CAMA Change Detection
with
Light Detection and Ranging

By KEITH W. CUNNINGHAM, Ph.D.

This session paper was presented at the
International Association of Assessing Officers (IAAO) Conference
in Atlanta, Georgia on September 10, 2007.

Appraisal data stored in computer-assisted mass appraisal (CAMA) systems have been successfully analyzed for change utilizing Light Detection and Ranging (LiDAR). Spatial information derived from LiDAR data is used to create a structure footprint. Injecting the scale-accurate sketch vectors from CAMA into the LiDAR-derived structure footprint allows a direct comparison of the sketch against the footprint. This comparison process to detect change in the CAMA data is highly automated and various degrees of change classification allow more efficient follow-on desktop reviews and field checks by assessors.

The IAAO Desktop Review

With the correct tools, the CAMA database can be analyzed for change. The recent IAAO Standard 3.3.5 approved October 7, 2006, states that as an "Alternative to periodic on-site inspections, jurisdictions may employ a set of digital image technology tools to replace routine cyclical field inspection with a computer-assisted office review."

The intent of the IAAO standard is to allow the use of high-resolution imagery to assess property grade, effective age, and condition. The imagery includes orthophotography, oblique imagery, and street-view images. Thus the IAAO is endorsing the use of these technologies to perform change detection from the office as part of a desktop review.

Change detection is not new to the assessor. The old-fashioned, but reliable approach is to go into the field with the property record, the property sketch, and a street-view photograph. On a cyclical schedule, the community is canvassed for updates and corrections.

It has been realized that the task of the assessor is simplified if the review can be performed from the office. But there will always be the need to visit the field to validate the office review, to measure new structures, and to reconcile ambiguous information.

Desktop Review Issues

There are a number of issues affecting the desktop review. The obvious one is the amount of interpretation and skill that is involved with analyzing imagery for change. There are also more subtle issues understanding the limitations of image interpretation, as well as a limited vocabulary for quantifying any change found.

The *stare-and-compare* process is how imagery is traditionally studied for change. But stare-and-compare is time-consuming, and error prone due to *occlusion* and *change blindness*. Change blindness is a condition of fatigue where the eyes and brain can longer find differences in images.

Occlusion describes how a closer object blocks or masks the view of an object further away. Occlusions in imagery are most commonly caused by trees and shadows obscuring the view of the structure you want to see and measure. Even street view images can have occlusions from parked cars and vans, hedges, trees, fences, and neighboring homes. Obliques have the advantage of multiple views and being able to see under the occluding roof eaves, but still have features obscured by trees and shadows.

Overcoming these image analysis problems is a set of quantitative tools that perform change detection utilizing data from airborne laser scanner called LiDAR.

LiDAR

LiDAR stands for Light Detection and Ranging. This airborne laser scanner is capable of collecting millions and millions of ground measurements. As a plane flies overhead, an infrared laser sensor continuously sweeps the land and maps the terrain.

LiDAR is now the preferred tool for preparing detailed surface elevation models. In the past five years, LiDAR technology has matured eliminating all of the skepticism that met the technology when it was first introduced. Today, there are industry-standard approaches[1] for collecting LiDAR, but we are only at the frontier of its capability.

LiDAR Basics

The LiDAR sensor utilizes a pulsed infrared laser. The laser pulse is bounced off of the ground. Range from the sensor to the ground is computed by measuring the reflected time-of-flight. The time of flight is calculated with a very-accurate clock. Pulse-time measurement are converted to ranges using the formula R=(T*C)/2. R is the range, T is the time of flight, and C is the speed of light.

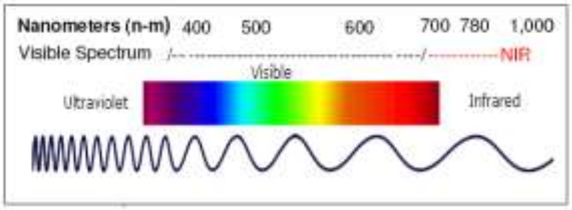

The graphic above shows the relative wavelength of the infrared beam compared to the visible spectrum. There is research into using other frequencies, such as ultraviolet to penetrate water to depths of 25 meters, being able to map the floors of water reservoirs and coral reefs. The military already uses the higher microwave frequencies to see through clouds and vegetation, with some rumors of being able to even see through building walls.[2]

Lidar Pulses

The pulses from the LiDAR are generated by a spinning mirror "sweeping" the infrared beam across the ground. The sweep is perpendicular to the aircraft's flight path.

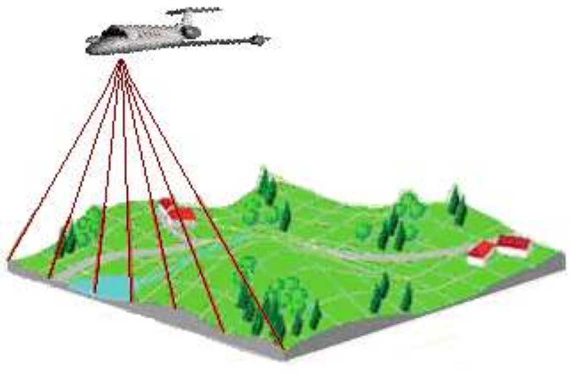

The LiDAR sensor emits several thousand pulses each second. For example, a 100 kHz sensor can produce 100,000 pulses per second. A typical county may have several million up to a billion or more lidar measurements.

As the laser pulse leaves the sensor on the aircraft, its width is about 0.1 centimeter (half an inch). The beam spreads as the pulse flies to the ground. The amount of spread is proportional to the distance between the sensor and the ground. Thus an aircraft flying 1,000 meters above the ground can have a pulse spread about 0.25 meters. Flying 2,000 meters above the ground would cause a doubling of the spread to half a meter.

Not all of the pulses strike the ground. Some are reflected from buildings, cars, trees, etc. Some pulses never return, being absorbed by the surface of a feature such as still, clear water.

---

[1] See the FEMA Guidelines and Specifications for Flood Hazard Mapping Partners Appendix A: Guidance for Aerial Surveying and Mapping

[2] With the use of a substantial amount of power.

Lidar Returns

Portions of the spreading laser pulse may strike a portion of a feature and bounce back to the sensor, which is called a *return*. Another portion of the same laser pulse may strike the ground, slightly later and this is called a *last return*.

These multiple returns represent variable and permeable features being scanned by the laser. For example, the first return from the pulse could be the crown of a tree, the second return may be a branch of the tree, a third return could be the fender of parked car under the tree, and finally the last return could be the curb.

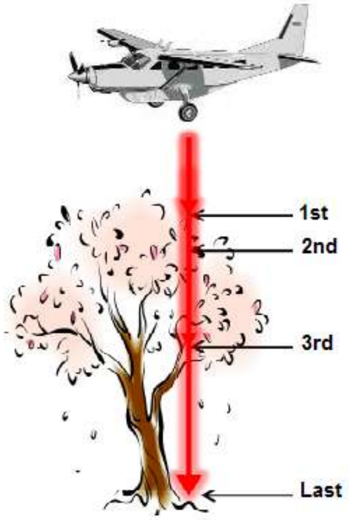

It is not unusual for most pulses to have more than one return. Many LiDAR sensors can track five returns per pulse. There have been considerations that future LiDAR sensors could detect an infinite number of returns but these sensors won't actually track returns, but the changed waveform of the reflected pulse.

LiDAR Coordinates

The LiDAR sensor measures both the distance and the angle of the return. The distance is computed from the laser pulse's time of flight. The angle is captured from the position of the spinning mirror that sweeps the laser across the ground.

Additional methods are used to refine every return's angle and distance measurements into an accurate XYZ coordinate of the ground strike.

Onboard the aircraft are two more sensors. They are the airborne global positioning system (GPS) and the inertial reference system (IRS). The GPS is used to initialize the IRS, which keeps track of the aircraft's flight path, including its crab, tilt & yaw. Sometimes the IRS is also called an inertial navigation system, an inertial positioning system, or an inertial measurement unit.

With the GPS and IRS, the aircraft flies a planned route of travel. This flight path defines the strips of LiDAR data to be collected. Flight paths are planned and followed to ensure that there no ground is missed in the aerial scanning.

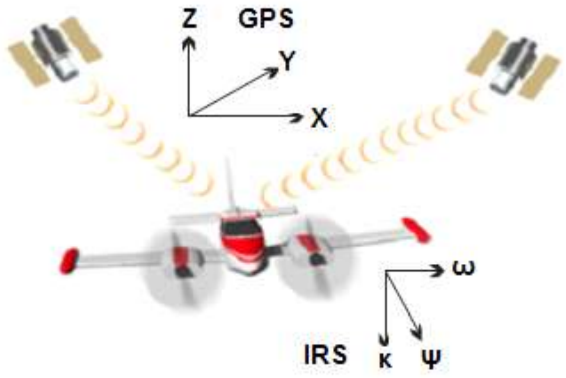

Each swath of LiDAR data is "stitched" together via a process called analytic triangulation (AT). The AT is a mathematical model that yields a complete blanket of laser returns from all of the swaths by tiling the swaths together.

This blanket of LiDAR data is then "tied" to the ground with traditional ground survey control. So traditional survey/photogrammetry techniques are used to ensure the accuracy of all of the LiDAR data collected on a project.

LiDAR Accuracy

In general each LiDAR return is about 1-foot accurate in the horizontal X-Y coordinates and about 6-inch accurate in the vertical Z coordinate. Note that these accuracies are expressed as standard deviations from a known value, as all map accuracies are.

Horizontal accuracy is generally worse because of the geometry involved with the sensor. This is because of the perspective error caused the angle of the pulses and the ground sloping away from pulse.

Important to a LiDAR project is the *error budget*. The error budget determines the accuracy of the LiDAR data being delivered. Part of this error budget includes the sample rate of the sensor and spacing of scan lines. This error budget also factors in the error from the scan pattern, sensor response times (latency), sensor (thermal) noise, and atmospheric error (fog, humidity, temperature, and refraction). Atmospheric error is why LiDAR missions are usually flown at night.

The greatest source of error is not the LiDAR sensor, but the GPS and INS on the aircraft. The exact position of the aircraft needs to be known in order to calculate the XYZ coordinate for every return. Any drift in the GPS calculations or with the INS will add error to the return's ground coordinate. Thus wind direction and speed can affect the aircraft's orientation.

Attention to the vertical datum is important to the final accuracy of the vertical coordinate, and to some degree the horizontal coordinates. Elevations are mapped in relation to either an imaginary curved surface of the earth called an ellipsoid or against an imaginary flat surface called a vertical datum.

Generally, NAD83 (North American Datum 1983) is sufficient for the horizontal coordinate component. But to define the vertical coordinate, a local orthometric (as opposed to an ellipsoidal) vertical reverences or datum should be used. Typical vertical reference systems are the NAVD88 (North American Vertical Datum 1988) and the ITRF (International Terrestrial Reference System). Note that each state and community may also have its own local vertical datum.

Another coordinate complication is that the equipment and techniques used for the survey control, GPS, INS, and final LiDAR point cloud could all have different coordinate systems and datums. Incorrect use of these systems and their translations could result in several meters of error.

Accurate survey ground control, a robust analytic triangulation model, and data densification with overlap and cross flights are important in the final accuracy of millions of LiDAR returns. But the final resulting cloud of points yields accuracies generally better than 15 cm. And they have precisions better than 2 cm.

LiDAR Density

LiDAR vendors will say they fly their missions "low and slow" to create denser point clouds, or "high and fast" to generate less dense, but more economical data.

The density of the LiDAR returns determines the resolution of the features that can be mapped. Smaller features require more dense data for their resolution.

Point density is dependent on the LiDAR sensor and aircraft altitude and speed. Another aspect of density is the amount of side lap in each flight line. A third way of collecting more dense data is to have cross flights perpendicular to the first flight paths.

Many users may want their data processed to thin the data or aggregate it into a smaller mapping data set. When data is thinned or aggregated, important details are lost. This loss of information may not seem important to LiDAR data users, such as a flood-control district or when the data is used in orthophoto production. But discarding the data may prevent the LiDAR from being used in future change detection processes.

Lidar Data Sets

There are several lidar data sets. They include buzz words of DTM, DEM, DSM, TIN, bare earth, intensity, etc. All of these data sets are typically available in digital file formats such as the binary LAS format and ASCII XYZ format.

There are also various levels of LiDAR products and amounts of value-add that can be performed. Of course, the more value-add, the longer and more expensive the delivery.

Detail	Use & Description
All Points	A raw point cloud of georeferenced returns, with no filtering.
Filtered	Morphologic filtering to create bare earth data and limited classifications.
Cleaned	Manual editing to clean data for use as a DTM and later processing.
Feature Extraction	Automated & manual data classification with specific data extracts, such as mass points & break lines.
Fused	Refined data is "joined" with other GIS data, imagery, and hyper-spectral imagery to generate new data sets.
DEM	Bare Earth LiDAR is aggregated to create a grid elevation model.

Raw LiDAR Data

After LiDAR data is georeferenced, it is available in both a binary (LAS) and ASCII (XYZ) files. The binary file is much more compact than the ASCII, but you need a LiDAR program to read the data. The LAS file also contains more information on each of the data points such as the intensity of the return, angle of the pulse, and number of returns for the pulse.

The ASCII file is more simply opened with a word processor or spreadsheet program. Below is a sample of the raw file in ASCII format, with comma delimiters. The first field is the easting (X-coordinate), the northing (Y-coordinate), and then the elevation (Z-coordinate).

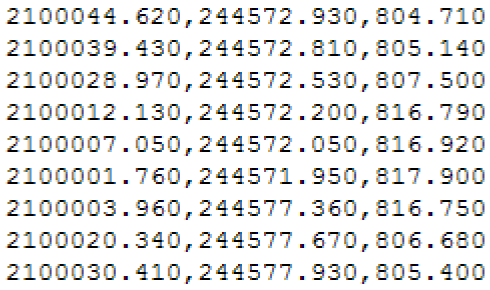

```
2100044.620,244572.930,804.710
2100039.430,244572.810,805.140
2100028.970,244572.530,807.500
2100012.130,244572.200,816.790
2100007.050,244572.050,816.920
2100001.760,244571.950,817.900
2100003.960,244577.360,816.750
2100020.340,244577.670,806.680
2100030.410,244577.930,805.400
```

These ASCII data sets can contain over a billion coordinates. Files of this size are hundreds of megabytes to many of gigabytes in size.

This same data in the example above can be geographically represented as a map showing each of the X & Y coordinates. When this geographic data is rendered so you can also visualize the Z coordinate, we have what is called a *point cloud.*

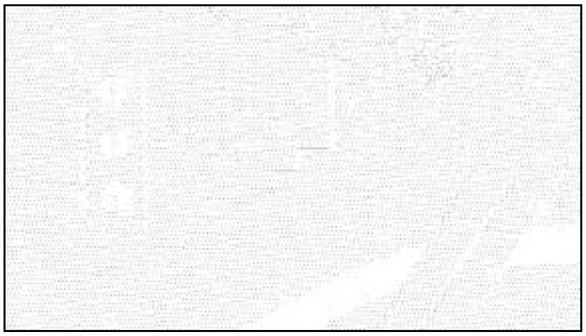

Lidar Processing

One significant issue with LiDAR is the amount of effort that goes into the processing of the data to create derivative products. This processing is measured in both computer and human processing time.

One of the initial post-processing steps is the removal of excess data. Sometimes called noise, this excess data is simply unwanted data that is not needed in the users' end application. Much research has gone into removing redundant information or "thinning" data. The motive for data thinning is to improve issues with data storage and to boost computational efficiencies of software that use the LiDAR data.

This post-processing can be completed with specialized software filters, but generally a person supervises the removal of the unwanted data. Eliminating this unwanted data can reduce file sizes by half, from many gigabytes to just a few.

A type of filter used to classify and remove data is the morphologic filter. This type of filter considers the patterns and texture created by a cluster of points in the cloud. For example, if the points indicate a smooth, sloping surface, then they are assumed to be a roof. If the points are random, then the feature may be trees or shrubs. This form of filtering is of course an intense area of research by academics and the military.

When the vegetation, cars, and other transient objects are removed leaving structures and the ground, the resulting file is called a digital surface model (DSM). The cleaned DSM is useful for

3D visualization and the user can visualize features such as buildings and even sloping roofs.

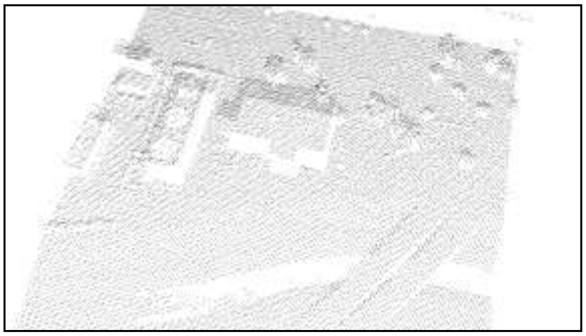

When the DSM is rendered as triangular polygons with each polygon creating an abstract surface, we have created a triangulated irregular network (TIN). The TIN is useful in visualization by creating relief maps and for view shed modeling. TIN models are the most common format for processing LiDAR data within a geographic information system (GIS).

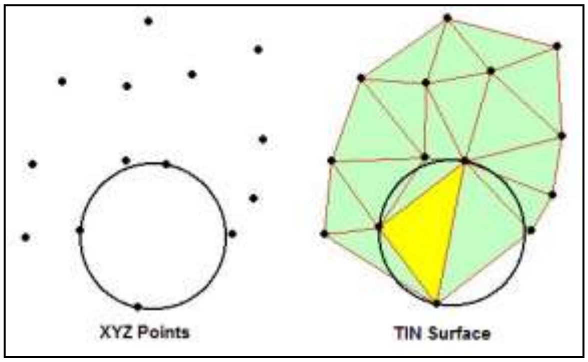

The Bare Earth model is a DSM with all vegetation and man-made structures removed. Examples of man-made structure are vehicles, bridges, dams, homes, buildings, rail cars, garages, etc. The bare earth model is of use to surveyors, engineers, and public works officials planning infrastructure development.

When structures are removed, and photogrammetry control features called *mass points* and *break lines* are added, the DSM is called a digital terrain model (DTM). The DTM is used in the generation of very accurate ortho imagery. Sometimes this is called *true* orthophotography. Some oblique products also *register* each oblique image to the DTM, ensuring very accurate vertical, horizontal, area, and façade measurements.

The DTM is also usually viewed as a TIN or in a simpler data extract known as the digital elevation model (DEM). The DEM is a method of averaging data and reducing the data content to create a very simple visualization product. The DEM is actually a grid into which several or even hundreds of LiDAR elevations are averaged. The averaged elevations are coded into each grid cell. Sometimes the DEM is created as a georeferenced raster file to improve file portability and for use in other applications.

An example of the file size savings was a DEM project in Douglas County, Kansas. The resulting raster file sampled 590 sets of LiDAR data of over 21 gigabytes into a single image of only 200 megabytes. Lidar Logic was able to embed the elevation data directly into the raster file using a proprietary process. This compressed image also preserved the elevation data in the visual rendering, so the user can intuitively see changes in the elevation. With a custom tool created in ArcView, users also can see the elevation value as they move their mouse cursor. Of course, this single image could be further compressed to only 10 megabytes, yielding a 2,000:1 compression, though there is the typical image quality loss incurred with image compression.

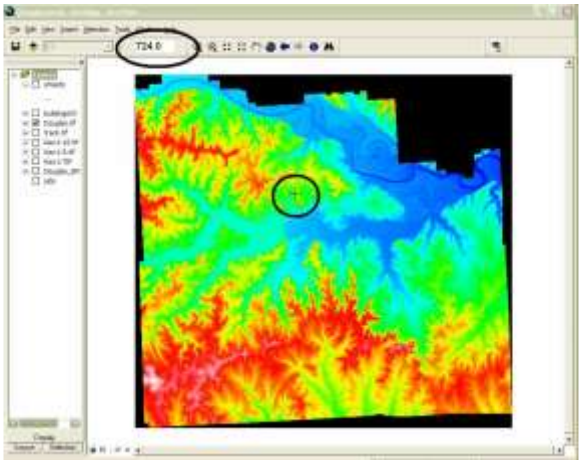

Note that the DEM is an isotropic model, implying the size of every cell has the same area and shape, whereas the TIN is an anisotropic model and each triangle is a different size and different shape. This has implications with future data processing if detailed features are to be extracted.

DEMs and TINs are often used in hydrologic and hydraulic models to determine flood risks. But these models are limited due to the density of the LiDAR ground strikes and should be augmented with break lines to yield accurate and reliable models.

LiDAR Ambiguity

Understanding the traditional approaches to LiDAR data processing helps explain how the data can be processed for CAMA change detection. This approach also considers how to manage ambiguity in an image, like the examples below.

Users can easily visualize the features in the LiDAR data. Looking at trees and structures in LiDAR point clouds makes the user aware of how these features are depicted with LiDAR. Surfaces of roofs can also be easily visualized. Part of the visualization process is how the brain comprehends the features and automatically fills in missing information.

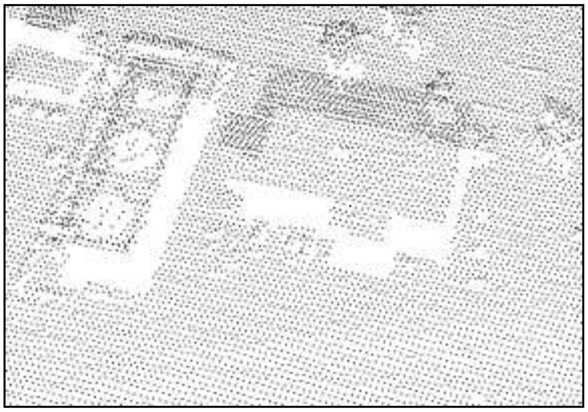

But a closer inspection of the point-could causes the brain to see "gaps" in the data, causing uncertainty in what was previously a clear image of features. This process is similar to how the brain comprehends optical illusions. At first glance, and image may appear one way, but with a closer look, the image changes.

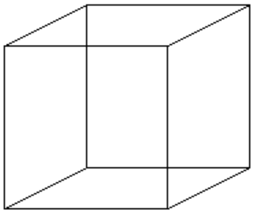

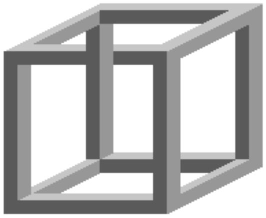

A practical example is the edge of this roof in the cloud of points. Note how with first glance you can clearly see the edge of the roof. But with closer study, the roof now appears ragged and poorly defined. This is especially true as you zoom in for more detailed study.

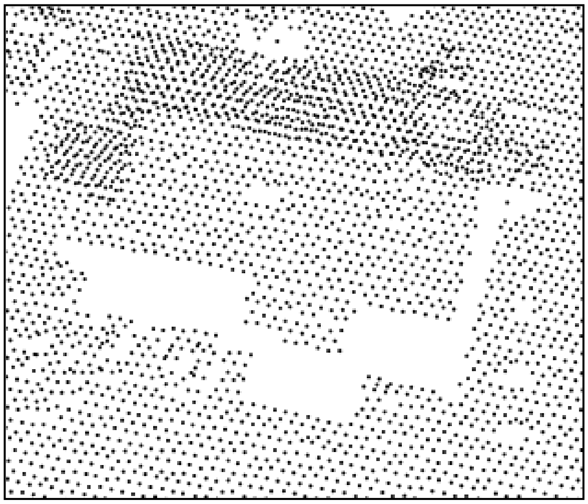

To overcome these fundamental limitations of the LiDAR data, several automated approaches were developed to make use of the noise and incomplete data. Understanding the mathematics of noise, the mind, perception, and data incompleteness helps us understand how to better extract features using rules of data sampling and resolution. The result the LiDAR mask.

LiDAR Mask

The goal of the mask is to identify features (structures) as part of the foreground while obscuring features (noise) in the background. This foreground/background processing of the data uses several concepts related to human cognition and how unwanted information can be reduced into the background, while preserving features of interest.

The lidar mask is a tool used to better define the edges of foreground features and to block the view of other unnecessary background information. The mask is used primarily to assist in the determination of feature breaks and later to locate their edges.

The first process of LiDAR masking is filtering out trees and other vegetation. We consider the trees and vegetation to be noise and will opt to remove it from our LiDAR data. However, we also utilize this noise in part to assist in determining features or parts of features that may be obscured visually by the noise. The resulting LiDAR data set will be the surface of the ground. This proprietary process is specifically designed to eliminate the change detection problem with vegetative and shadow occlusion experienced with aerial photography. Notice how in the example below you can see the trees in the ortho, see the trees in the raw LiDAR file, and then see the final LiDAR image with the trees removed.

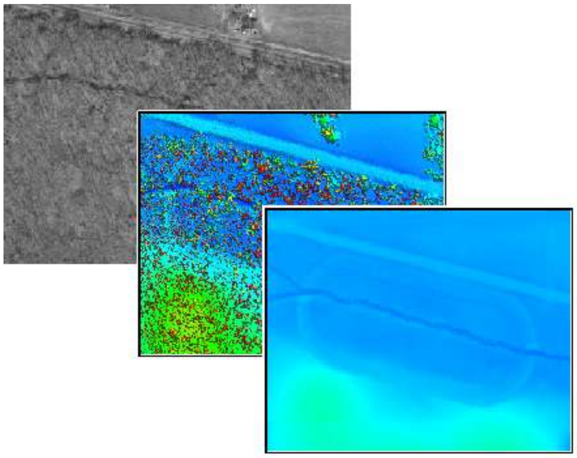

The second foreground/background processing component is the *dialing-in* of features that can be resolved. The size of features that can be resolved is dependent on the density of the LiDAR returns. In general, features of about 4 by 4 meters can be resolved from LiDAR data collected according to FEMA specifications. Smaller features can be resolved better with denser LiDAR data.

Part of the dialing-in process is also used to remove other ground clutter. For example, parked cars are about the same dimension as backyard storage shed. It is easy to remove both the parked cars and the small outbuilding, but the programming logic required is much more rigorous to only drop parked cars compared to sheds.

Sketch Change Detection

The LiDAR masking filters are very important in preparing the final mask for the assessor. The information in the lidar mask includes homes, buildings, and other man-made features relevant to an assessor such as gazebos and above ground swimming pools. This filtering or masking of the background helps eliminate the noise that could interfere with later automated change detection with the CAMA sketch as well as the final manual desktop review.

Change Detection Inputs

There are three important inputs in the CAMA sketch change detection process. They are the LiDAR Mask, the sketch vectors from CAMA, and the digital parcel fabric.

The sketch in CAMA must be digital and retrievable in some non-proprietary format. If sketches are drawings on a property card, then the property card must first be digitized. The example below shows a sketch from an AS/400[3] in a DB2[4] database. The sketch vector data was exported into an ASCII file, shown below.

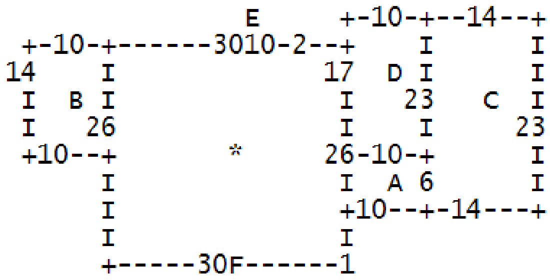

```
               E     +-10-+--14--+
 +-10-+------3010-2--+    I      I
14    I             17  D I      I
 I  B I              I   23   C  I
 I   26              I    I     23
 +10--+       *     26-10-+      I
      I              I  A 6      I
      I              +10--+-14---+
      I              I
      +-----30F------1
```

The sketches are all then converted into a set of digital scale-accurate polygons. In the case of this

---

[3] AS/400 is a trademark of IBM

[4] DB2 is a trademark of IBM

demonstration area, the polygons were created as an ESRI Shapefile[5].

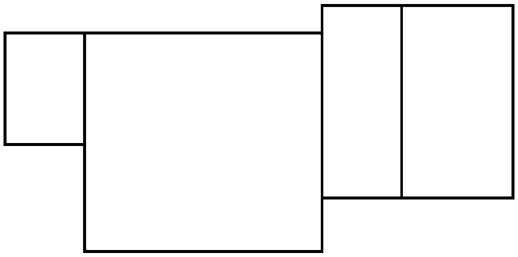

The parcel fabric should be a single shapefile. Its coordinate system should be the same as the LiDAR data. If not, it is easier to re-project the parcels into the same coordinate system and vertical datum as the LiDAR. Finally the parcel data needs to have unique parcel identifiers linking the sketch from CAMA.

Sketch Georeferencing & Analysis

The mask is an efficient tool in the automated georeferencing of the CAMA sketch. The mask represents the structure as a void in the background data. Software automatically determines the best fit of the sketch based on the geometry of the sketch into the void in the mask.

Compared to the LiDAR products discussed earlier, we are fusing the CAMA sketch to the appropriate void in the LiDAR mask. This fusing process creates new value added information. The first value-add is the georeferencing of the sketch from CAMA. The second value-add is the analysis of fit.

The change detection process is a quantitative measure of a qualitative fit. This means the wellness of fit is graded by the analysis software as each sketch from CAMA is fit and compared to the corresponding parcel void of the LiDAR mask.

The graphic below shows the LiDAR mask with voids as transparent, thus representing the missing elevation data. The parcel fabric is the bold line and the sketch before georeferencing is the fine outline.

Note that the mask color represents the elevation data. Thus the elevation for each sketch can be automatically determined from the mask.

The actual software performing the change detection completes its task in an unglamorous, non-visual, batch mode. The results are simply georeferenced sketches from CAMA and an attribute, or score, in the sketch database describing how well the sketch fits.

The next graphic demonstrates how the sketch is moved and rotated to fit the void in the mask. Note how the sketch typically fits into the void. But where the sketch doesn't fit, it is graded and flagged for desktop review.

Remember that the void in the mask represents the roof top of the structure. Therefore there is a certain percentage of overlap of the void into which the sketch will fit. The sketch georeferencing software anticipates the overlap. So a georeferenced sketch should be slightly

---

[5] Shapefile is a trademarks of ESRI

smaller than the mask void and the geometry of the sketch and void in the mask should be similar.

We can assume there is change if there is either too much or too little overlap between the sketch and the void. Also if the geometry of the sketch and void are different, we can assume there is change. Four logical outcomes are possible with each comparison. They are:

1) True Positive - Change detected correctly.
2) False Positive - Change detected incorrectly.
3) True Negative - No Change detected correctly.
4) False Negative - No Change detected incorrectly.

The software was developed to primarily achieve condition 3, when there is a "true negative," while minimizing the occurrence of condition 4, a "false negative." Condition 3 is also where most of the human time would be spent if this process were done manually.

- Condition 1 is easily determined, provided the change is the size of a deck or larger.
- Condition 2, which is an incorrect change identified is not necessarily a problem because this condition is flagged for Desktop Review and would be corrected during the review.
- Condition 3 is the most common result of the analysis, which is to be expected.
- Condition 4 is the most problematic because we don't know when the condition of "No Change" is incorrectly identified.

With respect to false negatives, the only solution for this condition is confidence in the performance of the software. We have confirmed that Condition 4 does not appear to be a problem on our first projects because of a 100% desktop review to evaluate the accuracy of the software. But the world is infinitely variable and software cannot handle all of exceptions the real world has.

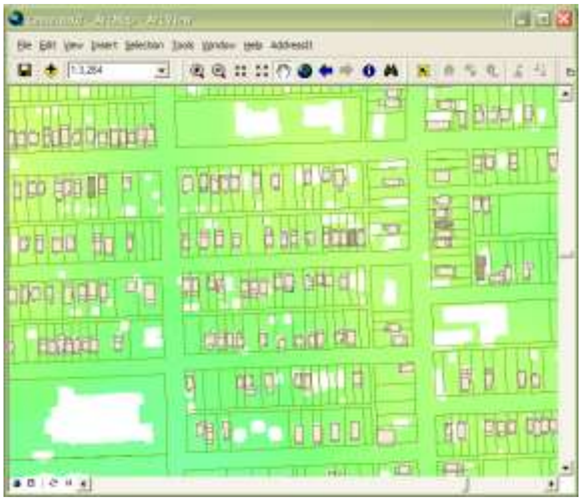

In the image above, note the voids in the mask where there is not a sketch. Some of these voids are outbuildings on the property that are not sketched and not assessed.

But there are notable exceptions when there are not sketches to be georeferenced, and every assessor will understand and appreciate these. Examples include missing sketches for mask voids which are for tax exempt organizations, such as government, schools, and churches. Another example of exceptions are multi-storey building sketches as we can only georeference and test CAMA fit on the ground floor.

Desktop Review

Analysis of the CAMA sketch is key to the desktop review. With LiDAR, the sketch is automatically georeferenced, analyzed for change, and *scored* with a quantitative value.

The scoring includes values for missing outbuildings, missing decks, minor change and major change. There are also scores for missing sketches, which may voids in the mask for tax exempt parcels. Another unusual score is for a void and sketch with the same dimensions, but with a different geometry, like the example below.

With specific sketches now flagged and scored, the assessor can more efficiently perform the desktop review. Or if necessary, go into the field to take measurements and collect updated street view imaging.

At the point of the desktop review, the LiDAR mask is no longer needed. But it is available as a GIS layer. The LiDAR mask is of value during the desktop review when it is rendered transparent so you can see through it to the orthophoto underneath. The approach helps minimize the image background information containing yards, shrubs, trees and other distracting information. This allows the assessor to focus their attention on the sketch being inspected.

The next graphic is the same LiDAR mask in the previous example, with the ortho toggled on.

Note there has been a fundamental change in the desktop review process. Because we are comparing sketch to LiDAR, the CAMA database is being analyzed for change. We are not comparing ortho to ortho, or oblique to oblique.

Determining the grade and condition of the change is another, separate task from the LiDAR. If obliques are available, then comparing the sketch with the obliques can provide information on the nature of new construction, grade, and perhaps even effective age and condition. Mega-pixel street-view imagery also is a very useful tool in the process.

Unfortunately, in some cases, a field visit to the parcel is necessary to determine the change.

Summary

LiDAR is now a proven tool for automating the fusion of many data sets, including the sketch from CAMA into real-world space. The LiDAR mask also enables a significant degree of automated change detection of the sketch against current LiDAR data.

Individuals performing the review appraisal from the office can now conduct their job more efficiently, achieve higher productivities, and be more confident that the changes they are identifying are correct.

The LiDAR mask approach to CAMA change detection is a proactive one. All communities have been vexed with having to continually bring their tax roles up to date in a reactive process. With LiDAR, the assessor can now ratchet up their CAMA valuations across the board, not in a piecemeal, cyclical approach.

The importance of lidar in the field of change detection will grow, especially in updating CAMA data. In communities with a reactive approach to database maintenance, the lidar mask is an effective tool for georeferencing the sketch in CAMA and for masking the ortho and oblique to more easily and accurately determine changes.

Other uses of LiDAR are the determination of building elevations and flood plain mapping. This will lead to 3D electronic sketches, including ceiling heights. Also, the CAMA data could become the nexus for better determining insurance loss claims during disaster recovery events.

On a final note, once a LiDAR baseline is created, change detection utilizing future LiDAR data will actually make the change detection process simpler. In the future, change detection for CAMA sketch measurements will be possible in 3D, as well as from LiDAR on LiDAR data processing.

About the Author

At Lidar Logic, Dr. Cunningham leads the exploration of new concepts and technologies for automated feature extraction utilizing LiDAR. Current work includes assisting several counties with CAMA change detection, impervious surface modeling, and systems process re-engineering.

Over the past twenty years, Keith has provided GIS, GPS & cadastral consulting on more than 150 county-level projects. This includes serving as a technical advisor to four states, several national organizations, and numerous projects overseas. He also consults with large spatial corporations where his knowledge of the human factors and process engineering is used to design better software and systems. Keith maintains his industry credentials with articles, classes, and editorial work for photogrammetry, GPS, and economic development. Keith also lectured the United Nations on GIS/GPS for infrastructure development.

Dr. Cunnigham received his PhD from Kansas University in 1997, with empahsis in artificial neural networks for automated feature extraction. While at the University of Kansas, he helped teach graduate seminars in advanced GIS, Urban GIS, and GPS for GIS. He also holds a BS in Geology and an MA in Geography from the University of Missouri.

Keith Cunningham can be reached at kwc@mobile-map.com