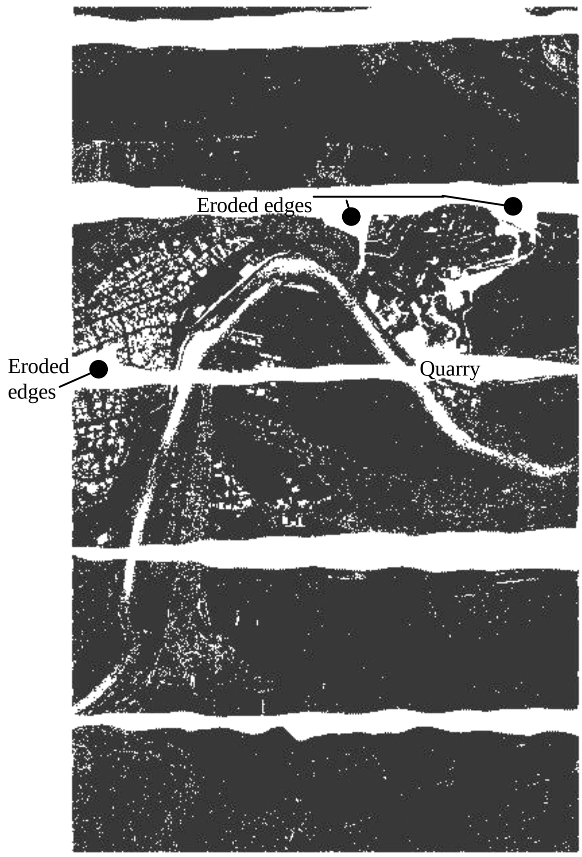

(a) Iteration angle = 2

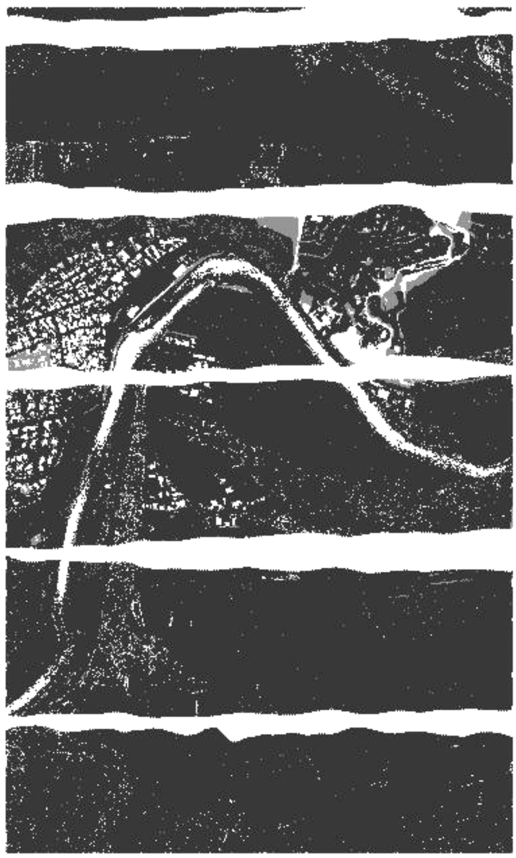

(b) Iteration angle = 8

Figure 8 Difference images generated using the result of the Axelsonn filter.

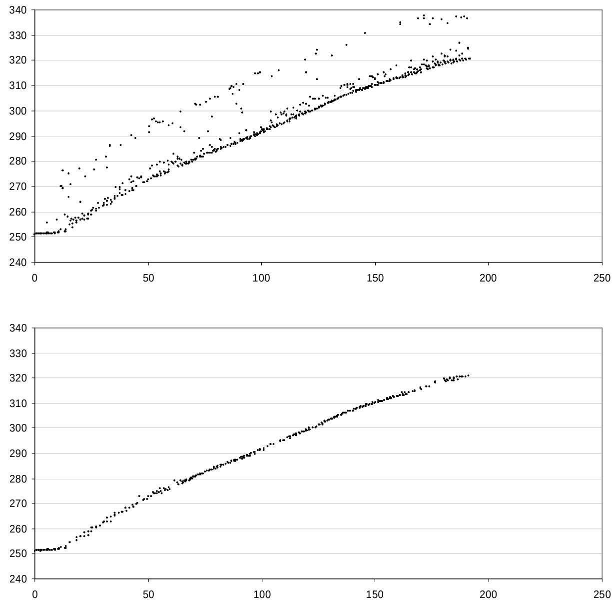

(a)

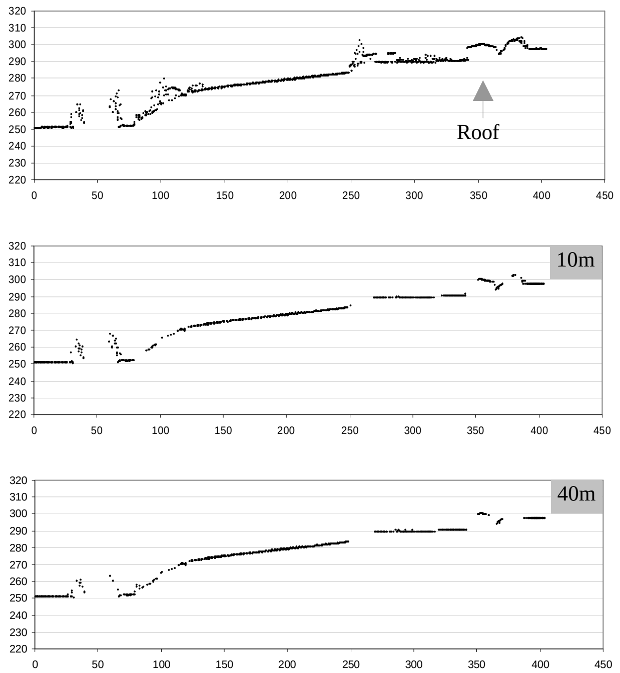

(b)

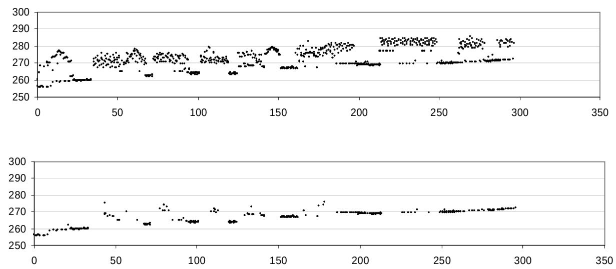

(c)

Figure 9 Unfiltered vs. filtered using slope adaptive filter

Reducing the resolution of the slopemap would result in buildings being filtered off (Figure 9 filtered using a slopemap with a pixel size of 10m and 40). However, this would result in over generalized slopes that may cause the incorrect rejection of terrain points. It can also be seen in Figure 9(b) that vegetation on riverbanks is not effectively filtered. Also, in Figure 9(c) is shown data for an area where there is low vegetation penetration. Here the filter does poorly. In both cases, the low vegetation penetration results in a slopemap that is adapted to the lowest points in the measured vegetation rather than the terrain.

Figure 10 shows more terrain profiles. Here an unfiltered section is followed by one filtered using the Terrascan filter (iteration angle = 6 degrees) and another using the slope adaptive filter (minimum slope = 0.15, slope factor = 1.0). The Terrascan filter completely removes all the vegetation if Figure 10(a). However, in the process of filtering the Terrascan filter thins the terrain points. The slope adaptive filter is not entirely successful, it removes most of the vegetation but it still leaves a few behind. A test on a few samples showed the average ratio between the slope adaptive filtrate and that from the Terrascan filter to be around 6:1.

Figure 10(b) shows the profile of a quarry. In this profile, there is very little vegetation and no buildings. What is of interest is the performance of the filters in very steep terrain. Both filters appear to capture most of the characteristics of the terrain. However, the Terrascan filter, filters off the steepest slope at the highest point in the profile. Moreover, the Terrascan filter still thins the data. Because the slope is not vertical, it is partially captured in the slopemap, and the dilation and multiplication ensures that the steep slope is not filtered off by the slope adaptive filter.

4.5 Summary

Pros. The filter does not remove steeps slopes (unlike the slope-based filter). While it is not able to correctly filter all data points, it does not thin the terrain points. This is good because it means that slope adaptive filter can be supplemented by another classifier to overcome the filter's shortcomings. If the operating range of the filter is kept low (relative to the resolution of the laser scanner data) mounds, hills, etc., are not filtered off (a problem with some filters, as noted by Huising and Gomes Pereira, 1998).

Cons. Classification is point to point. Systematic errors in a point (e.g., low points) can cause the incorrect rejection of valid terrain points. Furthermore, because the slopemap is discrete, it gives rise to side effects, which in turn have to be corrected themselves. The filter still faces problems in filtering off large buildings and areas with low vegetation penetration (forests, vegetation on riverbanks, etc.).

5 CONCLUSION

Modifications made to the slope-based filter correct the problem of Type I errors in steep sloped terrain. However, the price for this has been a small increase in the number of Type II errors and an increase in the number of filter parameters. This is unavoidable. A side effect of the increased number of parameters is that fine tuning the parameters becomes difficult, especially since it is not known if the parameters of the filter will respond similarly under different terrain conditions. Urban and vegetated areas were found to responded differently to the filter. This suggests that during

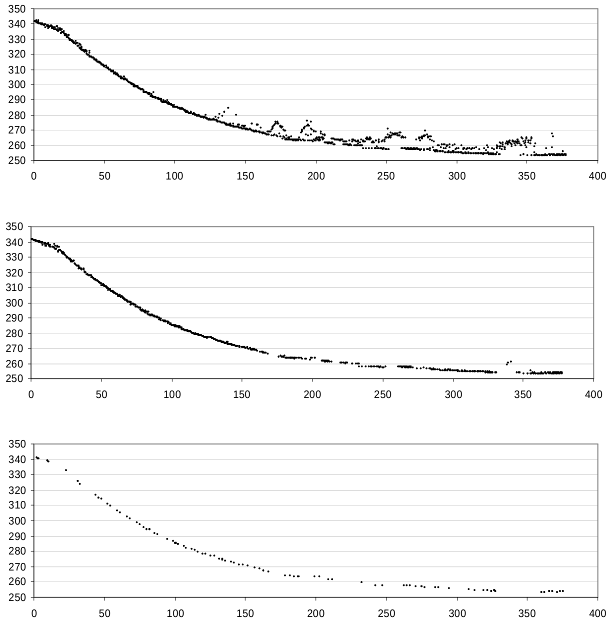

(a)

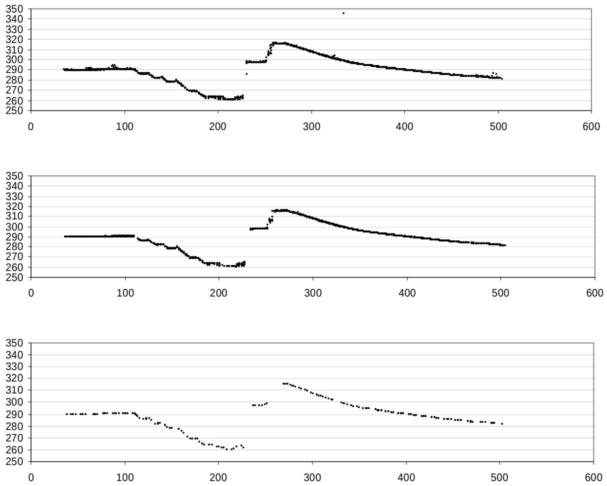

(b)

Figure 10 Unfiltered/ Slope Adaptive filter/ Terrascan filter

filtering urban and vegetated areas need to be separated (e.g. Oude Elberink and Maas, 2000) and different parameters applied for urban and vegetated areas. In the long term a filtering strategy based on using different classifiers for different terrain, coverage's may make more sense.

Another aspect of the filtering process that still needs attention is the measurement or prediction of the accuracy of filters. Currently the performance of filters is mostly reported using the rms of the filtrate or the number of points correctly classified by a filter (using a test site). The numerical and graphical comparisons presented here show that while useful, these statistics are not fully representative of the performance of the filter unless the reference data that is used is also representative of the terrain. Because of this, the approach that will be taken for assessing the accuracy of filters in future will be to determine in each data set where the filter will likely fail based on the characteristics of the filter.

REFERENCES

Axelsson, P., 2000: "DEM Generation from Laser Scanner Data Using Adaptive TIN Models. In IAPRS. Vol. 33. part B4/1, pp. 110-117.

Axelsson, P., 2001: "Ground estimation of laser data using adaptive TIN-models". Proceedings of OEEPE workshop on airborne laserscanning and interferometric SAR for detailed digital elevation models 1-3 March 2001, paper 5 (11 pages). Royal Institute of Technology Department of Geodesy and Photogrammetry 100 Stockholm, Sweden.

Elmqvist, M, 2001: "Ground Estimation of Lasar Radar Data using Active Shape Models". Paper presented at the OEEPE workshop on airborne laserscanning and interferometric SAR for detailed digital elevation models 1-3 March 2001, paper 5 (8 pages). Royal Institute of Technology Department of Geodesy and Photogrammetry 100 Stockholm, Sweden.

Huising E.J. and Gomes Pereira, L. M., 1998: "Errors and accuracy estimates of laser altimetry data acquired by various laser scanning systems for topographic applications". ISPRS JPRS, Vol. 53, no. 5. pp.245-261.

Kraus, K., Pfeifer, N., 1998: "Determination of terrain models in wooded areas with airborne laser scanner data. ISPRS JPRS. Vol. 53, pp. 193-203.

Oude Elberink, S., Maas, H, 2000: "The use of anisotropic height texture measures for the segmentation of airborne laser scanner data." IAPRS. WG III/3., Vol XXXIII, Part B3. Amsterdam, 2000. pp.678-684.

Schickler, W., Thorpe, A., 2001: "Surface Estimation Based on LIDAR". Proceedings ASPRS conference April 23-27, 20001. St. Louis Missouri. CD-ROM, 11 pages.

Vosselman, G. and Maas, H., 2001: "Adjustment and Filtering of Raw Laser Altimetry Data". Proceedings of OEEPE workshop on airborne laserscanning and interferometric SAR for detailed digital elevation models 1-3 March 2001, paper 5 (11 pages). Royal Institute of Technology Department of Geodesy and Photogrammetry 100 Stockholm, Sweden.

Vosselman, G., 2000: "Slope based filtering of laser altimetry data". IAPRS, Vol XXXIII, Part B3, pages 935-942. Amsterdam, The Netherlands.