Journal of Applied Remote Sensing

Automated bare earth extraction technique for complex topography in light detection and ranging surveys

Terry H. Stevenson
Lori A. Magruder
Amy L. Neuenschwander
Brian Bradford

Automated bare earth extraction technique for complex topography in light detection and ranging surveys

Terry H. Stevenson, Lori A. Magruder, Amy L. Neuenschwander, and
Brian Bradford

University of Texas at Austin, Geospatial Laser Applications and Measurements,
Applied Research Laboratories, P.O. Box 8029, Austin, Texas 78713-8029
magruder@arl.utexas.edu

Abstract. Bare earth extraction is an important component to light detection and ranging (LiDAR) data analysis in terms of terrain classification. The challenge in providing accurate digital surface models is augmented when there is diverse topography within the data set or complex combinations of vegetation and built structures. Few existing algorithms can handle substantial terrain diversity without significant editing or user interaction. This effort presents a newly developed methodology that provides a flexible, adaptable tool capable of integrating multiple LiDAR data attributes for an accurate terrain assessment. The terrain extraction and segmentation (TEXAS) approach uses a third-order spatial derivative for each point in the digital surface model to determine the curvature of the terrain rather than rely solely on the slope. The utilization of the curvature has shown to successfully preserve ground points in areas of steep terrain as they typically exhibit low curvature. Within the framework of TEXAS, the contiguous sets of points with low curvatures are grouped into regions using an edge-based segmentation method. The process does not require any user inputs and is completely data driven. This technique was tested on a variety of existing LiDAR surveys, each with varying levels of topographic complexity. © The Authors. Published by SPIE under a Creative Commons Attribution 3.0 Unported License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI. [DOI: [10.1117/1.JRS.7.073560](https://doi.org/10.1117/1.JRS.7.073560)]

Keywords: LiDAR; bare earth extraction; classification; terrain; DTM; DEM; TEXAS.

Paper 12119P received May 1, 2012; revised manuscript received Apr. 17, 2013; accepted for publication Apr. 19, 2013; published online May 22, 2013.

1 Introduction

Light detection and ranging (LiDAR) has become one of the most significant contributors to the production of usable and accurate surface data primarily due to the technology's ability to produce digital elevation models with centimeter-level accuracy. As an active system, an additional advantage of LiDAR is that it can be used to acquire data over wide areas regardless of time of day. Ultimately, LiDAR provides not only measurements of the topography but also produces the capability to discern between natural and man-made structures based on reflective properties that affect the nature of the returned laser energy.¹⁻⁴ In addition to buildings, other man-made features such as power lines and bridges can also be identified within a LiDAR survey using clustering and segmentation techniques while the precise characterization of the underlying ground surface provides for advancements in geodesy, geomorphology, forestry, urban planning, and natural hazard monitoring. The success of surface classification is undoubtedly dependent most on the accuracy of the digital terrain model (DTM), or bare earth component. The DTM is the fundamental layer of underlying topography which not only allows for the characterization/classification of vegetation and man-made structures, but also can assist in the derivation of human activity layers, obstruction determination, trafficability, floodplain mapping, and natural resource allocation.

Since the quality of the DTM product directly affects the usability, it is important to develop a flexible methodology (adaptable to topography and/or environment) capable of integrating multiple LiDAR point cloud data attributes for classification. The flexibility of the approach indicates that it can be used within different ecosystems, despite potential regional challenges.

Challenges to bare earth extraction can manifest themselves in the natural dynamics of the terrain, heavy vegetation that obscures or limits the number of ground detections, or non-standard man-made structures that might confuse the classification criteria. Many existing bare earth techniques struggle to maintain the topographic integrity of the terrain as the algorithm will either clip features or over-filter the surface in an effort to distinguish between geometrically similar terrain and non-terrain features. Clipping features can result in flattened mountain tops and ridges while the surface smoothing reduces the detail of the topography and produces incorrect elevations. Even in the absence of clipping and smoothing issues, the complete removal of non-terrain features (e.g., vegetation or buildings) is an added challenge in producing accurate terrain models.

Methods for classifying or filtering LiDAR data either work with the raw point cloud of individual LiDAR returns or with a gridded surface model. A benefit of filtering data based on a digital surface model (DSM) is an improvement in processing speed and performance as the more manageable raster DSM image can represent a data set comprised of millions of individual LiDAR returns. Typical methods for surface classification rely on height differences within local (neighboring) statistics as the discriminating parameter for the determination of ground points from non-ground points. The premise behind using this type of approach is that non-ground objects will generate discontinuities in elevation which will distinguish them from the bare earth product.^{5,6} Often, this process can be implemented on a pixel-by-pixel basis or by using homogenous pixel clusters produced through a segmentation scheme. Some segmentation methods, however, are not fully adaptive and rely on user interaction (often iteratively) for the selection of thresholds.⁷ Certainly, efficiency and autonomy are two of the most fundamental objectives of any LiDAR filtering algorithm. Furthermore, additional aspects of a successful terrain classification approach include no requirement for manual editing and no data dependencies. The adaptive lower envelope follower algorithm (ALEF) derived for bare earth extraction and described in Neuenschwander et al.,⁸ laid the groundwork for a fully automated, robust, and accurate methodology; however, the final results did not eliminate all post-classification editing required in order to have a clean digital terrain model. In particular, the ALEF method had difficulties in urban areas and did not automatically identify all the buildings without user interaction. The terrain extraction and segmentation (TEXAS) approach described here improves upon the results previously obtained by the ALEF.

2 Data

Airborne LiDAR data were collected by the University of Texas' Center for Space Research over the Rio Yuna watershed in the Dominican Republic to map flooding hazards in response to severe flooding stemming from Tropical Storm Noel in 2007. Within the Rio Yuna watershed, the topography varies from flat floodplains to steep mountains with elevations ranging from 40 to 600 m. The vegetation within the upper watershed is dominated by agriculture for cacao, tobacco, and other food crops. The floodplains within the lower Rio Yuna watershed are used primarily for rice cultivation with some trees providing demarcation of agricultural field boundaries. Two of the Holdridge lifezones (subtropical dry forest and subtropical moist forest) account for over 68% of the Dominican Republic landcover.⁹ Vegetation within these forests consists of dense broadleaf forest with patches of open, conifer forest. The LiDAR data were acquired in March/April of 2009 using a narrow beam divergence (0.2 mrad) at an average altitude of 1200 m above ground resulting in a footprint diameter of approximately 20 cm. The LiDAR survey was flown such that the point density of the data was ~ 1 laser shot per m^2 . Due to the relatively low point density, the data were gridded to a 2×2 m cell posting in order to generate 2-dimensional data images. In all, over 1000 km^2 were flown in the Rio Yuna watershed, however, for the purposes of this paper four small (approximately 2×2 km) subsets were selected and analyzed. Each subset provides a different land use and land cover in order to address many of the potential challenges to a terrain classification algorithm.

The first subset is situated within and around the town of Bonao, which has a population of approximately 125,000 people. Bonao is located in the middle of the Dominican Republic and is

prone to flash flooding due to its proximity to the Rio Yuna and the adjacent mountains. The building structures in Bonao are primarily small, single story homes built in close proximity to one another. The topographic extent within the Bonao urban subset ranges from 120 to 145 m and the subset extent is 2.5×2.5 km. The second subset contains the Rio Yuna floodplain, downstream from Bonao, with a topographic relief range of 100 to 180 m. This particular subset is 1.8×1.9 km in size. The vegetation in the floodplain is primarily agriculture and small patches of woody vegetation, however, small trees and brush line the alluvial features in the primary channel. The third subset is centered on a mountain adjacent to the Rio Yuna where elevations range from 65 to 215 m. This subset is also 1.8×1.9 km in size. Vegetation on this mountain is comprised of dense tropical forest with canopy heights up to 20 m. The fourth subset contains farmland to the east of the town, and was included to show an area that does not need a large amount of filtering. This subset is smaller than the others, at 1.1×1.4 km, with elevations ranging from 58 to 64 m. The terrain is relatively simple, and the only major features are two bridges across the river.

3 Methodology

In order to produce a clean, bare earth digital terrain model (DTM) from LiDAR data, several processing steps are required. Given a de-noised point cloud, a gridding process is performed to create a rasterized digital elevation model (DEM). The TEXAS algorithm processes the gridded surface image rather than a point cloud to decrease the processing speed. TEXAS can be executed on any gridded DEM; it is unbiased to a particular sensor or detector system, which makes it applicable to linear mode, Geiger mode, airborne, or space-based LIDAR. Figure 1 shows a flow chart of the TEXAS methodology, which will be described in detail in this section.

The general goal of TEXAS is to identify and remove any cells in the gridded input that do not represent bare ground. The concept behind this approach is that the underlying terrain should be relatively free of sharp vertical corners. Such corners are typically produced by artificial structures, like building walls and vehicles or by trees and other vegetation. There are limitations to this approach, since not all terrain follows this rule. Sharp cliffs, for example, are often smoothed by TEXAS, despite the steps implemented to reduce this effect. In addition, certain artificial structures are not surrounded by sharp corners, such as buildings with exterior roof-access staircases, or partially collapsed buildings. In these cases, the building might be misidentified as a hill or other natural feature, and not properly removed. Notwithstanding these limitations, TEXAS generally produces a very clean bare earth product based on subjective analysis as well as the quantitative analysis presented herein.

In terms of the specific steps required by TEXAS, the first task is to transform the input data point cloud into a gridded surface model. A subset of the point cloud taken from the town of Bonao is shown in Fig. 2, as an example. The image in Fig. 2 is at an elevation angle of approximately 45 degrees to give a 3-D perspective. The grid spacing used in the gridding process is determined by the density of the point cloud, and is a parameter chosen by the user. Choosing a relatively small grid spacing can be advantageous if the point cloud is very dense, but lower point densities will frequently create voids in the gridded image if it is not compatible with the selected spacing scale. Larger grid spacings will alleviate voids in the image, but might simply lower the resolution unnecessarily. The gridding process sorts the point cloud into vertical data columns for each grid cell. The process then selects the minimum elevation value for each column as the representative elevation for the corresponding grid cell. This methodology for creating a minimum elevation image ensures that ground returns below a canopy are included in the gridded product, but also allows for a greater vulnerability to noise if the noise points have an elevation below that of the true ground surface. An example of a minimum elevation image produced by the gridding process is shown in Fig. 3; this particular product was created using the subset point cloud from Fig. 2.

The effect of the gridded resolution on the final DTM quality depends heavily on both the data density and the composition of the scene. If the data density is low, increasing the resolution reduces the ratio of terrain points to vegetation points, since the minimum elevation is taken for

Fig. 1 Flow-chart of the methodology used in TEXAS to identify the terrain features in the digital surface model (DSM).

each cell. If, however, the data density is high, it will increase the number of available ground points and improve the final result. Urban scenes are nearly always improved by increasing the resolution, as long as the point cloud is dense enough to support the higher resolution. This is because urban features are typically not porous (like vegetation), so high resolution DEMs more accurately describe the sharp corners of buildings, cars, and other non-terrain features. The larger number of points also improves region comparisons, since the individual regions will have more elevations associated with them, making inter-region comparisons more accurate. The digital elevation product used as an input for TEXAS need not be produced with this gridding method, since TEXAS can accept any rasterized terrain model. However, including more ground points in the input image will improve the quality of the output, since less of the terrain will be interpolated.

The TEXAS algorithm repeats the work flow shown in Fig. 1 multiple times on each scene, using the results of the previously completed cycle to assist with the determination of newly selected parameter settings. The curvature threshold used in the segmentation process is the primary contributor to how aggressively the algorithm removes features within the scene. An aggressive setting will successfully remove non-terrain features, but can lead to clipping of valid terrain. A more lenient curvature threshold setting preserves terrain very well, but can also allow for misrepresentation of non-terrain features as terrain within the final DTM. In general, the majority of the terrain clipping occurs when natural features, such as hills,

Fig. 2 Subset of the point cloud for an area in the town of Bonao, showing buildings (left) and vegetation (right).

Fig. 3 Gridded minimum elevation image (2 m grid size) of the same subset, lighter colors correspond to higher elevations.

are covered by dense vegetation. Since the vegetation cover greatly reduces the number of ground points available, those ground points that do exist are often misclassified when the more aggressive curvature threshold settings are used. To mitigate this problem, the program is first run with lenient settings, removing a large amount of the vegetation. This initial feature removal allows TEXAS to more easily detect the ground during subsequent runs with more aggressive threshold settings. These later runs will then remove the remaining vegetation typically without clipping the valid bare terrain.

The TEXAS algorithm determines the number of iterations and ultimately the aggressiveness of the curvature threshold with an evaluation parameter called the terrain roughness factor (TRF). The TRF is the standard deviation of all cell elevations in the scene divided by the standard deviation of the difference between the original scene and a low pass filtered DEM. This low-pass filtered DEM is generated by convolving the input image with a 5×5 meter Gaussian kernel. The 5×5 m kernel size is an optimized setting that will allow the capture of trees and small buildings, but preserve gradual terrain features, like hills. Higher TRF values correspond to scenes where low frequency changes in elevation dominate (e.g., mountains) while lower TRF values correspond to high frequency elevation changes (e.g., buildings and vegetation). Despite the presence of flat/planar surfaces, buildings generally present as rough features, as the structure edges create abrupt elevation differences that correlate to high frequency changes much greater than those experienced by most terrain features and vegetation. If the algorithm determines a high TRF value for a particular scene, the program will select a more lenient curvature threshold setting in the first iteration to prevent clipping. Conversely, if the surface evaluation determines a lower TRF value, the curvature threshold setting is consistent with a more aggressive approach for feature removal. These decisions and performance settings are all

determined within the algorithm based on the input data.

The segmentation process within TEXAS begins by using a 4×1 finite difference equation to approximate the third spatial derivative at the edge of each cell in the input image. A sample edge map is shown in Fig. 4, this is the same area depicted in Fig. 3, with high curvature edges highlighted. As mentioned previously, the curvature threshold used to differentiate between “high” and “low” curvature is set by the TRF. Numerically, this threshold is set between 0.5 and 1.5 for each iteration, these bounds were determined empirically. A threshold much lower than 0.5 causes TEXAS to wrongly classify patches of uneven ground as vegetation. A threshold much higher than 1.5 begins to reduce TEXAS’s ability to classify buildings by producing low curvature areas along building edges. The curvature threshold, which varies inversely with the TRF value, is decreased between subsequent iterations but remains within the established numerical bounds. So, an image that rapidly becomes smoother (lowering the TRF) will not lower the curvature threshold value as much as an image that retains much of its original roughness (high TRF). Once the high curvature areas of the image have been defined, the image is segmented into contiguous regions of low curvature, surrounded by areas of high curvature; an example would be a rooftop surrounded by a sharp drop to the ground. Vegetation is rarely assigned a region as it is rarely produces areas of low curvature so typically it is excluded immediately as actual terrain. This concept is illustrated in Fig. 5, where the gray points indicate regions identified as terrain and the white and black points were excluded. Black points represent those points that could be immediately discarded due to high curvature while the white points are points assigned to a region but later determined not to represent terrain. Note that the regions are slightly larger than the low curvature areas shown in Fig. 4, this is because the regions are expanded slightly to include nearby points with similar elevations.

Following the segmentation, each region is individually validated to determine if it represents valid terrain. Large area regions are automatically accepted as terrain using a spatial threshold of 118,000 m². Although the roof area of the Pentagon is equal to this threshold value, in practice, very few buildings approach this size so in most cases regions of this magnitude (or larger) would, in fact represent a ground surface. Regions with elevations significantly above their neighbors (especially if the neighbors are confirmed terrain regions) are also rejected. Given that there are numerous regions, large and small, within the scene, TEXAS attempts to build connections between regions that are not immediately adjacent, and then combines them into larger regions. A downward flood fill is performed after the classification step to add points in the high curvature area back in, this serves a dual purpose. First, it adds legitimate ground points that were previously excluded because of nearby high curvature, points immediately adjacent to buildings and ground points underneath vegetation are the two most common examples. Secondly, if a building was mistakenly left in the image, it ensures that the sides of the building are properly preserved for the next iteration. If this method were not used, the building sides

Fig. 4 Edge map of the subset. White edges correspond to edges with high curvature (above the threshold), shading of the cell illustrates elevation.

Fig. 5 Region map of the subset. Gray indicates points in an accepted terrain region, black indicates points not segmented into a region, and white represents points in regions classified as non-terrain.

would be eliminated and interpolated for the next iteration, smoothing them and reducing the ability of the next iteration to correctly identify the building. By preserving the sharp edges, more aggressive iterations can correctly identify a building and remove it entirely.

The output image produced by TEXAS typically contains a large number of voids where non-terrain features were removed. Because large discontinuities in the surface will reduce the ability of TEXAS to properly segment the scene from one iteration to another, a bilinear interpolation is implemented to fill the voids within the image. The interpolation process is actually performed twice in order to better filter the data. After the initial interpolation, any interpolated points with an elevation greater than the corresponding elevation in the original DEM are replaced with that original elevation and the interpolation is performed again. This two-step process is useful for retaining features like ravines and small streams, which can be erroneously removed; however, this step increases the program's vulnerability to noise. Bilinear interpolation was chosen for its speed and quality, but TEXAS is not dependent on any one interpolation algorithm.

In some scenes, very large areas are interpolated to produce a continuous output surface with no voids. Since there were no ground returns from these areas (or the ground returns were undetected), the interpolation is nothing more than a guess at the terrain's actual shape. Interpolation necessarily produces very smooth surfaces, and there are many cases where this may not represent the actual terrain. Because of this potential misrepresentation, every output image is saved with a mask image that indicates which cells are true ground points and which cells were interpolated. With the mask applied, only the detected ground points are shown. However, since the underlying terrain is often relatively smooth, the interpolation provides a good approximation for the shape of the terrain. Overall, the segmentation and interpolation processes are run a total of three to five times, depending on the initial roughness. Scenes with low TRF scores are run fewer times with more aggressive settings, while scenes with high TRF scores are run more times with more lenient initial settings. While this approach increases processing time, the increased resistance to terrain clipping outweighs the longer run-time. After all of the iterations are complete, a small-window median filter is used to remove any interpolation artifacts.

4 Results

In order to provide consistency with regard to the performance of TEXAS, all of the results presented for the selected subsets are DTMs directly produced through a single, automated run of the software. In addition, none of these results have been manually edited in any way. Again, the strength of this algorithm is that it does not require user input other than the gridding resolution; all thresholds and decision criteria are determined automatically for each data set through the process described in the previous section. The TEXAS output images of the terrain for each selected subset are shown in Figs. 6 to 9. The image format in Figs. 6 to 9

show the input DSM on the left and the TEXAS output DTM on the right, they also consist of some elevation analysis transects along the bottom of the image.

Figure 6 presents the results for the town of Bonao. From a qualitative standpoint, the observed performance of TEXAS is good in terms of overall removal of built structures and vegetation with few artifacts present from man-made features that were missed by the classifier and left within the DTM. Unfortunately, no high resolution ground truth exists for these specific areas, which limits the amount of quantitative analysis available for the determination of the quality and accuracy of the output DTM. As such, to qualitatively address the capability of TEXAS, four transects from each output DTM are plotted against the input point cloud and shown at the bottom of each figure. For each transect plot the profile for the point cloud points are shown in green and the calculated DTM surface points are shown in red. By observing the DTM surface against the point cloud, errors in the terrain surface become apparent and provide an indication as to the program performance based on the input data. Transects A to B and C to D indicate apparent agreement between the point cloud and DTM. Transect E to F; however, shows a few spots where the terrain surface sits above (~ 30 cm) the point cloud. These artifacts seem to occur in the dense building areas where some pixels close to the ground were not fully eliminated. In spite of these few artifacts, TEXAS was successful in preserving the topography with no obvious clipping, particularly along the edge of the floodplain (transect G to H).

The second subset processed by TEXAS, the Bonao floodplain, is shown in Fig. 7. Again, this DTM product has not had any post-classification editing, it is simply the automated output from TEXAS. Examination of the four transects in Fig. 7 shows that there is no obvious terrain clipping in this image and the vegetation appears to be sufficiently removed, in fact it appears that all of the vegetation was removed throughout the image beyond the transects. These results are encouraging, as accurate representations of floodplains are critical for hydrological modeling applications; particularly for this region, which is prone to severe flash flooding events.

The third subset selected and analyzed by TEXAS is situated on a mountain adjacent to the Rio Yuna floodplain (Fig. 8). Forested mountains are excellent tests for LiDAR classifiers due to

Fig. 6 (a) Digital surface model of Bonao (urban) subset and (b) DTM of Bonao (urban) subset.

Fig. 7 (a) Digital surface model of floodplain subset and (b) DTM of floodplain subset.

the challenge of the canopy and dynamic topography combination. The canopy height of the trees within this particular scene ranges from 10 to 20 m. The selected transects were chosen such that a variety of geomorphological and vegetation configurations would be represented. Inspection of the four transects in Fig. 8 indicates that based on the comparison to the input point cloud the TEXAS DTM is an accurate representation of the underlying topography. Transect A to B at the 50 m mark and again at the 175 m mark, however, does reveal that the DTM estimate sits approximately 1 m above the clustering of points within the ravine. Transect G to H in Fig. 8 illustrates the importance of the pixel add-back aspect within the algorithm, since often the algorithm will miss terrain points associated with riverbeds and other low-lying features. A specific case of the pixel add-back advantage occurs at the 80 to 90 m mark on transect G to H where the ranging points consist primarily of returns from vegetation over a ravine. However, a few LiDAR returns do penetrate through gaps in the canopy and reach the ground. These few low pixels would normally be rejected as terrain points due to the high curvature surrounding them. By implementing the pixel add-back step into the algorithm, these ground points are included into the final terrain mask, producing a more accurate DTM. And although a few non-terrain pixels were not fully removed, there were no obvious signs of terrain clipping. All ridges, peaks, and drainage creeks appear to be fully maintained and represented within the final DTM.

Another subset used to determine the capability of TEXAS is an area located to the east of the town of Bonao, and down the river from the floodplain scene. This subset primarily contains farmland, with a small amount of vegetation. The only man-made features within the subset are two bridges in the north and southeast regions of the scene. The terrain is quite simple, as the only noteworthy feature is the river. Given the lack of topography and built structures, this scene was not expected to present a challenge for TEXAS but was included in this analysis to demonstrate that TEXAS is capable of processing scenes that require very little filtering. Figure 9 shows the farmland scene before and after processing and also includes four transects for comparison. Transect A to B crosses a small dip in the terrain that, despite vegetation cover, is successfully preserved. Transects C to D and E to F show relatively flat areas with some tree cover,

Fig. 8 (a) Digital surface model of mountain subset and (b) DTM of mountain subset.

in both cases, it is apparent that the trees have been removed, and the underlying terrain is preserved. Transect G to H shows the performance of TEXAS over an elevated roadway leading to the north bridge as the vegetation beside the road is successfully removed, while the roadway itself is preserved. For the purposes of this algorithm, bridges are treated as non-terrain objects so this particular subset highlights TEXAS's ability to fully remove bridges while leaving the nearby roadways intact.

The percent of required interpolation for a given area is an important statistic. In general terms, this percentage provides a measure of how confidently the DTM represents the actual terrain. Very high interpolation percentages indicate that the terrain is not well known based on the number of true ground data points within the input data set. Table 1 shows the interpolation percentages for the four scenes analyzed herein. The most complex terrain, in the Bonao Mountain scene, actually produces a much lower interpolation percentage than the simpler terrains. This is because the mountain scene actually has much more exposed terrain (i.e., true ground points) than either the heavily vegetated floodplain, or the urban area which has a large number of buildings. So, while the mountain scene presents a larger challenge to an automated bare earth algorithm in terms of its complex topography, the resulting output DTM product is actually more representative of the true terrain than some of the other cases where there are fewer actual ground returns. However, higher interpolation percentages do not preclude producing a DTM with reasonable accuracy as long as there are enough actual terrain points within the input data.

In order to provide a quantitative assessment of the TEXAS performance, ground truth is required. Since ground truth is not available for any of these cases analyzed, one of the scenes was manually filtered in order to produce a surface product for comparison to the automated DTM output. This manual filtering was performed on the minimum elevation image, the same starting input required by TEXAS. While this is not a perfect or ideal ground truth, it does provide a good basis for quantitative comparison, and has been a precedent for past bare earth algorithm evaluation.⁵ Figure 10 shows the minimum elevation image input (left) and the manually filtered result (right). Note the series of vertical lines in the upper half of both images; these

Fig. 9 (a) Digital surface model of farlands subset and (b) DTM of farlands subset.

Table 1 Interpolation percentages for each scene.

Scene	Interpolation Percentage
Bonao: Mountain	14.74%
Bonao: Urban	33.17%
Bonao: Plain	34.32%
Bonao: Farmland	10.22%

are errors incurred during data collection and the initial processing of the LiDAR data, not from actual terrain. These lines are 30 to 40 cm above the surrounding terrain, and present as high frequency features that TEXAS filters. Because of these spurious features, any DTM point within 40 cm of the manually produced ground will be considered accurate during the comparison between the derived ground truth and the automated product.

Figure 11 shows the result of subtracting the TEXAS DTM from the manual ground truth surface. The blue regions are areas where the TEXAS DTM is below the manual DTM by more than the predetermined threshold (40 cm); the red regions represent areas above the manual DTM by more than 40 cm. Following the terminology used by Sithole and Vosselman,⁵ missed terrain points (blue) are considered Type I errors, and accepted non-terrain points (red) are Type II errors. Many times these are also referred to as commission errors (red) and omission errors (blue). The results shown in Fig. 11 indicate that Type I errors are much more prevalent than errors classified as Type II. Numerically, 6.33% of the image in Fig. 11 is represented as Type I errors and Type II errors comprise less than 0.1%. This distribution of error types is consistent

Fig. 10 Subset of the Bonao urban scene (a) and the manually filtered result (b).

Fig. 11 Image change map of the manually filtered urban scene and the TEXAS DTM, with Type I (blue) and Type II (red) errors highlighted.

with the performance evaluation of other algorithms in 2004 by Sithole and Vosselman,⁵ of the eight they evaluated, six produced higher rates of Type I errors than Type II.

5 Conclusions

The TEXAS approach was developed to classify and filter digital surface models from LiDAR data. This method computes the third derivative of the input DEM and recursively isolates regions of low curvature. The elevations within each region are evaluated against neighboring regions, and are either validated or rejected as terrain points. The TEXAS software was run on four different surface types from LiDAR data collected in the Dominican Republic. In these selected cases, the software successfully eliminated vegetation and man-made features, with a few noted exceptions, while preserving the underlying topography. The quality assessment of the final DTMs was implemented by comparing transects of the terrain surface against the input point cloud and calculating the elevation difference as the error in the surface product. This approach for performance evaluation was derived in the absence of ground truth

measurements. Clearly, if the input data is of poor quality, the corresponding output DTM will be inaccurate regardless of close numerical agreement between the two product elevations. As such the intention of this error analysis was not to provide a quantitative analysis of the input data itself, but to evaluate the TEXAS algorithm based on the input data given. Toward that end, this comparison, although not equal to using ground truth, does provide some evidence as to the performance capabilities of the technique. Future studies of TEXAS will include areas with known ground elevations so that a true comparison can be made. Despite the limitations of the validation procedure, these results are encouraging because the TEXAS software does not require any parameters or thresholds from the user and does not require any post-classification filtering or clean-up. In addition, this methodology does not require (or incorporate) any general smoothing of the landscape. The valid terrain points are completely incorporated in the final DTM without losing information due to filtering/smoothing. In addition to the TEXAS advantage, this program is independent of sensor hardware or acquisition scenarios, as it only requires a point cloud as input to the algorithm and is not affected by type of instrumentation. Future evaluations of the TEXAS performance will be done on additional ecosystems and topography to ensure that the process can develop a robust and consistent capability regardless of the composition of the scene.

Acknowledgements

The authors would like to gratefully acknowledge the funding support from MIT Lincoln Laboratories for this research. Specifically, we would like to thank Dale Fried, and Bob Knowlton and their entire LIDAR technical team.

References

1. P. Gamba and B. Houshmand, "Digital surface models and building extraction: a comparison of IFSAR and LiDAR data," *IEEE Trans. Geosci. Rem. Sens.* **38**(4), 1959–1968 (2000), <http://dx.doi.org/10.1109/36.851777>.
2. H. Maas and G. Vosselman, "Two algorithms for extracting building models from raw laser altimetry data," *ISPRS J. Photogramm. Rem. Sens.* **54**(2–3), 153–163 (1999), [http://dx.doi.org/10.1016/S0924-2716\(99\)00004-0](http://dx.doi.org/10.1016/S0924-2716(99)00004-0).
3. K. Zhang et al., "A progressive morphological filter for removing non-ground measurements from airborne LiDAR data," *IEEE Trans. Geosci. Rem. Sens.* **41**(4), 872–882 (2003), <http://dx.doi.org/10.1109/TGRS.2003.810682>.
4. K. Zhang, J. Yan, and S.-C. Chen, "Automatic construction of building footprints from airborne LiDAR data," *Trans. Geosci. Rem. Sens.* **44**(9), 2423–2533 (2006), <http://dx.doi.org/10.1109/TGRS.2006.874137>.
5. G. Sithole and G. Vosselman, "Experimental comparison of filter algorithms for bare earth extraction from airborne laser scanning point clouds," *ISPRS J. Photogramm. Rem. Sens.* **59**(1–2), 85–101 (2004), <http://dx.doi.org/10.1016/j.isprsjprs.2004.05.004>.
6. K. Zhang and D. Whitman, "Comparison of three algorithms for filtering airborne LiDAR data," *Photogramm. Eng. Rem. Sens.* **71**(3), 313–324 (2005).
7. G. Miliareisis and N. Kokkas, "Segmentation and object-based classification of the building class from LiDAR DEMs," *Comput. Geosci.* **33**(8), 1076–1087 (2007), <http://dx.doi.org/10.1016/j.cageo.2006.11.012>.
8. A. L. Neuenschwander et al., "Classification of LADAR data over Haitian urban environments using a lower envelope follower and adaptive gradient operator," *Proc. SPIE* **7684**, 768408 (2010), <http://dx.doi.org/10.1117/12.866033>.
9. D. Knudsen, W. Chaney, and F. Reynoso, "Fuelwood and charcoal research in the Dominican Republic," Purdue University, Department of Forestry and Natural Resources (1988).

Biographies and photographs of the authors are not available.