An Overview of LIDAR for Urban Applications

A group of planners, engineers and concerned citizens look down a city street, discussing how a new development on vacant land will integrate with the existing fabric of people, structures, and infrastructure. Their words carry weight but communication is limited by different understandings of the problems faced during development, and so they turn to a laptop placed on a nearby truck hood.

Geographic Information System software on the laptop shows plan maps for the area drawn by planners and engineers, rough three-dimensional models of the street as it is now and as it may be in the future. A visualization tool shows photo-quality models of the same street with fly-overs. A discussion evolves, and changes are noted. The design lead passes around a card with a Web address that links to Google Earth, where the citizens can look at a lower-resolution plan for the area complete with real and simulated photographs.

Where does the data for such models come from? Does someone go out and measure each and every feature on the street manually? Do we have such data for our infrastructure at present?

LIDAR is one key technology that makes the construction of city-wide data sets of this type feasible. This document provides background and context for understanding how LIDAR supports present and future urban modeling and planning efforts.

Overview

LIDAR is a technology for measuring positions of things rapidly - in the urban case, for measuring where everything on a street is. LIDAR stands for LIght Detection And Ranging; LIDAR uses a laser pulse to measure how far an object is from the tool. If we know where to tool is and how it is oriented, we can then work out where everything else is.

LIDAR is useful not only because it can provide accurate positions over large areas but also because it is fast: LIDAR can collect tens to hundreds of thousands of positions in a second. Collecting urban data at this level of detail manually would take years - the buildings would likely fall down before the task was done! LIDAR is thus a viable solution to the massive task of mapping our urban infrastructure to support maintenance, modeling, and visioning exercises.

LIDAR model of Miller Hall, Queen's University, collected with an ILRIS time-of-flight LIDAR in October, 2006. It looks like a grainy black and white photograph, but is actually tens of thousands of three dimensional locations, shown here using a software viewer.

Rob Harrap and Matt Lato
Queen's University
harrap@geol.queensu.ca

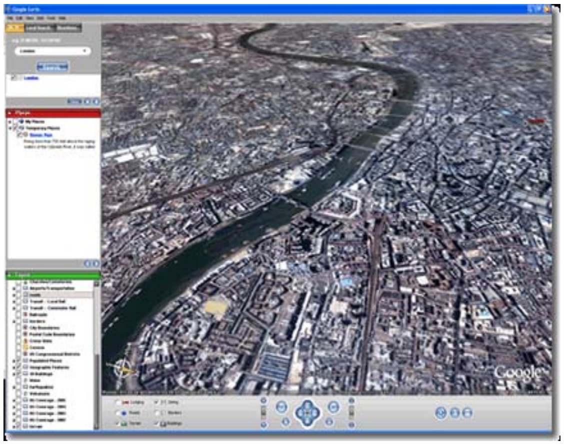

Google Earth - Three-dimensional GIS visualization, free on the Internet. Where does the 3d data come from?

LIDAR data can be collected from an airborne or terrestrial vehicle, or from a fixed position, usually on a tripod. Airborne LIDAR has been used for some time as a source of models of the landform for engineering, disaster management, and other visualization tasks; fixed LIDAR has been used in infrastructure mapping of detailed sites such as chemical plants for a number of years.

In recent years the use of LIDAR is growing rapidly, both in terms of the number of application domains and in the prevalence of the method in real-world practice. This document provides background on LIDAR and points out where LIDAR appears to be heading as a tool to support urban planning and engineering.

How It Works

LIDAR relies on two sets of measurements to generate a cloud of point locations for features around the known location of the scanner. First, the pointing direction of the laser must be known for each measurement. Depending on the physical mechanism of the scanner the points may be evenly or unevenly distributed on the target, and because systems normally operate on an angular offset between successive measurements, targets closer to the device will have a higher point density than those farther away.

The second piece of information needed is the distance. There are two approaches to measuring this: time-of-flight and phase-based. Time of flight LIDAR sends a pulse, waits, and measures the time of arrival of the return pulse(s). Given the travel time (the speed of light) and very precise time measurement, a distance can be derived. Time-of-flight systems are limited only by the need for a return signal, and so higher powered systems can see out to kilometers or more as needed. In practice this is limited by the requirement that the system not harm persons in the target area and by the gradual spread of the beam with distance.

LIDAR scan of buildings on the campus of the Royal Military College, Kingston, acquired by the TITAN system in May, 2007. In LIDAR, the point of view of a scene such as this is not the point from which the scanner collected the data - in this case, along the roadways.

Phase-based LIDAR sends a continuous beam with known phases. When these interact with a target, the phases are shifted, and the returned shifted signal is processed to derive distance. This is both faster and more accurate than time-of-flight methods but has the limitation that there is a finite distance beyond which offsets cannot be converted into distances. For example, the Leica HDS6000 unit we operate cannot ‘see’ past 70m. It can, however, collect 500,000 points per second!

In both cases the combination of direction and distance are converted into a position offset between the LIDAR device and the target or targets. Knowing where the LIDAR device is is thus key to building data sets that can be integrated with other geographic information.

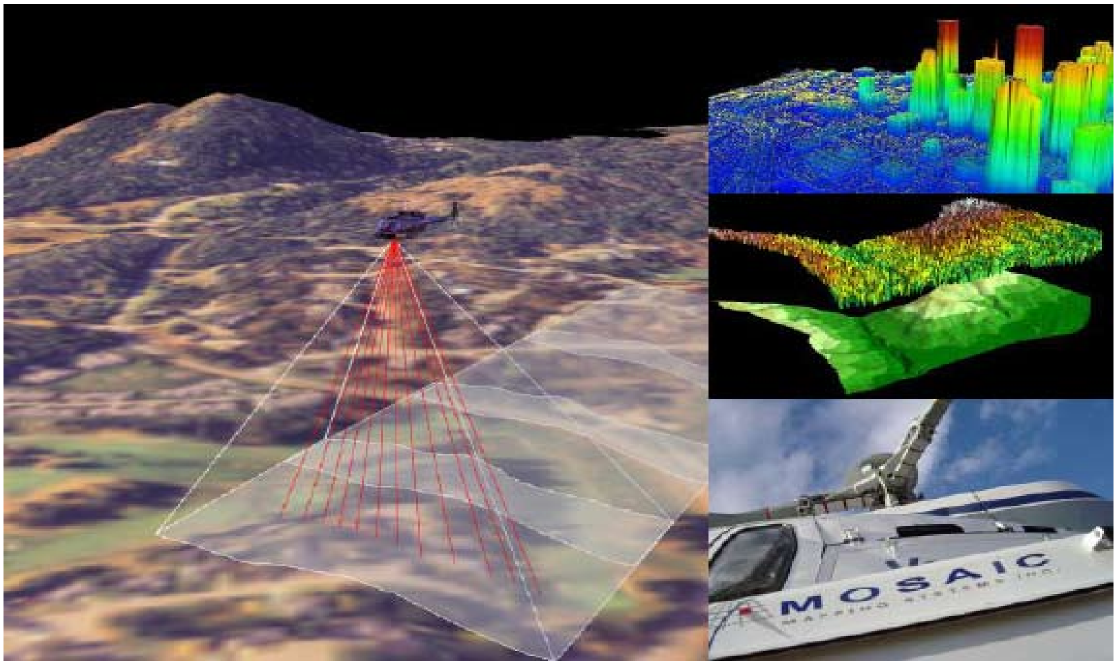

LIDAR airborne data collection and products for surface mapping. Mosaic Mapping, now Terrapoint Inc., is a commercial vendor with a range of LIDAR and GPS services. Figure courtesy Terrapoint Inc.

Navigation and Location

The accuracy of a LIDAR is limited by the accuracy to which we know its location. For a moving system we must track the position of the sensor for each pulse; this location is then used in combination with the LIDAR angle and distance information to place a position in three dimensions. For a static system the problem is somewhat simpler as all of the positions for one scan will be relative to the position of the sensor, which is not moving. In both cases, though, the overall accuracy of a scan relative to other geographic data (for example, existing GIS data for an area) is limited strongly by how well we can position the sensor.

Positioning is based on the Global Positioning System (GPS), augmented in the case of mobile systems by Inertial Measurement Units (IMU). Together these constrain the location of the sensor at every instant, and if we accurately correlate times of LIDAR acquisition to positions then the overall scan will be accurately located in space.

GPS is based on the measurement of position by the time of flight of radio waves. A number of satellites (at least 4, but usually more) contribute position offsets that, together, constrain where a GPS receiver is at any point in time. Practically this is limited by the availability of satellites, by clear views of the sky, and by the fundamental physics of the satellites and receivers.

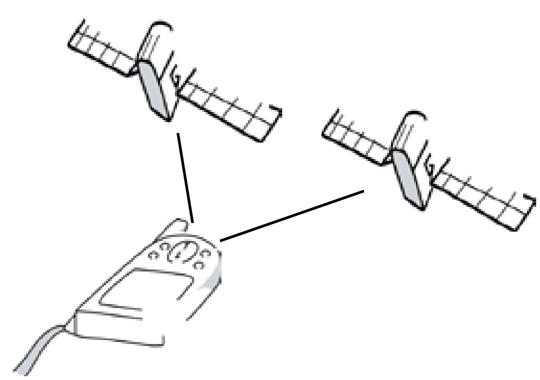

GPS is a core technology for LIDAR - if you don't know where the sensor is at time of data acquistion, everything else falls apart! GPS used in LIDAR is much more precise than systems used by consumers, but the principleas are the same.

GPS locations can be made much more accurate by a number of steps, including:

1) Ensuring satellite geometry at time of survey is adequate

2) Integration of fixed position data from base stations into the

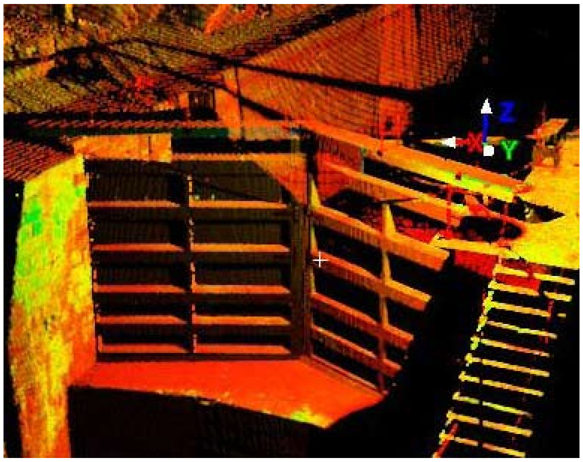

Scan data for the Kingston Mills locks from a Leica HDS6000 ultra-high resolution scanner. The scanner was at the point shown by the x,y,z axes. Scanning the entire region from this point took less than 5 minutes and generated a point cloud of more than 100 million locations. The color indicates the strength of the signal return from the target.

position determination process, and

3) Avoiding obstructed regions such as 'urban canyons' or tunnels.

The use of base stations, whose apparent motion can be used as a correction signal to offset clock and atmospheric physics limits on accuracy, is called 'differential GPS.' For mobile LIDAR surveys at least one and usually more base stations are placed at fixed locations in the survey area and provide correction signals for the position of the mobile unit; this typically takes GPS from meters to tens of meters accuracy down to sub-meter accuracy. With advanced GPS this can result in centimeter level accuracy.

Avoiding obstructed areas is an issue since these are of interest in urban scanning. Inertial Measurement Units address this issue to a degree.

Inertial Measurement Units (IMUs) rely on physical sensors to work out movement: very accurate determination of attitude and acceleration can be combined to provide a path through space which requires no external signal. IMU on airborne or terrestrial LIDAR allows the interpolation of positions between accurate GPS position fixes. In practice the accuracy of IMU-based position determination drops off quite quickly, and so it is the combination of GPS and IMU together that allow accurate sensor-position determination for LIDAR surveys.

Note that since an IMU is capable of accurately determining position during a 'gap' in GPS signal acquisition, a mobile terrestrial scanner equipped with an IMU can drive through urban canyons and tunnels for a significant distance with gradual loss of position accuracy. Once the position from GPS is re-acquired the position during the gap can be forward and back-corrected. In practice gaps of minutes in duration are permissible though not desirable.

Airborne Acquisition

In the case of airborne acquisition, the LIDAR is placed in an unobstructed location in a fixed wing aircraft or helicopter. This generally involves fixing a sensor pod to the bottom of the aircraft and putting in-flight controls inside where the operator will sit.

The aircraft is then flown to the target area and flies a series of paths across the area in a grid pattern. The density of the measurements is determined by the LIDAR data collection rate, the elevation of the aircraft, and speed. For dense surveys several overflights may be required.

Given the requirement of eye safety for people on the ground, higher power systems must be operated from greater height. As a result, systems are often divided into 'high range' and 'low range' systems for different applications.

Data from airborne scans is corrected for position and then processed for various output products. One typical product is the generation of a bare-earth elevation model; in this case the 'farthest' return for each orientation is used and features such as buildings are removed by an operator using specialized software. Such bare-earth models are useful for flood mapping, construction planning, and other visualization needs.

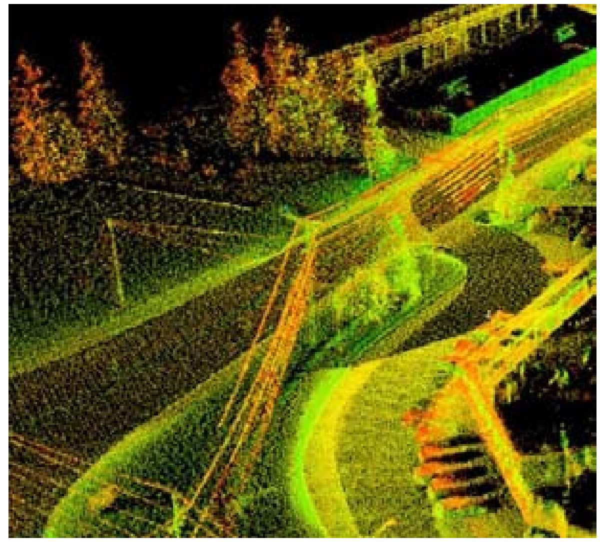

This urban data from the Terrapoint TITAN system shows how drive-by LIDAR can map infrastructure, buildings, and vegetation rapidly. The color here indicates the strength of signal return from the target.

Terrestrial Acquisition

In terrestrial acquisition, the LIDAR unit is mounted either on a tripod or on a vehicle. Since the range to target is dramatically less than with airborne systems, the point density and accuracy will be much higher, as much as a point per square millimeter. This is pragmatically limited by the diameter of the beam on incidence – if the intersection is 2mm across, collecting a point every millimeter is of limited use.

With terrestrial mobile systems, a coupled GPS-IMU solution is needed to track where the device is during data acquisition. This is in principle similar to airborne systems but has the added complication that it is much more likely that GPS visibility will be reduced so that reliance on IMU control will be larger. Ground control GPS locations are also used with mobile systems both for calibration and quality assessment.

With tripod mounted systems, GPS control can be either from a GPS on the unit, plus one or more control points on the ground to provide geometry, or from multiple GPS targets on the ground. Multiple scans can be combined as long as three or more common and distinct points exist between the scans. It is common practice to place targets in the scene that are highly reflective and have precise scannable markers in order to guarantee a minimum of scan combination points. These points, ideally, would be collected using GPS as well.

The Terrapoint TITAN scanning solution mounted on a truck. The box mounted on the scissor-lift contains multiple LIDAR and IMU units as well as a GPS receiver. The operator sits in the passenger seat in the truck.

Photos collected from axial cameras on a mobile terrestrial unit or collected with a digital camera from the position of the LIDAR device can be combined to create image domes that can be integrated with LIDAR data to provide colorization for cloud points.

The most mature example of a terrestrial mobile scanning solution is the TITAN system from Terrapoint, Inc. TITAN consists of multiple LIDAR scanners, IMU's and cameras in a boom (see photo) mounted on the back of a truck. As the illustrations herein show, TITAN can collect tens of thousands of locations accurate to approximately 2cm during normal driving, and so can collect data without interfering with the flow of traffic. TITAN post processing allows collection for several minutes without GPS and TI-

Kingston Whig-Standard building and area point-cloud, scanned November 2007 with a Leica HDS-6000 phase-based, tripod-mounted LIDAR

TAN has been used in tunnels and urban canyons to good effect. As part of our research we have collected large parts of Kingston, Ontario's downtown using TITAN and are working to enhance data processing methods for such data.

Processing : Basic Models

Once collected, LIDAR data is post-processed for geometric correction as needed. This can be simple, as in the case of a fixed scan from a tripod-mounted scanner, or highly complex, as in the case of IMU position correction for a mobile scanner. The result of this stage is a geometrically accurate collection of points, or a 'point cloud,' typically coded with intensity of return and in some cases with the normal vector (roughly, the vector back to the scanner).

Multiple scans for an area can be combined, resulting in even larger data sets to increase either the scan area or the scan density. In the case of scanning using targets on the ground, these targets are the control points for combination; in the case of data collected with GPS and IMU, the data are inherently spatial and so can be merged, although with an eye to evident data quality and obvious errors.

Processing to spatial products then proceeds. Since LIDAR will produce complex signals, often including multiple returns for one pulse, processing can produce a 'first' or 'last' return product. In the case of airborne surveys, the last return from many of the pulses will be the ground or built infrastructure such as buildings. The first return may be vegetation. A skilled operator can produce a 'bare-earth' model from last returns, using editing tools to remove buildings if desired.

Currently there are two daunting problems with LIDAR point cloud processing. First, there is no one software tool that does all of the necessary steps from input to model creation, and so files must be transferred between tools and formats. Second, the data volumes are so large – often in the tens of gigabyte range – that even the fastest workstations are hard-pressed to process the data in reasonable times. Typically a large data collection project will involve division of an area into zones, or 'tiles,' simply to provide smaller working targets that are of manageable size.

Although this resembles a LIDAR point cloud, this is a surface representation of the edge of the Kingston Whig Standard building constructed using surfacing tools in Leica's CYCLONE LIDAR processing package.

Processing: To Features

Engineers, urban planners, and other spatial data users do not want point clouds in most cases. They have neither the tools nor the interest in processing data in this format – they want spatial features such as 'roads' and 'trees' and 'buildings.' In LIDAR, each of these is found as part of an overall point cloud, and there

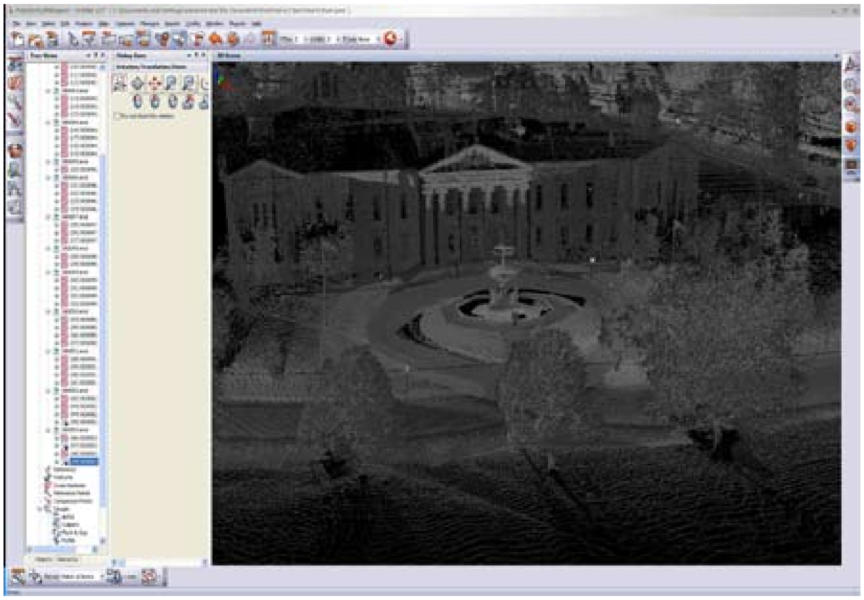

Processing of LIDAR point clouds requires specialized software and very fast computers. This point cloud of the region around the Kingston Courthouse, shown here in PolyWorks from Innovmetrics, consists of hundreds of thousands of points. To use it in a traditional visualization tool it would have to be reduced to geometric features. To use it in Google Earth it would need to be reduced to a few dozen shapes at most.

will usually be a correspondence between the material type and the intensity of signal return. Ideally processing should replace many points with few objects, and these objects should be from a standardized library (such as standard park benches, or standard hydrants) rather than being unique.

To date the process of converting point clouds to attributed spatial features is operator driven. There is abundant computer science research on this general topic but it has not yet trickled down into truly automatic tools that can 'look' at a point cloud and 'recognize' common features.

Two general approaches can be taken. First, points can be converted to surfaces with geometry and then corrected, with the source points being removed from the point cloud. As this process continues the point cloud will eventually consist of only features that are non-geometric (like trees, grass,…) or are unrecognizable. Alternatively a matching tool can look through the point cloud and try to recognize features by comparison to a library of known features, but this is largely a research method at the moment.

Even the process of building simple surfaces is problematic. How smooth should a surface be? Should the tool leave in texture – like bricks – or generate the overall wall shape instead? Can we tell instrument precision issues apart from feature texture? To date, LIDAR processing tools are reasonably well suited to surface building and object tracing for highly regular features like pipes, but less well suited to complex features such as building fronts.

Output To GIS and Visualization Tools

Ideally geometric output from LIDAR processing, as features where possible, should be passed to more general mapping and visualization tools for further work. For example, building shapes from a TITAN scan might be processed to building features at a desired resolution, brought into Autodesk's 3dStudioMax for refinement, and then exported as components for use in movie making. TITAN data might also be used to build out a large urban area at meter-level of detail for export to an open visualization platform such as Google Earth.

The real issue encountered here is of data set size. LIDAR data sets in the gigabyte range just do not transfer into 3d visualization tools effectively, and for Google Earth, even textured geometric features are an issue. For example, Google Earth engineers recommend representing a large building as images (up to about 200 kb in size total) and features (a few dozen flat shapes total) for a total size of about 250 kb whereas a LIDAR scan of the same building might have 500 million points or correspond to a few tens of megabytes of detailed building geometry. Clearly the workflow from LIDAR to Google Earth needs some research and pragmatic choices!

Conclusions and the Future

LIDAR scans, whether mobile or static, can be used to build urban models and to map urban infrastructure accurately and relatively cheaply. There is simply no alternative technology that allows thousands of measurements on geometry per second and can scan all features in line-of-sight to roadways. The combination of LIDAR with high resolution photography allows accurate and compelling 3d representations of urban areas to be built rapidly; to date the processing from data to features is not so rapid!

Research on LIDAR processing, especially on feature extraction and spatial analysis, is ongoing. As computers get faster and cheaper – one thing we can count on! – handling the large data volumes seen in LIDAR will be less daunting and the penetration of this technology into urban planning, architecture, and civil engineering will accelerate. While it may be some time before LIDAR-level models are used in online tools like Google Earth, in the near future features in online environments, games, and urban visualization tools will be at least based on LIDAR.

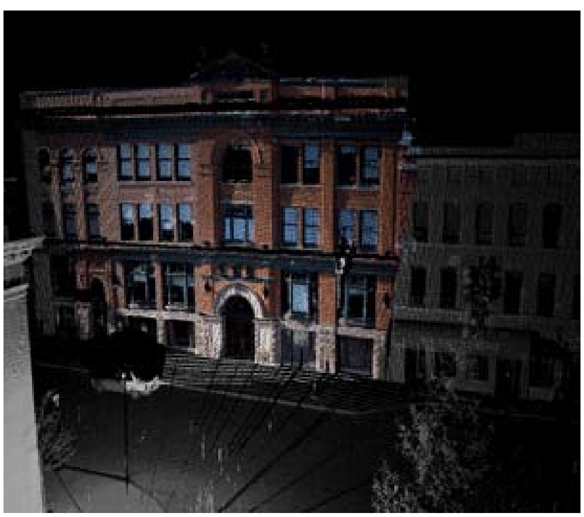

The addition of a little bit of color can make a big difference in how we 'see' LIDAR data. In this case, we've taken the Whig-Standard Building in Kingston, as seen in previous images, and added color information from a 35mm camera image shot from the same location as the LIDAR.

Related Subjects

For more information on LIDAR, see the references below. If you find this technology interesting, you may also want to look at related materials on:

GIS – tools for mapping, dominantly in plan view – with an emphasis on spatial analysis.

3d Visualization – tools for building and manipulating 3d data sets, making animations, and visualizing different scenarios.

GPS – fundamental geographic positioning technology

There are 'What Is' notes in this series on all of these!

References

LIDAR is new and much of the science and practice is not yet found in texts or generally available reference works. Some places to start include:

Maune, D.F. 2001 Digital Elevation Model Technologies and Applications: The DEM Users Manual. Published by the American Society for Photogrammetry and Remote Sensing, Bethesda, Maryland

Leica Geosystems website: http://www.leica-geosystems.com

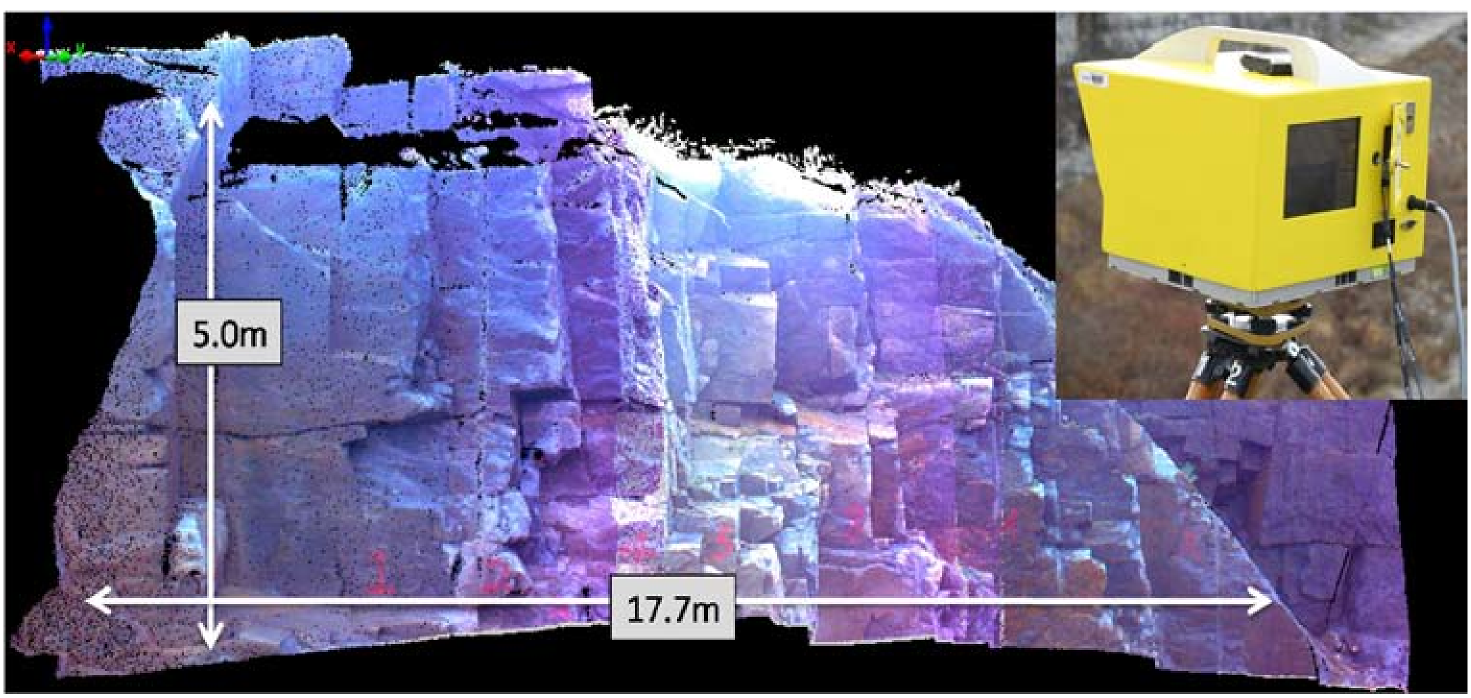

LIDAR can be used to map foundations, roadcuts, and other natural features. In this example, from research by Matt Lato, an Optech IlRIS scanner was used to scan a roadcut which was then colorized from imagery from an internal digital camera. This process is known as 'pixel mapping.'