An ESRI® White Paper • June 2010

Lidar Analysis in ArcGIS® 9.3.1 for Forestry Applications

ESRI 380 New York St., Redlands, CA 92373-8100 USA
TEL 909-793-2853 • FAX 909-793-5953 • E-MAIL info@esri.com • WEB www.esri.com

Copyright © 2010 ESRI
All rights reserved.
Printed in the United States of America.

The information contained in this document is the exclusive property of ESRI. This work is protected under United States copyright law and other international copyright treaties and conventions. No part of this work may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or by any information storage or retrieval system, except as expressly permitted in writing by ESRI. All requests should be sent to Attention: Contracts and Legal Services Manager, ESRI, 380 New York Street, Redlands, CA 92373-8100 USA.

The information contained in this document is subject to change without notice.

ESRI, the ESRI globe logo, ArcGIS, 3D Analyst, ArcMap, ArcEditor, ArcToolbox, ArcCatalog, ArcInfo, ArcGlobe, www.esri.com, and @esri.com are trademarks, registered trademarks, or service marks of ESRI in the United States, the European Community, or certain other jurisdictions. Other companies and products mentioned herein may be trademarks or registered trademarks of their respective trademark owners.

Lidar Analysis in ArcGIS 9.3.1 for Forestry Applications

An ESRI White Paper

Contents **Page**

Executive Summary 1

Keywords 1

Author 1

Introduction 2

What Is Lidar? 2

Advantages to the Forest Industry 3

Managing and Understanding Lidar Data 4
- Understanding Raw Lidar Data 4
  - Point File Information Tool 5
  - Lidar Classification in ArcGIS 7
- Loading the Lidar Files to ArcGIS 8
  - LAS To Multipoint Tool 9

Visualizing and Storing Lidar Data with ArcGIS 12
- Visualizing Lidar Data 12
  - Advantages of a Raster 13
  - Advantages of a Geodatabase Terrain 13

Building and Delivering DEMs and DSMs from Lidar 14
- The Workflow to Create a Terrain and Deliver to Clients 14
- Building a Geodatabase Terrain 15
- Building a Raster DEM 20

Analyzing Lidar Data for Foresters 22
- Calculating Vegetation Characteristics from Lidar Data 22
  - Tree Height Estimation 22

Contents **Page**

Biomass Density Calculation 23
Point to Raster 24
Replacing NoData Values as Zero Vegetation Density 24
Merging the Aboveground and Ground Results 25
Creating a Floating Point Raster File 25
Calculating Density 26

Distributing Large Lidar Datasets 27
Preparing Raster DEM for Serving with the ArcGIS Server Image Extension 27
Serving an Elevation Service through the ArcGIS Server Image Extension 28
Creating an Elevation Image Service 30
Visualizing an Elevation Service 33
Estimating Tree Height Using Elevation Services 37
Preparing Image Service Data 37
Creating the Height Estimation Service 38
Adding .ISRef Files to the Height Estimation Service 39
Adding Algebraic Process to Service to Return Height Estimation 41
Applying the Image Algebra Process 42

Conclusion 45

Acknowledgments 46

J-9884

Lidar Analysis in ArcGIS 9.3.1 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 raw lidar data with ArcGIS before any extensive processing occurs
- Storing and managing millions of lidar returns 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 an excellent tool for managing, storing, and analyzing lidar data. Coupling ArcGIS with the ArcGIS Server Image extension, the forestry professional is able to access large amounts of lidar data quickly and efficiently without the need to produce additional resultant datasets.

Keywords Lidar, ArcGIS, terrains, geodatabase, ArcGIS Server Image extension

Author Gordon Sumerling, ESRI Australia Pty. Ltd., Adelaide, South Australia

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 data into ground and canopy returns.
- Pass the resultant point clouds to a terrain that creates a viewable and displayable surface.
- Perform analysis on the terrain 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.

Also captured by the lidar sensors is the intensity of each return. 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 pseudoimage to provide context of the lidar acquisition area. See the image below:

J-9884

Figure 1
Lidar Intensity Image

Lighter areas represent strong returns. Darker areas represent weaker 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 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

http://www.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.

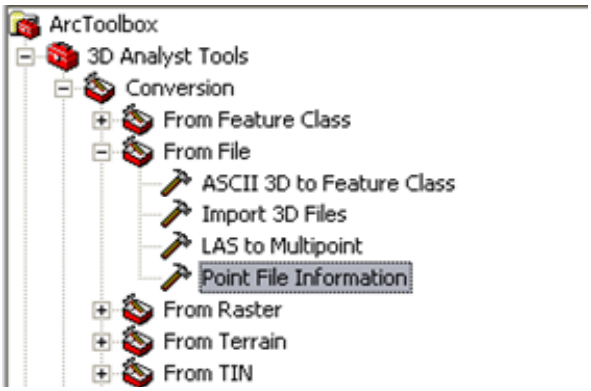

The tool is designed to read the headers of LAS or scan ASCII files and summarize the file contents. As a single lidar file often contains millions of points and many lidar datasets contain more than one file, the Point File Information tool can accommodate reading one or more files by specifying either individual lidar files or folders.

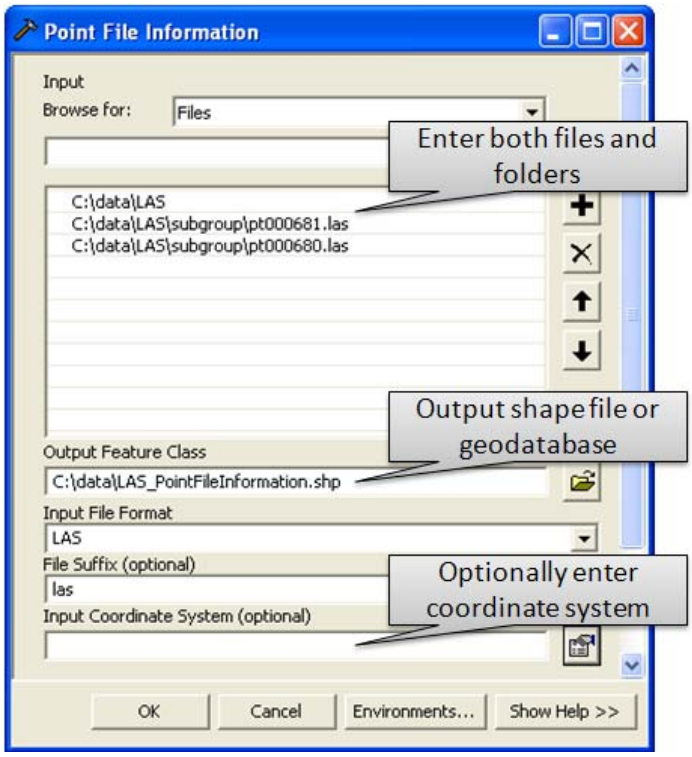

The result from this tool is a feature class that shows the

- 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, ArcMap will flag 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.

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, min and max 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.

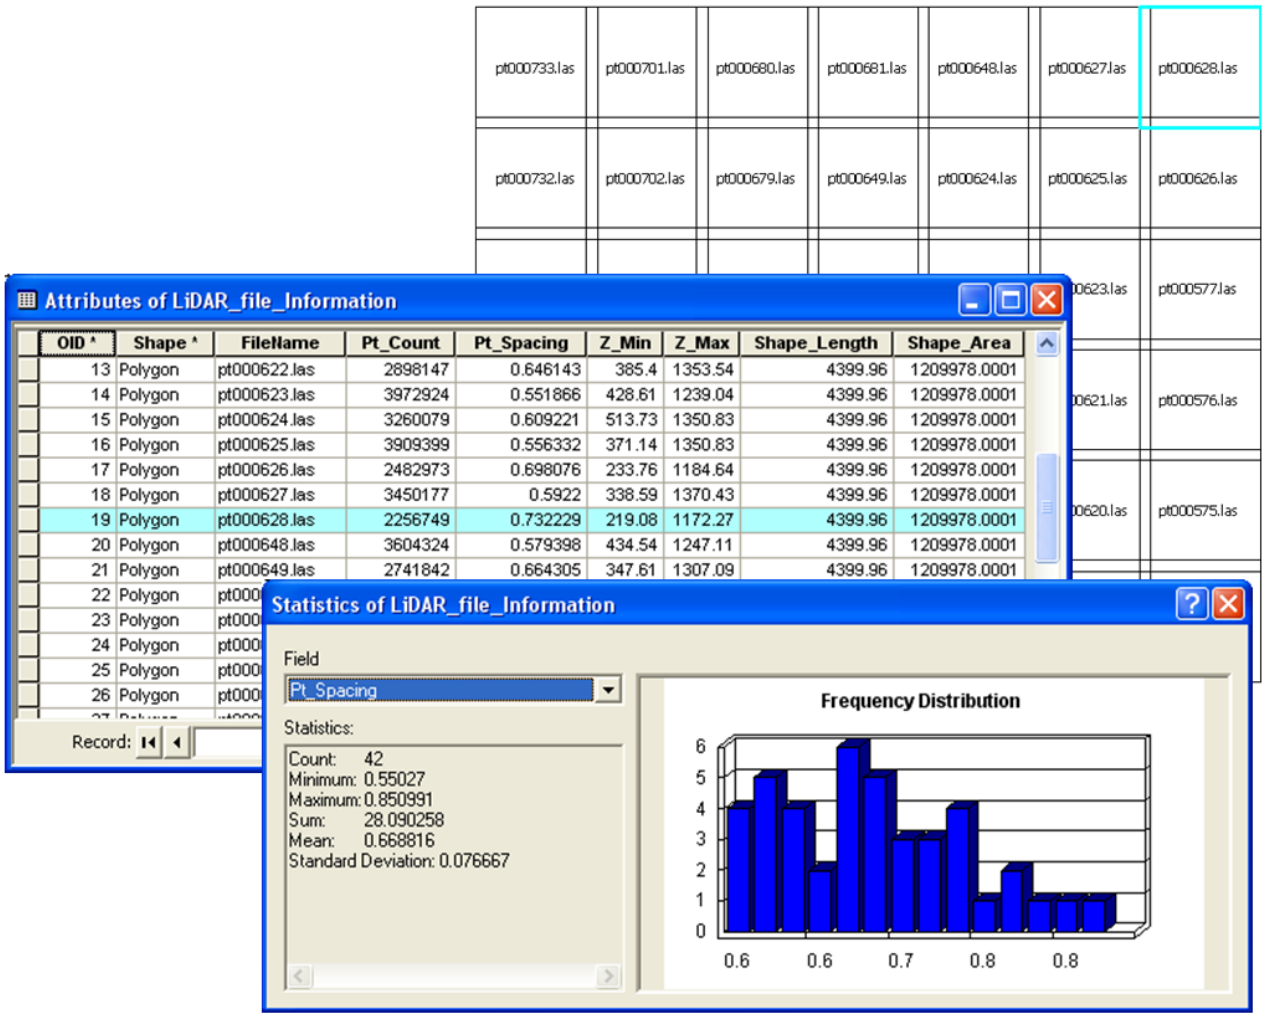

Lidar Classification in ArcGIS

LAS files contain a classification field that identifies each point's return type. 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 and stores it as a binary large object (BLOB) in the geodatabase. This is not exposed through software, so users will require prior knowledge of their data to separate it by the classification values.

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.

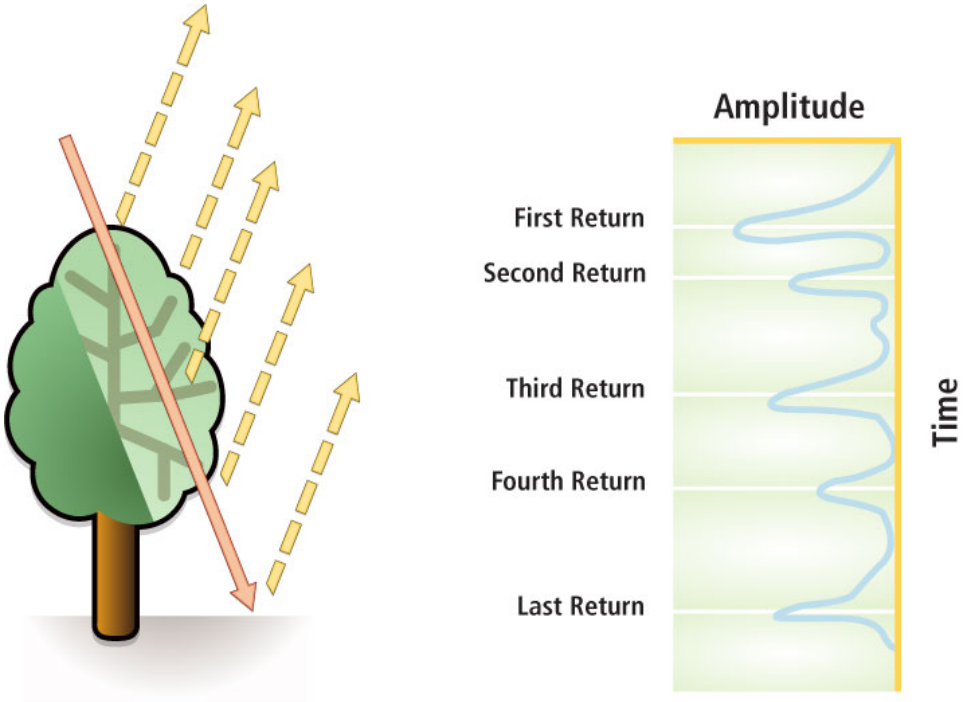

These multiple returns create a data management challenge. A single lidar file can typically be 60 MB to 100 MB in size and can contain several million points. If this data is loaded directly into a table, it creates many millions of records, which results in a large data file that becomes difficult 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 dramatically increases database efficiency by storing the data in an ESRI BLOB field. A BLOB is a collection of binary data stored as a single entity in a database management system and enables data compression. Any tool written to exploit these fields needs to understand the ESRI BLOB field structure.

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. The ability to store the data as a BLOB structure helps with data management by reducing the space used to store the data.

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 sometimes 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

http://www.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 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 result shown in the example below is a multipoint feature class of ground returns.

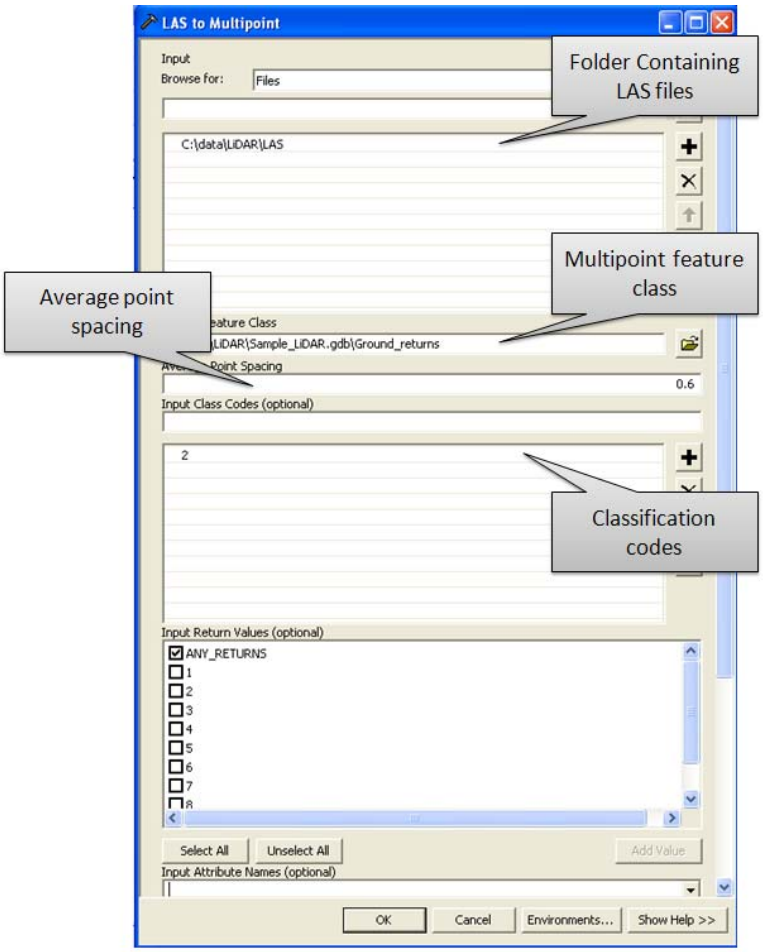

The screen capture below represents ground returns using a series of points. In its raw form, a multipoint feature class is not designed to be displayed; it is designed as an efficient storage medium for the many millions of points found in a lidar dataset.

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

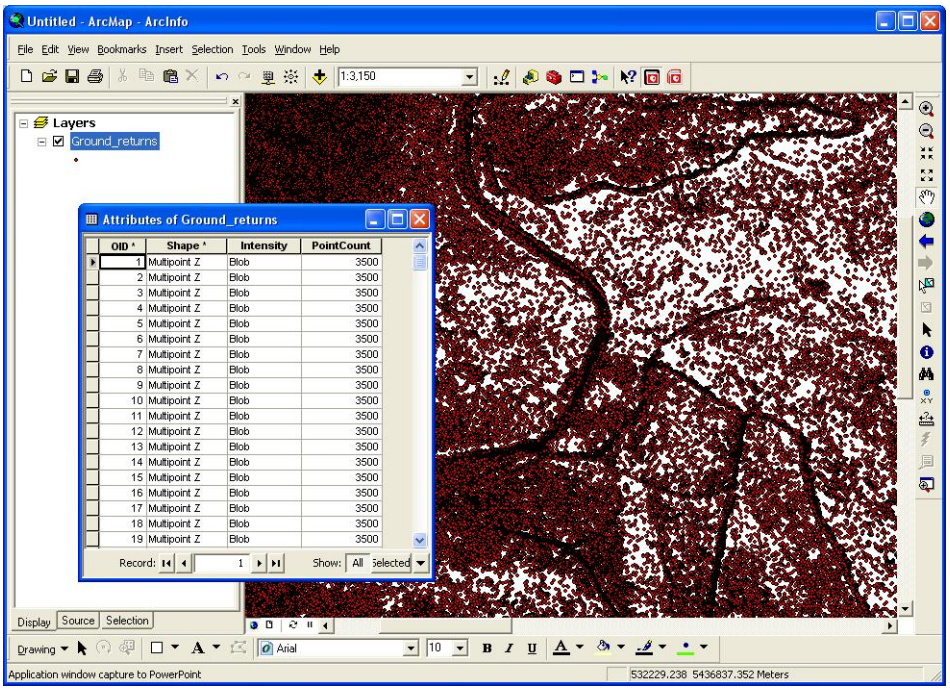

Ground return multipoint feature class shows Shape, Intensity, and PointCount fields.

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.

J-9884

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. 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.

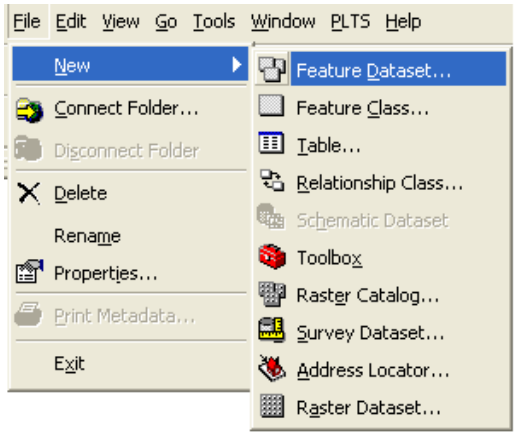

It is important when creating a feature dataset that the coordinate system be the same as the data to be used in the terrain dataset. All features that reside in a feature dataset, including the terrain, have the same coordinate system.

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.

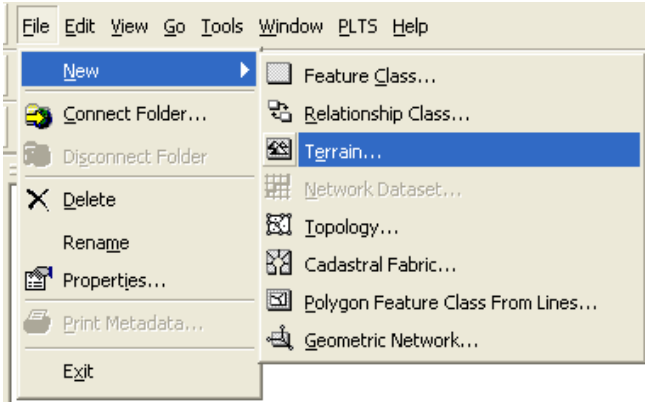

Note: The ArcGIS 3D Analyst extension will need to be active to perform this operation.

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.

J-9884

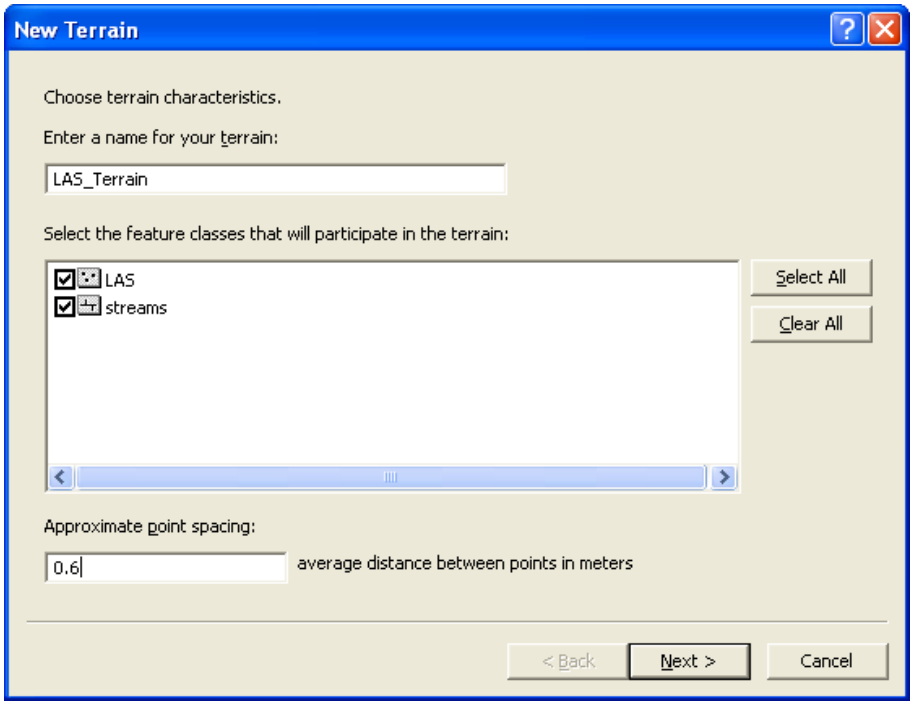

On the next wizard form, select feature class characteristics and decide how the feature classes affect the interpretation of the terrain and whether they are viewable at all levels of the terrain. It is important to ensure that if there are breaklines or areas of interest, the field containing the z-values is used in the height source.

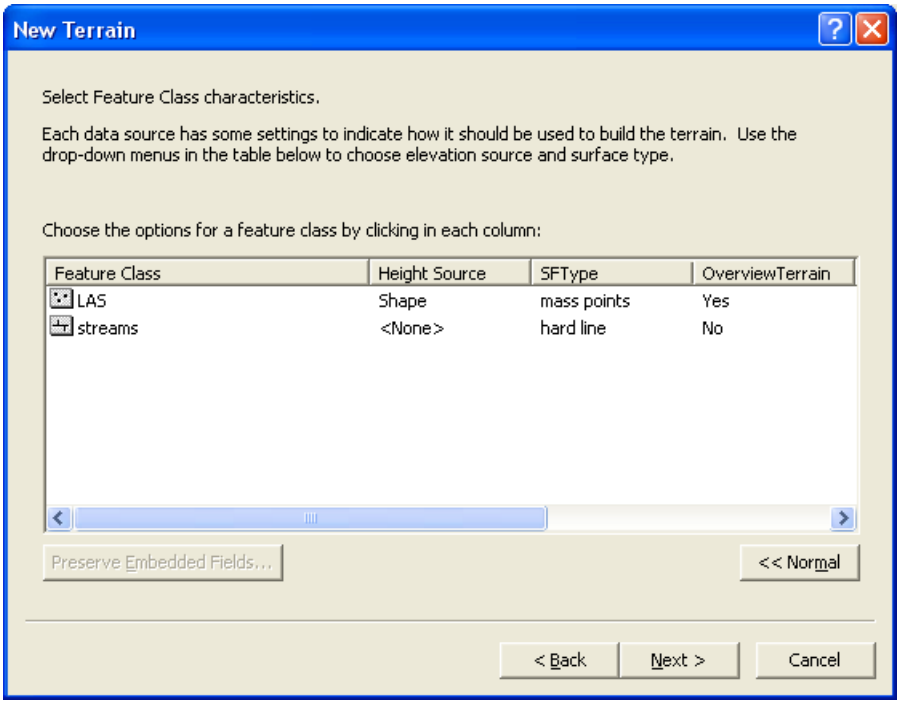

Terrain datasets can be pyramided using one of two point-thinning filters: z-tolerance and window size. Pyramids are reduced resolution versions of the original underlying data, which are displayed when working at small scales. This reduces display times.

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 because it is significantly faster. If thinned data can be used for analysis, which is reasonable if the lidar is oversampled for user needs, the z-tolerance filter should 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 Mean 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.

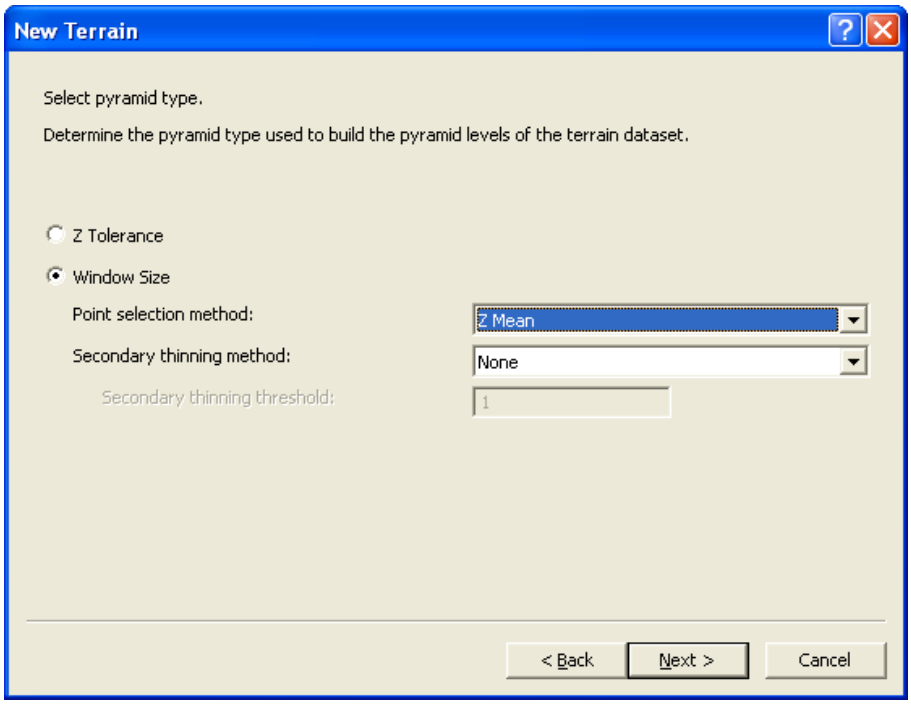

Finally, generate the terrain pyramids.

Pyramids are levels of detail generated for a terrain dataset to improve efficiency for some applications. They are used as a form of scale-dependent generalization. Pyramid levels take advantage of the fact that accuracy requirements diminish with scale. They are similar in concept and purpose to raster pyramids, but their implementation is different.

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.

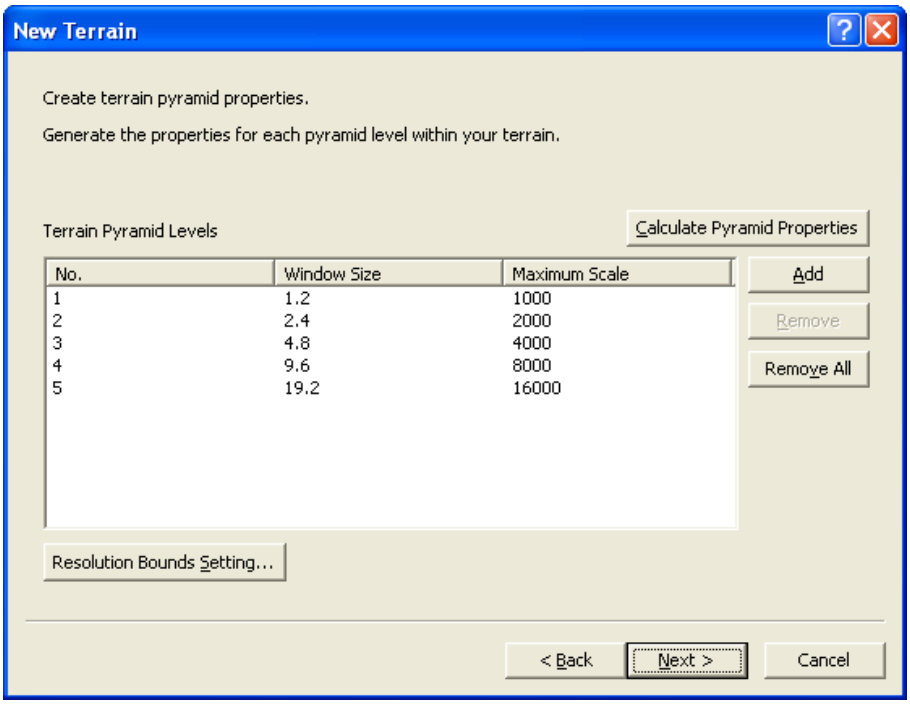

Once the process is finished, the wizard will prompt the user to build the terrain. The number of pyramids chosen to be built will dictate how long it will take to build the terrain.

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

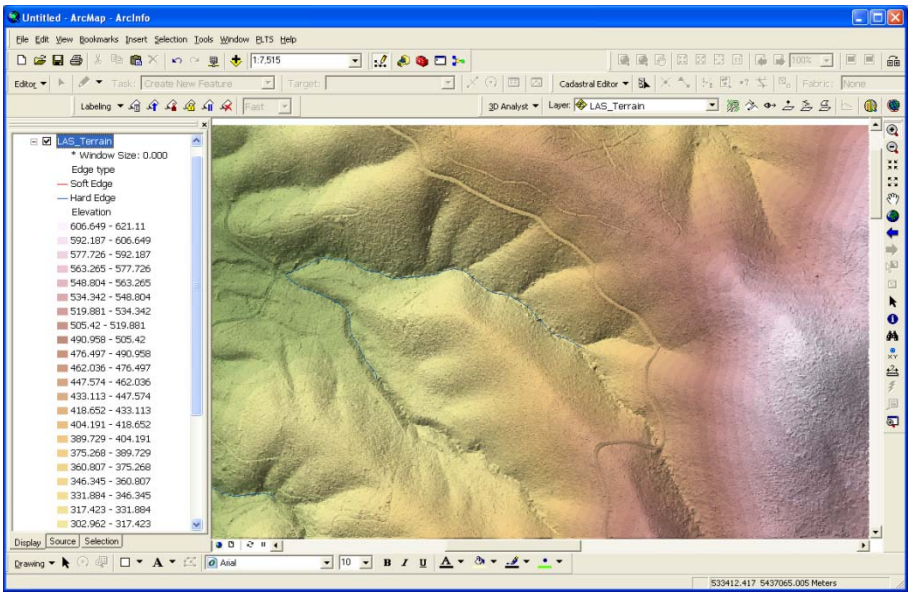

Building a Raster DEM

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.

The export formats supported by the Point To Raster tool include, but are not limited to, .tif, .jpg, 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, while 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

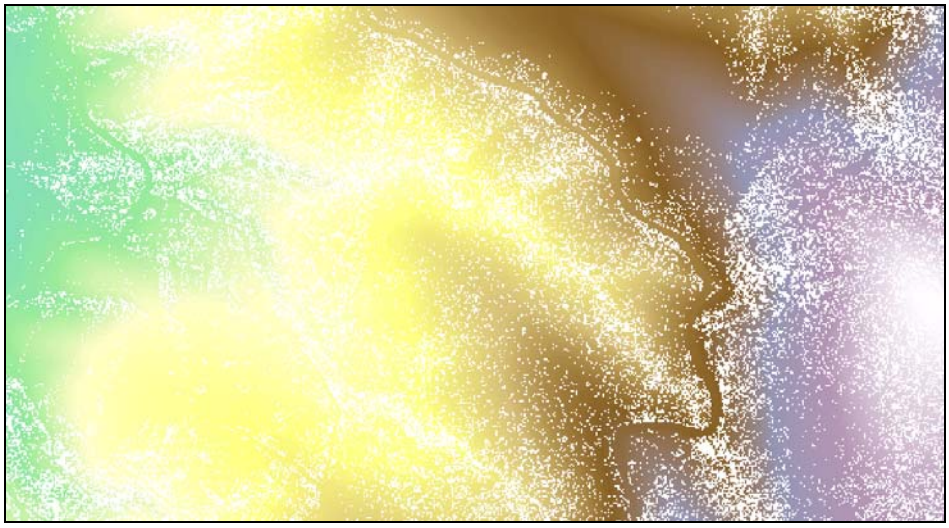

It is possible to reduce this effect by postprocessing the raster DEM with the ArcGIS Spatial Analyst Raster Calculator tool. Using the Conditional evalution function, each cell in the raster DEM is evaluated for the NoData value. If the evaluation is true, then a floating filter is used to gain the 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 statement for removing the NoData values:

```
Con(IsNULL([inputraster]), FOCALMEAN ([inputraster],
RECTANGLE, 3, 3, data), [inputraster])
```

In this example, the condition applied is IsNULL.

- If it returns true, then a focal mean 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

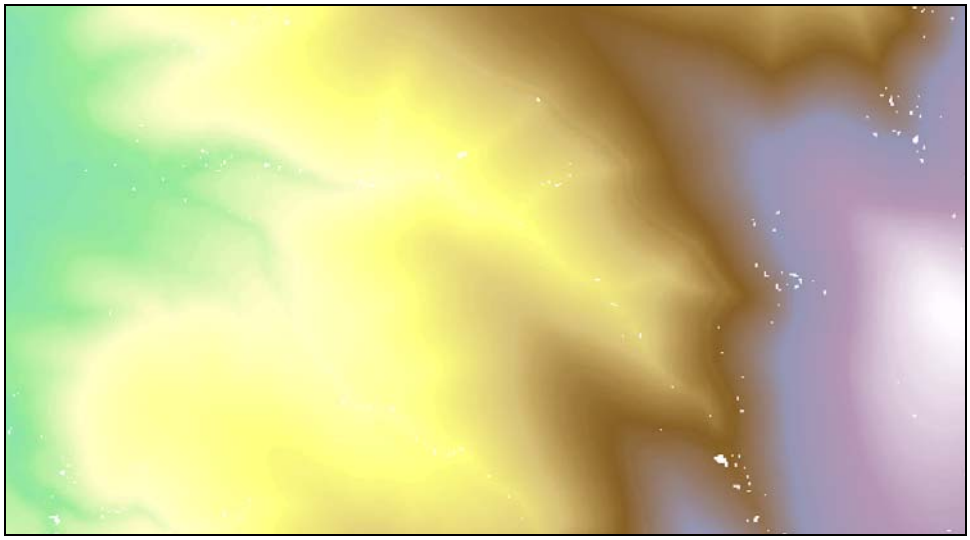

Analyzing Lidar Data for Foresters

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

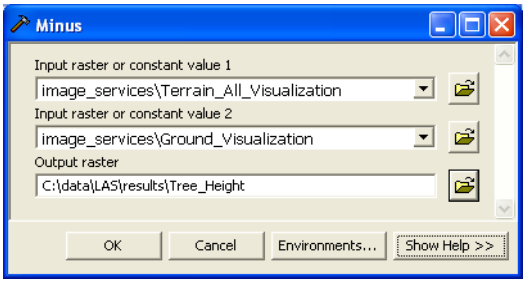

In the results below, blue portrays low vegetation, and red portrays high vegetation. In this example, a break in tree growth can be clearly seen in the area highlighted as a hard line between tall and low growth.

Figure 6
Height Estimation Raster

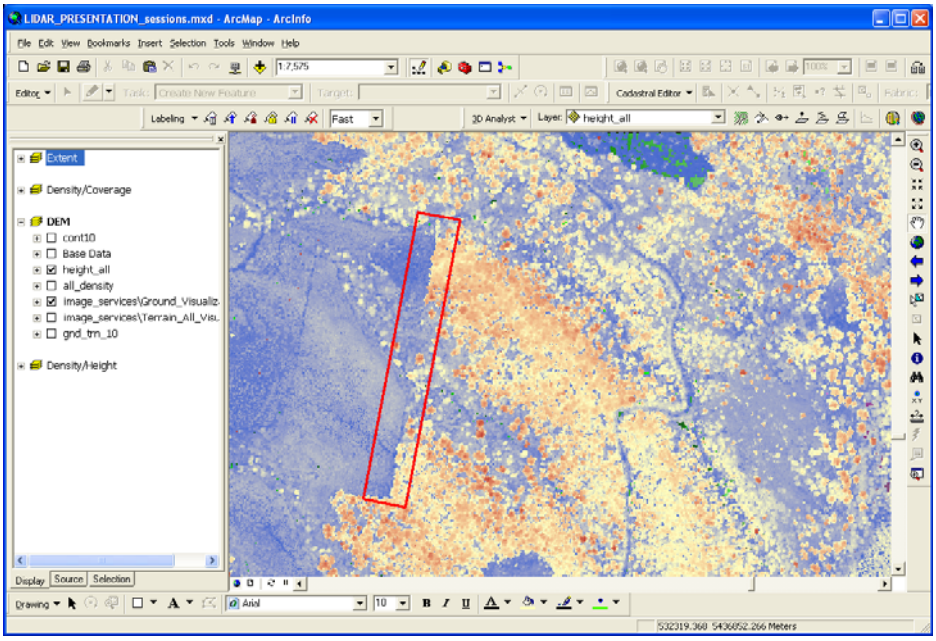

Biomass Density Calculation

The biomass density will give an indication of tree vigor and growth. Where a forest is of the same species, areas of poor soil nutrients are easily identified by the proportion of biomass. Lower biomass readings indicate poorer soil conditions. Higher biomass readings indicate ideal soil conditions for tree growth.

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

The first stage of this process is to convert the multipoint feature classes to raster files. When calculating the DEM and DSM, interest is in the height values of the lidar data, so the Shape field of the multipoint feature class is used to provide the height of the surface. In this process, the focus is density, so the height is not appropriate. In this case, use COUNT for the cell assignment, which gives an approximate density.

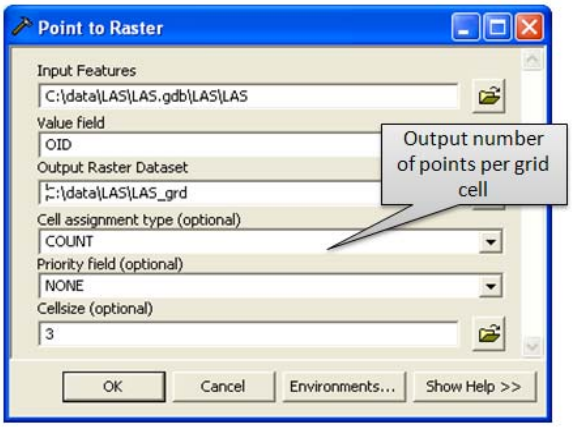

Replacing NoData Values as Zero Vegetation Density

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.

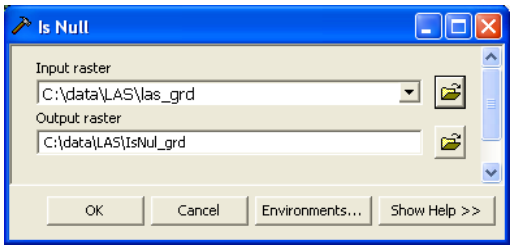

The second process then merges the original raster file with the NULL raster file so that the resultant raster file has a complete range of values from 0 upward. There are no cells that have an unassigned value. The tool used here is the Con tool in ArcGIS Spatial Analyst (ArcToolbox\Spatial Analyst 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.

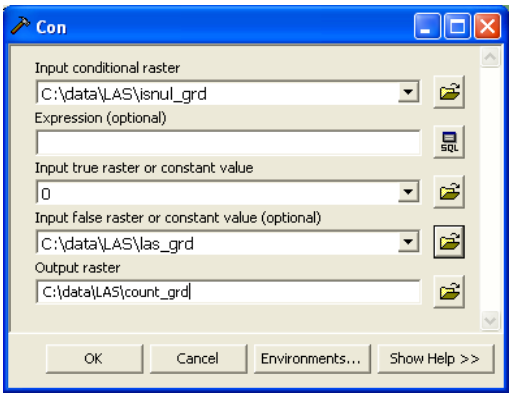

Repeat the above three processes for the aboveground or canopy grids.

Merging the Aboveground and Ground Results

Merge the above ground 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).

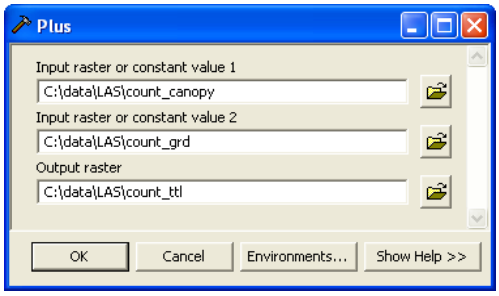

Creating a Floating Point Raster File

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).

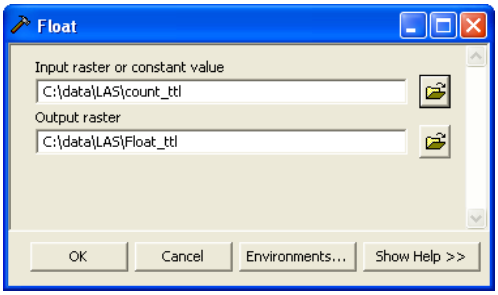

Calculating Density

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.

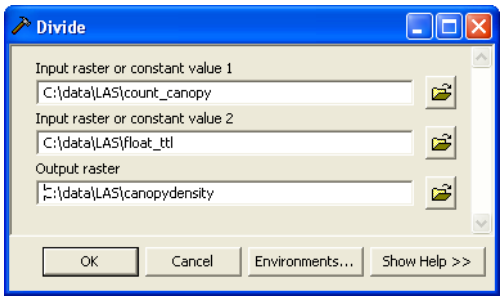

The canopy density raster is depicted in the screen capture below. Yellow represents low-density canopy coverage, and dark blue represents high-density canopy coverage.

Figure 7
Canopy Density Raster

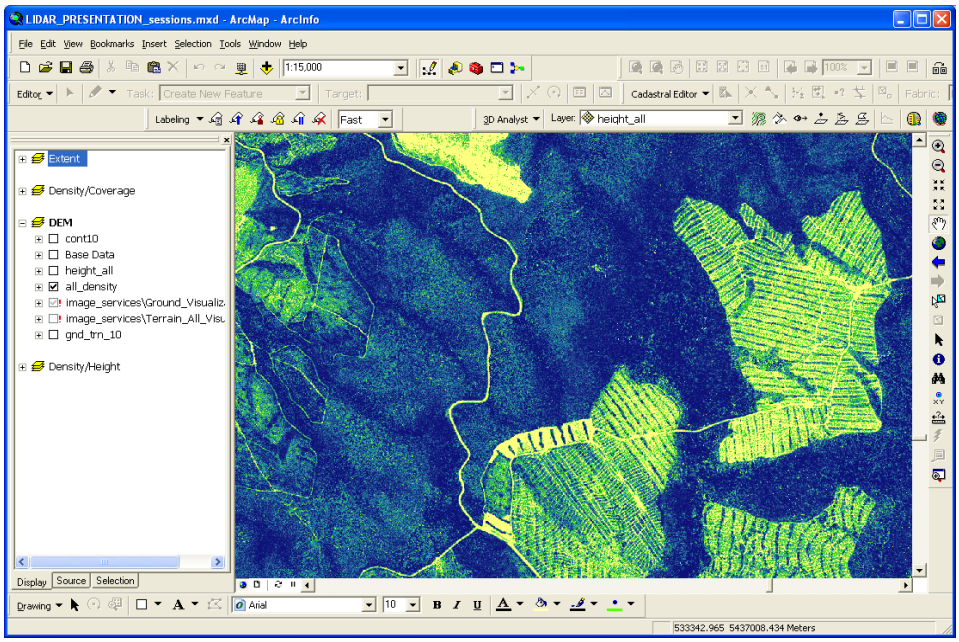

Distributing Large Lidar Datasets

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 or a single terrain. 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, the ArcGIS Server Image extension is the optimum solution 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 data, it is necessary to output geodatabase terrain 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 producing a raster for further analysis, then it is recommended that an ESRI Grid file be produced. If it is for including in an image service, then a TIFF file is more efficient.

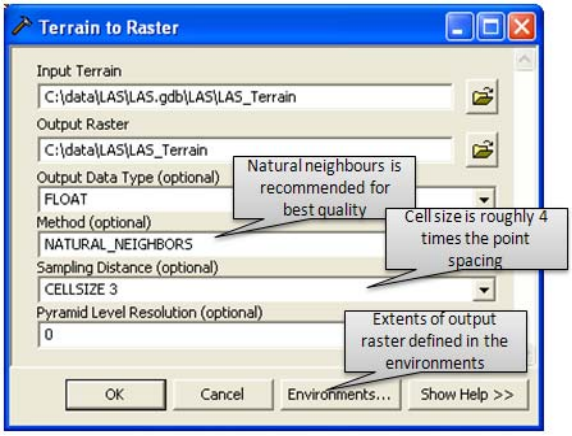

For an interpolation method, natural neighbors provides the smoothest result. Although it is not as fast to produce as linear interpolation, it is more accurate.

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 the raster format. 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 primarily used to serve and analyze image datasets seamlessly such as aerial photographs and satellite images. It also has the capability to serve terrain data and visualize the terrain data when it is stored in a raster dataset. This is done 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.

When the elevation service is visualized via the ArcGIS Server Image extension, any of the following interpretations can be seen:

- Elevation coded
- Hillshade
- Shaded relief
- Slope
- Aspect
- Curvature

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 the user can view the entire lidar dataset rather than viewing each lidar file. This enables the user to quickly see the DEM and DSM and perform comparisons.

Figure 8
Example of a Hillshade

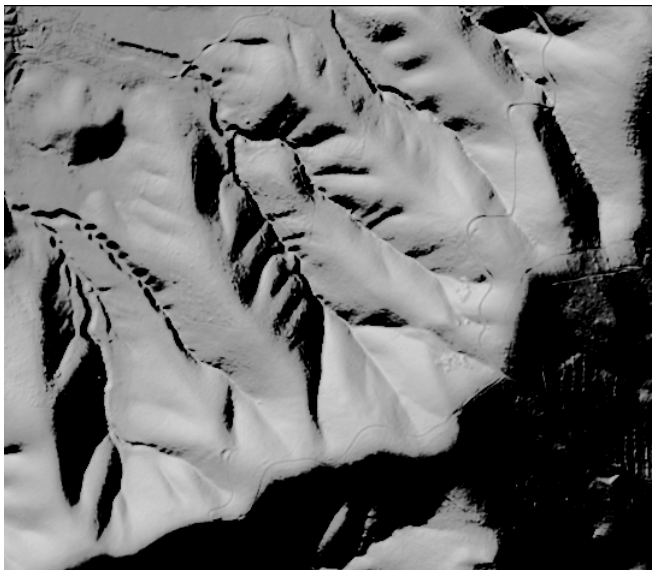

Figure 9
Example of a Shaded Relief

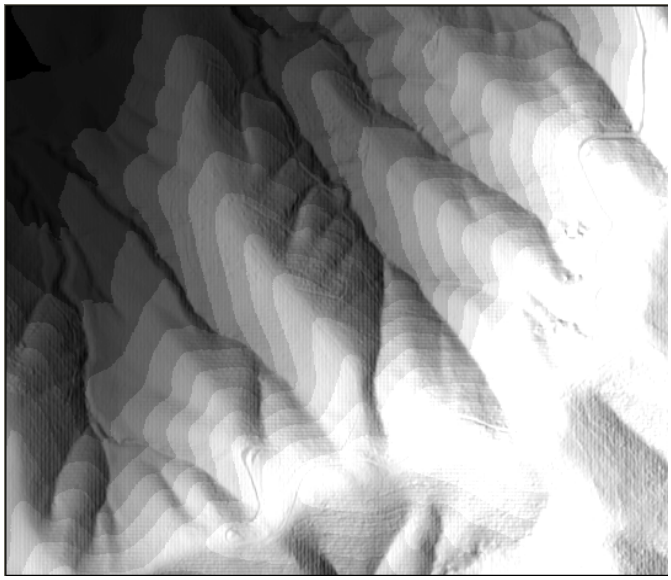

Creating an Elevation Image Service

Creating an elevation service in the ArcGIS Server Image extension is essentially the same process as developing a standard image service, 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 image service. It uses the New Image Service Wizard from the Image Service Editor toolbar in ArcMap.

- From the Image Service Editor toolbar in ArcMap, select Image Service > New Image Service.
- The New Image Service Wizard appears.
- Navigate through the wizard to the Please specify the following parameters to define your image service form and set the following parameters:
  - New image service location
  - Spatial reference system (coordinate system of the service)
  - Service type (elevation)

This service type defines the image service.

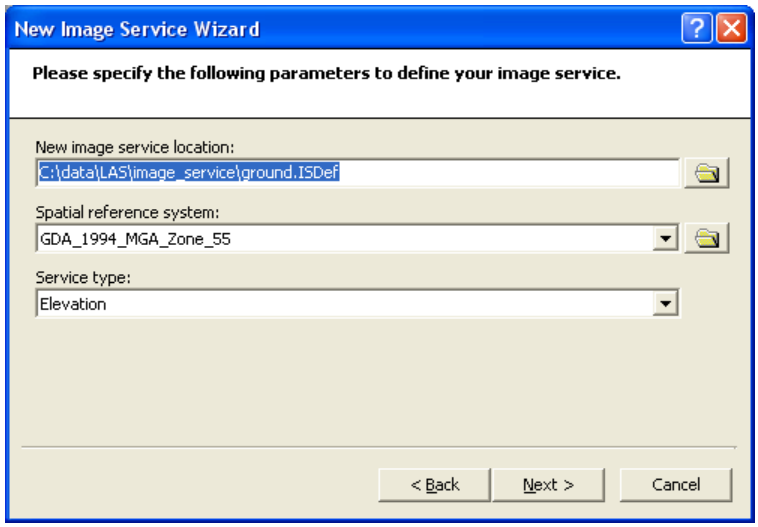

- Navigate through the wizard to the form How is the elevation image service going to be used?
  - Click the button to select the option, As an elevation image service.

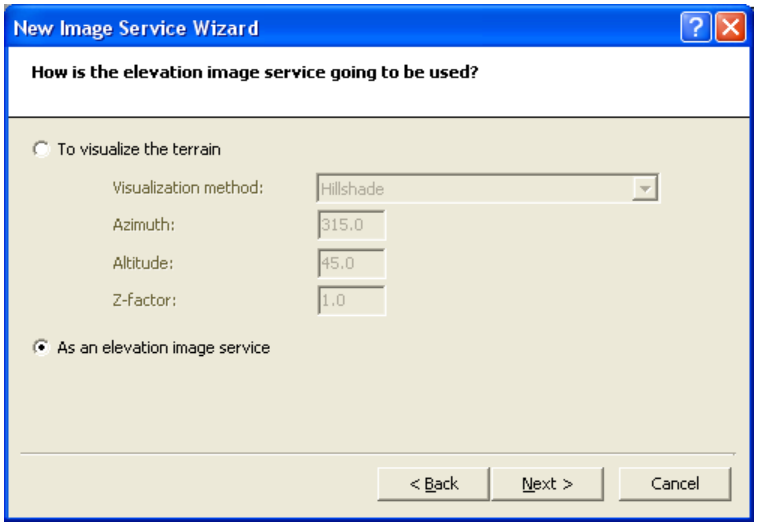

- Navigate to the Input data for the image service form and select the following parameters:
  - Raster type (In this example, Elevation/TIFF is the image elevation image source.)
  - Input (for the data folder)
  - Use this spatial reference for all input data (the coordinate system of the elevation data)

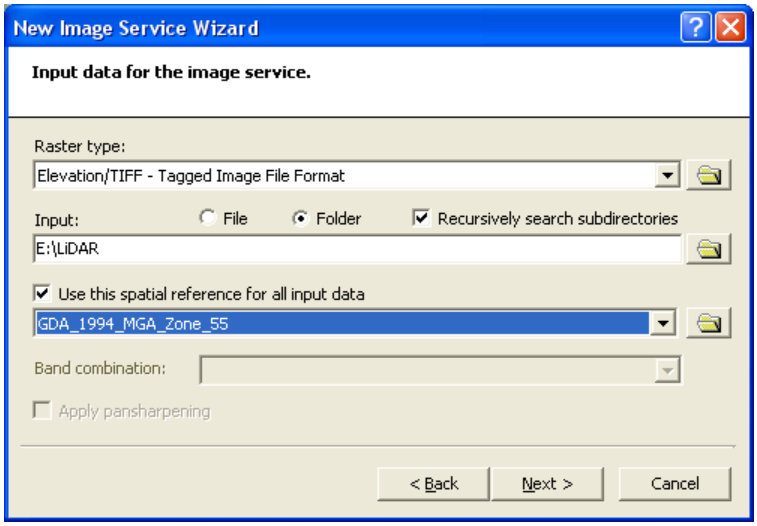

Note: If the images are to be located on a system separate from the ArcGIS Server Image extension, the directory mappings should be by UNC path, for example, \\server\share\data. Do not use drive mappings such as J:\data.

- Navigate to the form Add more data to your image service? Select the following settings:
  - Click No to Do you want to add more data?
  - Select both the Generate overviews and Compile the service for publishing check boxes.

- Navigate to the Enter the information about the image service form. This form is where the metadata about the image service is stored. Enter information for Service name, Title, Geographic region, Pixel unit, Pixel source, Publisher, and Contact organization parameters. This metadata information is transmitted with the image service and can be accessed by client systems.

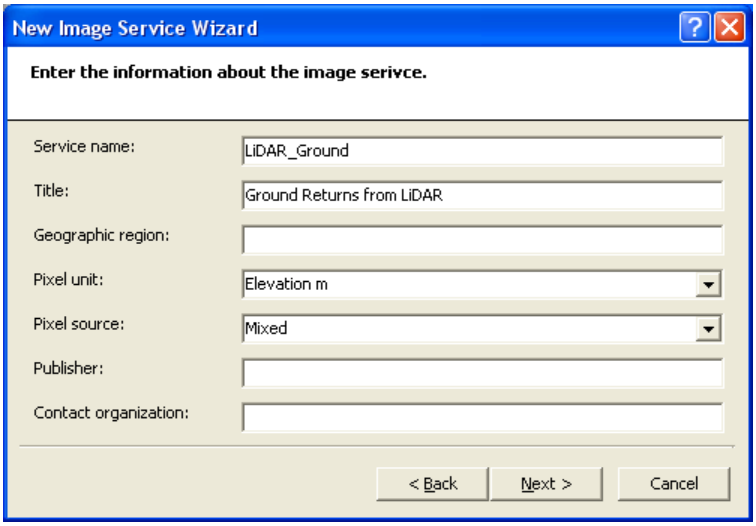

- Navigate to the last form and select Finish to create the service.

The image server process will now generate the Elevation Image service. It will create the service overviews and compile the service ready for publishing as an image service. The service overviews are smaller, lower-resolution images based on the originals but provide the ArcGIS Server Image extension with its performance. The resultant image will look similar to the screen capture below.

Figure 10
Elevation Image Service

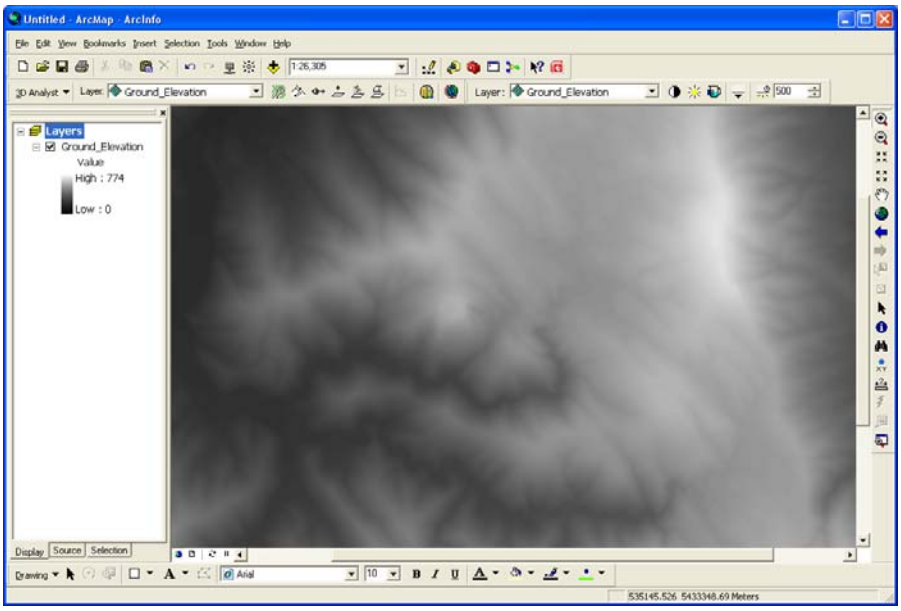

This is the completed stage of an elevation service creation. The elevation service can be used as an elevation source for terrain visualization in ArcGlobe or as input to ArcGIS 3D Analyst or ArcGIS Spatial Analyst geoprocessing tools such as contour or viewshed.

Visualizing an Elevation Service

The next stage shows how to apply a visualization to the elevation service. The visualization in this example is colored shaded relief.

With the image service group layer loaded in ArcMap, do the following:

- Right-click the image service group and select Properties.

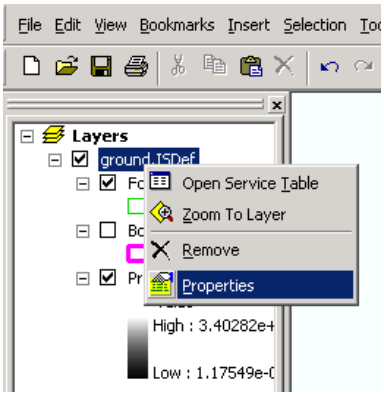

- The Image Service Properties—Service Processes dialog box appears. Select Service Processes and add the Visualize Elevation process to the Processes Selected pane.

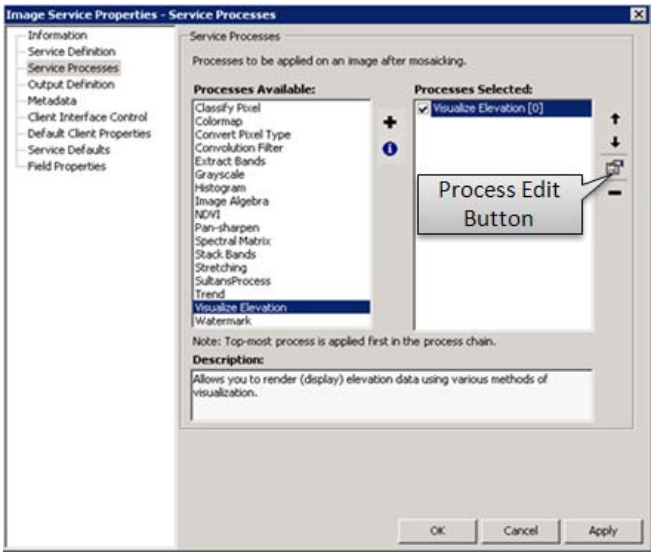

- To configure the Visualize Elevation dialog box, select the process, then click the Process Edit button. This will cause the properties of the process to appear.

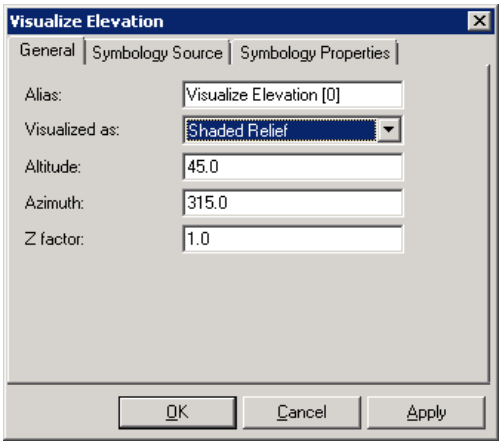

- Ensure that Shaded Relief is selected, then click the Symbology Source tab. For the initial visualization, there is no need to alter any of the default settings.

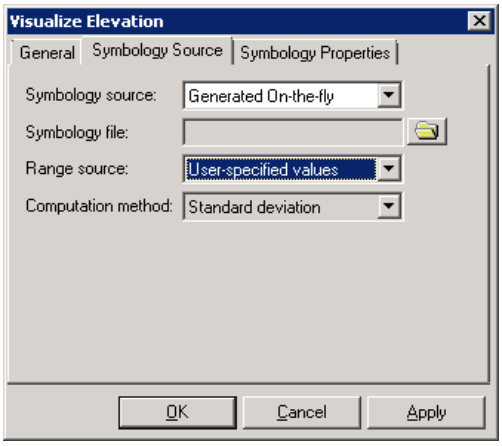

- Make sure Range source is set to User-specified values and Symbology source is set to Generated On-the-fly.

Note: The range source is important for fixing the colors for the terrain visualization. The default for the range source is computed from the area of interest. If this is configured, the color visualization will change as the map is zoomed and panned, while the elevation ranges will change for the map currently being viewed.

With these configured, select the Symbology Properties tab.

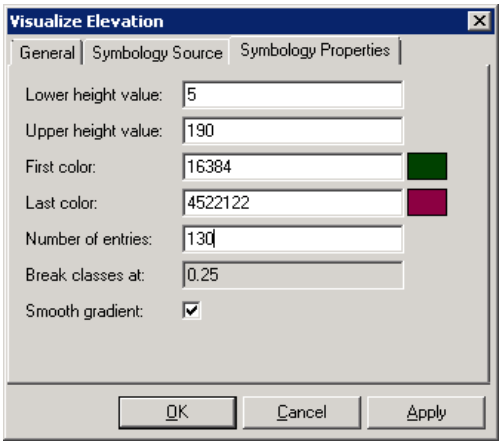

This tab is data dependent. It may take several attempts before the best values are established. It is imperative to understand the minimum and maximum height in the dataset to set the Lower height value and Upper height value parameters. In this service example, Lower height value is set to 5 meters, and Upper height value is set to 190 meters.

- The First color and Last color parameters are personal preferences; however, the values of 16384 and 4522122 provide a good representation.

- The color gradient is defined by Number of entries. Using a figure that is 90 percent of the total range of values is suggested.

Click OK on the Visualize Elevation dialog box and the Image Service Properties form.

The last step is to build the service by specifying

- Compute output pixel properties.
- Compile Service.

The resultant visualization below is of the ground returns from the lidar data.

Figure 11
Visualized Shaded Relief Image Service

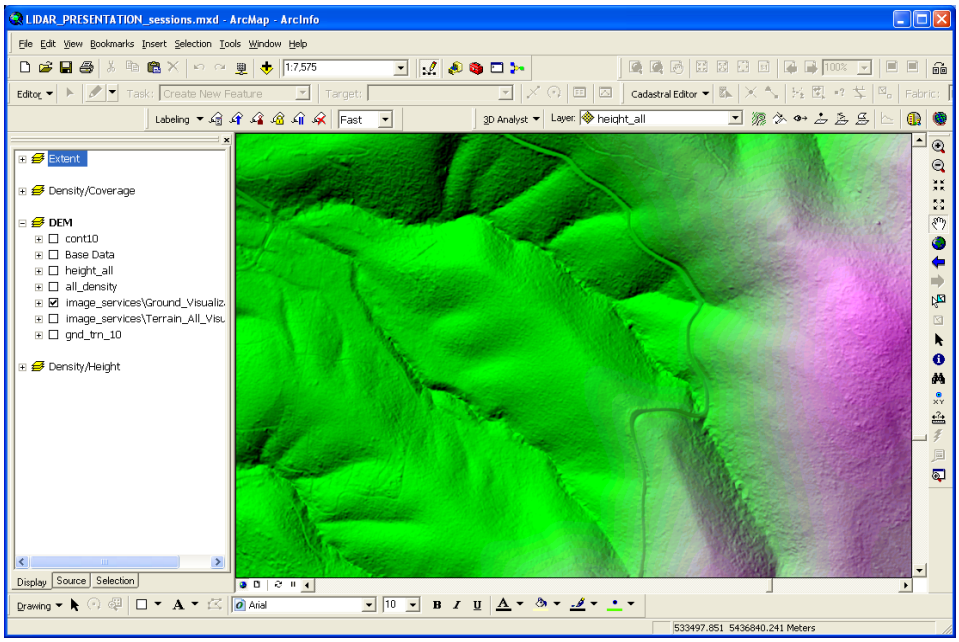

This can now be published by the ArcGIS Server Image extension for consumption by many different clients.

Estimating Tree Height Using Elevation Services

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 ArcGIS Server Image extension, similar services can be created that return the same results. The advantage of using the ArcGIS Server Image extension is there is no need to create a height estimation raster as the ArcGIS Server Image extension creates the height estimation on the fly for the area of interest on the screen, removing the need to create extra raster datasets.

This next section describes how such a service can be created.

Preparing Image Service Data

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

The height estimation image service takes the input of the two services and performs an algebraic operation between them to produce the third service. The source data for this service is two direct image server connections or .ISRef files. The .ISRef file defines an image service connection and the image service properties.

The .ISRef files are created from inside ArcMap and need to be generated before the height estimation service can be created. To create the .ISRef files, perform the following tasks:

- Load each of the DSM and DEM image services into ArcMap.
- Right-click each of the layers and choose Save As ISRef File.
- Save them as Canopy_Elevation.ISRef and Ground_Elevation.ISRef.

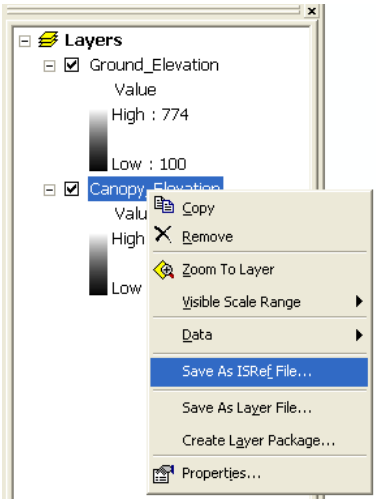

Creating the Height Estimation Service

When creating the height estimation service, it is necessary to create the shell of the elevation image service first, then add the two .ISRef files as two distinct steps. The .ISRef files are classified as a georeferenced image source and cannot be loaded into an elevation service via the New Image Service Wizard.

The image service is created from the Image Service Definition Editor toolbar in ArcMap by selecting Image Service > Advanced > New Service Definition.

When this is selected, the Image Service Definition dialog box appears. The following settings need to be made for the height estimation service:

- Service definition (name)
- Spatial reference
- Service type (Choose Elevation.)

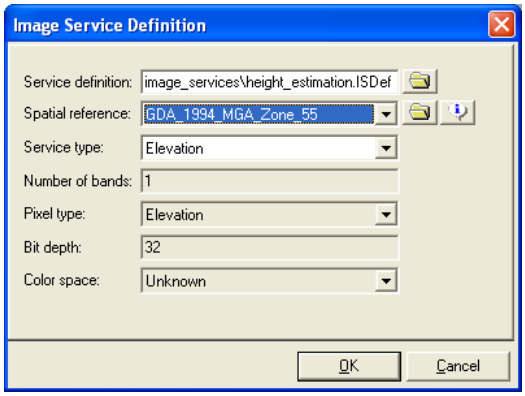

The image service is now ready to load data.

Adding .ISRef Files to the Height Estimation Service

The two .ISRef files are added to the height estimation service from the Image Service Definition Editor toolbar in ArcMap by selecting Image Service > Advanced > Add Raster Dataset.

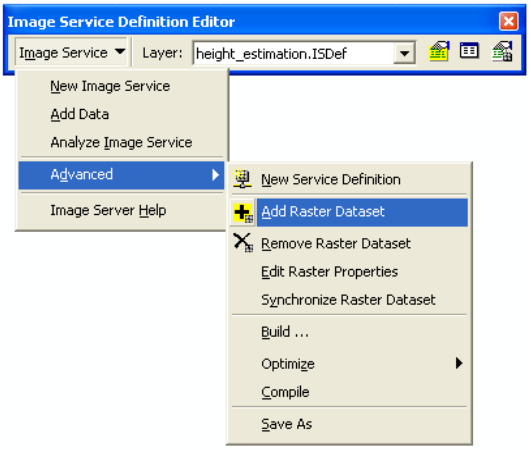

Select the raster type to be input as Direct Image Server Connection, which is found under the Georeferenced Imagery folder.

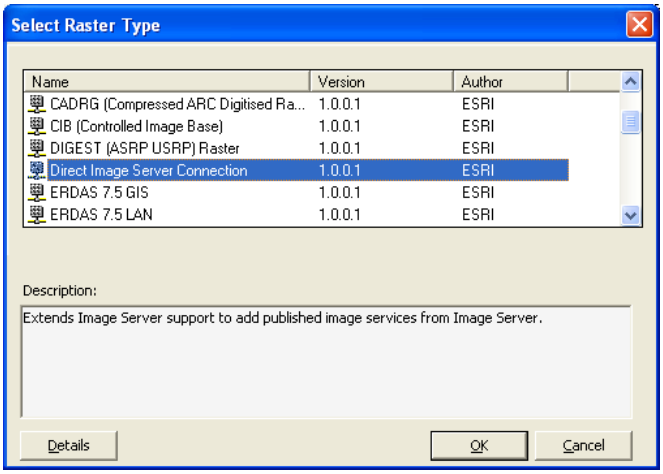

From the Add an Image Service dialog box, choose the two .ISRef files you saved earlier and click OK.

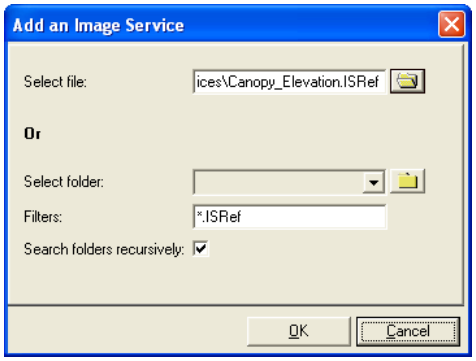

With the image sources loaded, the service now needs to be built. Select Image Service > Advanced > Build. The Build Options dialog box appears.

On the form, check the following check boxes:

- Compute pixel size ranges.
- Create service boundary.
- Compute output pixel properties.
- Load preview.

Click OK, and the initial build occurs.

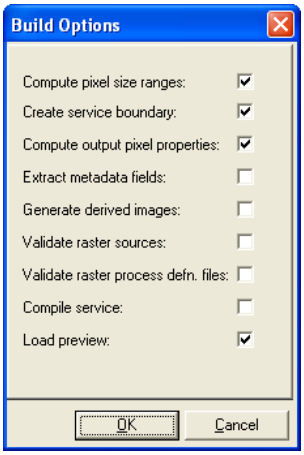

The height estimation service is now ready for the algebraic operations.

Adding Algebraic Process to Service to Return Height Estimation

The height estimation service uses the Image Algebra process from the ArcGIS Server Image extension. The Image Algebra process allows you to perform algebra on the spatially overlapping pixels from two raster bands. There are three rasters involved in this process: primary, auxiliary, and output. The primary and auxiliary rasters are the processed inputs, and the result is the output raster.

The Image Algebra process is applied directly to the primary image in the service rather than the service as a whole as described in the Visualizing an Elevation Service section. In this service, the primary image is the DSM, and the auxiliary is the DEM.

To apply the process, it is necessary to disable the auxiliary image (DEM) from all mosaic operations. To disable the DEM from the service

- Open the service table and select the DEM service.

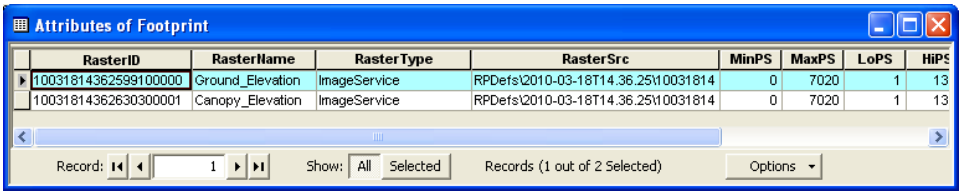

- From the Image Service Definition Editor toolbar, click the Raster Properties button.

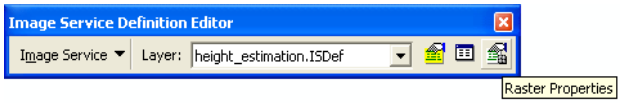

- The Raster Properties—Rasters dialog box appears.
- Select the Rasters node and uncheck the Enabled check box.

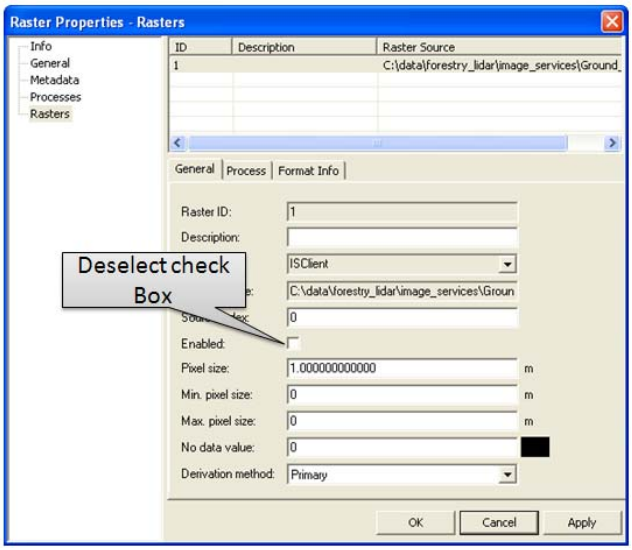

Applying the Image Algebra Process

To apply the Image Algebra process to the primary image, in this case, the DSM

- Open the service table and select the DSM service.

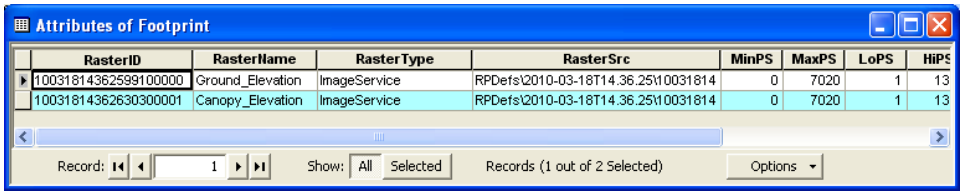

- From the Image Service Definition Editor toolbar, click the Raster Properties button.
- The Raster Properties—Processes dialog box appears.
- Select the Processes node.
- From the Processes Available pane, select Image Algebra and add it to the Processes Selected pane.

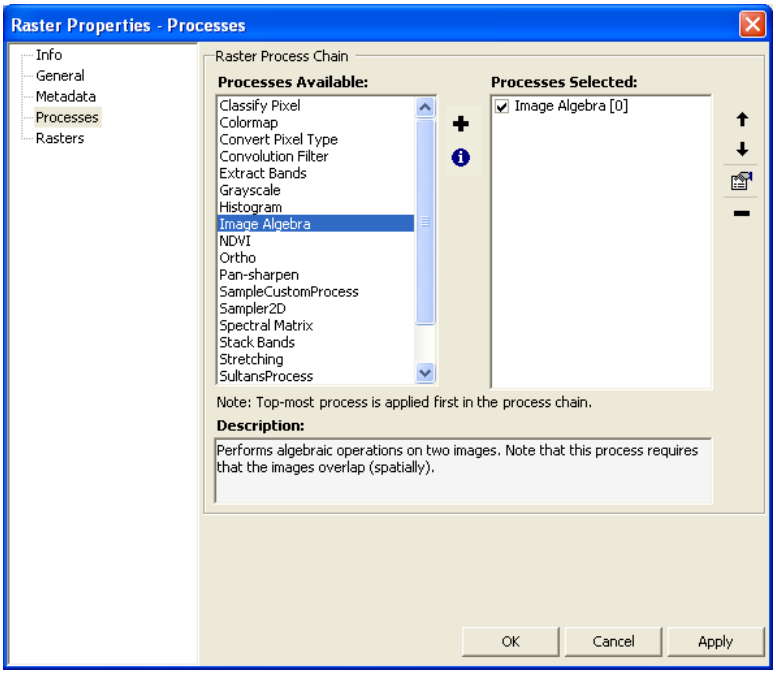

- Click the Image Algebra Properties Edit button.

- In this case, subtract the auxiliary image from the primary image to obtain a result. Ensure the Method parameter on the Primary tab is set to Subtract.

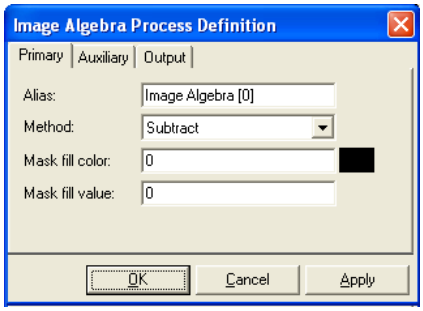

- Switch to the Auxiliary tab and enter the auxiliary raster ID. The auxiliary raster ID is a combination of two identifier values: <RasterID> and <RasterIDinRPDef>. The RasterID value is found in the RasterID column of the service table. The RasterIDinRPDef is the value in the Raster Properties—Rasters dialog box. When you open this dialog box, click the Rasters node and identify the correct ID value from the displayed table. In this case, the RasterIDinRPDef is 1. The input looks similar to the following:

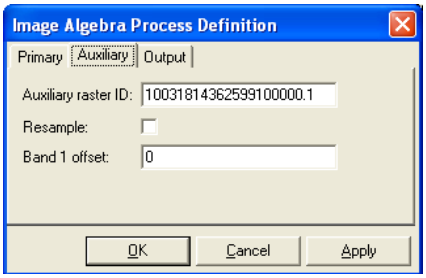

Note: When using Elevation .ISRef files as the source imagery, the RasterIDinRPDef is always 1. Only when multiple band sources are used will the value change.

- Click OK and close the Raster Properties dialog box.

It is necessary to build the image service to take into account the algebraic process being applied to the service.

- Select Image Service > Advanced > Build. The Build Options dialog box appears.
- On the Build Options dialog box, check the following check box:
  - Compute output pixel properties.
- This will now recalculate the output image for the service.

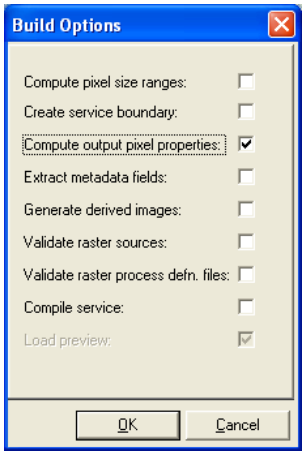

The result will appear as in the image below. In this image, the light green color indicates tall vegetation and the dark blue color indicates low or no vegetation. In the example, roads are clearly seen as dark blue lines winding through the tall surrounding vegetation.

J-9884

Figure 12
Height Estimation Image Service

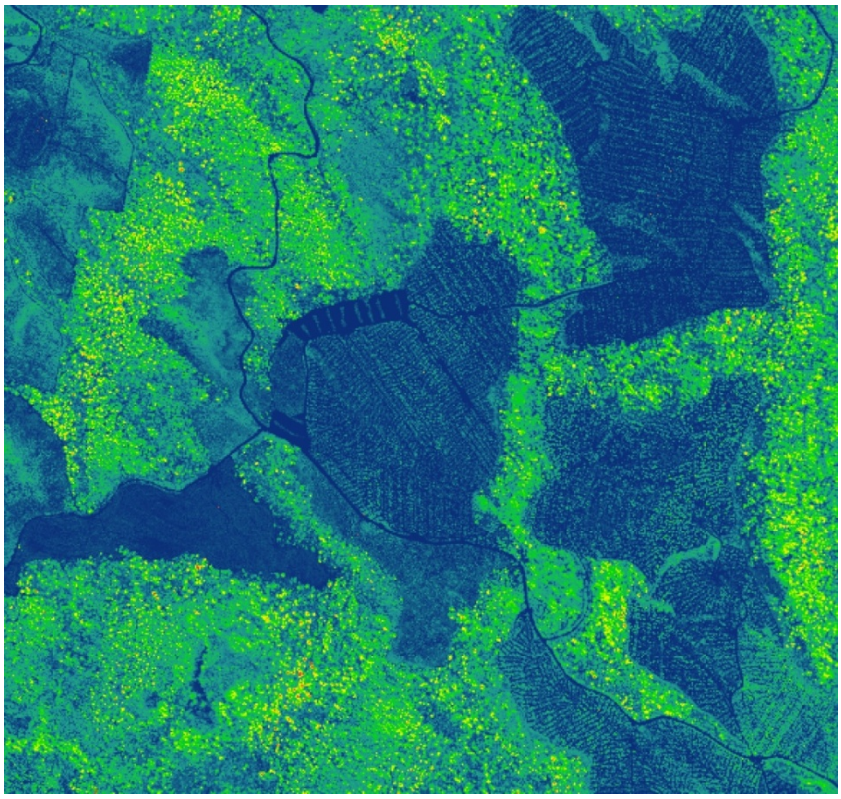

Light green areas represent tall vegetation. Dark blue areas represent low vegetation.

This height estimation service returns real height values. Thus, any pixel can be interrogated to return the tree height. The results from this can then be used by foresters to get an indication of growth and tree vigor.

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. Results indicate 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, and logging operations.
- 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 an excellent tool 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
www.esri.com

Offices worldwide
www.esri.com/locations

ESRI®
380 New York Street
Redlands, CA 92373-8100 USA

Printed in USA