Lidar Analysis in ArcGIS® 10 for Forestry Applications

Lidar Analysis in ArcGIS 10 for Forestry Applications

An Esri White Paper

Contents   **Page**

Lidar Analysis in ArcGIS 10 for Forestry Applications

Executive Summary

Foresters use light detection and ranging (lidar) data to understand the forest canopy and terrain, which helps them with forest management and operational activities. Combining lidar data with Esri® ArcGIS® helps analysts assess forest health, calculate forest biomass, classify terrain, identify drainage patterns, and plan forest management activities such as fertilization, harvesting programs, development activities, and more.

This paper will step through processes to convert lidar data into a format ArcGIS can process, explain methods to interpret the lidar data, and show how ArcGIS can disseminate the data to those who are not geospatial analysts. It will present methods for reading raw classified lidar data and demonstrate methods for

- Analyzing and validating lidar data with ArcGIS before any extensive processing occurs

- Storing and managing millions or billions of lidar points within the geodatabase in a seamless dataset, regardless of the number of original lidar files

- Processing to extract digital elevation models (DEMs) and digital surface models (DSMs) from the lidar data and store them as terrains in a geodatabase or as raster elevation files

- Extracting vegetation density estimates and tree height estimates from lidar, which aid in growth analysis, fertilization regimes, and logging operations

- Serving and analyzing large amounts of lidar data as a seamless dataset to geographic information system (GIS) clients

In all areas, ArcGIS is a complete system for managing, storing, and analyzing lidar data. Coupling ArcGIS Desktop with ArcGIS Server, the forestry professional is able to access large amounts of lidar data quickly and efficiently without the need to produce additional resultant datasets.

Introduction   ArcGIS can be used to analyze and manipulate lidar data to provide useful results for the end user. This paper provides the processes to analyze and manipulate lidar data and details how to

■   Check the supplied data.

■   Read and separate the point cloud data into ground and canopy returns.

■   Pass the resultant point clouds to a terrain dataset that creates a viewable and displayable surface.

■   Perform analysis on the terrain dataset for tree height delineation and canopy density.

■   Pass the terrain data to the ArcGIS Server Image extension for dissemination to a wider audience as a seamless viewable surface that can be accessed from GIS technology.

What Is Lidar?   Lidar stands for light detection and ranging. In its most common form, it is an airborne optical remote-sensing technology that measures scattered light to find range and other information on a distant target. Similar to radar technology, which uses radio waves, the range to an object is determined by measuring the time delay between transmission of a pulse and detection of a reflected signal. Instead of radio waves, lidar uses much shorter wavelengths of the electromagnetic spectrum, typically in the ultraviolet, visible, or near-infrared range.

This technology allows the direct measurement of three-dimensional structures and the underlying terrain. Depending on the methodology used to capture the data, the resultant data can be very dense, for example, five points per meter. Such high resolution gives higher accuracy for the measurement of the height of features on the ground and above the ground. The ability to capture the height at such high resolution is lidar's principal advantage over conventional optical instruments, such as digital cameras, for elevation model creation.

In addition to the height attribute being captured, each return can capture other attributes such as the return intensity, return number, and number of returns. The existence of these attributes is dependent on the lidar data supplier and the supply order. These additional attributes can all be used in the analysis of lidar data.

The intensity value is a measure of the return signal strength. It measures the peak amplitude of return pulses as they are reflected back from the target to the detector of the lidar system. Intensity is often used as an aid in feature detection and, where conventional aerial photography is not available, can be used as a pseudo-image to provide the context of the lidar acquisition area. See the image below:

Figure 1
Lidar Intensity Image

Lighter areas represent strong returns. Darker areas represent weaker or partial returns.

In forestry, lidar can be used to measure the three-dimensional structure of a forest stand and produce a model of the underlying terrain. The structure of the forest will typically generate a first return from the uppermost limit of the canopy, followed by less intense returns through the canopy, down to the underlying terrain. Returns are classified into ground and aboveground sources. The ground returns can generate a detailed terrain of the area of interest, while the canopy returns can be filtered to provide forest structure at the canopy and middle level of the forest.

Advantages to the Forest Industry

The ability to simultaneously visualize the ground and model the canopy structure provides significant advantages to the forest industry. Traditionally, foresters and land managers have relied on topographic maps for terrain classification and field-based surveys to obtain tree volumes and height information. Lidar data provides significant improvements over both these techniques.

Existing topographic maps depict contours and rivers, which have been, for the most part, captured from aerial photography using stereographic terrain generation techniques. In areas where the tree canopy obscures the underlying terrain, interpretive methods are used to depict where streams and contours occur. Terrains generated from lidar data more accurately represent these geographic features. Lidar penetrates the tree canopy to return a more accurate interpretation of the ground surface. This increases the accuracy of terrain classification and thereby the resultant interpretation and analysis of the geographic features.

Lidar has provided significant benefits for forest development and engineering operations including locating roads, harvest planning, forest regeneration, and more. The ability to identify suitable creek crossings, determine optimal routes, and locate previously unmapped historic roads aids in reducing costs and creating operational efficiencies.

Lidar has also offered an improvement to existing forest inventory methods and procedures. Traditional field-based timber inventory methods are based on measurements derived from systematically sampling plots in forest stands. This statistical sampling method is most often used in forests where measuring every tree is impractical. Tree volumes and heights are calculated in each sample plot, then generalized throughout a forest stand that shares similar characteristics. Estimated results help describe stand characteristics but are inaccurate due to variability in growing conditions throughout the forest, sampling bias, and lack of precision. In addition, the time to collect such measurements is both lengthy and expensive, as many sample plots may be required to provide a reliable representation. Lidar can overcome these limitations.

An increasing number of forestry and land management organizations are using lidar for forest inventory measurements. A wide range of information can be directly obtained from lidar including

- Digital elevation models
- Tree heights and digital surface models
- Crown cover
- Forest structure
- Crown canopy profile

Postprocessing of lidar data can reveal

- Volume—Canopy geometric volume
- Biomass—Canopy cover
- Density—Height-scaled crown openness index and counts of delineated crowns
- Foliage projected cover—Crown dimensions

The forest industry is requiring increasingly precise inventories to guide forest management activities. Using lidar data, forest inventories can be conducted at nearly the single tree level, offering more accurate representations of the true forest stand structure.

For forest inventory activities, lidar has been used primarily to retrieve basic structural tree attributes including height, canopy cover, and vertical profiles. These attributes can be used to derive other critical forestry measurements including basal area and timber volume, as well as biomass for alternative energy and carbon sequestration analysis. This paper will address these attributes.

Managing and Understanding Lidar Data

Understanding Raw Lidar Data

Before any analysis is performed with lidar data, the data received must be checked for any inconsistencies. Lidar data can be delivered in either binary .las format or ASCII .xyz files. The LAS file format is a public binary file format, developed by the American Society for Photogrammetry and Remote Sensing (ASPRS), that is an alternative to proprietary systems or a generic ASCII file interchange system used by many data providers. Details on the format can be found at asprs.org/society/committees/lidar/.

Although a data provider will endeavor to provide the best quality data to its clients, there is always a chance a client will encounter anomalies in the data. These can be in the form of irregular minimum bounding shapes or holes in the sampling. It is therefore necessary to check the quality of the data before performing any analysis.

The Point File Information tool in Esri's ArcGIS Desktop 3D Analyst™ assists in performing data quality assurance checks.

Point File Information Tool

The Point File Information tool, found in the 3D Analyst toolbox in ArcGIS (ArcToolbox\3D Analyst Tools\Conversion\From File\Point File Information), reports important statistics about the raw lidar files.
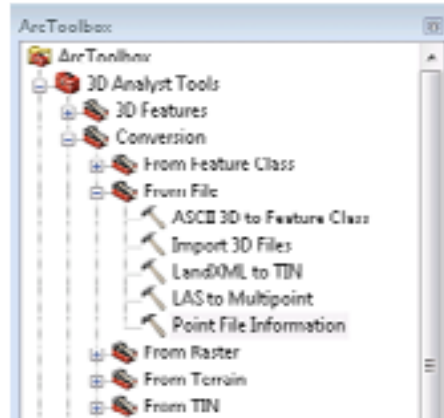
The tool is designed to read the headers of LAS or scan ASCII files and summarize the file contents. As a single lidar file often contains millions of points and many lidar datasets contain more than one file, the Point File Information tool can accommodate reading one or more files by specifying either individual lidar files or folders.
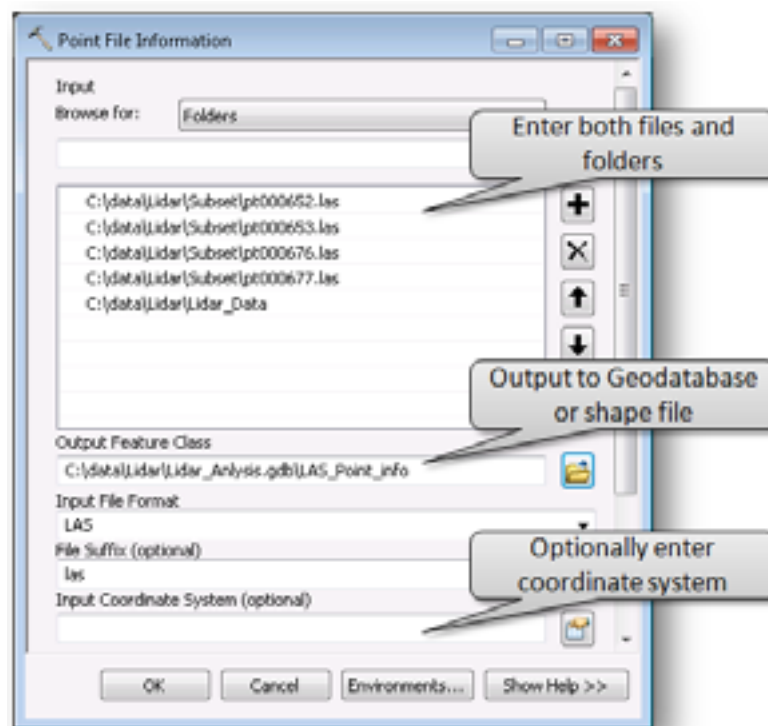
The result from this tool is a feature class that shows the

■ Minimum bounding rectangle for each file
■ Number of points recorded
■ Average point spacing
■ Minimum/Maximum z-values

When the feature class is loaded into ArcMap™, the minimum bounding rectangle of each lidar file is drawn. Lidar data files are usually uniform in size, so if any of the feature shapes appear large or irregular compared to the majority of features from the feature class, they will appear different in ArcMap, which will alert an analyst, who can then refer to the corresponding lidar data file for further investigation.

The average point spacing is important and should be uniform throughout the data files. If any of the files have an average point spacing that is significantly larger than other files, this may indicate incorrect sampling. In addition, average point spacing is important when building geodatabase terrains and converting lidar files to feature classes.

The average point spacing is a product of the total number of points divided by the area of the lidar data file. In cases where a lidar data file is only partially covered by points, such as along a coastline, the average point spacing will be calculated to be greater than the point spacing of the area sampled. These anomalous files would still be used in the dataset, but their calculations would be excluded from further processing.

J-9999

The images that follow show three of the elements as reported by the Point File Information tool, including

■   A uniform grid showing the extents of each lidar file.

■   The attribute table associated with the lidar extents, showing the average point spacing, point count, minimum and maximum z-values, and originating file names.

■   The average point spacing as indicated by the statistics from the Pt_Spacing column. In this example, the average point spacing tends to be approximately 0.6 meter. The lidar dataset used in this paper was captured at a sampling density of two returns per square meter; thus, 0.6 meter gives a good approximation to the ordered capture rate. Again, if there were any significant outliers in the files, these would be highlighted for further inspection.
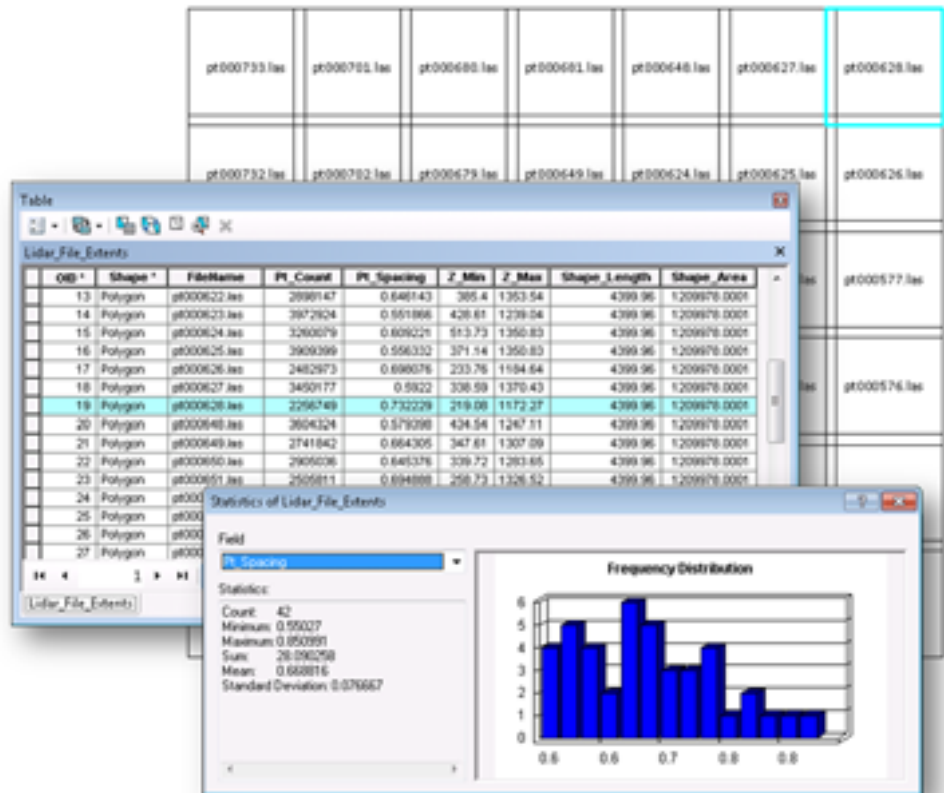

| Lidar Classification in ArcGIS | LAS files contain a classification field that identifies each point's return type. This is known as the class code. A classification describes the point return as Ground returns, Canopy returns, Building returns, or Unclassified. This is useful in determining the content of each lidar file. |

Having the classification field available as part of the tool immediately identifies whether the lidar file has been classified and whether it can be used for interpreting the terrain or forest structure. If no classifications exist, either there is a problem with the file, which may need to be updated, or the lidar file has not been classified at all. The field also helps the analyst when interpreting data where no documentation exists. It provides a good understanding of the file's content and how it has been classified.

ArcGIS reads the classification field with the Point File Information tool. Toggling the Summarize by class code causes the tool to scan through the LAS files and analyze the class code values. The attribute table of the output feature class will contain statistical information for each class code encountered.
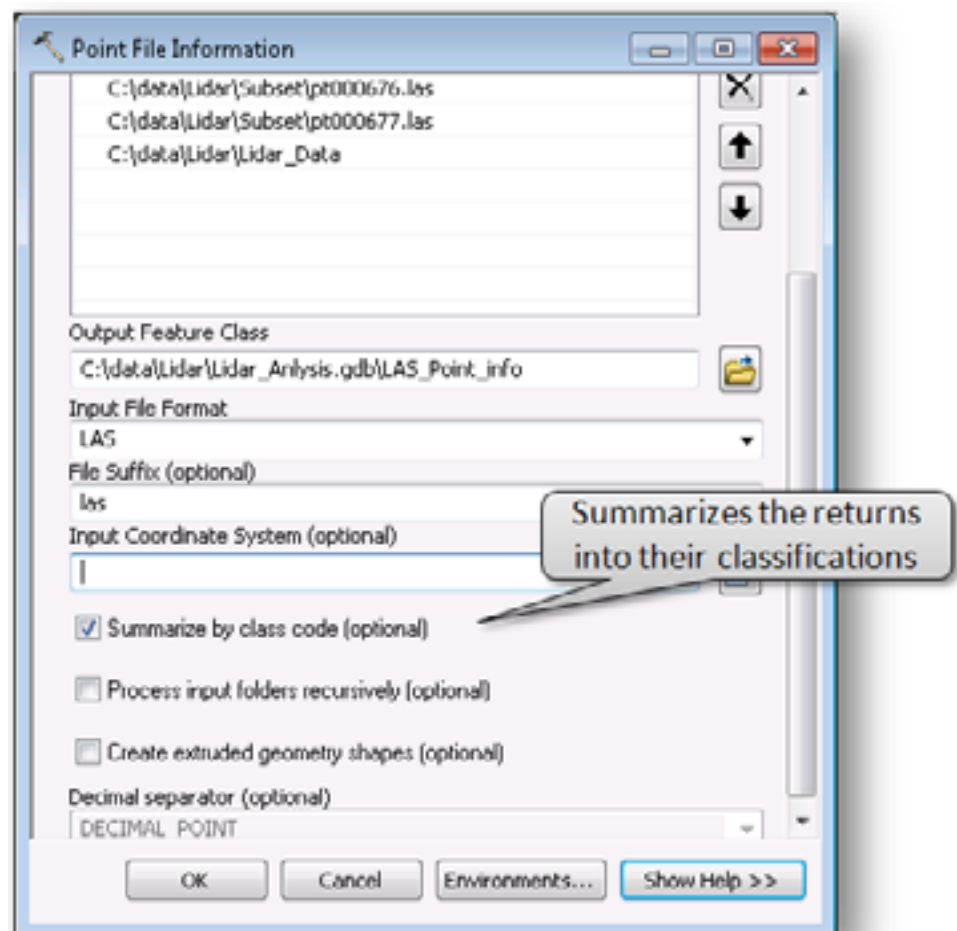
Loading the Lidar Files to ArcGIS

Lidar data is characterized by very dense collections of points over an area, known as point clouds. One laser pulse can be returned many times to the airborne sensor. A pulse can be reflected off a tree's trunk, branches, and foliage as well as reflected off the ground. The diagram below provides a visual example of this process.
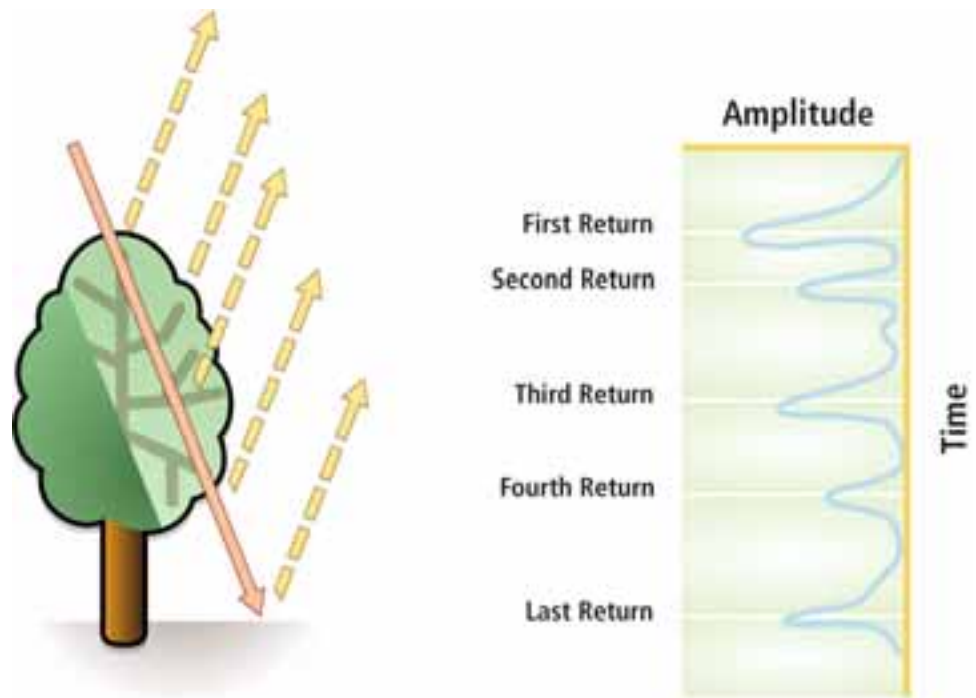
These multiple returns create a data management challenge. A single lidar file can typically be 60 MB to 100 MB in size and can contain several million points. If this data is loaded directly into a table, it creates many millions of records, which results in a large data file that becomes difficult to manage. This challenge is overcome by loading the points into the geodatabase feature type known as multipoint.

Multipoints are used to store thousands of points per database row. This is achieved by storing the geometries in a Shape field and optional attributes in an Esri binary large object (BLOB) field. The Shape field and the BLOB field are collections of binary data stored as a single entity in a database. Storing the data as multipoints allows optimized compression of shapes, reduces storage requirements, and improves database performance. Any tool written to exploit these fields needs to understand the Esri BLOB and Shape fields structures.

The tool to load lidar files into the geodatabase is called LAS to Multipoint. It is part of the 3D Analyst toolset in ArcGIS (ArcToolbox\3D Analyst Tools\Conversion\From File\LAS to Multipoint).

LAS to Multipoint Tool
The LAS to Multipoint tool enables the user to read the lidar data files and load them into the geodatabase. Many lidar analysis applications on the market today can perform detailed analyses against lidar files, but only on individual files. Loading the lidar files in a geodatabase allows a seamless mosaic of the entire lidar dataset, which then can be analyzed by ArcGIS tools.

When using the LAS to Multipoint tool, all the points can be loaded into the geodatabase. This is useful for producing a point density map; however, this is not useful for canopy and ground analysis. For this type of analysis, it is better to separate data into unique classifications.

LAS files captured since September 2008 should conform to the LAS 1.2 specification. This specification allows the separation of lidar data into ground returns and nonground returns by the classification field. A full description of the specification can be found at asprs.org/society/committees/standards/LiDAR_exchange_format.html.

Lidar datasets captured prior to this will often contain the classification in the metadata that is associated with lidar data.

The table below is an extract from the specification and describes the classification codes. When lidar data is provided as part of a data order, the classifications would normally be provided as part of the delivered documentation.

Classification Value	Description
0	Created (never classified)
1	Unclassified
2	Ground
3	Low Vegetation
4	Medium Vegetation
5	High Vegetation
6	Building
7	Low Point (noise)
8	Model Key Point (mass point)
9	Water

When the LAS data files are read by the LAS to Multipoint tool, it can accommodate these classifications and separate them into unique feature classes.

The LAS 1.2 specification also defines how to separate returns from the first return through the last return. The return value is stored in the LAS file with the point information. Having these values enables the extraction of the upper canopy data based on the first reflected values. Midcanopy and ground values are reflected at any time.

The LAS to Multipoint tool needs certain specifications depending on the project. In the following screen capture

■ A folder is specified for the data source. (Individual files can be specified, but when large amounts of lidar data are being read, then it is a best practice to specify the folder.)

■ The Output Feature Class is specified in a file geodatabase. (Multipoint feature classes exist in geodatabases and shapefiles, but geodatabases are preferred due to the extended capabilities of the geodatabase and the size restrictions of a shapefile.)

■ The ground spacing is specified. (This was acquired from the Point File Information tool or from the supplied metadata.)

■ The Input Class Code to be extracted is specified. (The specific code entered will vary depending on the classification being analyzed.)

■ The returns are focused on the ground returns.

■ The LAS file extension is designated as a .las file.

■ The coordinate system is specified.

Other specifications may also be considered:

■ If the goal is to create a canopy surface, then the Input Class Codes need to be specified as 5 and the Input Return Values as first returns only. If the supplied lidar data has only been classified as ground and unclassified, the Input Class Code is 1 for Unclassified, and the Input Return Value is first.

■ If the entire canopy is to be modeled, then the Input Class Codes need to be specified as 3, 4, and 5 as these contain the upper, middle, and lower canopy returns, respectively.

■ If the goal is to produce a ground surface, then the Input Class Codes need to be specified as 2 and the Input Return Values as ANY_RETURNS.
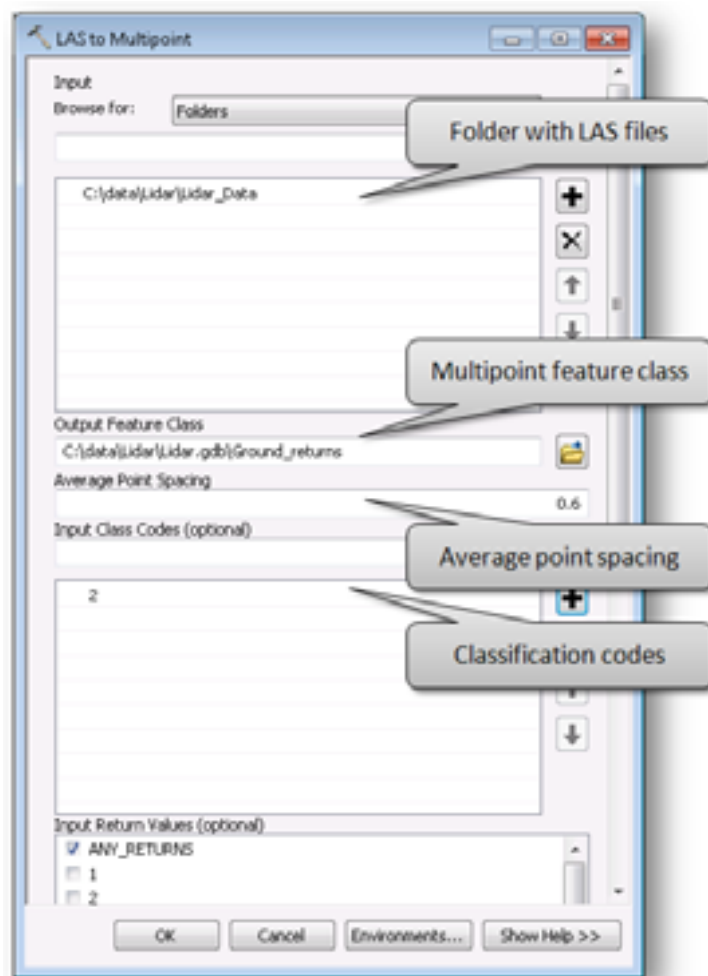
The screen capture below represents ground returns using a series of points. In its raw form, a multipoint feature class is not designed to be displayed; it is designed as an efficient storage medium for the many millions of points found in a lidar dataset.

In this example, the points appear to be merging into a single dense mass of points, as there are so many points contained in the feature class.

The ArcMap screen capture also shows the attribute table with the Shape field as a Multipoint Z, the Intensity field is a BLOB, and the PointCount field shows how many points are stored per record. The PointCount in this case is showing 3,500 records per row. This is 3,500 points per record in the geodatabase. In a normal point feature class, there is one point per record. Having many points per record enables the feature class to be highly compressed, thus making the geodatabase an efficient method to store and manage lidar data. The intensity value is a BLOB record. This means that each intensity value is linked to each point via a specialized method in the Intensity field. To access this BLOB field, the program needs to understand how to interpret the BLOB field.

Figure 2
Ground Return Multipoint Feature Class
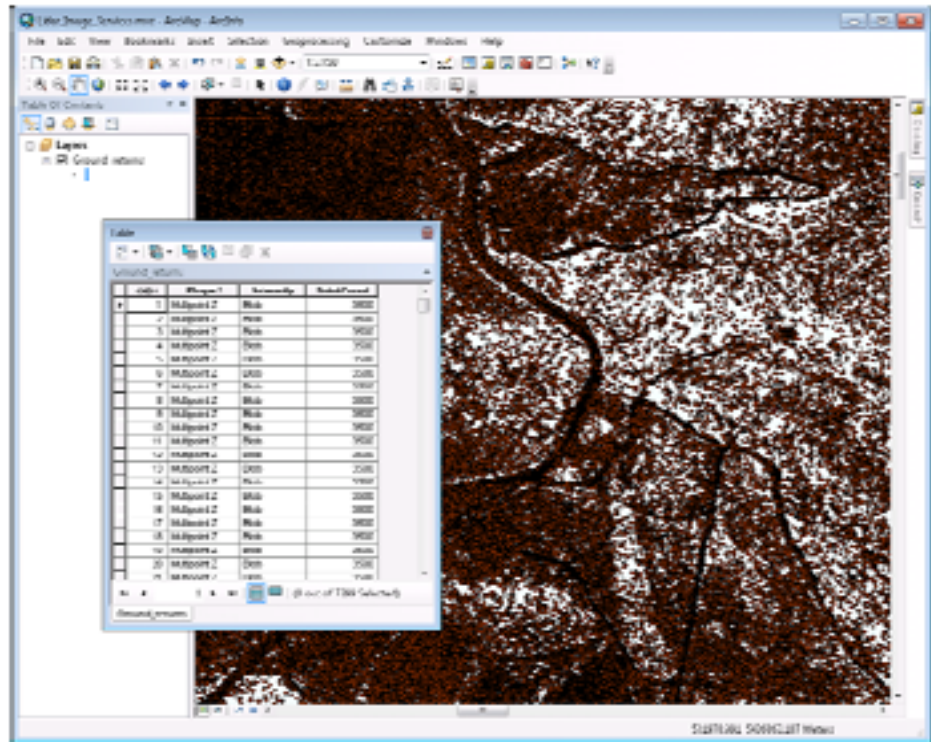
Ground return multipoint feature class shows Shape, Intensity, and PointCount fields.

Visualizing and Storing Lidar Data with ArcGIS

Visualizing Lidar Data

Storing, mosaicking, and separating data in a multipoint feature class in a geodatabase is the first stage of managing lidar data. The next stage is analyzing and visualizing the data.

A forester or land manager may want to visualize the data to enable understanding of

- Ground terrain
- Canopy structure
- Forest type and/or species

A forester may want to analyze the data to enable understanding of

- Tree heights
- Vegetation biomass/density
- Creek and river lines
- Locations of road networks for both existing and new road planning
- Existing terrains for the location of new plantations

When analyzing and visualizing the data, a decision needs to be made whether to convert raw elevation point data to a geodatabase terrain or to an elevation raster file.

Although both these formats are useful for analysis in the forest application, selecting the best format depends on the application.

Advantages of a Raster

If the only source of data is lidar, then an elevation raster can be ideal as there is no blending of additional data sources required. An elevation raster can be quickly produced and created at any resolution. The raster often does not produce the highest-quality results, but lidar data tends to be so dense that for many applications, the reduced accuracy may be sufficient.

When working with lidar raster datasets, there will be situations where no returns are recorded. With ground returns, this can be exaggerated where dense canopy exists and the lidar cannot penetrate to the ground. Where no returns occur, holes or NULL values will appear in the raster. This is the main disadvantage of using the Point To Raster tool. These holes can be reduced with postprocessing techniques that will be discussed later in this paper.

Advantages of a Geodatabase Terrain

A geodatabase terrain is the optimal format to use if elevation data sources include lidar (mass points); breaklines such as roads, water bodies, or rivers; and spot heights. The geodatabase terrain is capable of blending these multiple data sources into one uniform surface with a simple, easy-to-use wizard. A geodatabase terrain resides inside a feature dataset in the geodatabase with the corresponding feature classes that were used to build it.

A geodatabase terrain references the original feature classes. It does not store a surface as a raster or a triangulated irregular network (TIN). Rather, it organizes the data for fast retrieval and derives a TIN surface on the fly based on the feature classes it resides with.

A terrain is not static and can be edited and updated as required. Local edits can be performed, and rebuilding of the terrain is only required for the area of editing. If new data is acquired, it can be easily added to the existing terrain and, again, only rebuilt for the area of new data.

Supported data types for a terrain include

- Points
- Multipoints
- Polylines
- Polygons

In forestry applications, data types can be derived from

- Bare earth lidar data
- First/All return lidar data
- Breaklines representing water body shorelines, rivers, culverts, and roads

Building and Delivering DEMs and DSMs from Lidar

In forest applications, DEMs are useful for planning and operational activities. Terrain beneath the tree canopy provides important information needed by silviculturists, engineers, and equipment operators.

DSMs delineate aboveground vegetation and are therefore useful for understanding the forest structure. They identify stands with similar characteristics, and when used in conjunction with a DEM, use it to calculate tree heights.

Lidar data provides the user with the ability to make two distinct high-resolution surfaces: a first return, or canopy surface, and a ground surface. Typically, the DSM will contain tree canopy and buildings, and the DEM will contain bare earth or ground returns. With the data loaded into a multipoint feature class in a geodatabase, it becomes necessary to consider the workflow for DEM analysis.

Deciding whether to build a geodatabase terrain or a raster grid model will depend on the requirements. Processing the data into a geodatabase terrain will be the most efficient method for maintaining the data, but delivering to clients for consumption will require the conversion of the final geodatabase terrain to a raster DEM format. This presents the user with a problem—trying to process a single terrain into a single file with billions of points will create a file too big to work with and process with most DEM applications. It is therefore necessary to divide these raster files into smaller workable files. Again, the problem presents itself of how these can then be served as a single terrain to clients.

Esri's ArcGIS Server Image extension solves this problem. It can consume the raster DEM files and serve them through ArcGIS Server as an image service. The image service can be consumed by ArcGIS clients as a visualized terrain or an elevation service. The elevation service can then be utilized in ArcGIS extensions such as ArcGIS 3D Analyst or ArcGIS Spatial Analyst for further terrain analysis.

The Workflow to Create a Terrain and Deliver to Clients

Typically, the workflow to get the data from the raw lidar files to a format that can be consumed by client applications is as follows:

- Convert the raw lidar data files to a multipoint feature class in a geodatabase.
- Incorporate the multipoint feature class into a geodatabase terrain.

At this point, the terrain can be visualized and consumed by ArcGIS and geoprocessing tools. If the datasets are very large, they can be served to GIS clients by the ArcGIS Server Image extension. The workflow to move a geodatabase terrain to the image service is to

- Convert the geodatabase terrain to a series of DEM rasters.

■ Consume the multiple raster DEMs in the ArcGIS Server Image extension and serve to clients as an elevation service or Web Coverage Service (WCS).

This workflow will be addressed in the ArcGIS Server Image extension section of this paper.

Building a Geodatabase Terrain

Geodatabase terrains reside in a feature dataset in the geodatabase. All features used by the geodatabase terrain also reside in this feature dataset. This ensures that all the contributing features have the same spatial reference. A terrain dataset is a multiresolution, TIN-based surface built from measurements stored as features in a geodatabase.

As all terrain datasets reside in a feature dataset in the geodatabase, the feature classes used to construct the terrain must also reside in the feature dataset.

Here is how to generate a terrain dataset:

Initially, create a feature dataset in the geodatabase. From the File menu in ArcCatalog™, select New > Feature Dataset.
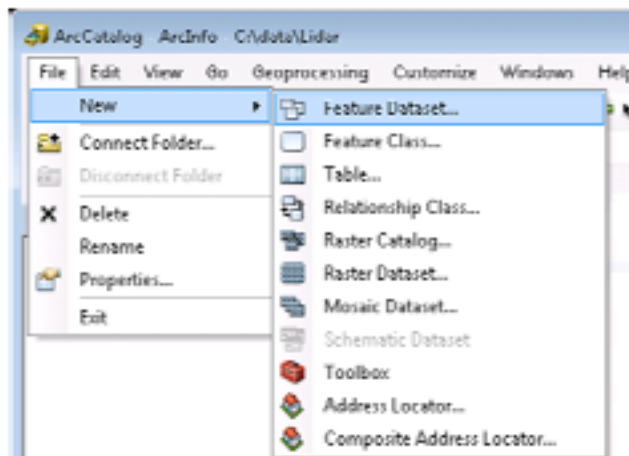

It is important when creating a feature dataset that the coordinate system be the same as the data to be used in the terrain dataset. All features that reside in a feature dataset, including the terrain, have the same coordinate system.

With the feature dataset created, load the raw lidar data into a multipoint feature class in the feature dataset using the LAS to Multipoint tool. If other elevation data sources are available, such as breaklines and spot heights, copy them into the feature dataset.

Note: The feature class copy will fail if there is a mismatch between the coordinate system of the feature dataset and the feature class.

With all the source elevation data in the feature dataset, create the terrain in the feature dataset. From the File menu in ArcCatalog, select New > Terrain. This can also be achieved by right-clicking on the feature dataset and enabling the context menu.
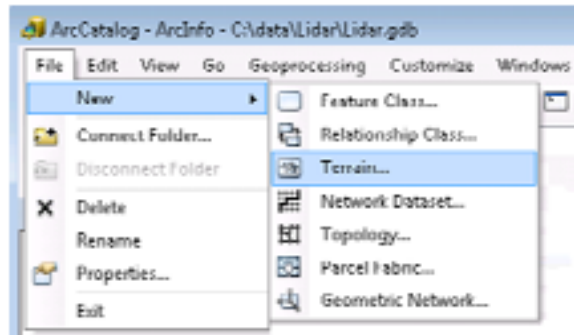

**Note:** The ArcGIS 3D Analyst extension will need to be active to perform this operation.

This initializes the New Terrain wizard, which will lead you through the terrain generation process.

The first form presented contains terrain characteristics. On this initial form

■   Specify the terrain name.

■   The feature dataset is scanned, and the wizard displays all feature classes that can participate in the terrain. From the supplied list, select the multipoint feature class containing lidar points with shape geometry and stream data to be used as breaklines.

■   Specify the average point spacing for the multipoint feature class. If the point spacing is not known, the Point File Information geoprocessing tool can provide point spacing for the supplied data files, as discussed earlier in this paper.
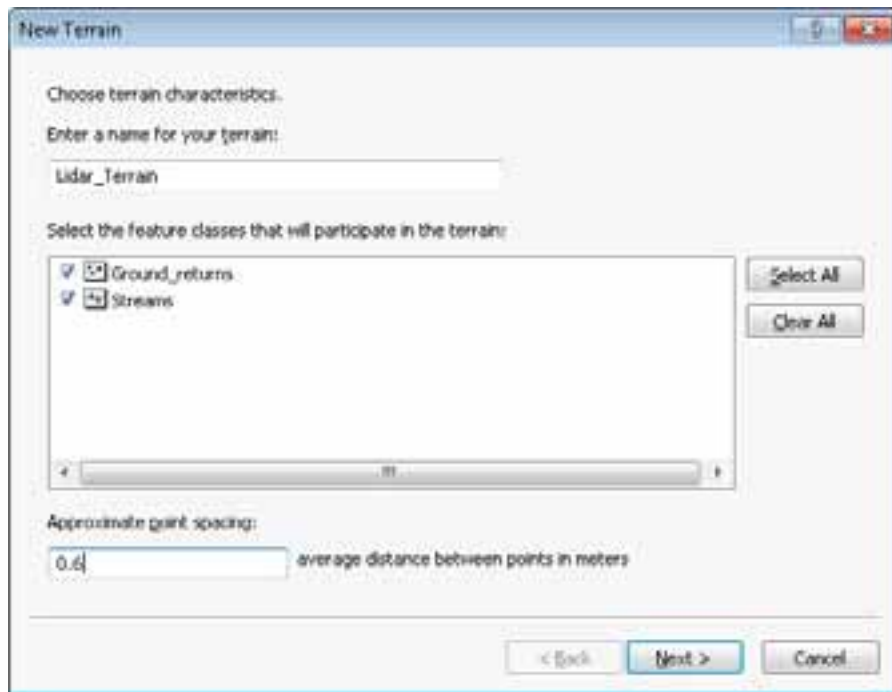
On the next wizard form, select feature class characteristics and decide how the feature classes affect the interpretation of the surface and whether they are viewable at all pyramid levels of the terrain. It is important to ensure that if there are breaklines or areas of interest, the field containing the z-values is used in the height source.
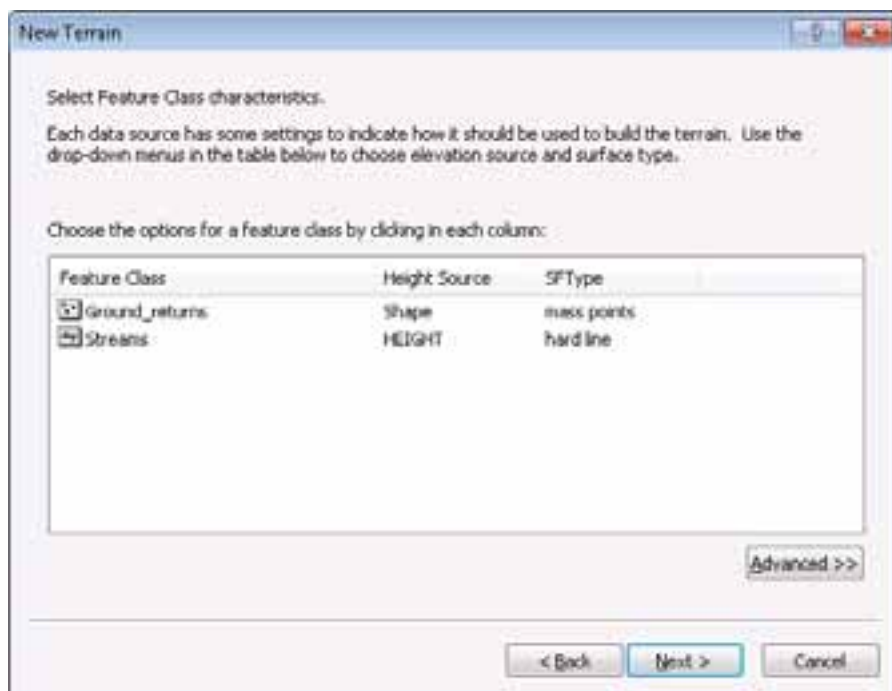
Terrain datasets can be pyramided using one of two point-thinning filters: z-tolerance and window size. Pyramids are reduced resolution versions of the original underlying data, which are displayed when working at small scales.

For DEM production, either method for pyramid production can be used. If rasterizing from the full resolution point set, use the window size filter for terrain construction as it is faster. If thinned data can be used for analysis, which is reasonable if the lidar is oversampled for user needs, the z-tolerance filter can be used. Although more time consuming, this method is most appropriate because it provides an estimate of vertical accuracy of the thinned representation.

For DSM production, use the window size filter with the Z Max option. For DEM production, use the window size filter with the Z Minimum option. It is not necessary to perform secondary thinning unless the terrain is over relatively flat areas. In these regions, performance will be improved by implementing secondary thinning.
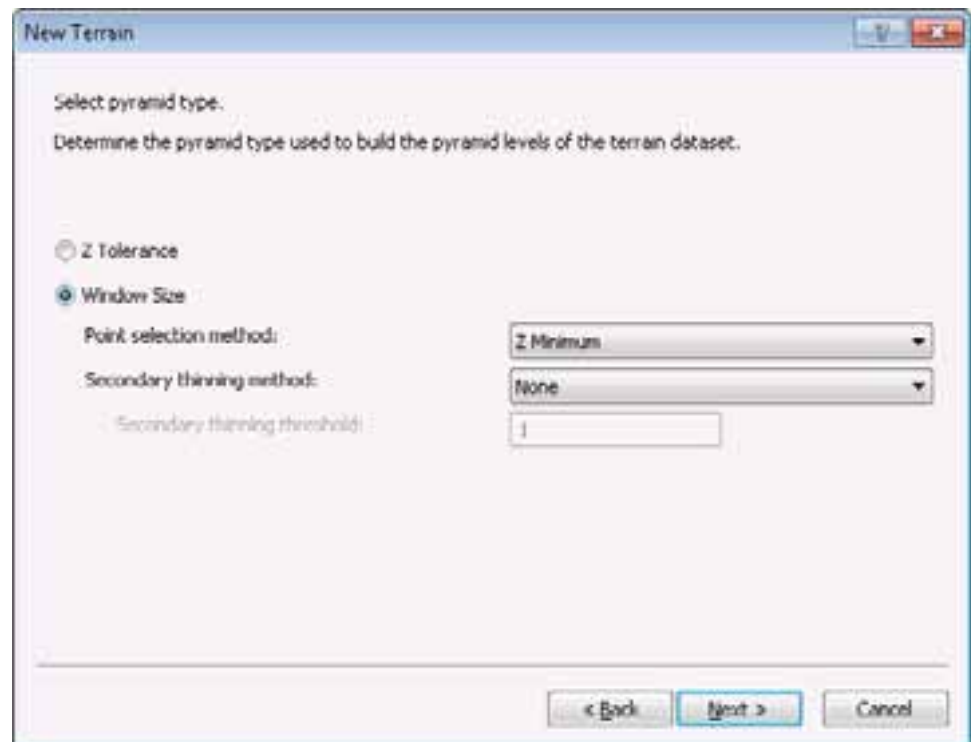
Finally, generate the terrain pyramids.

Terrain pyramids are generated through the process of point reduction, also known as point thinning. This reduces the number of measurements needed to represent a surface for a given area. For each successive pyramid level, fewer measurements are used, and the accuracy requirements necessary to display the surface drops accordingly. The original source measurements are still used in coarser pyramids, but there are fewer of them. No resampling or derivative data is used for pyramids. It takes time to produce pyramids, so you need to consider how best to use them to your advantage.
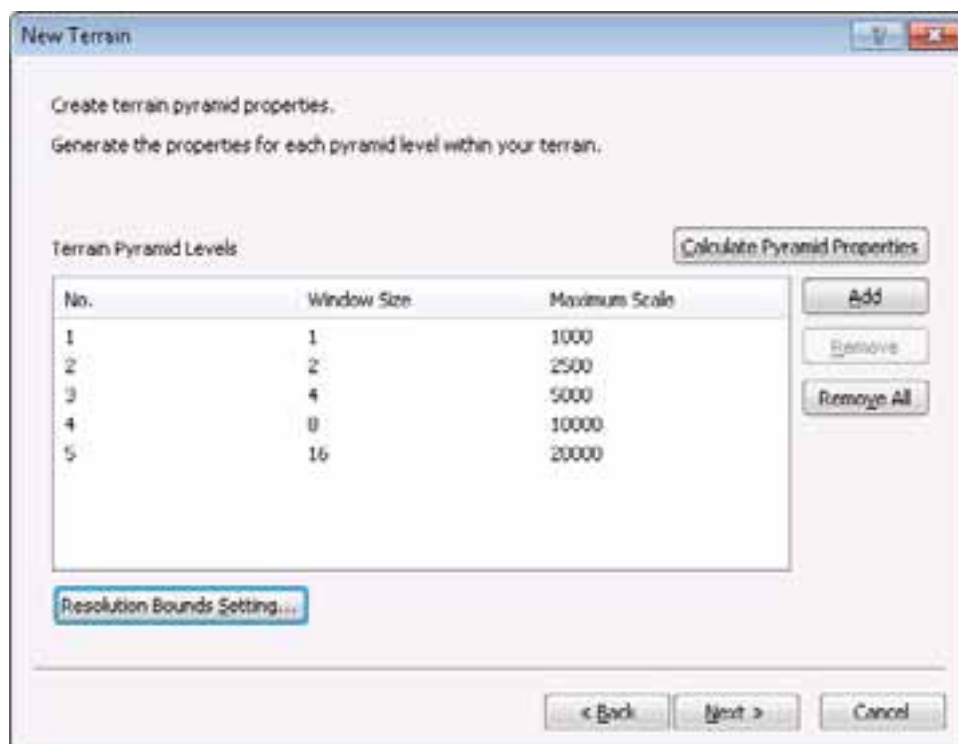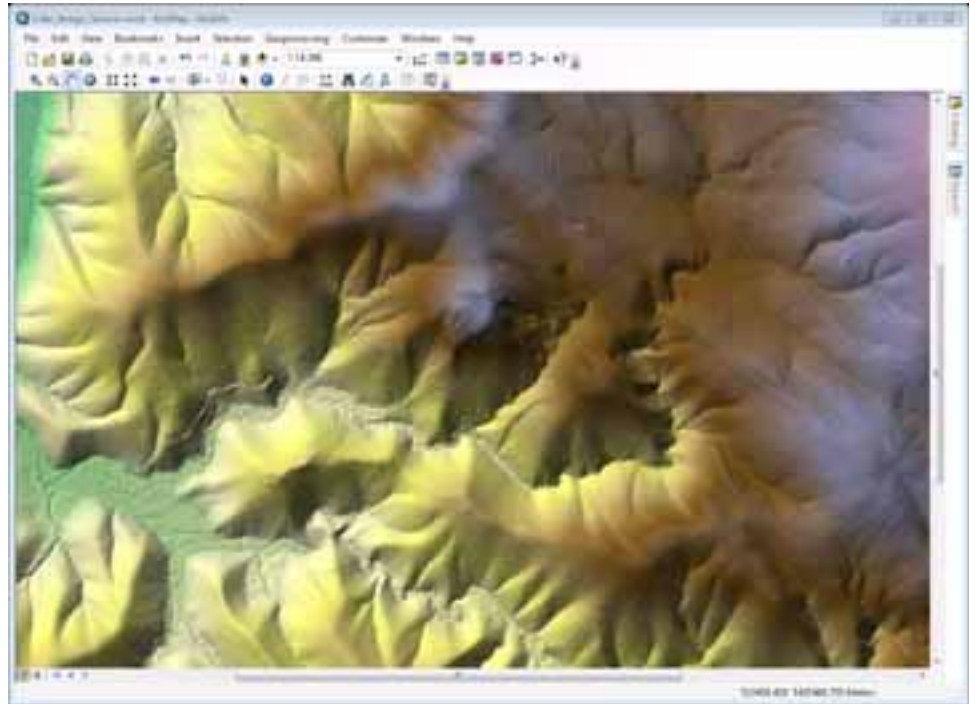
Once the process is finished, the wizard will prompt the user to build the terrain. The number of pyramids chosen to be built will dictate how long it will take to build the terrain.

Below is an example of a terrain view. This example is a DEM derived from the level 2 classification in the raw lidar dataset.

Figure 3
Terrain Derived from Ground Returns


Building a Raster DEM

A DEM or DSM can be built directly from the multipoint feature class using the Point to Raster tool. The Point to Raster tool is ideal if the only data source is lidar data. The tool takes a single feature class as input, is fast, and gives results suitable for most forestry applications. As mentioned earlier, the disadvantage with this tool is that NULL value cells will appear in the raster where no returns are encountered.

The export formats supported by the Point to Raster tool include, but are not limited to, .tif, .img, and Esri GRID.

The Point to Raster tool is initialized from the Conversion Tools toolbox (ArcToolbox\Conversion Tools\To Raster\Point To Raster). On the form, select

■ The input multipoint feature class

■ The value field, which is the shape field from the multipoint feature class (This contains the z-heights for the feature.)

■ The output raster (If the intended raster is an Esri GRID file, do not place an extension on the file name. If the output is TIFF, terminate the file name with .tif.)

■ A cell assignment of MEAN or MAX (Set the MEAN values for an average surface height. Setting the value to MAX is useful when producing a first return, canopy surface.)

■ A cell size (This is important. Too fine, and the surface will have many NoData cells; too coarse, and the surface will lose detail. A good rule here is four times the average point spacing. In this example, the average point spacing is 0.6, so a cell size of 3 meters is optimal [2.4 rounded up to the highest whole meter].)
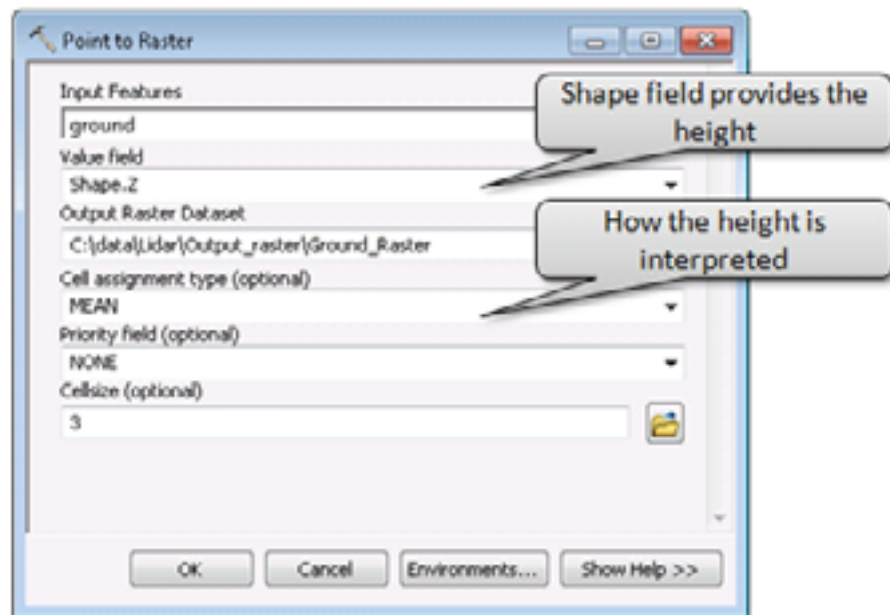
The results from this tool are quick to generate, but as discussed, the frequency of the NoData cells may make the raster DEM appear noisy. This problem can be further magnified where vegetation cover is so dense that it has obscured the ground returns. The diagram below shows the output from ground returns at a three-meter pixel size.

Figure 4
Raster Returned from the Point to Raster Tool
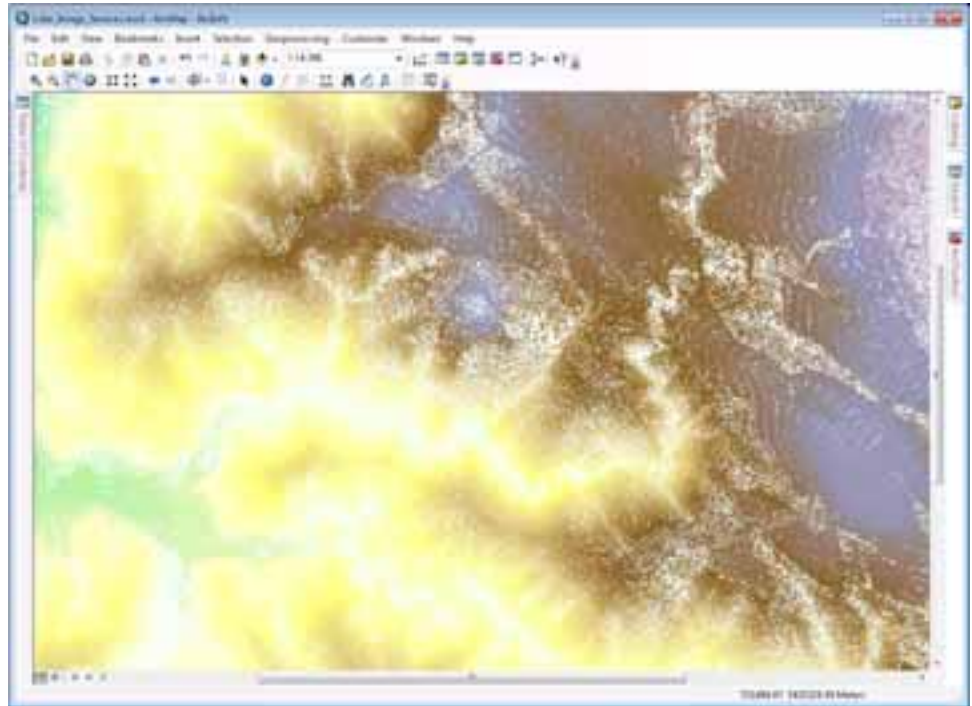
It is possible to reduce this effect by postprocessing the raster DEM with a Python script that incorporates the conditional function (con). When using the Conditional evalution function, each cell in the raster DEM is evaluated for the NoData value. If the evaluation is true, then a floating filter is used to gain the average values of the surrounding cells and applied to the NoData cell. If the evaluation is false, the original raster is used. The conditional (Con) expression looks like the following:

```
Con(<condition>, <true_expression>, <false_expression>)
```

Following is a sample Python statement for removing the NoData values:

```
from arcpy.sa import *
outputfile = Con(IsNull("point2ras"),
FocalStatistics("point2ras", NbrRectangle(3, 3, "CELL"),
"MEAN", "DATA"), "point2ras")
```

In this example, the condition applied is IsNULL.

■ If it returns true, then a FocalStatistics filter is applied. This finds the mean of the values for each cell location on an input raster within a specified neighborhood and sends it to the corresponding cell location on the output raster.

■ If it returns false, then the original cell value is returned.

Here is the result of such a filter.

Figure 5
Raster Returned after Postprocessing Using Conditional Evaluation Function
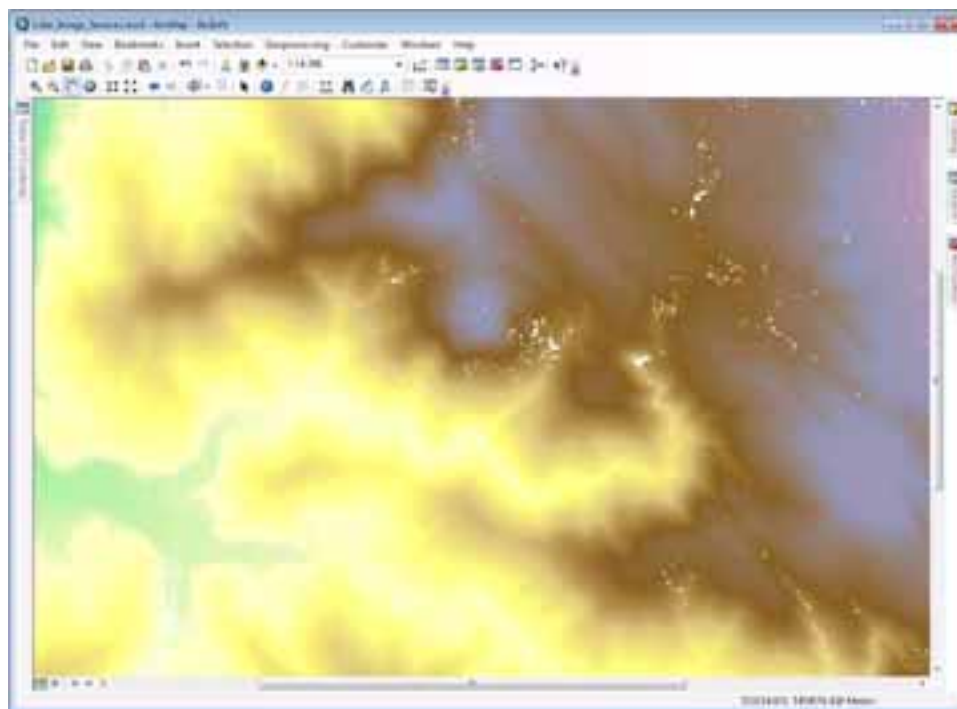
Analyzing Lidar Data for Foresters

Calculating Vegetation Characteristics from Lidar Data

With geoprocessing, lidar data can be made to reveal characteristics of a forest. The forest height is calculated by analyzing the difference between the canopy surface and ground surface. The vegetation density, or biomass, is calculated by analyzing the density and frequency of returns for a given area. The following two sections outline these processes.
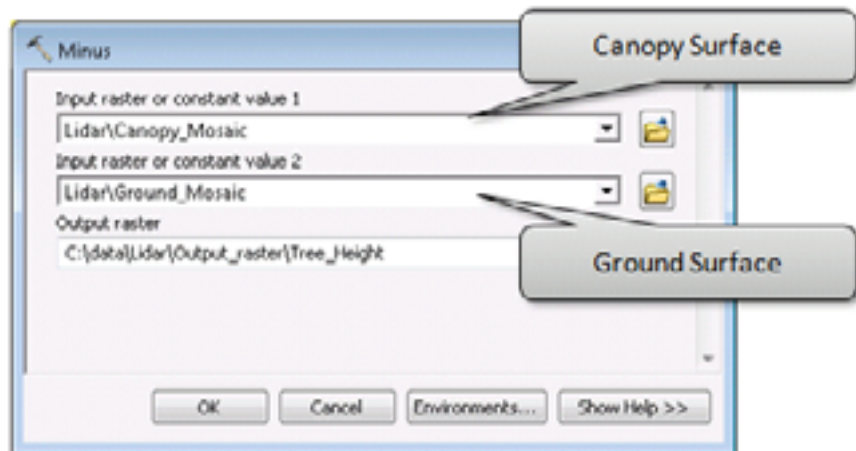
Tree Height Estimation

Tree height estimation is useful for growth analysis and approximating timber volume. Areas of fast and slow growth can be quickly identified and fertilization schemes developed based on these growth statistics.

With both the DEM and DSM generated from the lidar data, it is possible to estimate the canopy height above the ground. To calculate the canopy height, simply subtract one surface from the other by using the Minus tool found in the ArcGIS Spatial Analyst toolbox (ArcToolbox\Spatial Analyst Tools\Math\Minus).
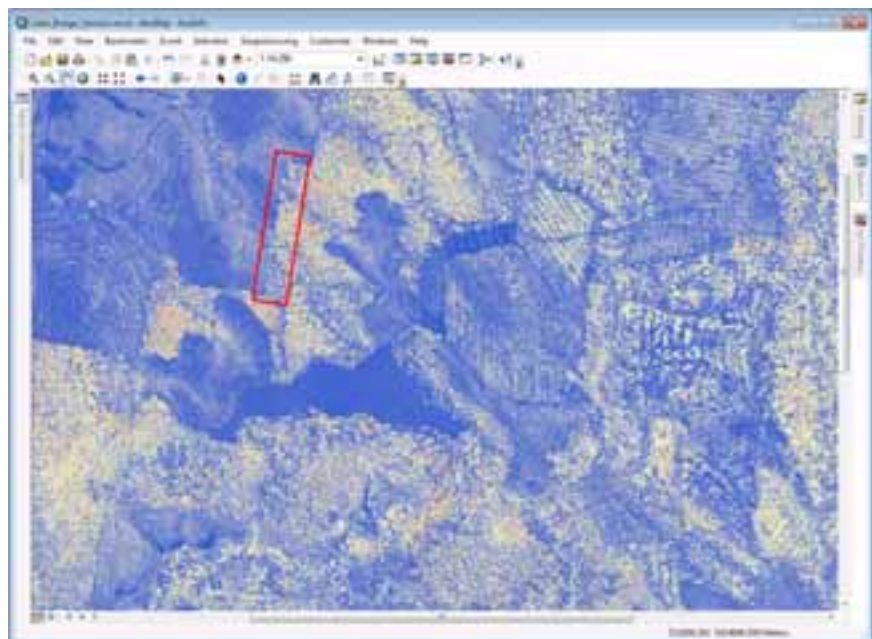
The inputs for this tool are

- Canopy surface (value 1)
- Ground surface (value 2)
- Output height raster

In the results below, blue portrays low vegetation, and red portrays high vegetation. In this example, a break in tree growth can be clearly seen in the area highlighted as a hard line between tall and low growth.

Figure 6
**Height Estimation Raster**

Biomass Density
Calculation

The biomass density will give an indication of tree vigor and growth. Where a forest is of the same species, areas of low vigor can be quickly separated from areas of high vigor and the causes investigated.

To calculate biomass density, it is necessary to have bare earth multipoints in one feature class and all the aboveground points in another feature class. When creating the aboveground feature class from the raw lidar data, it is necessary to include all vegetation returns. According to the LAS 1.2 specification, vegetation classifications are 3, 4, and 5, although some data suppliers may place vegetation in class 1 due to the expense of classifying to the three vegetation classes. Ultimately, the lidar metadata will provide the correct classification.

The key to determining biomass density is to calculate the raster file to be used with the correct cell size. Normally, cell sizes four times the size of the average point spacing should be used. This allows pixel averaging and removal of NULL cells. If smaller pixel sizes are used, the frequency of the NULL cells increases and can bias the results. In the examples used here, the average point spacing is 0.6, so the cell size is 3 meters (2.4 rounded up to the nearest whole meter).

Point to Raster Tool

The first stage in determining biomass density is to convert the multipoint feature classes to raster files. This is done with the Point to Raster tool. The Point to Raster tool is initialized from the Conversion Tools toolbox (ArcToolbox\Conversion Tools\To Raster\Point To Raster). When calculating density, the number of returns in a given area is important rather than the elevation values returned, so the value field on the Point to Raster form is irrelevant. As we are interested in the density of the returns, use the COUNT qualifier for the cell assignment type. This provides an approximate density of the lidar pulse returns.
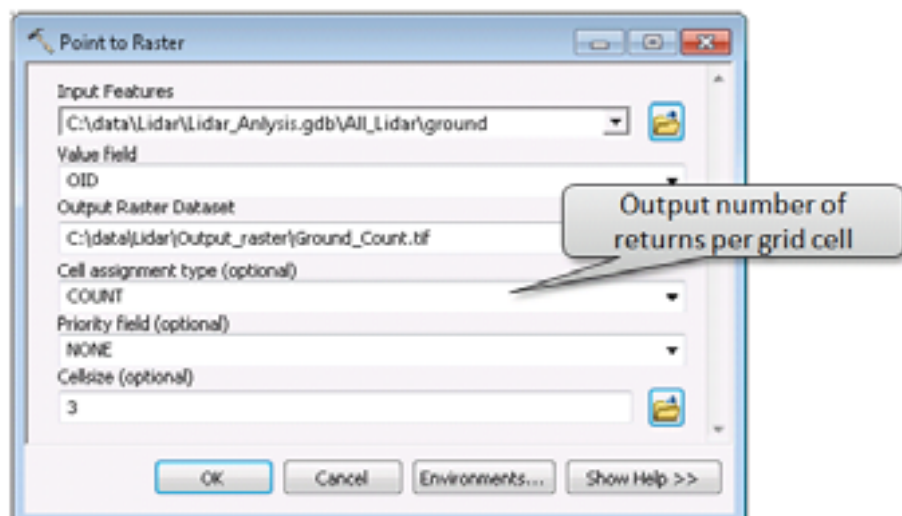
*Replacing NoData Values as Zero Vegetation Density*

In the following two steps, the user takes all cells that have NULL values or cells of NoDATA in them and assigns them a value of 0 to indicate the vegetation density is zero. This is done so that all subsequent operations treat the NULLs as zeros (i.e, no vegetation density) and real data values are returned.

The first of the two steps uses the Is NULL tool in ArcGIS Spatial Analyst (ArcToolbox\Spatial Analyst Tools\Math\Logical\Is Null). The tool designates all NULL values in a raster file as zero. It reads the original grid file and writes out a binary file of 0 and 1. A 1 is assigned to values that are not NULL.
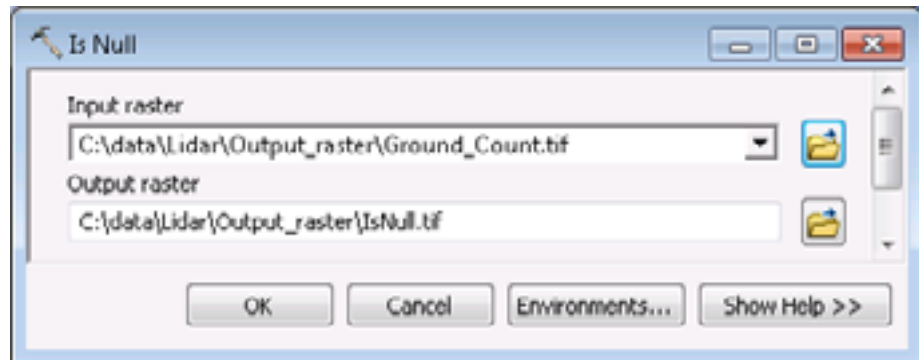
The second process then merges the original raster file with the NULL raster file so that the resultant raster file has a complete range of values from 0 upward. There are no cells that have an unassigned value. The tool used here is the Con tool in ArcGIS Spatial Analyst (ArcToolbox\Spatial Analyst Tools\Conditional\Con).

When the Con tool is run, if a value of 0 is encountered, it is accepted as a true value. If a value of 1 is encountered, the tool pulls the value from the original raster file. This results in a final raster file without NULL values.
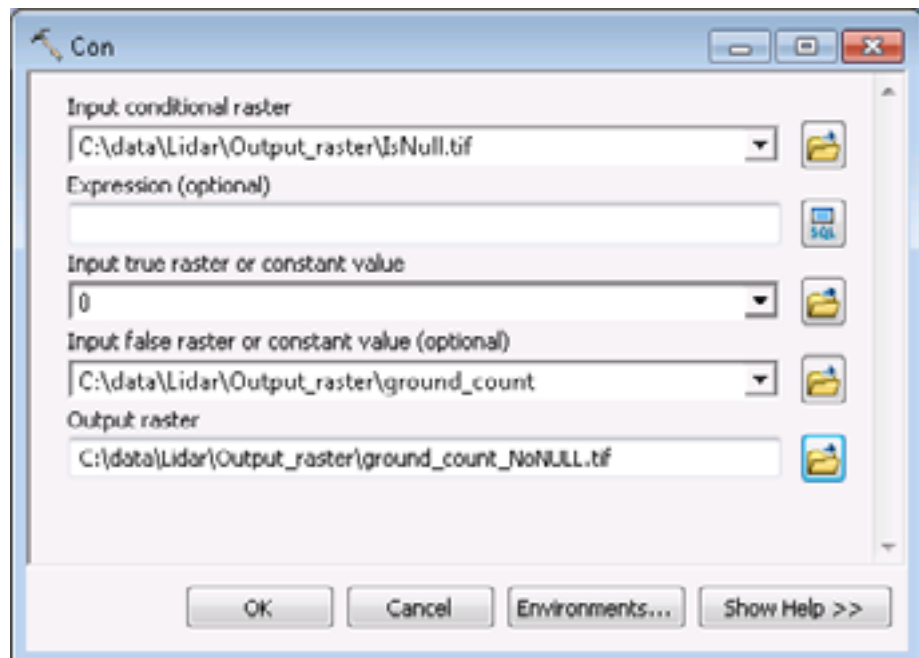
Repeat the above three processes for the aboveground or canopy raster.

Merging the Aboveground and Ground Results

Merge the aboveground density raster with the ground density raster to derive overall density of returns. To do this, use the Plus tool in ArcGIS 3D Analyst (ArcToolbox\Spatial Analyst Tools\Math\Plus).
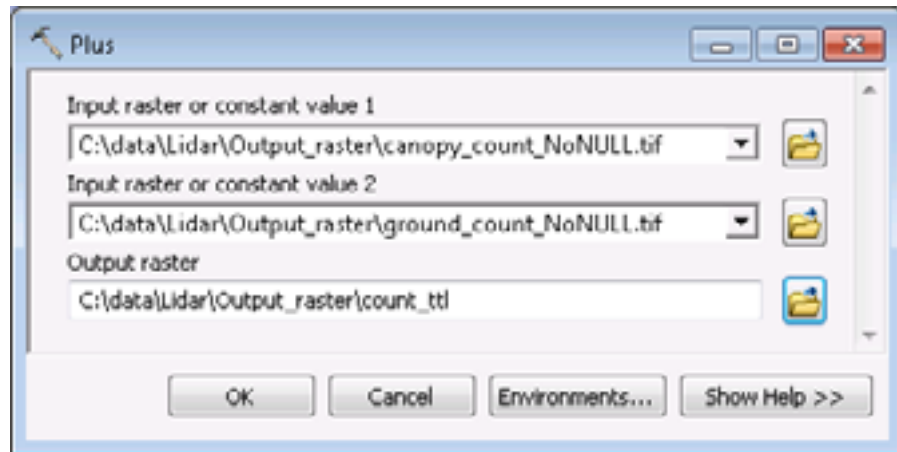
Creating a Floating Point Raster File

All rasters used have been integer value rasters. Each pixel in the raster is a whole number. To calculate the density, use the Divide tool. The results from the Divide tool range from zero to one. Using two integer rasters will result in an integer raster, which will provide whole numbers for each cell and not a true representation of the result from the Divide tool. To change the result of the output raster type, one of the input rasters needs to be of the data type float. To transform a grid from data type integer to data type float, use the Float tool in ArcGIS Spatial Analyst (ArcToolbox\Spatial Analyst Tools\Math\Float).
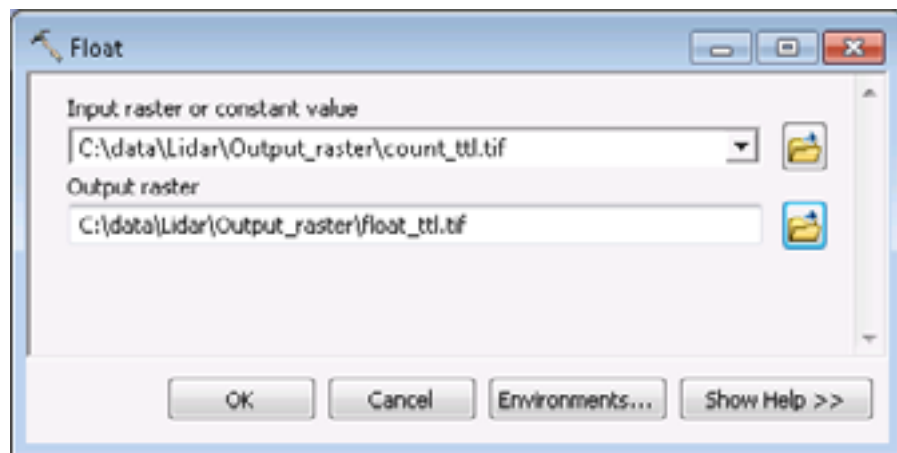

Calculating Density

To calculate the density, use the Divide tool in ArcGIS Spatial Analyst (ArcToolbox\Spatial Analyst Tools\Math\Divide). The result from the Divide tool is a raster with a range between 0.0 and 1.0—hence the need earlier to create a float raster. Dense canopy is represented by a value of 1.0, and no canopy is represented by a value of 0.0. In this case, the canopy returns are divided by the total returns.

Note: Where data has been classified into ground (2) and unclassified (1), results may vary from what is seen in this paper. The unclassified class may contain data other than vegetation returns, which will influence the final result.
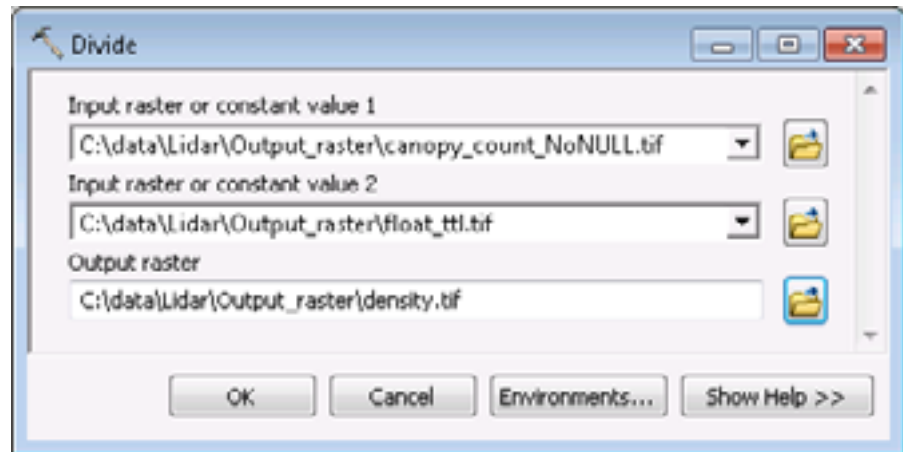
The canopy density raster is depicted in the screen capture below. Yellow represents low density, and dark blue represents high density.

Figure 7
Canopy Density Raster
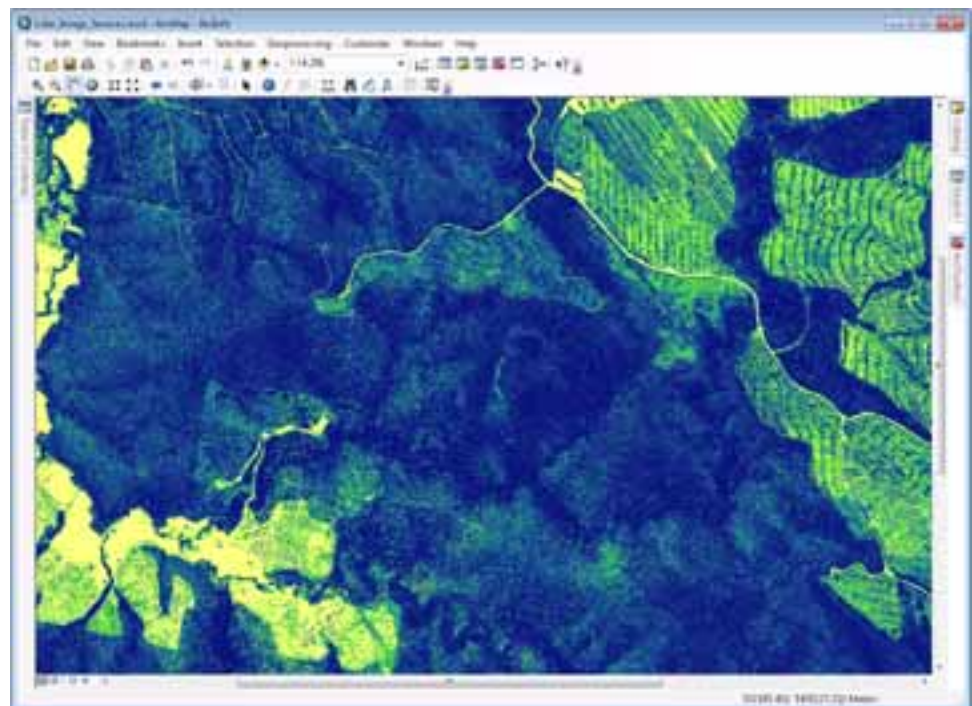
Distributing Large Lidar Datasets

Lidar data is characterized by large volumes of data that can be stored in a geodatabase terrain or a raster file. There are some situations where lidar data is so large that it should not be represented by one raster file. Additionally, in some cases, multiple raster files may also be slow to display, especially in situations where the geodatabase may have to be transferred across a network.

In many situations, a mosaic dataset is the optimum solution for storing anad managing large raster datasets. New at version 10, a mosaic raster dataset is a data model within the geodatabase used to manage a collection of raster datasets (images) stored as a catalog and viewed as a mosaicked image. Mosaic datasets have advanced raster querying capabilities and processing functions and can also be used as a source for serving image services through the ArcGIS Server Image extension. The ArcGIS Server Image extension is used for delivering many large raster files to a client system quickly and efficiently. It delivers images at only the current scale resolution and for the current view extent.

Preparing Raster DEM for Serving with the ArcGIS Server Image Extension

As the ArcGIS Server Image extension only serves raster datasets, it is necessary to output geodatabase terrain as raster files.

Use the Terrain to Raster tool (ArcToolbox\3D Analyst Tools\Conversion\From Terrain\ Terrain To Raster) to produce the rasterized version of the terrain. This tool provides methods for interpolation of the terrain to grid cells, cell size, and the pyramid level to use when producing the terrain. The extents for the output raster can be checked by clicking the Environments button. If the output image is going to be used in further analysis or inclusion in a mosaic dataset, Esri recommends the TIFF format. ArcGIS 10 includes support for the new BigTIFF format, which means TIFF images greater than 4 GB can be created.
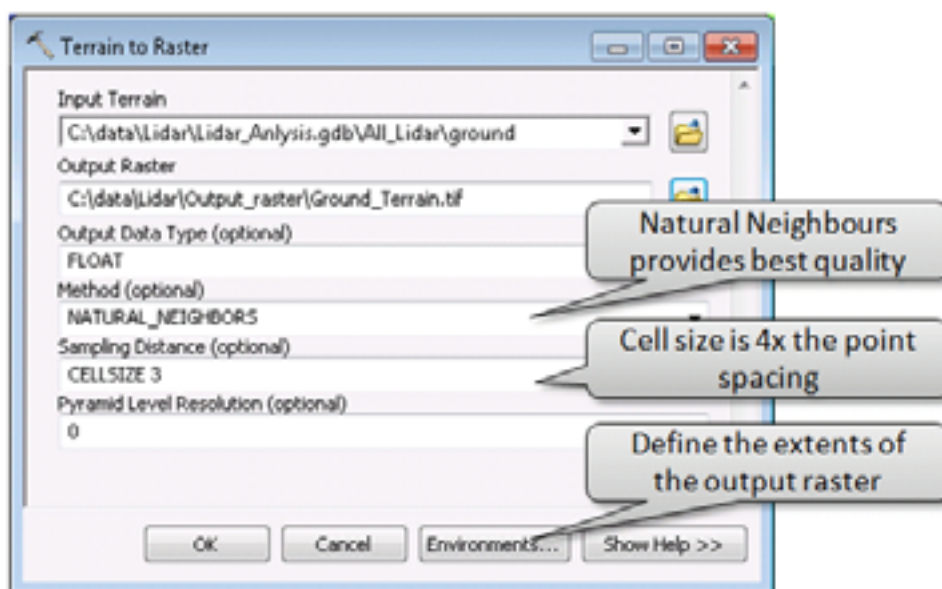
For an interpolation method, natural neighbors provides the smoothest result. Although it is not as fast to produce as linear interpolation, it is more accurate.

Set the cell size to four times the density of the lidar data. This provides a smoothed average. In this case, this is 0.6 meters, so 3 meters is a good size to work with (four times the average spacing rounded up to the nearest whole meter). In forest applications, there is no advantage to using a smaller size than this.

A significant concern here is the size of the terrain being exported. If the terrain is large, the exported raster can exceed the maximum size of some raster formats. To overcome this, a terrain can be divided into small resultant raster files. This is achieved by setting the extents on the Environments tab. The extents can be defined by a group of features in a feature class. These can then be used as input to a model that loops through the feature class to export the final grid layers. The ArcGIS Server Image extension can then consume and serve them as a seamless surface.

Serving an Elevation Service through the ArcGIS Server Image Extension

The ArcGIS Server Image extension is used to serve and analyze mosaic datasets seamlessly, such as aerial photographs and satellite images. It also has the capability to serve terrain data through a specialized image service called an elevation service.

When the elevation service is serving raw elevation information, it can be consumed by ArcGlobe™. It generates a surface to depict elevation over which other imagery can be draped. It also can be used as an elevation input into 3D Analyst and Spatial Analyst tools for complex models and terrain analysis.

A function can be applied to a mosaic dataset, which allows the elevation data to be visualized. The function is applied to the mosaic dataset on the fly, reducing the need to create multiple versions of the same dataset, thereby reducing dataset redundancy. Any of the following functions can be applied to an elevation mosaic dataset and served through the ArcGIS Server Image extension on the fly:

- Hillshade
- Shaded relief
- Slope
- Aspect

Typically, the best methods for visualizing the surface are hillshade and shaded relief. Examples of these follow.

The advantage of using an elevation service is that an entire lidar dataset can be viewed as a seamless image rather than a series of separate lidar files. This enables quick viewing and interpretation of the entire DEM or DSM and performance of comparisons using such tools as the Swipe tool or Flicker tool found on the Effects toolbar in ArcMap.

Figure 8
Example of a Hillshade
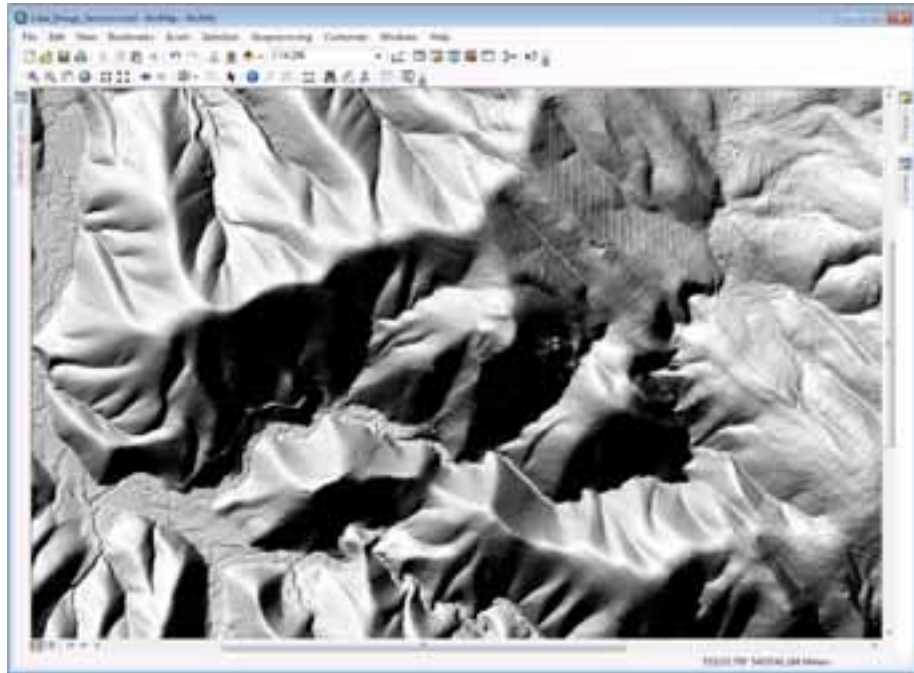
Figure 9
Example of a Shaded Relief
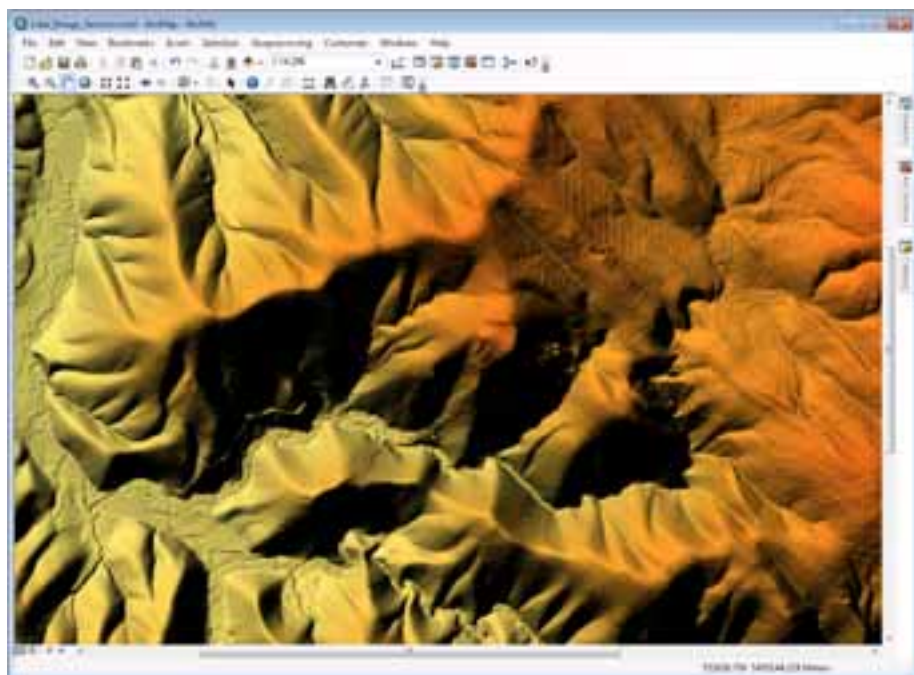
Creating an Elevation Mosaic Dataset

Creating an elevation mosaic dataset in a geodatabase is the same process as developing a standard image mosaic dataset, with the exception that it uses elevation sources rather than standard image data. Although data accessed is raster data, it is interpreted as an elevation source.

The process outlined below is used to generate an elevation mosaic dataset. It uses the ArcGIS 10 mosaic dataset geoprocessing tools. The Mosaic Dataset tools are located in ArcToolbox\Data Management Tools\Raster\Mosaic Dataset.

To create a new mosaic dataset

■   Select the Create Mosaic Dataset geoprocessing tool. The Create Mosaic Dataset tool has three primary entry fields:

   ●   Output Location—This is the geodatabase that will contain the mosaic dataset.

   ●   Mosaic Dataset Name

   ●   Coordinate System—This is the coordinate system of the mosaic dataset, which may differ from the underlying raster sources.

   There is no need to specify the type of service, as this is interpreted from the pixel source when the raster data is added to the mosaic dataset.
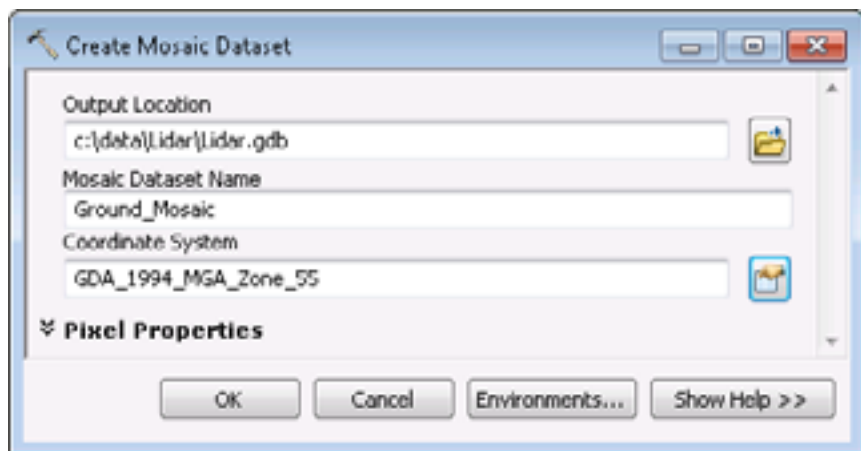
■   Select OK; the mosaic dataset will be created.

The mosaic dataset can be considered a container to which the raster data now needs to be added.

To add data to the mosaic dataset

■   Select the geoprocessing tool Add Rasters To Mosaic Dataset. Input to this form can be divided into two distinct groups—information about storage of the rasters and advanced information about the rasters.

The first group of inputs includes

- Selecting the mosaic dataset created earlier

- The Raster Type drop-down list—This drop-down list can specify different image sources/sensors such as Landsat or QuickBird. In this sample set, the elevation source is a 32-bit .tiff file, so Raster Dataset is specified.

- The Input for this raster source is Workspace. All the exported terrain data files from earlier processing are stored in a folder, which here is referred to as the Workspace.

- Update Overviews—This generates the pyramids over the entire mosaic dataset. The mosaic dataset considers the entire raster set as a seamless mosaic and builds lower-resolution pyramid files over the single seamless mosaic.

J-9999
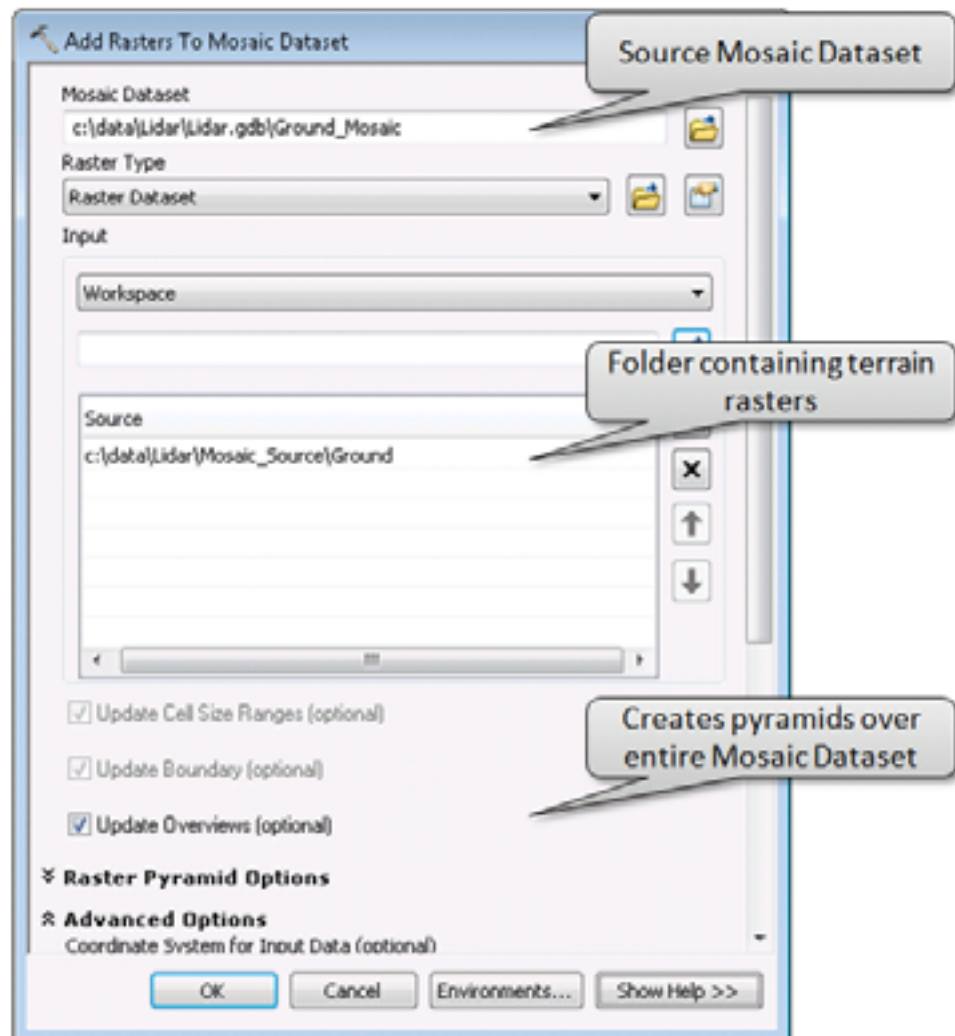
The second group of inputs is selected from Advanced Options. This group of inputs describes the characteristics of the rasters and how to work with them. Generally, the following settings are made:

- The coordinate system of the source rasters

- The option to recursively search subfolders if the data resides in more than one folder

- The option to include or exclude duplicate rasters or overwrite existing rasters

  This option can be useful if some raster data is found to be corrupt and needs to be updated.

- Build Raster Pyramids option for each of the source rasters if they do not already exist

  Selecting this option increases the performance of the display of the raster dataset when served through the ArcGIS Server Image extension.

- Calculate Statistics option to provide feedback to the system of each raster's pixel range
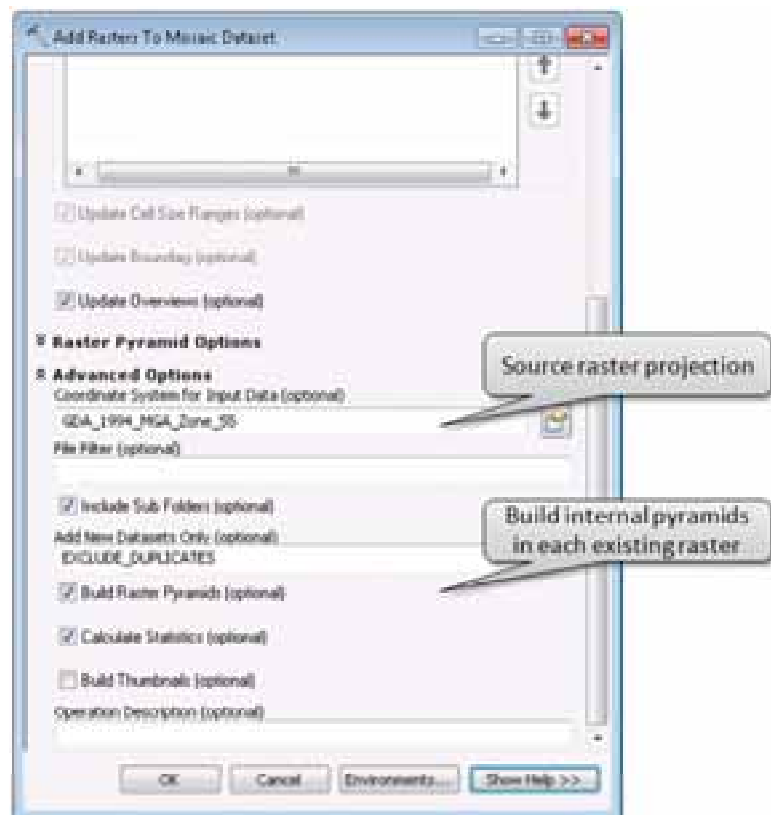
■ Select OK to populate the mosaic dataset with select raster datasets and generate the overviews of the mosaic dataset.

The resultant mosaic dataset can now be loaded into ArcMap. A mosaic dataset is composed of three principle components:

● A Boundary layer, which defines the extent of the mosaic dataset

● A raster Footprint layer, which is the extent for each of the rasters in the mosaic dataset

● The Image layer, which represents the entire image layer being displayed at the required resolution for the current display scale

Figure 10
Elevation Mosaic Dataset
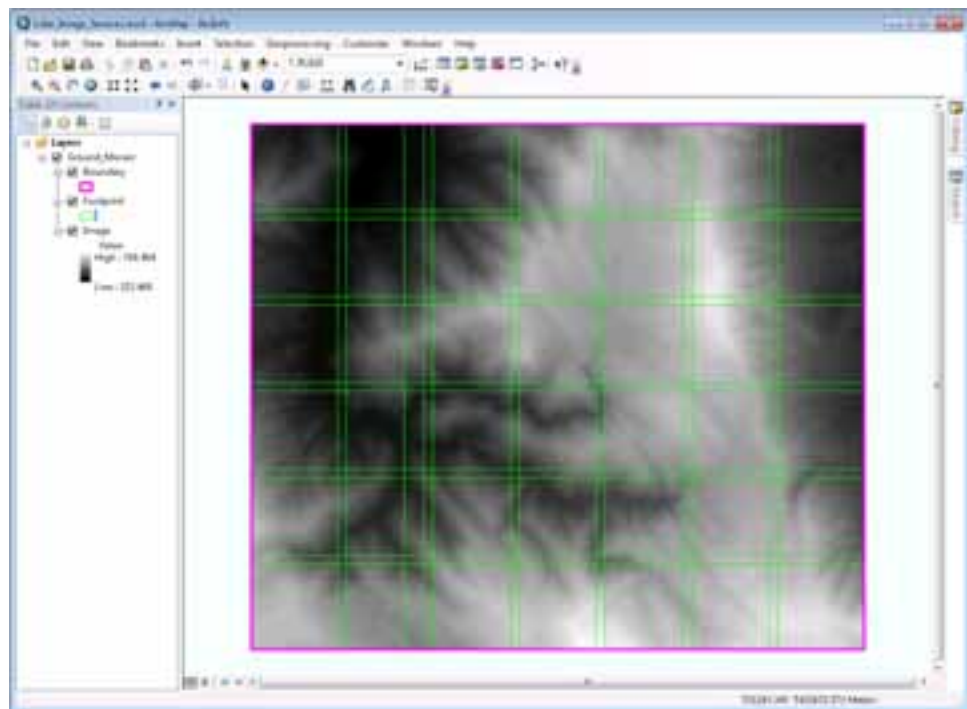
In figure 10, the Footprint layer is composed of the original raster data sources and the generated overviews. The overviews are smaller; lower-resolution images based on the original raster data provide the mosaic dataset with its performance.

This mosaic dataset can now be published as an image service using the ArcGIS Server Image extension. To publish a mosaic dataset, right-click on the mosaic dataset from within ArcCatalog and select Publish to ArcGIS Server.
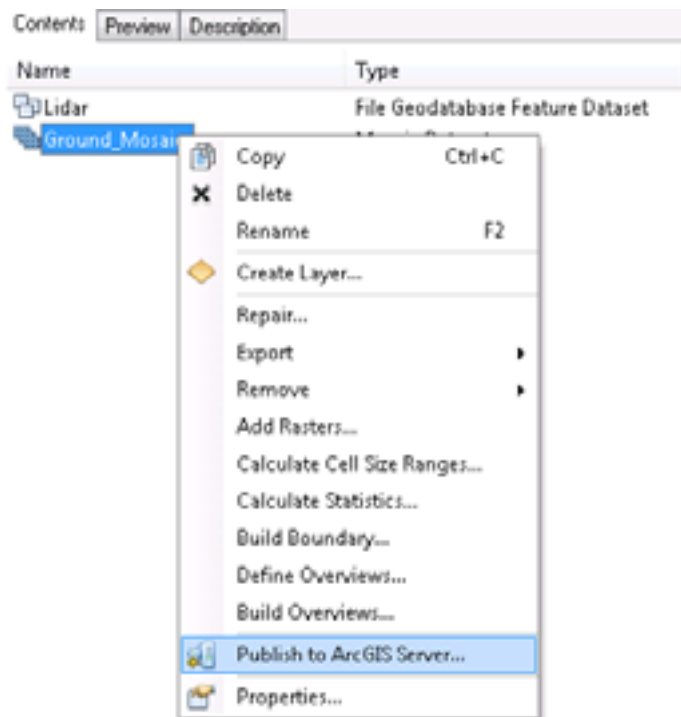
Publishing the mosaic dataset is the final stage in the elevation service creation. The elevation service can be used as an elevation source for terrain visualization in ArcGlobe or as input to ArcGIS 3D Analyst or ArcGIS Spatial Analyst geoprocessing tools such as contour or viewshed.

Visualizing an Elevation Service

The published elevation service is useful to gain an understanding of the terrain but often does not provide the visual context if the service is to be used as a background. The next section shows how to apply visualization to the mosaic data source through the use of a mosaic dataset function. Mosaic dataset functions are applied to the raster data on the fly as the mosaicked image is accessed and viewed.

While a function can be applied to a mosaic dataset directly, it is often not desirable, as it affects the final output and changes the values returned from the primary mosaic dataset. A referenced mosaic dataset can be created that references all the original parameters from the primary mosaic. A unique function can then be applied to the referenced mosaic dataset, which provides a different visualization of the lidar data from the source mosaic dataset, and the visualization is created on the fly for the area of interest. Multiple referenced mosaic datasets can be created for the original source mosaic dataset.

In this example, a referenced mosaic dataset is created, and the function used is the Shaded Relief function.

From within ArcCatalog, perform the following steps to apply a Shaded Relief function to a referenced mosaic dataset.

*Creating a
Referenced Mosaic
Dataset*

A referenced mosaic dataset is created by launching the Create Referenced Mosaic Dataset geoprocessing tool from the Mosaic Dataset toolbox found at ArcToolbox\Data Management Tools\Raster\Mosaic Dataset.

When the dialog box appears, enter

■   The Input Mosaic Dataset as the Ground Mosaic created earlier
■   The Output Mosaic Dataset as the name for the New Mosaic Dataset

The Coordinate System should be automatically populated from the Input Mosaic. Check the Build Boundary option, then select OK to create the referenced mosaic dataset.
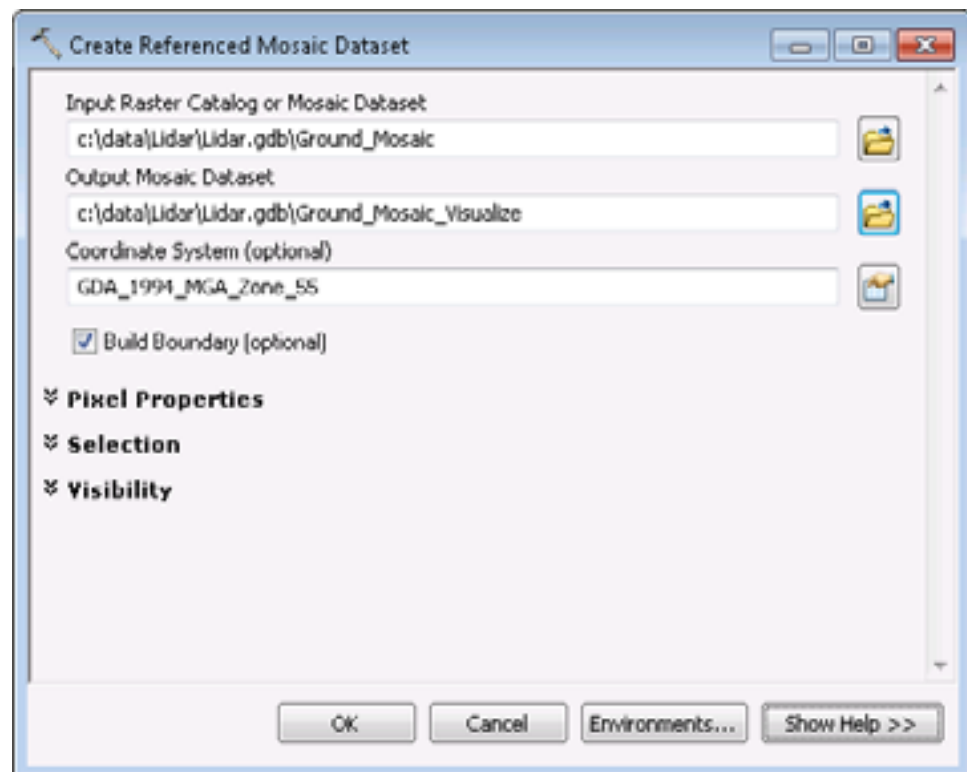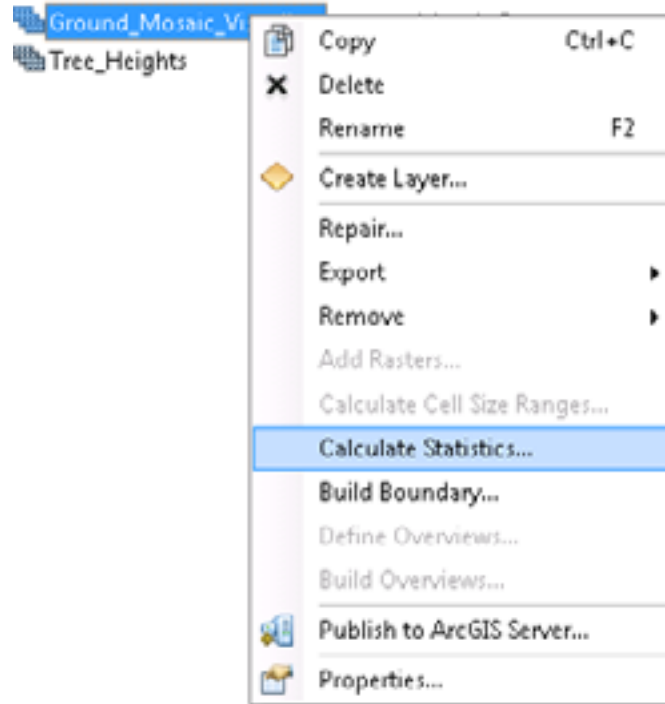
Before the function can be applied to the referenced mosaic dataset, the statistics of the dataset need to be calculated.
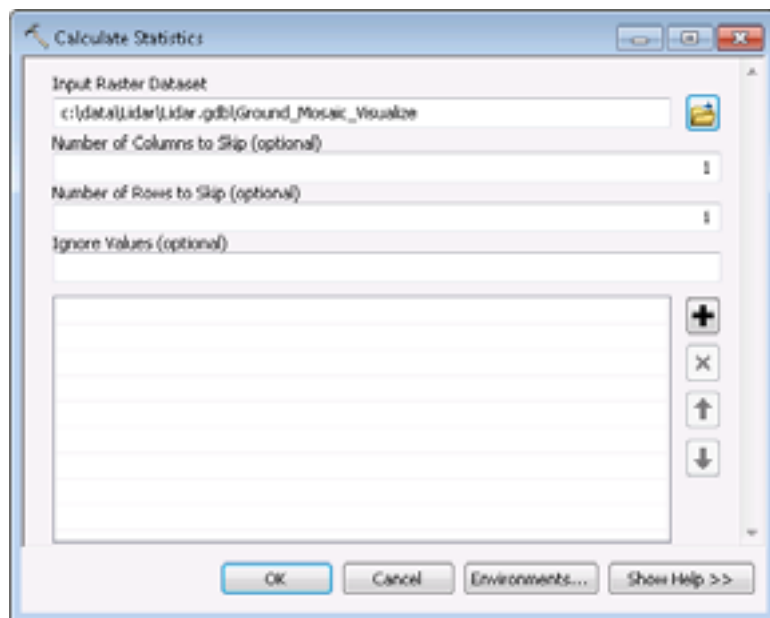
To calculate the statistics

■   Select the referenced mosaic dataset and right-click to enable the context menu.

■ Select the Calculate Statistics option.
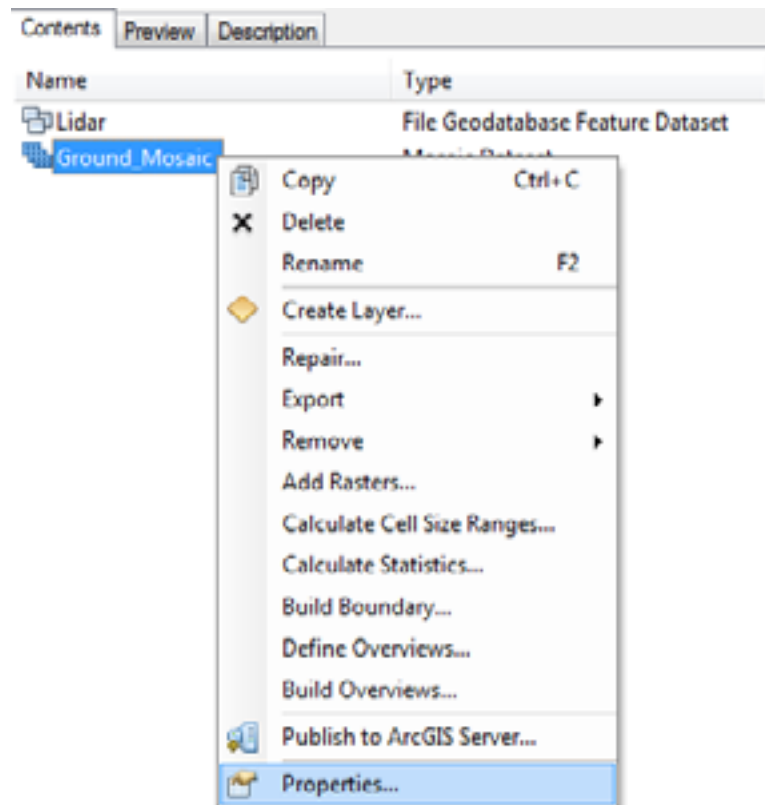

The Calculate Statistics dialog box appears.

■ Select OK to have the statistics calculated.

***Applying a Mosaic Function***

With the referenced mosaic dataset created, a function now needs to be added to the mosaic dataset.

To add the Shaded Relief function

■ Right-click the referenced mosaic dataset and select Properties.

- Select the Functions tab on the Mosaic Dataset Properties dialog box. The Functions tab is where the functions are configured for the mosaic dataset. One or many functions can be applied to the mosaic dataset to provide on-the-fly interpretations of the original mosaic dataset.
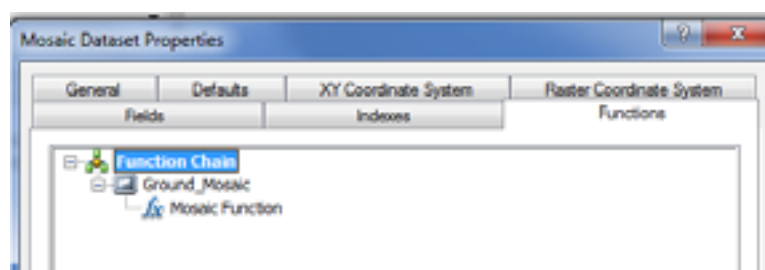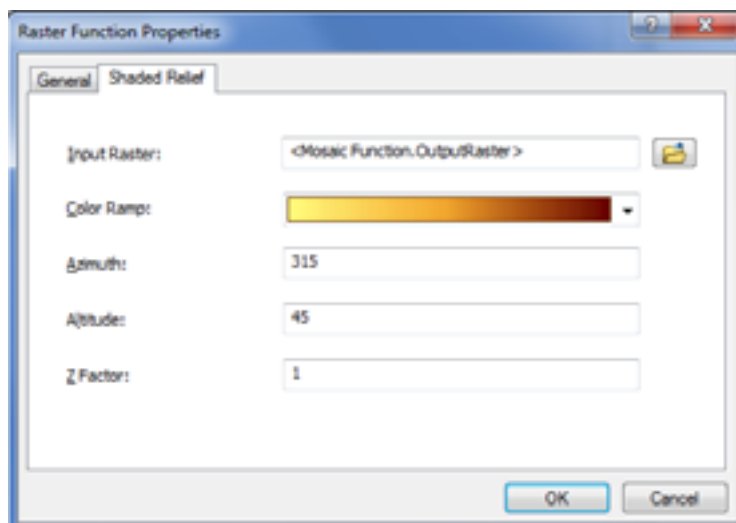

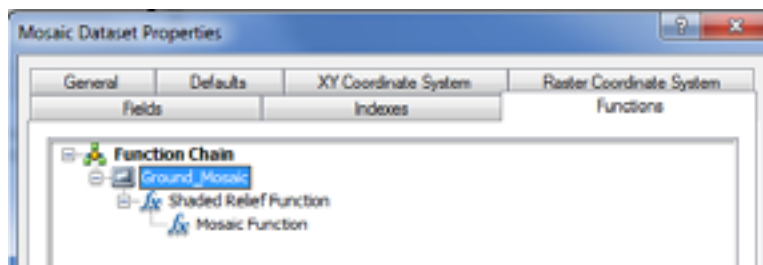- Right-click Mosaic Function; the functions menu appears.

J-9999

- Select Insert > Shaded Relief Function.
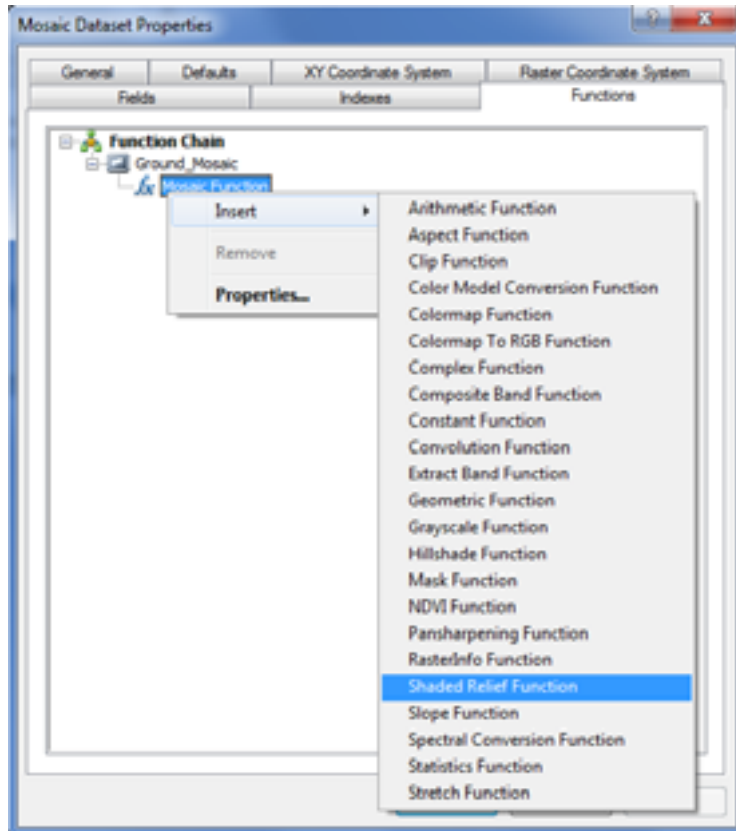
The parameters for the Shaded Relief form appear. In this example, the yellow to dark red color ramp is used, and the default settings are used for Azimuth, Altitude, and Z Factor.
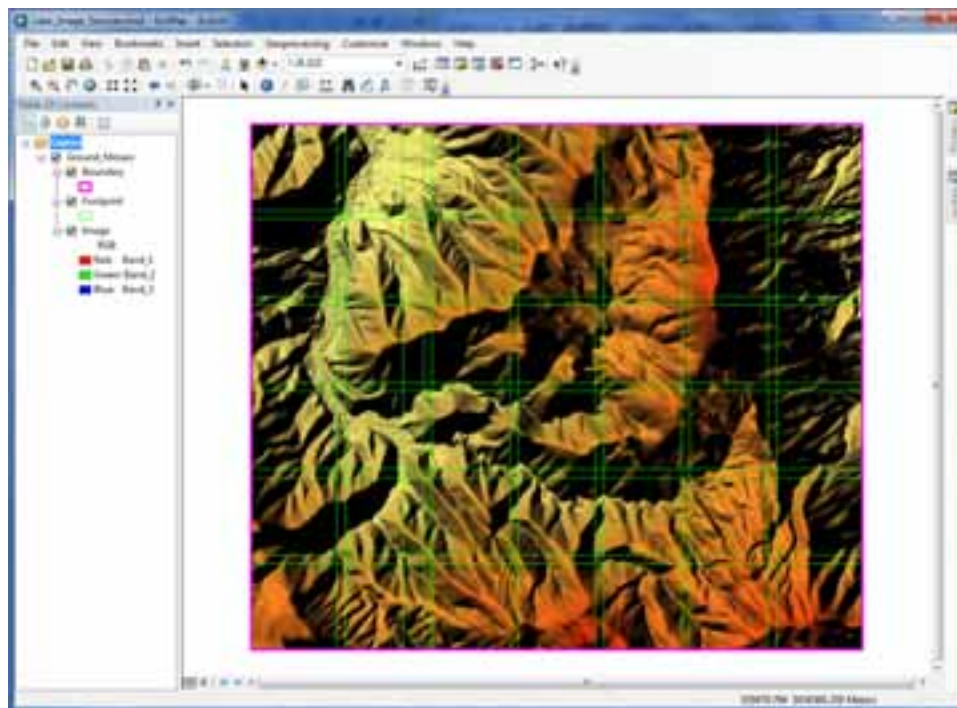
- Select OK; the Shaded Relief function is added to the Function Chain.

- Now, select OK to apply the function to the mosaic dataset.

The resultant visualization in figure 11 is of the referenced mosaic dataset of the ground returns.

Figure 11
Visualized Shaded Relief Mosaic Dataset

This mosaic dataset can now be published as an image service using the ArcGIS Server Image extension. As before, to publish a mosaic dataset, right-click on the referenced mosaic dataset from within ArcCatalog and select Publish to ArcGIS Server.

Publishing the referenced mosaic dataset is the final stage in the visualized elevation service creation. It can now be consumed in ArcMap and ArcGlobe and published on the Web through any one of Esri's Web APIs.

Estimating Tree Height Using Mosaic Dataset Functions

As was detailed earlier in this paper, it is possible to estimate the height of vegetation from lidar by performing algebraic operations between the DSM and DEM. When using the mosaic datasets, functions can be applied that return the same results. The advantage of using the mosaic datasets is that there is no need to create a height estimation raster, as the Mosaic Dataset function creates the height estimation on the fly for the area of interest on the screen, removing the need to create extra raster datasets.

The creation of a height estimation mosaic dataset requires two elevation mosaic datasets as the source inputs: a DEM mosaic dataset and a DSM mosaic dataset. The authoring of the DEM mosaic dataset was described in the section Creating an Elevation Mosaic Dataset. To create a DSM service, follow this same method but use the DSM terrain as the source data.

The height estimation mosaic dataset is a referenced mosaic dataset that takes the input of the two mosaic datasets and performs an arithmetic function between them to produce a third service.

Creating the Height Estimation Mosaic Dataset

When creating the height estimation mosaic dataset, it is necessary to create a referenced mosaic dataset and add the two source mosaic datasets files as two distinct steps.
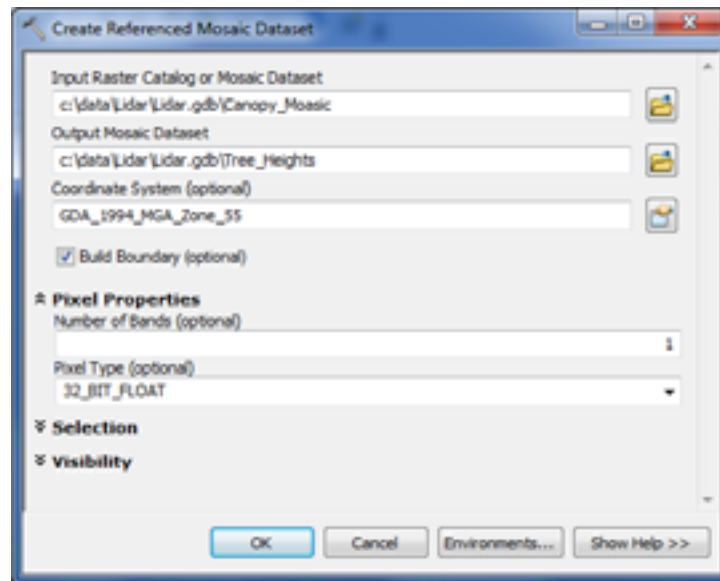
The referenced mosaic dataset is created by launching the Create Referenced Mosaic Dataset geoprocessing tool from the Mosaic Dataset toolbox found at ArcToolbox\Data Management Tools\Raster\Mosaic Dataset.

■ Select Input Mosaic Dataset as the mosaic dataset that displays the canopy data. It is important that the canopy be selected for the input mosaic dataset, as the function applied later will be against the input mosaic dataset.

■ Select Output Mosaic Dataset as the name for the new mosaic dataset.

■ Specify the Coordinate System for the mosaic dataset.

■ Check the Build Boundary option.

Under Pixel Properties,

■ In Number of Bands, enter 1.
■ In Pixel Type, choose 32_BIT_FLOAT.
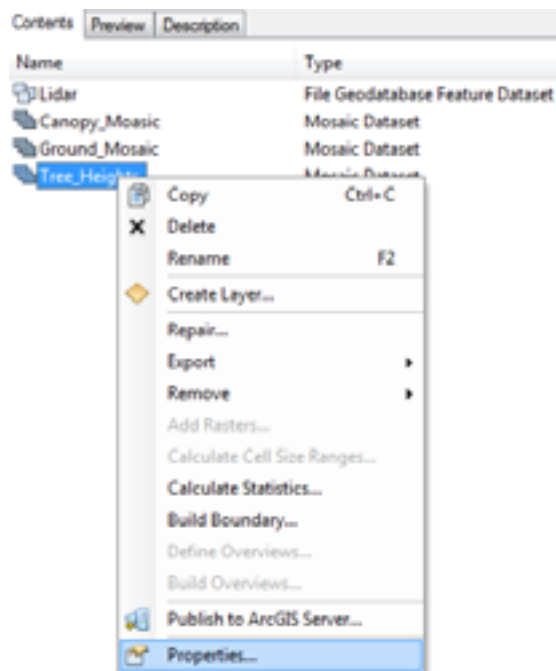
This is the same as the source mosaic datasets.

■   Select OK to create the referenced mosaic dataset.

The referenced mosaic dataset is created, and the arithmetic function now needs to be added.

***Applying the
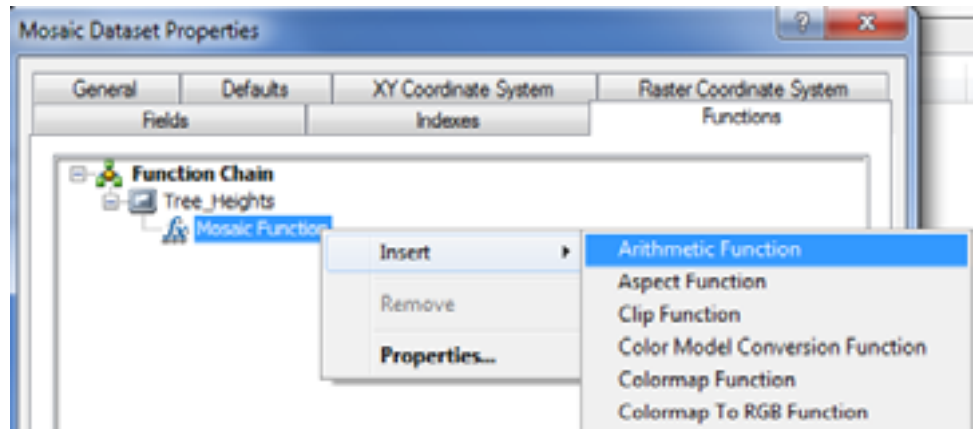Arithmetic Function***

To add the arithmetic function

■   Right-click on the referenced mosaic dataset and Select properties.

■ Select the Functions tab on the Mosaic Dataset Properties dialog box.

■ Right-click Mosaic Function and select Insert > Arithmetic Function.

■ On the Raster Function Properties dialog box, make the following settings:

● Input Raster 1 is left untouched, as this is the canopy mosaic dataset specified earlier.

● In Input Raster 2, click the Browse button, then choose Ground Mosaic Dataset.

● Operation is set to Minus. This is how Raster 2 interacts with Raster 1.
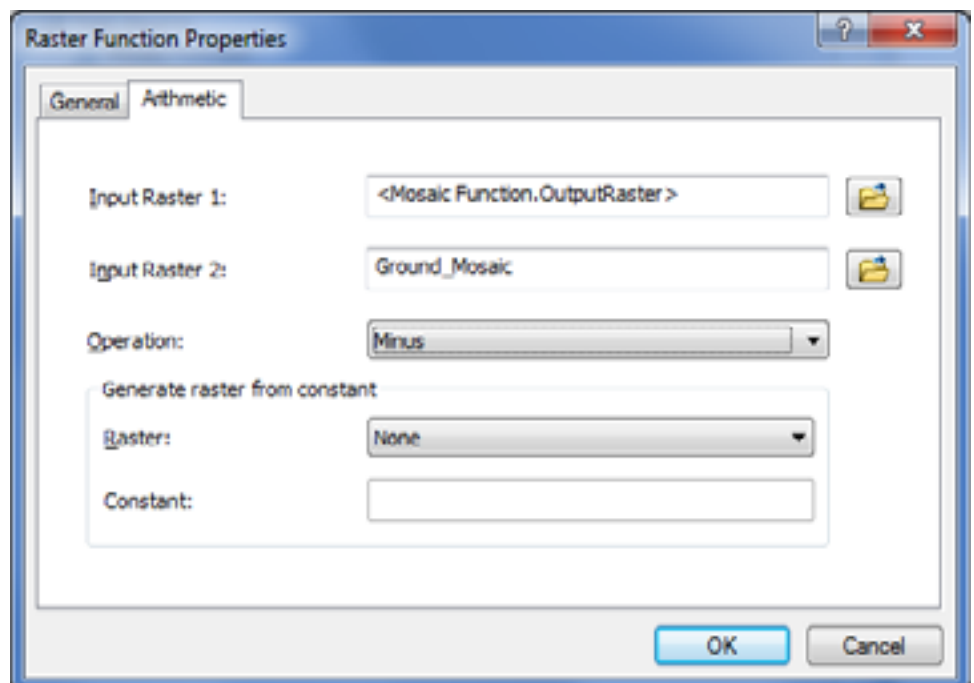
■ Select OK; the function will be applied to the mosaic dataset. As can be seen in the screen shot below, the ground mosaic now forms part of the Arithmetic Function.
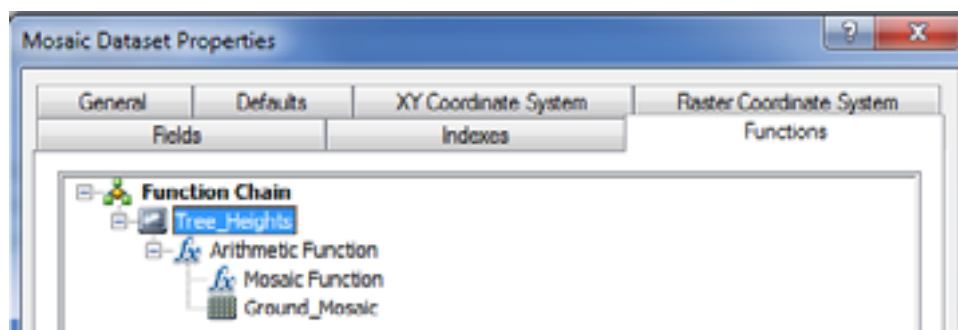
■ Select OK on the Mosaic Dataset Properties dialog box; the function will be applied to the referenced mosaic dataset.

For ArcMap to render the mosaic correctly, the statistics for the referenced mosaic dataset need to be calculated.

To calculate the statistics

■ Select the referenced mosaic dataset in ArcCatalog and right-click to enable the context menu.

■ Select the Calculate Statistics option. The Calculate Statistics dialog box appears.

■ Select OK to calculate the statistics.

The referenced mosaic dataset can now be added to ArcMap.

Note: By default, the height estimation mosaic dataset is in grayscale. To add context to the raster, a color ramp is applied. In figure 12, the color ramp is yellow to green to dark blue.

Figure 12
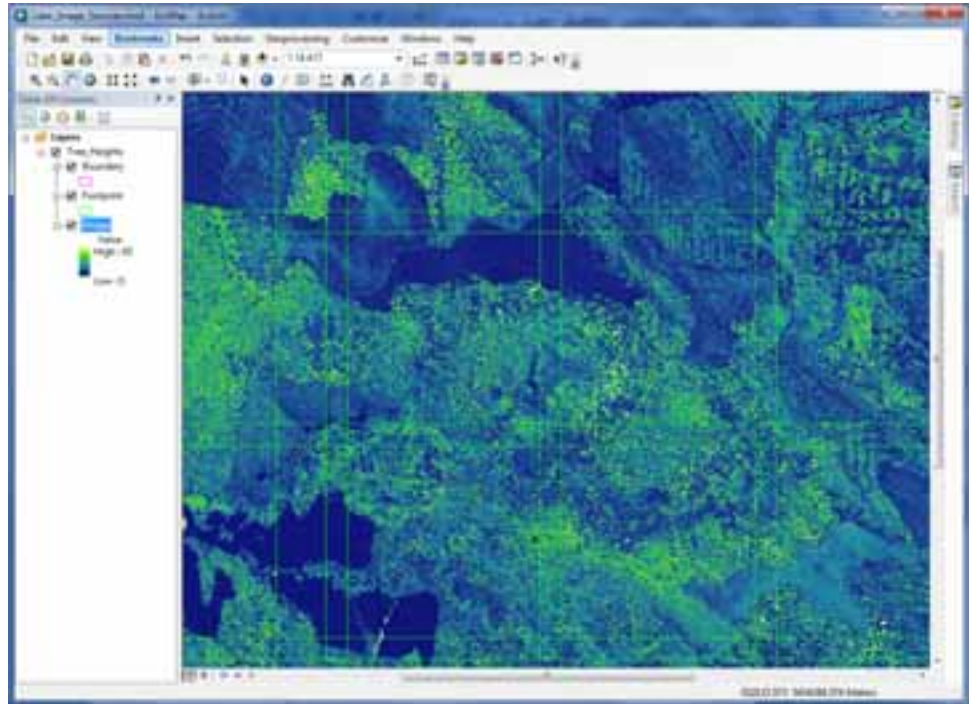**Height Estimation Image Service**

Figure 12 shows the results of the height estimation mosaic dataset. In this example, light green represents tall vegetation and dark blue areas represent low vegetation.

The real vegetation height values are returned as interpreted heights from the lidar data. Thus, any pixel can be interrogated to return the tree height. The results can be used by the foresters to aid in forest stand delineation and field inventory layout and to estimate size class and other forest inventory metrics.

This referenced mosaic dataset can now be published as an image service using the ArcGIS Server Image extension. To publish a mosaic dataset, right-click on the mosaic dataset from within ArcCatalog and select Publish to ArcGIS Server.

It is important to note here that the results displayed by the ArcGIS Server Image extension are only calculated on the fly for the extent of the ArcMap window and at the display scale of the map window. This means that broad acreage data sources can be used as the input for the elevation services, and the results are only calculated for the screen area. This reduces data duplication and returns results quickly.

Conclusion   This paper has demonstrated that the benefits to the forest industry for the use of lidar are wide and varied. These include methods to

■   Analyze and validate raw lidar data with ArcGIS before any extensive processing occurs.

■   Store and manage millions of lidar returns within the geodatabase in a seamless dataset, regardless of the number of original lidar files.

■   Extract DEMs and DSMs from the lidar data and store them as terrains in a geodatabase or as raster elevation files.

■   Extract vegetation density estimates and tree height estimates from lidar, which aid in growth analysis, fertilization regimes, logging operations, and supplementing and validating traditional field forest mensuration techniques.

■   Serve and analyze large amounts of lidar data as a seamless dataset to GIS clients, reducing the need to analyze each lidar file on a file-by-file basis, providing good overall analysis of the forest.

ArcGIS is a complete system for managing, storing, and analyzing lidar data. Coupling ArcGIS with the ArcGIS Server Image extension enables organizations to distribute and access large amounts of lidar data quickly and efficiently without the need to produce additional resultant datasets.

Acknowledgments   The author would like to thank Forestry Tasmania, Australia, for the sample lidar data used as examples throughout this document.

About Esri

Since 1969, Esri has been helping organizations map and model our world. Esri's GIS software tools and methodologies enable these organizations to effectively analyze and manage their geographic information and make better decisions. They are supported by our experienced and knowledgeable staff and extensive network of business partners and international distributors.

A full-service GIS company, Esri supports the implementation of GIS technology on desktops, servers, online services, and mobile devices. These GIS solutions are flexible, customizable, and easy to use.

Our Focus

Esri software is used by hundreds of thousands of organizations that apply GIS to solve problems and make our world a better place to live. We pay close attention to our users to ensure they have the best tools possible to accomplish their missions. A comprehensive suite of training options offered worldwide helps our users fully leverage their GIS applications.

Esri is a socially conscious business, actively supporting organizations involved in education, conservation, sustainable development, and humanitarian affairs.

Contact Esri

1-800-GIS-XPRT (1-800-447-9778)
Phone: 909-793-2853
Fax: 909-793-5953
info@esri.com
esri.com

Offices worldwide
esri.com/locations

380 New York Street
Redlands, CA 92373-8100 USA