LIDAR LAB 4: POINT INTERPOLATION

GROUND

1st RETURNS

DSM + RGB

TIN

The TIN method uses a Delaunay triangulation. The algorithm creates triangles by drawing lines between data points; in turn creating a patchwork that covers the extent of point distribution. This method results in closely matching results to the original inputs as the face of each triangle is calculated in reference to the elevation of, and tilt between, each contributing point of the triangle.

The above TIN hillshade shows how well it works in areas of evenly distributed points. In general the standard deviation of the surface to validation points is 0.1032, the second closest of the methods tested. However in areas of sparse point coverage, such as the slope visible in the south east portion of the image, the coverage becomes wider for each triangle to meet a point. The lack of points, such as the river or dense tree stands, also leaves empty areas and triangle artifacts. This leaves an overall “choppy” appearance to the end ground result, despite truer values.

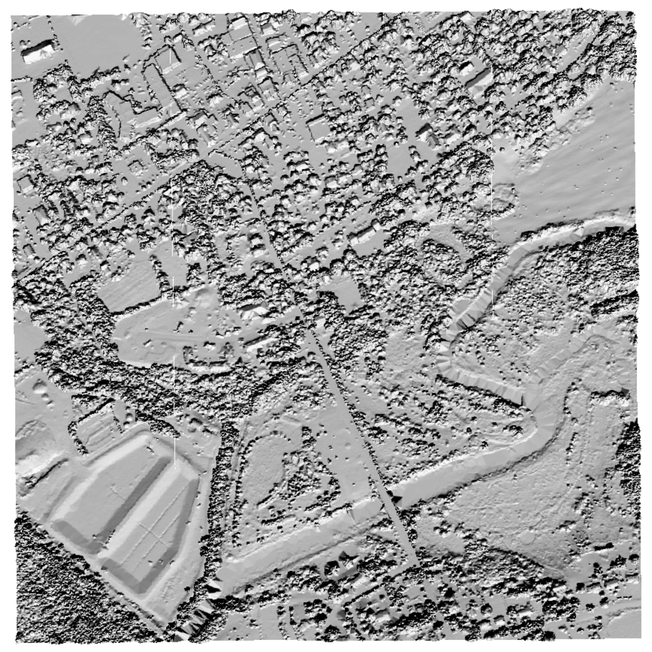

The TIN interpolator also covers the full extent of the first return points. In this view all features above the ground are represented. In general this eliminated the triangular, empty areas seen before in the ground TIN with additional point returns like tree canopy and buildings filling these gaps. However their inclusion also causes problems of their own.

For instance the central cluster of buildings in the above image at first glance do not appear to be buildings. Figure 1. shows an example of these buildings up close, and how the interpolation coverage of these differ from reality. This is due to the distribution of LiDAR points. While a scan might catch the roof top of the building, it can be a slight or drastic connection to the next point on the ground.

This effect can be amplified by the angle in which it was scanned; with more space between the last high point on a building to the ground below, from the roof blocking the ground directly next to it. This seems to distort other objects, such as one to the right of the bridge- most likely a lamp post. These effects become highlighted by placing an RGB image over this grid. Depending on the viewing angle this can be apparent or hidden by other features. Overall the appearance of the RGB image looks rigid and choppy with the elevation values of the TIN method. While this works well in areas with drastic changes (i.e. residential areas) it detracts from the whole, making it less than optimal with a wide range of features.

IDW

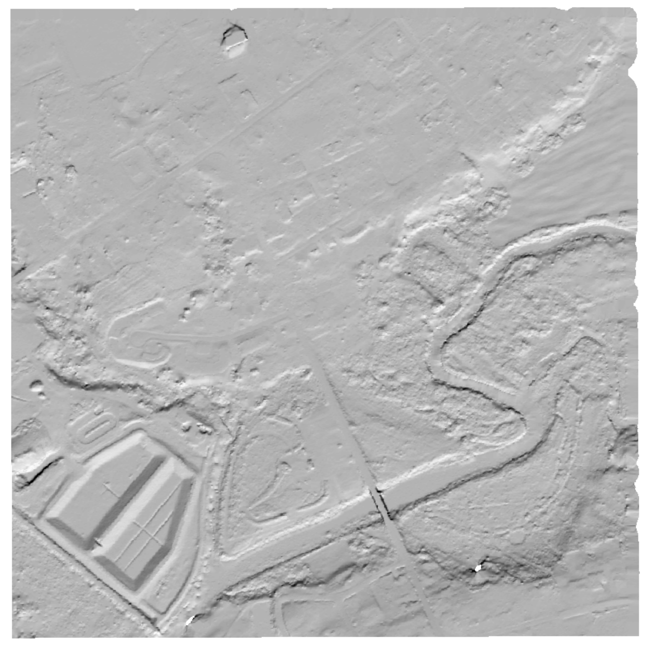

The IDW algorithm is a weighted average interpolator. When determining the value of a grid node a search radius is used to compare this location to the nearby point values. However the node takes on these values in regards to the weight placed on the point’s location. For instance the example above uses a weighted power of two and a search radius of 1. When determining the node value the points farther away, closest to the search radius edge, have less effect on the interpolated value. Thus nearest points tended to be used. However if a smaller power was used the weights would be more evenly distributed among the neighboring data points (within the search radius).

This choice resulted in a smoothed out appearance, as compared to the TIN. This was accomplished even with the smoothing parameters set to zero, an indication of how this method considers the location of grid nodes to points, verses a generalization between points. The concern with using this method was the balance of a search area large enough to avoid gaps, but without so much overlap that the appearance of “bulls-eye” circle effects do not appear at overlapping interpolations at solitary points. While no bulls-eye circles are apparent, there are three areas of no coverage, and the “dinged” effect along the top and right hand edges of the image.

Again, the first return points produce a full coverage with no warped spots from lack of ground data. However in this representation there are some differences in how the features have been represented. For instance the rows that are visible in farming fields are diminished and smoothed over compared to the previous TIN hillshade. This generalizing effect also extends to buildings, whose edges have become more gradual in comparison to other methods. This smoothing/generalizing of the grid would explain why the average and standard deviations for this method are the highest among the four processes attempted.

IDW was also the only hillshade in which points where “smeared” as they do on the top and left portions of the image. This may be a relic of meeting the required cell size or a possible error in running the hillshade function.

When applied to the RGB elevations the IDW proved quite adept at representing open fields and nearby tree stands from this grid’s gently rolling effects. However upon inspection of residential areas the gradual slope from the last roof points to the nearest ground points were exaggerated, producing a “melting” effect of buildings.

NATURAL NEIGHBOUR

The Natural Neighbour algorithm is kind of the catch-all selection of the test group. It works by using Voronoi polygons (created from TIN) and treats each node of the grid as a perspective new vertex of these polygons. In turn the overlap between the original polygons to the altered version is used as a weight, with the average distance from the centroid of the polygons to the node being the source of height information. This latter process is much like IDW since the height of the node is based on the contribution value. As a result of this constant comparison to the original polygons and change iterations this method can take longer to compute, but remains a popular choice of interpolation (Bater & Coops, 2009).

This method subsequently produced the best results when compared to the other outputs. All areas of the grid have been modeled (sans no data zones) without any strange artifacts. In problem areas, such as the forest stand which slopes the South East of the bridge, Natural Neighbour most resembles the TIN results with a large area that looks similar. Unlike the TIN result there are no resulting triangles that stretch over this expanse, smoothing out the result and a better quality grid. Not surprising that Natural Neighbour ranked the best out the residual report both in average and standard deviations.

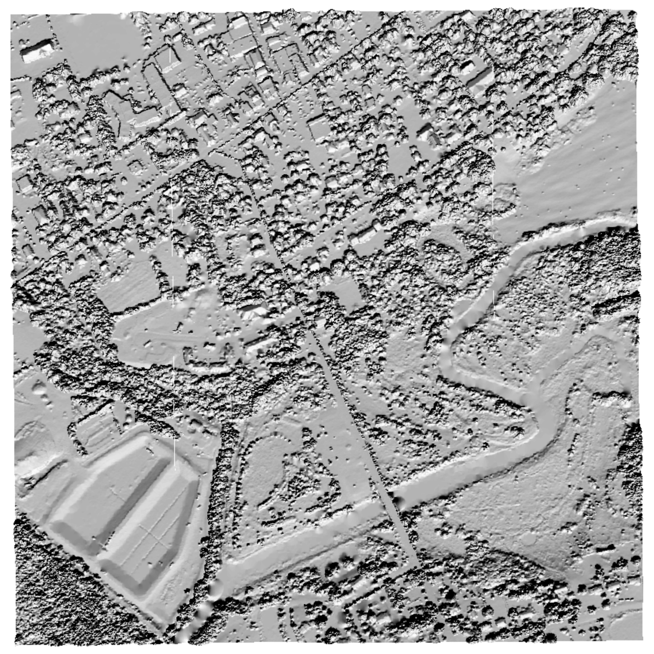

The Natural Neighbour interpolation proved to be adept at modeling first return points. In this representation features seem to be well defined while still employing a smoothing effect; this is observable at the bend in the Annapolis river. Compared to the IDW or TIN the water looks smoother with definition at the river banks. The success of capturing the terrain is from the point spacing of the LiDAR data which provides adequate point density. Despite this there are still sloping effects of the buildings, but are not exaggerated as IDW.

When examining the results in reference to the RGB image not much can be discerned as unique to this method. It is near identical to the TIN and very similar to Nearest Neighbour; minute differences exist between Natural and Nearest Neighbour and are discussed in the latter’s focus section.

Ultimately if choosing the best method for modeling the LiDAR collection subset the Natural Neighbour algorithm would be suggested. It performs the best when modeling ground points while also representing a sort of “average” result of the first return points. It contains no unique errors and looks near identical to two other, acceptable, method results. Considering the terrain features present in the study area it could be extrapolated that Natural Neighbour is a useful algorithm when a mixture of natural and artificial terrain exists.

NEAREST NEIGHBOUR

20 x 20 search range

10 x 10 seach range

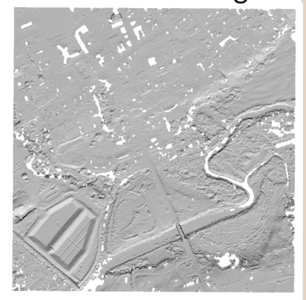

The Nearest Neighbour interpolation method works by finding the closest point to the node location to be determined, within a specified range. This is one of the quickest computations, so the estimated range needed for full ground coverage started at 10 X 10. It was apparent that this needed to be increased once large tracts of the river, tree stands and large buildings were not covered by the grid. Increased to a 20 X 20 search range the outputted grid covered all extents of the point coverage, though with some difficultly. For instance the points on the left edge of the image stretch sharply. Also there appears to be darker edges along the bridge and tree stands .If run again a different option would be to use a 15 x 15 range, to avoid missing areas but avoid the darkened locations.

The overall appearance is quite sharp as compared to the other methods. This is an example of how the values are interpolated. Instead of creating new values each node takes on the exact value of the nearest inputted point. Therefore it is possible that the “repetitiveness” of the values enhances the difference between unlike elevations. When compared to the validation points both sets of deviations were second highest, not unexpected considering this method of obtaining values.

When the Nearest Neighbor Interpolation was applied to the first return points their spacing provided the optimum coverage needed for successful use of this algorithm. In truth not much change is apparent between this method, TIN, and Natural Neighbour; though one example that was noted was the way a pond in the bottom left corner took shape. In this iteration the pond actually forms a more accurate account of the pond in question, whereas the Natural Neighbour has a rounded appearance.

When the RGB image is set with elevation from this grid method some additional differences tend to surface. These include what appears to be added “mass” to features along the extent edges. For instance another road can be viewed along the North West edge of the image below. A little further down from this an obvious error occurs where a higher elevation is continued to the edge, but the RGB appears grey like a road or structure, now appearing as tall as the neighbouring trees. Other locations with trees tend to be slightly larger as they extend further but look natural by corresponding to the RGB.

Therefore without much divergence this method is tied for optimal use with first return points. In some cases it works better in its incorporation of more features, but it’s use for ground elevations eliminates it as an all purpose method in these trials.

FIGURE 1

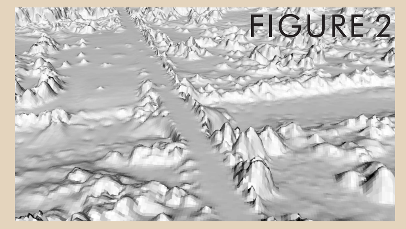

FIGURE 2

FIGURE 3

OVERVIEW

The purpose of this assignment was to compare different methods of point interpolation and evaluate each in relation to ground and first return LiDAR points.

To complete this task LiDAR points were loaded into Microstation TerraScan. A subset of 1km x 1km was extracted from this set and subsequently classified. The location chosen is south central Middleton, which includes a range of buildings (i.e. commercial, residential), construction zones, the Annapolis river, and vegetation (i.e. fields, tree stands). This would provide a variety of elevation features to create a Digital Terrain Model (DTM) with first returns.

Only two classes were specified after outlying points (isolated, low point, etc) were removed; ground and first returns. These sets of points were subsequently brought into Golden Surfer and placed through this program’s gridding function. In total eight grids were made: four types for each point type. These include Triangulation with Linear Interpolation (TIN), Inverse Distance to a Power (IDW), Nearest Neighbour interpolator, and Natural Neighbour interpolator. ArcScene was employed to visualize the results by creating hillshades. All grids were set to an azimuth of 165 degrees and an altitude of 45. This setting was used to imitate the shadows existing in the RGB imagery of the same area.

To assess the accuracy of the outputted grids, two methods were used. The first was the calculation of grid residuals which compared the grids to a validation las file; this contained a series of points from the same study area but was not used in the creation of the grids. Points were compared in reference to how similar or different they were in comparison to the validation set, providing qualitative statistics. To compare the grid types visually, they were placed through a grid math function which separated one from the other to output their difference. These areas indicate the locations of least correlating elevations and ultimately problem areas of the interpolation process.

GRID MATH

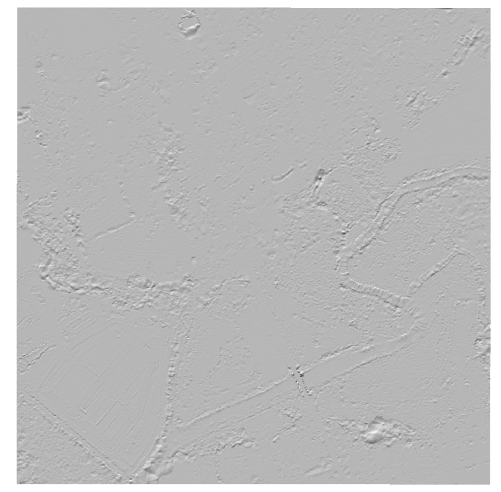

TIN - IDW

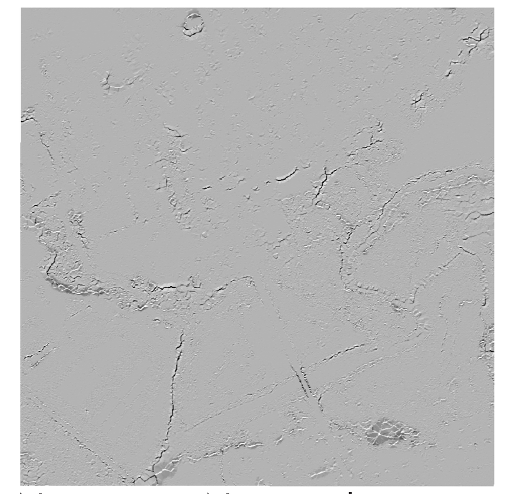

Nearest - Natural

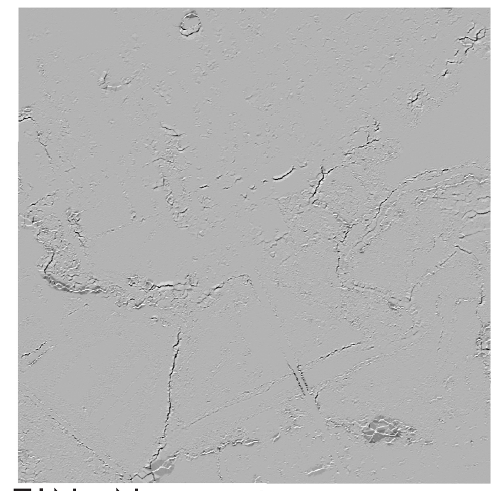

TIN - Nearest

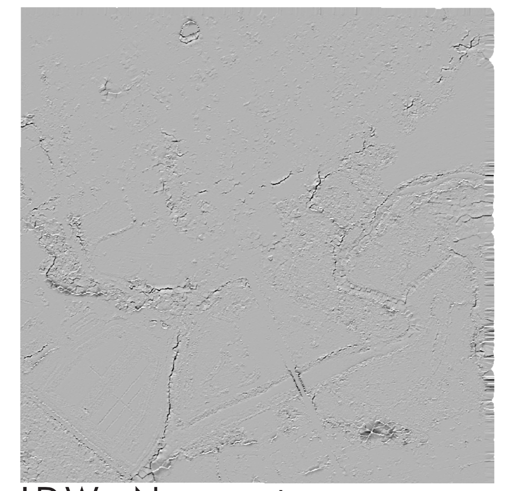

IDW - Nearest

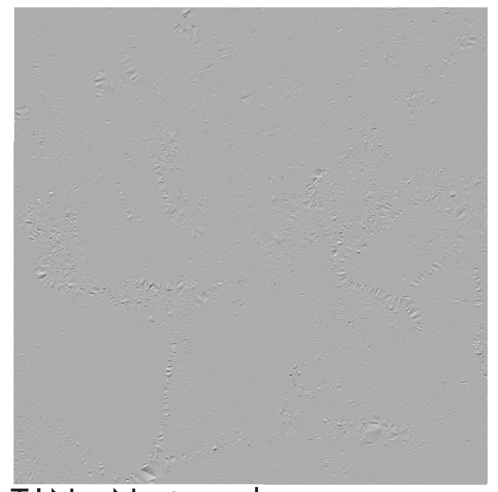

TIN - Natural

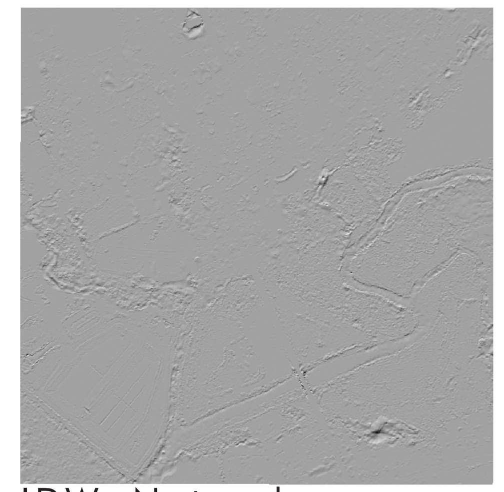

IDW - Natural

The main difference between any of the grid math equations seems to be the river banks and tree stands. This is usually more prevalent in those equations that included IDW or Nearest Neighbour grids, logical given their higher residuals from the validation points. Based on the issues seen between Natural and others, it becomes apparent that it is the most accurate of the four types. Natural and TIN are the most similar of the group, the difference being faint triangle artifacts (obviously contributed by the TIN). Considering the neutral property of the Natural grid, compared to Nearest it highlights some prevalent issues with the latter, including dark distinctions around tree stands, general areas of shading (problem areas), and some building footprints. IDW equations, while not as sharp and dark as Nearest, show deep and distinct differences especially around the river, tree stands and includes more building footprints.

GRID

RESIDUALS	TIN	IDW	Natural Neighbour	Nearest Neighbour
Number of values	21855	22807	21854	22321
Minimum	-2.45636671	-3.459079317	-2.449641528	-2.82
Maximum	1.90103067	2.23078922	1.803285318	2.574884001
Mean	-0.025970614	-0.028111048	-0.026242768	-0.025899474
Average deviation	0.048629065	0.064480784	0.048355757	0.056504411
Standard deviation	0.103210245	0.146727039	0.102122503	0.12496072

LiDAR Data
Acquired by the Applied Geomatics Research Group (AGRG), August 18, 2010 (Julian Day 230). Collected with Optech ALTM 3100 system over Middleton, NS. All data and derived results in coordinate system UTM 20 NAD83 CSRS98, vertical datum CGVD28.

Programs
Bentley Microstation v8 2004 and Terrasolid Suite 2008 for LiDAR point management. Golden Surfer 11.3 2013 for interpolation grids, grid math and residual statistics. ArcScene 10 for hillshades and manipulated imagery. CorelDRAW X6 for poster design. Microsoft Excel 2010 for table generation.

COGS
Centre of Geographic Sciences
nscc

Class: REMS 6082 LiDAR
Produced by: Samantha Cyr
Instructor: T. Milne
Date: April 2013