



GUIDANCE

Lidar Provisioning Guidance for the Digital Coast Data Access Viewer

January 2014

NOAA Coastal Services Center
www.csc.noaa.gov

Lidar provisioning in the Digital Coast Data Access Viewer (DAV) allows users to generate customized data sets, and this requires input on what the user wants as a product. The implications of these choices aren't always obvious. This document provides guidance on the choices available and the results of those choices. The focus is on the checkout procedure, after the area of interest and specific data sets have been chosen, when the screen looks similar to the figure below.

Checkout Notification **Outputs** Delivery Confirmation

Copy choices from: **Select**

Let me edit choices

Dataset #1 (Lidar)
2010 USACE Lidar: Southeast

Dataset #2 (Lidar)
2006 USACE Topo/Bathy Lidar

Projection & Datum Options:

Projection: Geographic(Lat/Lon) Zone: N/A

Horizontal Units: Decimal Degrees Horizontal Datum: NAD83

Vertical Units: Feet Vertical Datum: NAVD88

Output Options:

Output Product: Raster Output Format: Grid - GeoTiff 32-bit

Grid Method: Average Grid Units: Meters Grid Size: 2

Fill small gaps

Data Options: Use Advanced Options

Data Classes:
Ground

Back **Next**

The screen has all the selections grayed out to avoid overwhelming novice users. To make the choices to get the desired product, the first step is to check the **Let me edit choices** box in the upper right.

Projection and Datum Options

Once the editing of choices is enabled, the first set of choices relates to the projection and datum of the output. That dialog box looks appears below.

Projection & Datum Options:

Note: The first option provided is the recommended choice [? Help](#)

Projection:	Zone:
Geographic(Lat/Lon) ▼	N/A ▼
Horizontal Units:	Horizontal Datum:
Decimal Degrees ▼	NAD83 ▼
Vertical Units:	Vertical Datum:
Feet ▼	NAVD88 ▼

Projection

Choices for projection include geographic coordinates, Universal Transverse Mercator (UTM), the U.S. state plane system (1927 and 1983), and Albers Equal Area. For the creation of rasters or contours, the points are first converted to the desired projection, and then the grid is made. This reduces the loss of information resulting from reprojecting rasters. The projection desired is often dictated by the other data sets. UTM is commonly used for scientific studies, while state plane is often required for official work for states. For derived products such as raster digital elevation models, or DEMs, geographic is rarely a good choice, since it introduces the most distortion.

Zone

The zone selection is directly related to the projection. The zones are only applicable to UTM and state plane projections. The applicable zones will become available after a projection is chosen, with the zone that is the best choice for the center of your area of interest being chosen as the default.

Horizontal Units

Most of the horizontal units are self-explanatory and are often dictated by the projection (e.g. geographic data can only be in degrees). UTM is almost always in meters. State plane unit defaults will vary with the state plane zone. The most confusing units are the U.S. feet versus international feet. This is caused by two different adoptions of a conversion from meters to feet that vary slightly (find more information online from the National Geodetic Survey: www.ngs.noaa.gov/faq.shtml#Feet). Projection for the state plane 1927 system should always use survey feet.

Horizontal Datum

A datum can be thought of like a coordinate system, and different datums may differ by their scale, the orientation of their axes, and their origin. For some, such as the World Geodetic System of 1984 (WGS84) and the North American Datum of 1983 (NAD83), the differences are very small, while NAD83 and the North American Datum of 1927 (NAD27) have significant

differences. Pick the datum that is most applicable to the rest of your data. The official datum for the U.S. is NAD83, but many military and scientific applications will use WGS84. While horizontally, these two are often considered the same for mapping applications, you'll need to choose NAD83 if you want a vertical datum of NAVD88 (see the vertical datum section below). There are a few different versions of NAD83. Exactly which one was used for data measurement depends on the survey monuments used in calibration and processing. Versions of NAD83 after NAD83(86) are all very close to each other, but may be a few meters from NAD83(86). The lidar data are generally stored in the version of NAD83 that they were collected in, and no conversions are done to other versions.

Vertical Units

The choices here are feet or meters. There is no distinction made between survey feet and international feet because the differences are well below the noise level for any reasonable values of elevation. Choose whichever works best for your end analysis.

Vertical Datum

Most users will want to use the North American Vertical Datum of 1988 (NAVD88), the official U.S. vertical datum. This is an orthometric datum (i.e., water flows downhill). Other choices include the NAD83 ellipsoid and the National Geodetic Vertical Datum of 1929 (NGVD29). The data are stored in ellipsoid heights and converted to orthometric heights as needed using the latest model for the difference (i.e., geoid model). The original measurements are in the ellipsoid reference frame, so this allows us to retain the data accuracy as new geoid models become available. One thing to note, however, is that the same data set downloaded in an orthometric (NAVD88) datum a few years apart might produce different answers because of the possibility of a different geoid model being applied. The system does not currently allow users to select which geoid model to use, but could if there were user demand.

Note that NAVD88 is only available if NAD83 is chosen for the horizontal datum, even though WGS84 and NAD83 are very close together. This is because the conversion to WGS84 is done as a three-dimensional transform, and the vertical values are no longer appropriate for the NGS geoid models. While a meter shift in the horizontal may not be significant for most mapping, a meter shift in the vertical is. Those who need NAVD88 and WGS84 together should pick NAD83 and NAVD88 and then transform horizontally to WGS84 in a GIS or other package that supports it.

Output Options

The output options will initially look something like those in the figure below. However, the options will change considerably according to the desired output product.

Output Options:

Output Product:
 Raster ▼

Output Format:
 Grid - GeoTiff 32-bit ▼

Grid Method:
 Average ▼

Grid Units: Feet ▼

Grid Size: 6 ⚠

Fill small gaps

Output Products – Points, Rasters, or Contours

Users must choose the type of product they want. The lidar data are a collection of X, Y, Z points with attributes, and the other products are generated from them. Below is a look at all the product types and the output format choices that go with them.

Points

Users who need the full point cloud will be able to select from among a comma-separated ASCII text file of X, Y, Z values, the lidar [LAS format](#), or the [LAZ format](#) (a compressed LAS format). The text file will only contain the X, Y, and Z values, not the attributes of each point, but it can be imported into many different software packages. The LAS format can be read by most, possibly all, software targeted at lidar and many GIS software packages as well. LAS format can also retain information about each point, including pulse intensity, classification (e.g., ground point versus vegetation points), and return number (e.g., this is the first return of three from the laser pulse). The LAZ format is a lossless compression for LAS and can be uncompressed with free software from www.laszip.org. Compression ratios are generally about 7:1, so this can greatly decrease download times. Generally, points are needed by more advanced users who wish to create derived products through their own processes or need to be able to examine the input.

Rasters, Grids, Lattices, Images, DEMs, etc.

A raster is one of many terms used to describe a rectangular grid of cells with values. In most cases for DAV lidar rasters, the values are floating point numbers providing the elevation estimate for that cell. Described below are the options for the format of raster output and then the options for interpolating the points to get the elevation estimates for each cell.

Output Formats for Raster

Grid – Floating Pt. (*.flt)

The floating point grid has two files. The *.flt file contains the cell values as 32-bit floating point binary numbers, and the *.hdr (header) file contains text information needed to interpret the *.flt file, such as the number of rows and columns, the corner coordinates, and the byte order. This file format is the one expected by ArcGIS when doing a float to raster import. It can also be imported into many other packages, although it may require modifications to the header file. The header file lacks information about the projection, so users will need to remember the projection or refer to the provided metadata.

Grid – GeoTIFF 32-bit

The TIFF file format is an industry standard for image processing and can contain many different configurations of data. The 32-bit floating point TIFF contains the elevation values in a similar way to the *.flt file above, but also contains all the information about cell size and georeferencing (including projection) within the file. One lack in the TIFF format is the absence of a “no data” value for those areas where the elevation isn’t known. DAV uses the IEEE standard “Not A Number” value for empty cells. Some programs have difficulty recognizing these as no data, which may result in apparent minimum or maximum bounds with very high exponents. Assigning the min-max manually may correct this. You can find more information in the [Geozone blog post about TIFF](#).

Grid – GeoTIFF 8-bit

This is the 8-bit version of the TIFF format and is essentially a pretty picture. All the values have been scaled to fit in a range of 0 to 255, colored from blue to red, and a color bar provides the link between color and elevation. This format provides a look at the data without requiring sophisticated software but is not useful for analyses where the precision of continuous values are needed.

Grid – ASCII

This is a text format version of the grid designed for import into ArcGIS. It is generally not as space efficient as the other formats but does allow the user to look at the values in a text editor.

Grid Methods

Taking a set of random elevation points and generating a regular grid of values requires some sort of interpolation to determine the value to assign to each cell. Much of the time the differences in output between the methods is not significant, but for some data sets it can result in markedly different results.

Minimum

This is a very simple method where all points within the bounds of a cell are examined and the lowest value is assigned to the cell. This may be of particular value for data sets that haven’t had their points classified for type (e.g., bare earth) and you want to use the lowest point as a proxy for ground.

Maximum

This is the same basic operation as the minimum, except the highest value is kept. As an example, you might want to use the maximum to generate a raster of the top of a tree canopy. Generating both a minimum and a maximum raster and taking the difference could be used to generate an estimated canopy height raster.

Average

The average value of all points within the cell bounds is assigned to the cell. When points

are not classified, the interpolated value will represent the average of all points in a cell and may include vegetation, buildings, ground, and other features.

Smoothed (IDW)

Inverse distance weighting is used to assign the value of the cell. The distance from the center of the cell is calculated for the 12 closest points and a weighted average is used such that closer points have a greater influence on the resultant value. The algorithm is limited in search radius such that points must be within 10 meters to be considered. If there are not six points within 10 meters of the cell center, the cell is given a “no data” value. This is fairly slow to generate compared to the other grid methods.

Spline

The spline method uses the ANUDEM 5.2 program (<http://fennerschool.anu.edu.au/research/products/anudem-vrsn-53>; an older version is used in the ArcGIS 9x Topogrid tool). This is the only option that will interpolate to fill the grid regardless of how far away valid data are. While this can be very useful to fill in small gaps if a complete grid is required, it can also lead to clearly incorrect results at long distances from real data points. This is particularly evident if the program extrapolates into the water from a topographic lidar set, although it can work well to fill the missing surf zone area in a topobathy data set.

Grid-Cell Size

You can select the grid-cell size for the raster output. The default value has usually been chosen to provide a grid without too many empty cells, though it won't be finer resolution than 1 meter cells. You can change the value, but you can't select a finer grid than the nominal point spacing of the data. A larger cell size will result in a smaller output file and a smoother elevation model. A smaller cell size will result in more detail and a larger output file, but may also result in more holes in the output elevation model. Note that the density of lidar points on the ground will usually vary with the land cover. A grid-cell size that barely has the cells filled in open terrain will tend to have many holes in forested areas for the same data set.

Filling Small Gaps

Several of the interpolation methods only use the points within a grid cell to determine the value for that grid cell. There is an option to use inverse distance weighting to fill in cells that had no points as long as they are close enough to cells that did have points. The input for the calculation is solely from the raster, so it will be much faster to calculate an average raster and then fill the small gaps than it would be to do inverse distance weighting (smoothed method) for the entire raster from the points. The process will require three valid grid cells within a distance of three cells (i.e., a 7x7-cell box) before it will interpolate a given cell.

Contours

There are two output formats for contours. The shapefile is a common format used by GIS systems. The DXF format is typically used by computer-aided design (CAD)-oriented systems. In all cases of contour generation, the points are first converted to a DEM, and then the contours

are derived from the raster. Thus, all the options for a raster above are relevant for the contours. In general, picking a small grid-cell size will result in a less visually pleasing contour. The contour routine will not cross large areas with no data and will break the contours; thus, too small a grid-cell size will result in many very small segments. Interpolation choices are important here. The spline method will fill the grid and provide unbroken contours. However, it will not be clear where the obscured areas are that should have significantly less confidence in the location of the contours. When generating contours, the user will also have to select a contour interval or a single contour value. For example, provisioning contours at an interval of 2 feet will result in contours at 0, 2, 4, 6, n feet. Map accuracy standards provide guidance on contour intervals that can be supported for a given data accuracy. In general, the data should have a vertical accuracy of 9.25 centimeters root mean square error (RMSE) to support 1-foot contours.

Another option is to select a single value. For example, if the mean high water tidal datum was 1.5 feet above the North American Vertical Datum of 1988 (NAVD88) and the user wanted to generate the mean high water line, he or she could select a single contour at 1.5 feet. This would give a single elevation line instead of a series of spaced lines.

Shapefile Format

The shapefile format is documented at www.esri.com/library/whitepapers/pdfs/shapefile.pdf. It is a relatively simple format and can be read by most, if not all, GIS systems. However, one important limitation is a maximum 2 gigabyte (GB) file size. There is currently no method to estimate the size that a contour set will be prior to doing the contouring, and there is a potential for data sets with high variability to fail. Small cell size, small contour interval, and contouring all points instead of ground points are leading causes of failure.

DXF Format

The DXF format is documented at www.autodesk.com/techpubs/autocad/acad2000/dxf/dxf_format.htm. It is generally used by CAD-oriented programs such as AutoCAD or MicroStation.

Data Options

The system provides two methods to filter the points used to make your final product. This allows you to eliminate points that would be inappropriate for your particular needs.

Unfortunately, not all data sets can support the filtering methods, and some options may not be available for a specific set of data. Note that the filtering methods only select data points based on information already in the files; they do not classifying data on the fly. The data options section (with advanced options turned on) is shown below.

Data Options: Use Advanced Options

Adv. Data Classes:

Unclassified
Ground
Water Surface

Ctrl-Click to multi-select

Return Types:

Any ▼

Data Classification

The majority of data, particularly county-wide data, has been classified by the vendors. In general, the classifications are limited to ground, water, and unclassified, but some data sets will have additional types identified. People looking for the points in LAS format generally want all the points, since they will have software that can extract, and possibly reclassify, the points as needed. However, for the generation of rasters or contours, the type (classification) chosen will dramatically affect the product. A common product is a raster or contour set representing the ground elevations. In this case, the user would want to make sure that the unclassified points were not included in the calculations, since they would include trees, buildings, and so on. Generating contours can be particularly problematic if vegetation points are included and can exceed the file size limit of a shapefile. Contour generation may be restricted to only those data sets that have ground classifications. If the advanced options are turned on, the user will be able to multi-select the point classes desired from all the classes available in that data set. If the advanced options are not selected, at most the user will choose between ground and all classes. In this case, ground will include all the classes considered to be solid earth (ground, model-key points, bathymetric bottom).

Return Types

To access the return types option, users must check the advanced options box. Most lidar data are collected with multiple return systems such that a given laser pulse may be reflected off several surfaces, for example, first the tree top, then a branch, and finally the ground. The DAV system allows users to choose all the returns, only the first returns, or only the last returns. Generally users will want all the returns, but there are cases where the other options are helpful. For example, to make a map of the canopy height, the user could make two raster requests. The first would be for the ground points using the classifications filter to make a ground surface. The second would be using all first return points to generate a top of canopy surface. The final step is to subtract the two rasters in a GIS or remote sensing package to create the canopy heights. While more subtle filtering could be done (e.g., looking at points inside the canopy by ignoring the first and last returns), this is better done by downloading all the points in LAS format and working with software designed for lidar.

How Does the System Work?

Only the points are stored, and all products are generated from the points. The points are primarily stored in [LAZ format](#), geographic coordinates, and ellipsoid heights. We store the data in ellipsoid heights because this is the reference frame of the original measurements. As improved models become available to transform from ellipsoid heights to orthometric heights (e.g., to NAVD88 using a geoid model such as GEOID12a), the models can be incorporated into the system without having to transform all the underlying data. Vertical transformations are done as needed using published methods from the National Geodetic Survey (VERTCON, GEOID12a, VDatum). Horizontal transformations (projections) are done using programs employing the General Cartographic Transformation Package from the U.S. Geological Survey.

To generate raster output, interpolation programs were written in-house for minimum, maximum, average, and IDW. The spline method uses the ANUDEM program Version 5.2 (<http://fenner.school.anu.edu.au/research/products/anudem-vrsn-53>).

To generate contours, the Generic Mapping Tools program pscontour is used (<http://gmt.soest.hawaii.edu>).

Examples

Raster Generation – Resolution

To illustrate how the provisioning options affect an output raster, several rasters of the same area were generated. In the first set, the grid-cell size was varied in a series of 1, 2, 5, and 10 meters for an area in Charleston, South Carolina. The original data were intended to support a 2-meter raster. The interpolation routine was set to averaging, and only points classified as bare earth were used, hence the holes where buildings are. Filling has been turned on to fill small holes relative to the pixel size, but large holes remain unfilled where point spacing is below the interpolation method's threshold. The data are shown such that the 2-meter raster is at full resolution. The 1-meter raster is therefore not showing the true detail it would have but is here to illustrate the effects of choosing a cell size below what the data support.



Figure 1. Charleston, South Carolina, lidar at 1-meter resolution. Note the many holes where the point spacing is below the threshold.

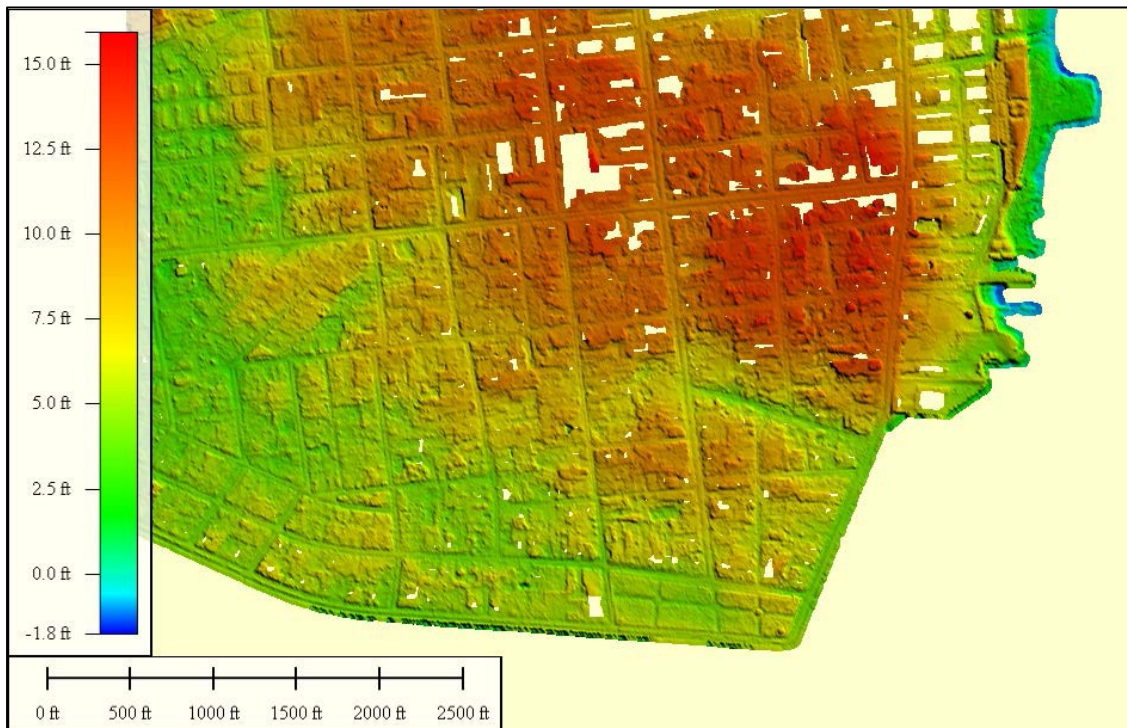


Figure 2. Charleston, South Carolina, lidar at 2-meter resolution.




Figure 3. Charleston, South Carolina, lidar at 5-meter resolution. Much of the detail has been lost, but many holes are filled. Filling is noticeably expanding into the water.




Figure 4. Charleston, South Carolina, lidar at 10-meter resolution. All holes have been filled, but filling has extrapolated into the water and the finer detail is gone.

Raster Generation – Interpolation Methods

The choice of interpolation method makes considerably less difference than the resolution in data that has been classified (filtered). Shown below are 2-meter resolution results from the same Charleston, South Carolina, area but using interpolation methods of minimum, inverse distance weighting, and spline instead of the average method used above. In general, they are very similar, with the exception of the spline. The spline fills the entire area of interest, even if this involves significant extrapolation and fills the Charleston harbor elevations based upon the Charleston peninsula elevations. The spline has many positive attributes, since it fills all the holes and maintains the detail; however, as with any method that fills the holes, we may no longer know where there is uncertainty due to filling.

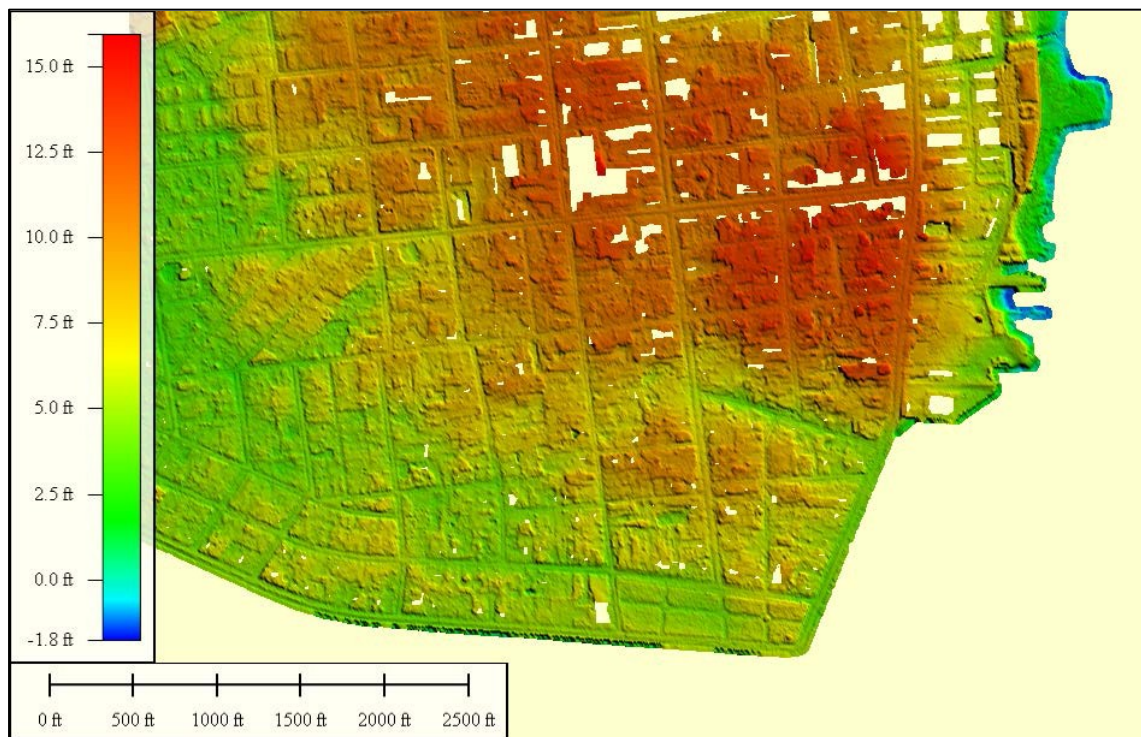


Figure 5. Charleston, South Carolina, lidar using the minimum interpolation method. Values represent the minimum point in each cell.

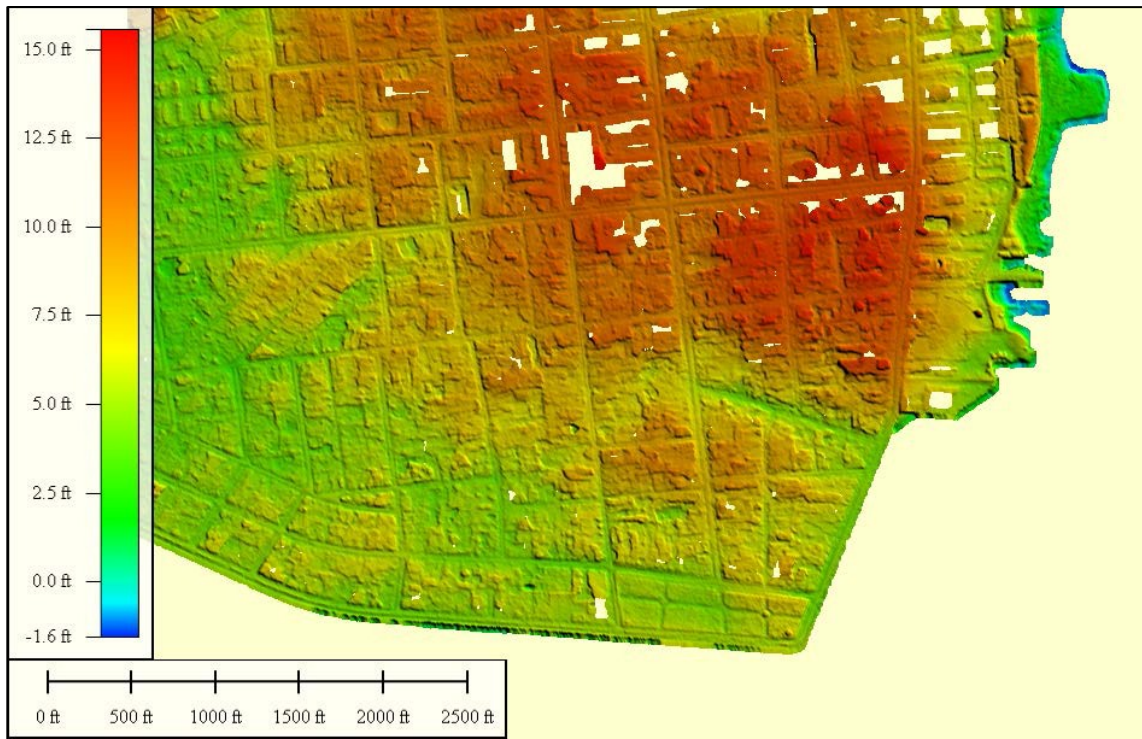


Figure 6. Charleston, South Carolina, lidar using inverse distance weighting (IDW) interpolation method.




Figure 7. Charleston, South Carolina, lidar using the spline method from the ANUDEM program. Detail is maintained and all holes filled, but significant extrapolation happens in the water.

Contour Generation

Contours are primarily a visual product meant for human interpretation, and we are often looking for smooth pleasing lines instead of the jagged ones that might more accurately fit the data.

Below are several examples of contours generated for the same area from the same data set in Charleston, South Carolina. The area includes a freeway overpass that has been classified as non-ground, and therefore its points are excluded from the contouring (an image of the general area is shown in Figure 8). DAV generates contours by first generating a raster grid. The alternate method of generating contours from a triangulated irregular network (TIN) is not currently available. This means the horizontal resolution selected for gridding will influence the contours. Generally, higher detail and noise will be seen as the grid-cell size is reduced. A 1-meter grid (Figure 9) has considerable noise in the contours. This grid was created using the average interpolation method and filling small gaps. The three empty strips are the overpasses of the I-26/I-526 interchange. These structures were too big for the gaps to be filled by a 1-meter grid because they are greater than 7 meters wide. Figure 10 also uses the average method and fills small gaps but on a 5-meter grid. This has smoother lines, and the grid size is large enough that the overpass gaps are filled. Figure 11 and Figure 12 show the spline interpolation method for the same 1- and 5-meter grid-cell sizes. These provide smoother contours than the average method at a given resolution, and the 1-meter grid has interpolated across the overpasses.



Figure 8. Image of general area used in contour illustrations.




Figure 9. Two-foot contours generated with 1-meter grid using the average interpolation method. Note the missing data from the overpass.




Figure 10. Two-foot contours generated from a 5-meter grid using the average method.



Figure 11. Two-foot contours generated on a 1-meter grid using the spline method.




Figure 12. Two-foot contours generated on a 5-meter grid using the spline method.