

PROBABILISTIC LINE-OF-SIGHT WITH LIDAR POINT CLOUDS

Peter L. Guth

Department of Oceanography
United States Naval Academy
572C Holloway Road
Annapolis, Maryland 21402
pguth@usna.edu

ABSTRACT

Intervisibility computed from a lidar point cloud offers several advantages compared to using a gridded digital surface model: (1) even with high resolution grids, the point cloud offers a better graphical depiction of the line-of-sight, and (2) the point cloud avoids interpolation problems that frequently result in the grid not accurately capturing the forest canopy. The algorithm for computing intervisibility uses the following steps: (1) extract the points in the vicinity of the line-of-sight; (2) rotate the points to along-profile and cross-profile coordinates; (3) extract the floor and ceiling points within the point cloud in small bins along the profile, which give the ground surface and canopy top; and (4) compute intervisibility. Several refinements allow improved accuracy: (1) the vegetation density at various heights, derived by counting returns in voxels along the profile, allows probabilistic estimates when only a few trees might affect the results; (2) use of LAS classification codes to pick out ground, vegetation, and buildings; (3) where the classification was overly general, buildings can be rapidly identified on the profiles from the planar roof reflections; and (4) using a larger bin size for the ground determination can compensate for a lower return rate compared to the canopy. Using the point cloud requires a significantly larger data set and more processing, but improving computer power makes the effort much more effective by improving results.

KEYWORDS: lidar, line-of-sight, intervisibility, algorithms, point clouds

INTRODUCTION

Intervisibility using line-of-sight profiles or viewsheds represents an important uses for digital topography for the military, communications engineers, landscape planners, and others. Guth (2009b) discussed options for incorporating vegetation in line-of-sight computations, but those existing vegetation databases typically have resolutions of 30 m or larger and can only provide estimates at those scales.

Lidar topography promises great improvements for a wide range of analyses, including intervisibility, although with an increase in both data storage and required processing as the data resolution approaches 1 m or even smaller. The data can use grids or the raw point clouds. Lidar grids require floating point elevations, with 4 bytes per elevation posting, so two complementary grids with the ground elevation and vegetation height require 8 bytes per posting, which typically use 1 m spacing. The LAS format, developed by ASPRS and the standard for most lidar data, typically uses 28 bytes per posting. This is 3.5 times the storage required for a point in the grid, and the point cloud frequently has a much higher density than 1 point/m², so that a point cloud could easily have 50 times the storage requirement compared to 2 grids. In addition to the storage costs, random access of specific point elevations is more expensive for the point clouds.

This paper describes algorithms to use vegetation for line-of-sight intervisibility for lidar grids and point clouds. The algorithm has been implemented in MICRODEM (Guth, 2009a), but could be adapted to any GIS software.

LIDAR PROCESSING AND DATA CHARACTERISTICS

While the primary focus of the work is developing algorithms to use lidar data, the work must consider the characteristics of the lidar data. What the data sets show constrains what can be done with them. Grids require less storage and allow much faster processing, so the advantages of the point clouds must be clearly weighed against their disadvantages.

This work will focus on three data sets from OpenTopography (2012) covering two areas: a deciduous forest in

eastern Pennsylvania with both leaf-on and leaf-off conditions, and a coniferous forest in western Nevada (Table 1). These data sets are among those with the highest point density that are freely available. For each data set, a point cloud with about 150 million points in LAS format and two 1 m grids in UTM projection were downloaded. All three point clouds have minimal point classification, with only a relatively conservative ground classification used. The grids correspond to a digital surface model (DSM, called the Zmax grid by OpenTopography, and also called the first return surface) and a digital terrain model (DTM, called the Zmin grid by OpenTopography, and also called the last return or bare earth surface). OpenTopography can produce three other grids (Zmean, Zidw, and TIN) but the Zmin and Zmax grids bound them and provide better estimates for intervisibility computations.

Table 1. Lidar data sets used

Data Set	Survey Date	Point Density	LAS Classification
Christina River Basin Critical Zone Observatory, PA	April 2010	13+ points/m ²	41% ground, 59% unclassified
Christina River Basin Critical Zone Observatory, PA	July 2010	13+ points/m ²	30% ground, 70% unclassified
Lake Tahoe Basin Lidar, CA/NV	August 2010	11+ points/m ²	20% ground, 80% unclassified

Preprocessing included tiling and indexing the lidar LAS files for rapid random access, and creating a height above ground (HAG, Stoker, 2010) grid from the DSM and DTM. While both areas have a few buildings, this work focuses on vegetation and the HAG surface should represent a good proxy for the vegetation height. For rapid computation, the analysis program loads the DTM and HAG grids into memory. The DSM, when needed for visualization or computations, can be created on the fly as the sum of the DTM and HAG.

Figures 1 shows subsets of the three grids for the Pennsylvania data set, with a hybrid display including shaded reflectance and elevation colors, at the 1 meter resolution of the grids. In the leaf-on data set, the DTM is very noisy in the forested parts of the map. This paper will focus on the area shown in this figure, and the profile indicated with the red line, and then briefly consider the Nevada region at the end of the paper. For this profile, the observer is on the left side of the map and on the profiles, and the target is on the right. Target and observer are both 2 m above the ground.




Figure 1. Three grids of the leaf-on Pennsylvania data set.

Figure 2 shows the leaf-on and leaf-off DTMs for the same area as Figure 1. The leaf-off DTM shows very few tree trunks and represents a much better estimate of the true bare earth surface, but does have a number of tree trunks along a lane and in the copse of trees in the southeast corner of the map. A number of buildings appear along the road at the southern edge of the map.




Figure 2. Leaf-off and leaf-on map comparisons of the DTM, Pennsylvania dataset.

Profiles (Figure 3) confirm that the leaf-off DTM has many fewer vegetation artifacts (there is one small spike at about 235 m along the profile) compared to the leaf-on DTM, and that the leaf-off DTM generally has a smooth, realistic surface as it does in the map view (Figure 2). The DSM is much less sensitive to leaf-on or leaf-off data collection, although the leaf-on DSM, collected three months after the leaf-off grid, is generally higher, reflecting the extra time for foliage growth and the increased chances for return reflections from leaves.

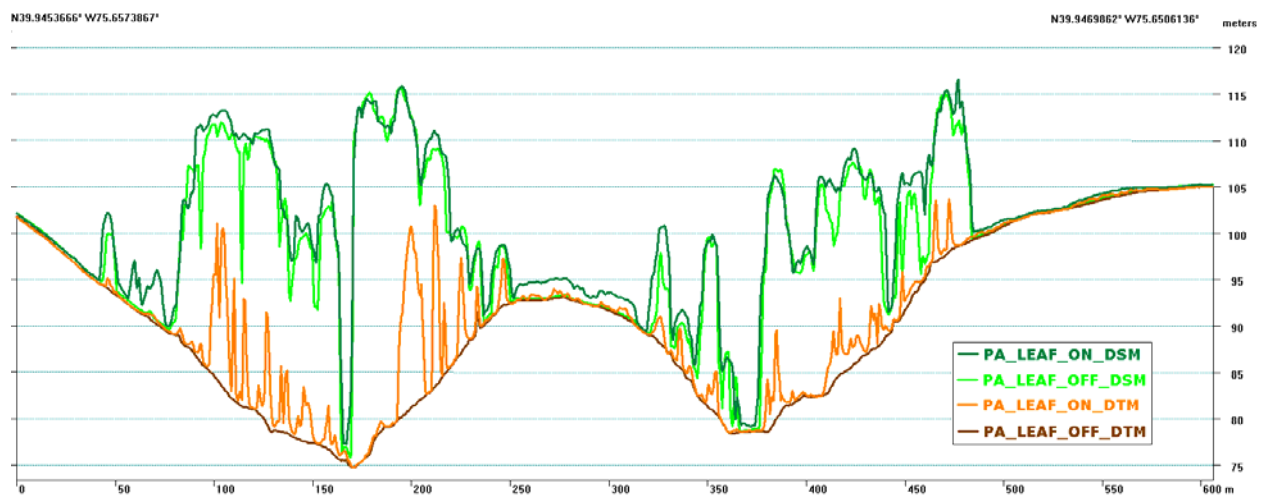


Figure 3. Profile comparing leaf-on and leaf-off DSMs and DTMs, Pennsylvania dataset.

Figure 4 shows the changes in the three months between collection of the leaf-off and leaf-on data sets. Positive numbers indicate that the July surface is higher, and they vastly outnumber negative values which indicate the April values were higher. For the DTM, the higher increases (over 8 meters) occur in the center of the forested regions. For

the DSM, the highest values tend to occur on the edges of the forested regions, where lateral growth (or filling out of the foliage, increasing the chance for reflections) expands the apparent forested areas laterally.




Figure 4. Maps showing differences between leaf-on and leaf-off DTMs and DSMs, Pennsylvania dataset.

Figure 5 shows the same profile as Figure 3, computed solely from the point cloud. At points every 0.5 m along the profile, the highest and lowest points in a small (1 m) box surrounding the point define the ceiling and the floor. In the bottom profile, the tree trunks appear, much as they do in the DTM. In the top profile, increasing the region size used for the floor, but not the ceiling, removes almost all of the spikes, and the profiles resembles the leaf-off DTM even though it used the leaf-on point cloud. This solution works well in flat areas, but creates stair-stepping where the terrain slopes.




Figure 5. Profiles through the point cloud extracting floor and ceiling points from the leaf-on dataset, showing the effect of changing the search radius used for the floor. A larger floor region removes or greatly reduces the three trunks remaining in the profiles.

LIDAR INTERVISIBILITY ALGORITHMS

Line-of-sight can use the DTM and HAG grids, and a database to hold the results (Table 2). Using a database allows standard GIS routines to plot the results on graphs and maps and reduces the requirements for custom programming. Table 3 shows part of the database with the results for the calculation in the Pennsylvania data set shown throughout this paper. In this part of the profile, 4 of the computation points show vegetation blocking intervisibility. Since this corresponds to only 2 m, if this were the only blockage, the point under consideration might be visible through the foliage.

Table 2. Database structure for grid-based intervisibility

Field	Use
LAT	Latitude of points along the profile. The first is the observer, and the last the target.
LONG	Longitude of points along the profile.
ELEV_M	Elevation of the ground surface along the profile, from bilinear interpolation in the DTM grid.
RANGE_KM	Distance in KM along profile. By default this is half the DEM grid spacing, but the user can adjust the spacing.
CURV_M	Earth curvature in meters.
LOS_HT_M	Height of the line-of-sight, including earth curvature.
VEG_HT	Vegetation height at the point, from bilinear interpolation in the HAG grid.
BLOCK_TERR	Binary field showing hard blockage by terrain. It will be "Y" if $LOS_HT_M < ELEV_M$.
BLOCK_VEG	Binary field showing potential blockage by vegetation. It will be "Y" if $LOS_HT_M < ELEV_M + VEG_HT$.

Table 3. Database showing computations for line-of-sight with DTM and HAG grids.

LAT	LONG	ELEV_M	RANGE_KM	CURV_M	LOS_HT_M	VEG_HT	BLOCK_TERR	BLOCK_VEG
39.946389	-75.653112	82.34	0.3824	0	105.79	23.06		
39.94639	-75.653106	82.3568	0.3829	0	105.79	23.33		
39.946392	-75.6531	81.9315	0.3834	0	105.79	24.01		Y
39.946393	-75.653095	83.9745	0.3839	0	105.8	22.25		Y
39.946394	-75.653089	88.9875	0.3844	0	105.8	17.22		Y
39.946396	-75.653084	89.6112	0.3849	0	105.8	16.43		Y
39.946397	-75.653078	86.2675	0.3854	0	105.8	19.53		
39.946398	-75.653073	85.242	0.3859	0	105.81	20.38		
39.9464	-75.653067	83.7889	0.3864	0	105.81	21.67		
39.946401	-75.653061	81.9709	0.3869	0	105.81	23.17		

The algorithm for computing intervisibility with the point cloud uses the following steps: (1) extract the points in the vicinity of the line-of-sight, which might be in multiple tiles, and load them into memory; (2) rotate the points to along-profile and cross-profile coordinates; (3) extract the floor and ceiling points within the point cloud in small bins along the profile, which give the ground surface and canopy top (see Figure 5); (4) compute intervisibility blocking; and (5) estimate how much vegetation the line-of-sight intersects.

Using the point cloud for intervisibility requires computing several different parameters for the database (Table 4). For testing the computations use the DTM and HAG, plus the point cloud, but ultimately they could use just the point cloud. The ground elevation, top of the canopy, and vegetation height could be taken from the grids, but the point cloud has several potential advantages: (1) it can use the greater resolution possible from the point cloud compared to the grids; (2) the search radius can be different for the ground compared to the canopy top (Figure 5); and (3) the vegetation density can be estimated from the number of returns both around the line-of-sight, and the number of returns above the line-of-sight.

Table 4. Additional fields for point cloud-based intervisibility

Field	Use
MINZ_PTCLD	Minimum (floor) return elevation in a small box around the computation point. Because of problems with returns around tree trunks on the ground, the box size for the ground can be larger than that for the maximum. This should be close to the value in the DTM.
MAXZ_PTCLD	Maximum (ceiling) return elevation in a small box around the computation point. This should be close to the value in the DSM.
VEGZ_PTCLD	Vegetation height from the point cloud (MAXZ_PTCLD – PMINZ_PTCLD). This should be close to the value in the HAG grid.
PTS_AROUND	Number of returns in a small voxel surrounding the LOS_HT_M above the computation point. If there is little penetration of the lidar pulse into the vegetation, this value may not give a good estimate of the vegetation density.
PTS_ABOVE	Number of returns above the LOS_HT_M above the computation point. This shows how far down into the vegetation the line-of-sight occurs.

Table 5 shows part of the table for the same line-of-sight shown in Table 3. The profile, 606 m long, uses 1213 points for the computation. Where the grids showed 4 points blocking intervisibility in this portion of the profile, the point cloud only shows 2. For these points, there are more returns around the computation point than there are above it, indicating the line-of-sight is near the edge of the vegetation and that some of the returns are actually below the line-of-sight.

Table 5. Database showing computations for line-of-sight with the lidar point cloud.

LAT	LONG	RANGE_KM	LOS_HT_M	MINZ_PTCLD	MAXZ_PTCLD	VEGZ_PTCLD	PTS_AROUND	PTS_ABOVE
39.946389	-75.653112	0.3824	105.79	85.35	104.22	18.87		
39.94639	-75.653106	0.3829	105.79	85.11	104.22	19.11		
39.946392	-75.6531	0.3834	105.79	85.11	105.99	20.88	4	3
39.946393	-75.653095	0.3839	105.8	94.76	105.99	11.23	6	3
39.946394	-75.653089	0.3844	105.8	92.31	105.63	13.32		
39.946396	-75.653084	0.3849	105.8	92.31	104.84	12.53		
39.946397	-75.653078	0.3854	105.8	81.69	104.57	22.88		
39.946398	-75.653073	0.3859	105.81	81.69	104.57	22.88		
39.9464	-75.653067	0.3864	105.81	91.54	103.88	12.34		
39.946401	-75.653061	0.3869	105.81	96.88	103.63	6.75		

Figure 6 compares the computations, both visually and statistically, for the DTM/HAG grids and the point cloud. For the grids, the program displays the vegetation blockages in a different color from the vegetation which does not block intervisibility. The statistical summary shows the length of the entire profile, the bin size used, the length of the portion of the profile blocked by vegetation, and the number of patches of vegetation encountered along the profile. A patch boundary has the profile going from a computation point within vegetation to one without vegetation. In this case the grids show a much longer portion (122 m) of the profile being blocked, along with fewer patches. For the point cloud, only 48 m are blocked by vegetation in the immediate vicinity, but 99 m have overlying vegetation and might be blocked if limited penetration by the laser results in underestimation of the foliage density. Overall the point cloud computations produce a finer resolution for the results.

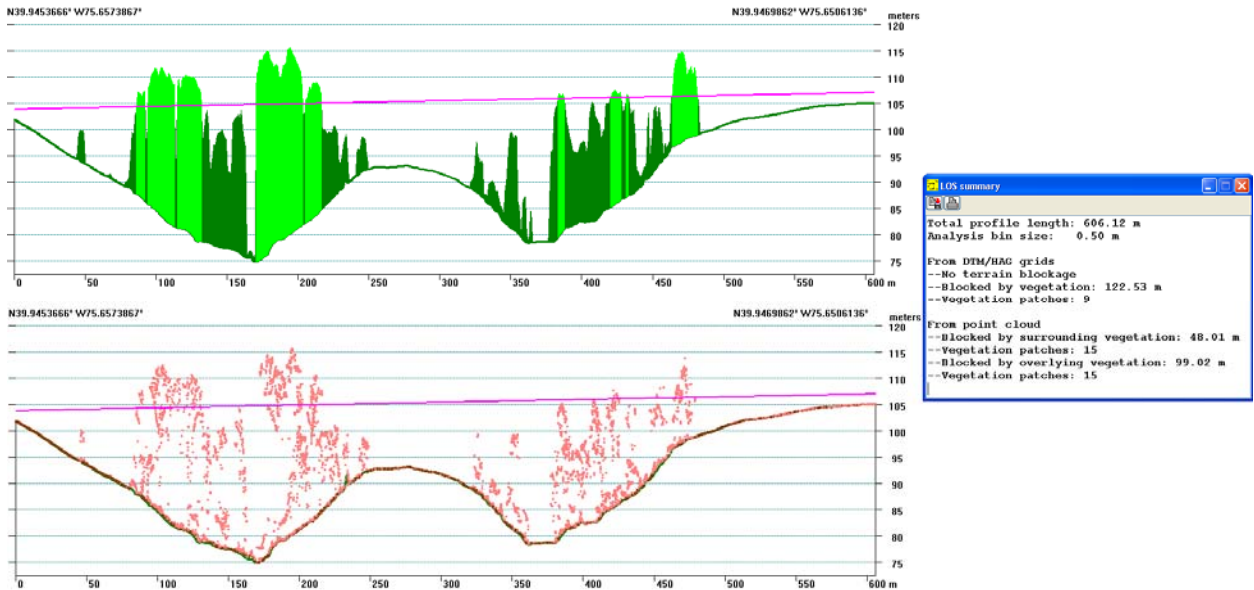


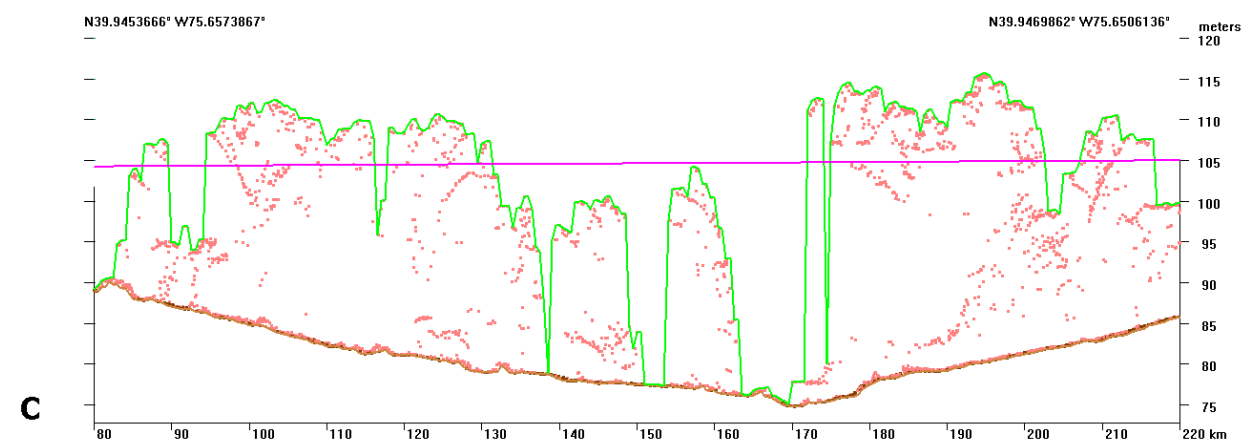


Figure 6. Line-of-sight through leaf-off DTM and HAG grids (top) and leaf-off point cloud (bottom) with statistical summary on the right.

Figure 7 shows a blowup of part of this calculation through the forested region on the left side of the profile. From about 0.09 to 0.13 km, there are 1-6 lidar returns per m³ in the vicinity of the line-of-sight, but 10-20 returns in the space above it; this portion should be blocked by vegetation. Between 0.15 and 0.16 km, there are about 8 lidar returns per m³ in the vicinity of the light of sight, but many fewer above it. The line of sight here is just skirting the top of the vegetation canopy, and if this were the only blockage along the entire profile, the two end points might in fact be intervisible depending on a number of factor such as the target's size, the type of vegetation, and the time of year.



Leaf-on



Leaf-off

Figure 7. Blowup of part of the line-of-sight. The profile in B shows the DTM surface in dark green, while the profile in C shows the point cloud ceiling (proxy for DSM) in light green.

The point clouds can be displayed and manipulated with 3D graphics, but often simple slices through the cloud better reveal the structure because they allow measurement, and restricting the display to a narrow slice without the distraction of all the other data points. Figure 8 shows a slice through the line-of-sight, looking from the observer

toward the target. Negative values on the x-axis are to the left, or north side, of the line-of-sight. The diagram also shows the first two Fresnel zones, graphically showing the degree of intrusion.

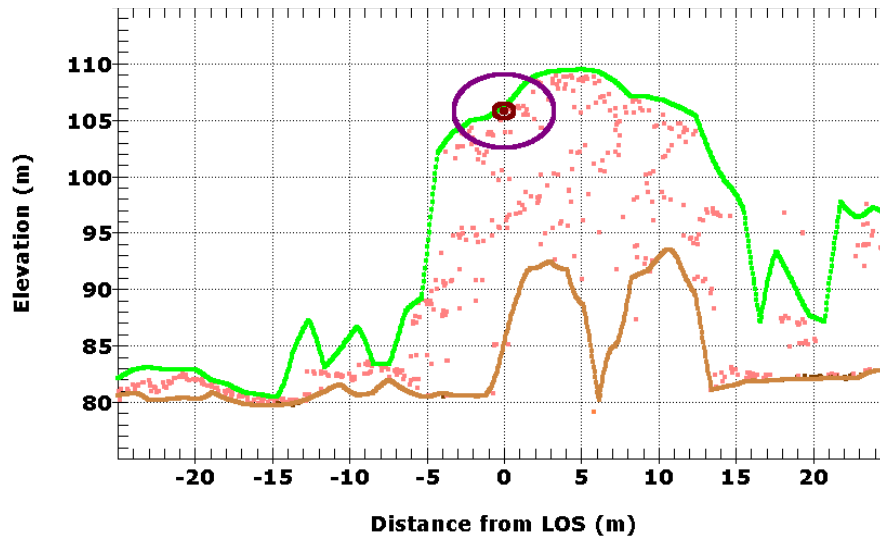


Figure 8. Section parallel to the line-of-sight, showing the interpolated surface from the DTM and HAG, the locations of lidar returns. The first and second Fresnel zones allow a detailed analysis for grazing line of sight. This is at 383 m along the profile, which Figures 5 and 6 show in along-profile view.

DISCUSSION

This work used extremely dense lidar point clouds, but as the technology matures this density will increasingly become the norm. While the displays colored the lidar returns by their LAS classification, except for some correctly classified ground returns this data had not been classified and creating a classification was beyond the scope of the work. Obviously a complete classification would improve the intervisibility computations, although that is probably not necessary since extraction of the ground from the point cloud is relatively easy.

Getting an accurate ground surface proved a challenge, which was worse for the leaf-on data set. Evergreen forests present this problem year round, and Figure 9 shows the spikes present in the Nevada lidar data. De-spiking the DTM is part of classification, and will require greater adoption of algorithms like those of Haugerud and Harding (2001) or Evans and Hudak (2007).

Llobera (2009) discussed the application of the Beer-Lambert's attenuation law to develop probabilistic line-of-sight, and used a synthetic DEM to show how an algorithm might work. The arrival of high density lidar clouds will now allow developing models and applying them with relatively easily obtained, high resolution data. This will also allow computation of intrusions into the Fresnel zones for predicting radio line-of-sight.

CONCLUSION

Using two grids, the DTM and the HAG, greatly improves on intervisibility results compared to using a single DSM. Obstructions can be broken into hard blockages by terrain, and potentially soft blockages by vegetation. Using the lidar point cloud can greatly improve both the visualization and quantitative computations of the vegetation blockage. Blockages could be interpreted with physical models or attenuation, or empirical results. As 64 bit computers become mainstream, efficient computations with the point cloud will reduce any penalties for the much larger data sets. Intervisibility will use full voxel computations (Stoker, 2009) along the line-of-sight.

Improved analysis will require better classification of Lidar data sets, especially to pick out the ground surface. This will create better DTM grids, and include all ground returns in the LAS data files. This is generally now done by proprietary algorithms from the data providers, and needs to be done more consistently to avoid analysis software having to perform some classification prior to use of the data. Because the results depend on the nature of the landscape, and the lidar survey density and scan angles, this will not be an easy fix to automate. But the potential for

improved analysis, with consistent input data, makes the work valuable.

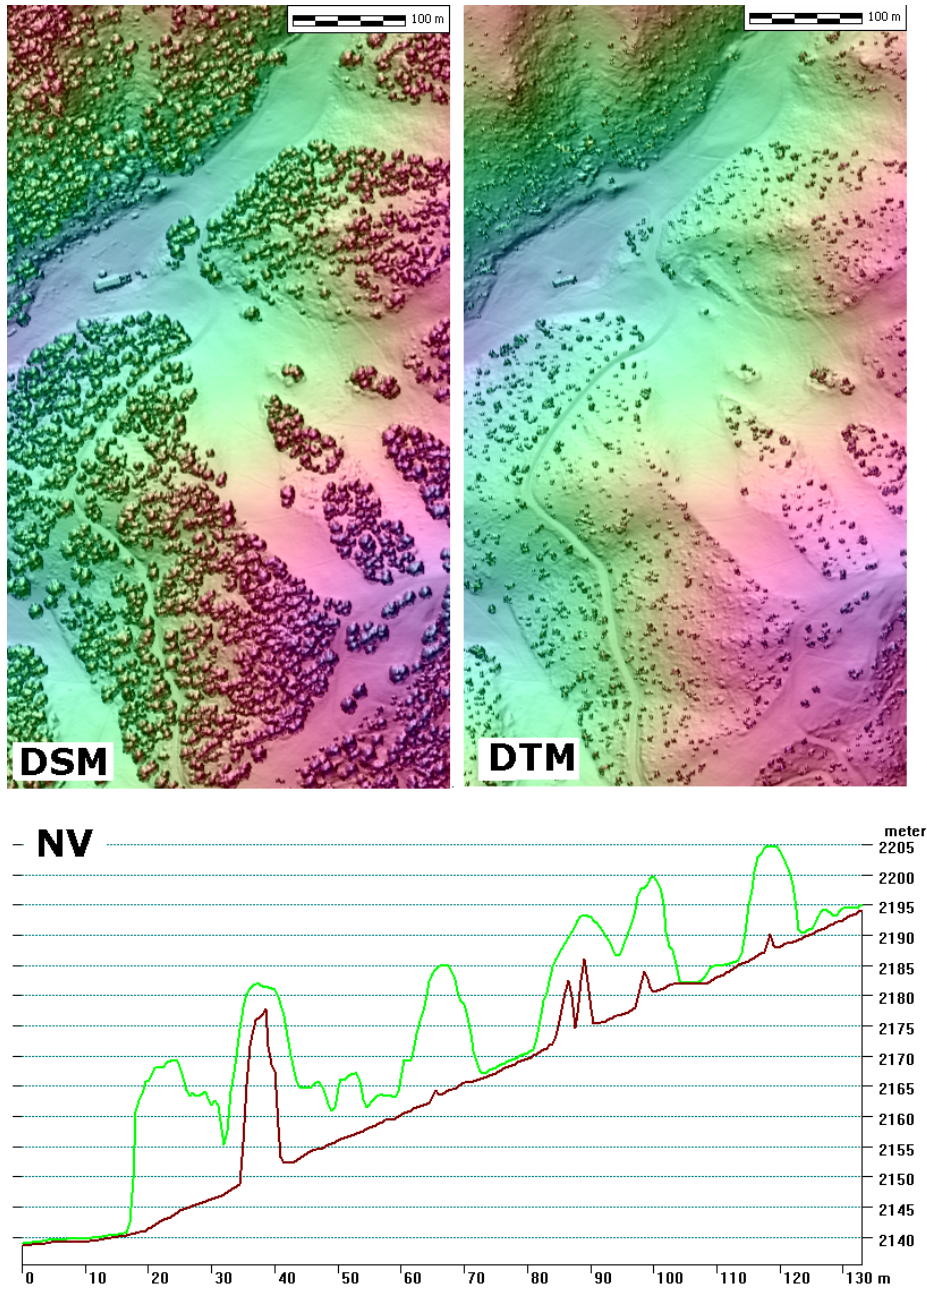


Figure 9. DSM and DTM for the Nevada data set, and a profile through both showing that coniferous trees have the same problems as deciduous trees in blocking lidar pulses from hitting the ground near the trunks.

ACKNOWLEDGMENTS

This work used the freeware MICRODEM program (Guth, 2009), available at <http://www.usna.edu/Users/oceano/pguth/website/microdemdown.htm>, but the algorithms described here could be applied to any GIS software. I thank OpenTopography (2012) for hosting the lidar data and creating the grids on demand. Data acquisition, ground-truthing, vegetation surveys and processing for the Pennsylvania data were funded

and coordinated by NSF Award EAR-0922307 to Qinghua Guo. The Tahoe Regional Planning Agency provided the Nevada data set.

REFERENCES

- Evans, J.S., and Hudak, A.T., 2007, A multiscale curvature algorithm for classifying discrete return lidar in forested environments: *IEEE Transactions on Geoscience and Remote Sensing*, 45(4):1029-1038.
- Guth, P.L., 2009a, Geomorphometry in MICRODEM, In Hengl, T., Reuter, H.I. (eds), *Geomorphometry: concepts, software, applications*. Developments in Soil Science Series, Elsevier, p.351-366.
- Guth, P.L., 2009b, Incorporating vegetation in viewshed and line-of-sight algorithms: Proceedings ASPRS/MAPPS 2009 Conference, San Antonio, Texas, November 16-19, 2009, 7 page paper on Conference CD-ROM.
- Haugerud, R.A., and Harding, D.J., 2001, Some algorithms for virtual deforestation (VDF) of lidar topographic survey data: *International Archives of Photogrammetry and Remote Sensing*, Graz, Austria, XXXIV-3/W4:211-217.
- Llobera, M., 2007, Modeling visibility through vegetation: *International Journal of Geographical Information Science*, 21 (7): 799-810.
- OpenTopography, 2012, OpenTopography: A portal to high-resolution topography data and tools: <http://www.opentopography.org/>, last accessed 16 January 2012.
- Stoker, J., 2009, Volumetric visualization of multiple-return lidar data: Using voxels: *Photogrammetric Engineering & Remote Sensing*, 75(2):109-112.
- Stoker, J., 2010, Making lidar more photogenic: Creating band combinations from lidar information: *Photogrammetric Engineering & Remote Sensing*, 76(3):216-220.