Bounding Methods for Heterogeneous Lidar-derived Navigational Geofences

Andrew J. Moore
Langley Research Center, Hampton, Virginia

Matthew Schubert
Analytical Mechanics Associates, Inc., Hampton, Virginia

Terry Fang
University of Pennsylvania, Philadelphia, Pennsylvania

Joshua Smith
University of Arkansas, Fayetteville, Arkansas

Nicholas Rymer
National Institute of Aerospace, Hampton, Virginia

NASA STI Program . . . in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA scientific and technical information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI program provides access to the NTRS Registered and its public interface, the NASA Technical Reports Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA Programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counter-part of peer-reviewed formal professional papers but has less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include organizing and publishing research results, distributing specialized research announcements and feeds, providing information desk and personal search support, and enabling data exchange services.

For more information about the NASA STI program, see the following:

- Access the NASA STI program home page at http://www.sti.nasa.gov

- E-mail your question to help@sti.nasa.gov

- Phone the NASA STI Information Desk at 757-864-9658

- Write to:
  NASA STI Information Desk
  Mail Stop 148
  NASA Langley Research Center
  Hampton, VA 23681-2199

NASA/TM–2019-220399

Bounding Methods for Heterogeneous Lidar-derived Navigational Geofences

Andrew J. Moore
Langley Research Center, Hampton, Virginia

Matthew Schubert
Analytical Mechanics Associates, Inc., Hampton, Virginia

Terry Fang
University of Pennsylvania, Philadelphia, Pennsylvania

Joshua Smith
University of Arkansas, Fayetteville, Arkansas

Nicholas Rymer
National Institute of Aerospace, Hampton, Virginia

September 2019

Acknowledgments

Available from:

NASA STI Program / Mail Stop 148
NASA Langley Research Center
Hampton, VA 23681-2199
Fax: 757-864-6500

Abstract

Safe Unmanned Aerial Vehicle (UAV) operations near the ground require navigation methods that avoid fixed obstacles such as buildings, power lines and trees. Aerial lidar surveys of ground structures are available with the precision and accuracy to geolocate obstacles, but the high volume of raw survey data can exceed the compute power of onboard processors and the rendering ability of ground-based flight planning maps. Representing ground structures with bounding polyhedra instead of point clouds greatly reduces the data size and can enable effective obstacle avoidance, as long as the bounding geometry envelopes the structures with high spatial fidelity. This report describes in detail four methods to compute bounding geometries of ground obstacles from lidar point clouds. The four methods are: 1) 2.5D Maximum Elevation Box, 2) 2.5D Ground Map Extrusion, 3) 3D Bounding Cylinder, and 4) 3D Bounding Box.

The methods are applied to five point cloud datasets from lidar surveys of UAV flight research sites in Georgia and Virginia with an average point spacing that ranges from 0.1m to 0.6m. The methods are assessed using survey areas with geometrically heterogeneous ground structures: buildings, vegetation, power lines, and sub-meter structures such as road signs and guy wires. The 2.5D Maximum Elevation Box method is useful for simple structures. The 2.5D Ground Map Extrusion method efficiently encloses vegetation, but requires hand-drawn ground footprints. The 3D Bounding Cylinder method excels at enclosing linear structures such as power lines and fences. The 3D Bounding Box method excels at enclosing planar structures such as buildings.

The methods are compared on the basis of data compression and boundary fidelity on selected areas. The 2.5D methods yield the highest data compression but the polyhedra produced by them enclose significant amounts of empty space. Boundary fidelity is superior for the 3D methods, though this fidelity comes at the cost of a roughly thirtyfold lower data compression ratio than the 2.5D Maximum Elevation Box method. A mix of these output geometries is proposed for autonomous UAV navigation with limited on-board computing.

Both the accuracy and spatial detail of emerging satellite-based survey technology lower than that of aerial lidar scanning survey technology. Sub-meter structures and thin linear structures are not reliably mapped at present by satellite-based surveys.

I.   Introduction

Buildings, power lines and trees are collision hazards in Unmanned Aerial Vehicle (UAV) operations altitudes just above ground level. Techniques to avoid these obstacles are needed for safe low altitude navigation. Avoidance can be accomplished while planning the flight (proactively) and during the flight (reactively) if the GPS coordinates of the boundaries of obstacles are known. Before a flight, a UAV operator can choose flight waypoints that are a safe distance from obstacles if the obstacle boundaries are displayed in a flight planning map. During flight, onboard autonomy can compare the vehicle coordinates to the boundary coordinates and adjust the flight path to avoid collision in the event that the vehicle strays off-course due to off-nominal conditions such as high winds.

There are two prerequisites to accomplishing these operational practices. First, accurate surveys of obstacle boundaries are required, and second, the obstacle boundaries must be represented compactly. In modern practice, surveys of ground structures combine traditional tripod-based human measurements (to establish known reference control points) with lidar point clouds from aerial [1] or ground [2] scans which are fitted to the control points [3]. Richly detailed surveys with decimeter level accuracy are commonplace with this method. Compact representation of obstacle boundaries is needed, because detailed three dimensional (3D) lidar surveys of a site generate enormous data sets. The raw data size will overwhelm both the visualization software used to plan flights and the onboard computing [4] used for inflight obstacle avoidance. Naïve data reduction such as downsampling prior to boundary determination risks the removal of critical obstacle features.

This report first describes in detail four methods to compute a compact representation: 1) 2.5D Maximum Elevation Box, 2) 2.5D Ground Map Extrusion, 3) 3D Bounding Cylinder, and 4) 3D Bounding Box. The

methods are applied to five point cloud datasets from lidar surveys of flight research sites in Georgia and Virginia. Research flights were conducted at these sites to test technologies for UAV inspection of high voltage electrical infrastructure [4]-[9] and for safe UAV urban operations [10]. The spatial sampling (average point spacing) of the datasets ranges from 0.1m to 0.6m.

A. *Brief summary of literature*

Simplified boundary representations of ground structures have been available for some years. Commonly referred to as digital elevation maps, these representations are used, for example, to map tree canopies in environmental and geological studies [11], buildings [12], and piping [13]. In early studies (see [14] for a review) digital elevation maps represented x and y dimensions with much more fidelity than the z dimension, and were also called 2.5D models [15].

As aerial lidar scan and computing technologies advanced, full 3D models were developed to classify ground obstacles. While the variety of ground structures under study and its partitioning into geometric classes differs from study to study (e.g., [16] [17]), the majority of researchers considered at least three distinct geometric groups: planar structures, linear structures, and vegetation. Man-made structures such as buildings and bridges are predominantly collections of planar facets, wires and poles are linear, and trees and other vegetation are neither (or quasi-fractal [18]). Planar models are almost universally employed for building facets, while linear structures are modeled either as cylinders [12][13][19], or, in the case of suspended lines such as cables and electrical conductors, as catenary curves [20][21][22]. Because the geometry of vegetation is so complex, diverse, and difficult to model [18][23][24], a classification pipeline is commonly used: first planar and linear structures are identified and enclosed with 3D bounding geometries, and then the enclosed points are removed from the point cloud, leaving only lidar points from vegetation [25]. The remainder is then modeled with a battery of techniques [25], fitted with a convex hull [26], or (analogous to draping it with an elastic sheet) simply enclosed with a 2.5D surface spline [27].

Surveys derived from satellite-based [28] and UAV-based [29] photogrammetric stereo imaging is rapidly progressing. While inadequate for low altitude UAV navigation (see Discussion), survey data from these technologies can be processed with the methods described in this report and used for low resolution applications such as topographical classification and forestry management.

B. *Outline of the report*

The performance of each method is described in Section II and the methods are compared in Section III using illustrative subsets of the five lidar surveys. The survey data includes ground structures with heterogeneous geometries: buildings, vegetation, power lines, and small structures such as road signs and guy wires. Both the boundary fidelity and the degree of data compression afforded by the method form the basis for comparison. In Section IV, the suitability for autonomous UAV navigation of the output geometries from the methods is discussed in the context of computing requirements and spatial fidelity, and the accuracy and spatial detail of emerging satellite-based survey technology is critically assessed.

II.   Lidar Data and Bounding Methods

A. *Lidar Data*

Five sets of aerial lidar survey data were used to assess and compare the methods. All data was received in the LAS (LASer) point cloud file format. When converting from the state plane coordinate system of the input survey to the WGS84 (World Geodetic System 1984) latitude/longitude/altitude coordinate system, a grid unit spacing of at least six decimal places of decimal latitude and longitude was used, which corresponds to approximately 0.1 meter spatial resolution at middle latitudes. The datasets are:

A. A survey of a 1.0 square mile area centered on the Southern Company Klondike Training Facility in Lithonia, Georgia collected in May, 2015. At an average point spacing of 0.3m, the full 282M point dataset size is 7.3GB.

B. A survey of a 0.3 mile by 0.2 mile area centered on the Dominion Energy Training Facility in Chester, Virginia collected in November, 2016. At an average point spacing of 0.33m, the full 8M point dataset size is 270MB.

C. A survey of a 29 mile transmission line corridor originating at a Dominion Energy facility in central Virginia collected in October, 2015. At an average point spacing of 0.6m, the full 117M point dataset size is 3.4GB.

D. A survey of a 1.5 square mile area centered on the NASA Langley Research Center in Hampton, Virginia collected in April, 2015. At an average point spacing of 0.2m, the full 117M point dataset size is 3.2GB.

E. A survey of a 1.0 square mile area centered on the NASA Langley Research Center in Hampton, Virginia collected in February, 2018. At an average point spacing of 0.1m, the full 195M point scan dataset size is 5.2GB.

B. Bounding Methods

Four methods for creating bounding geometries that enclose lidar-surveyed ground obstacles are described in this section:
1. 2.5D Maximum Elevation Box,
2. 2.5D Ground Map Extrusion,
3. 3D Bounding Cylinder, and
4. 3D Bounding Box.

Since the two 3D methods have similar processing steps, they are described together. Figure 1 schematically illustrates the four methods.
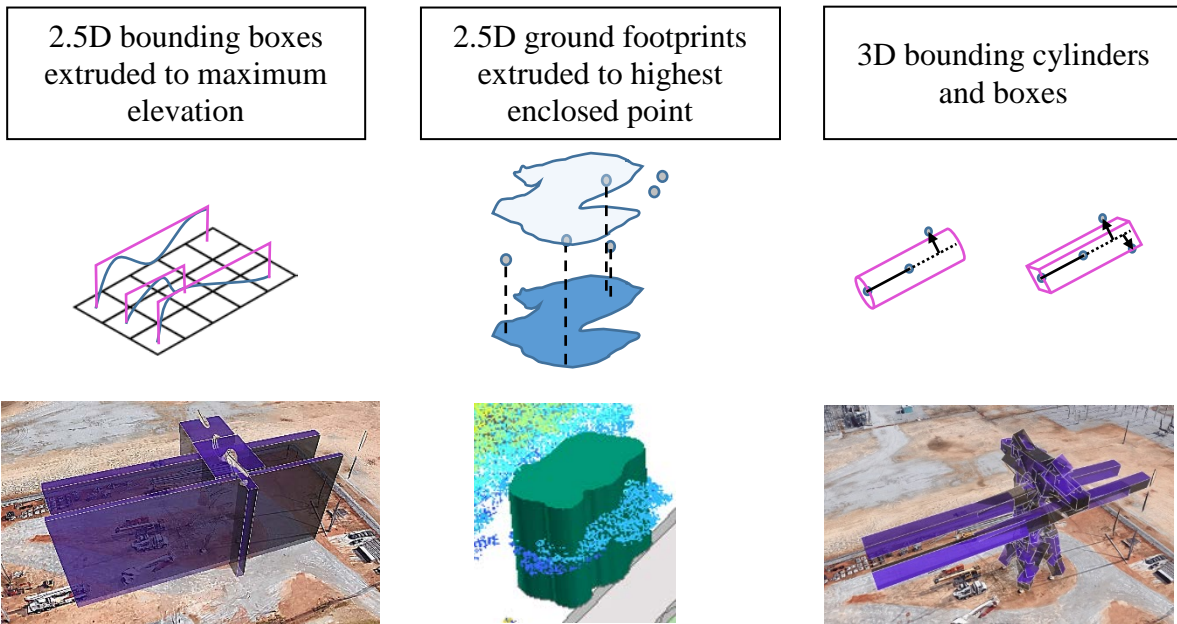
1) 2.5D Maximum Elevation Box

The 2.5D Maximum Elevation Box method, written in C++, constructs bounding rectangular polyhedra aligned with a latitude/longitude grid which extend perpendicularly from the ground plane to the maximum enclosed altitude. This method was applied to lidar dataset A, provided by Southern Company. The structures in this survey are largely aligned to lines of latitude and longitude; with suitable rotation in the ground plane, the method can be applied to other structure surveys as well. After ground point removal the lidar points are converted from the local state projection to the WGS84 (x,y,z: latitude, longitude, altitude) coordinate system.

After sorting based upon x and y values, the set of lidar points cloud is traversed, first along lines of latitude and then along lines of longitude, finding clusters of points within a radius set to twice the minimum point cloud spacing. A provisional polyhedron is constructed with a height equal to the maximum altitude within each cluster. As points are traversed and examined, they are marked as processed to avoid overlaps and double counting. After all points are traversed, provisional polyhedra are merged along lines of latitude such that grid locations with lidar points above ground level are set to the maximum value along the line. This simplification results in a striped morphology that was appropriate for the elongate high voltage structures in the flight test range (Fig 1. left).

2) 2.5D Ground Map Extrusion

The 2.5D Ground Map Extrusion method (Fig. 1, middle) merges two representations of ground structures: hand drawn footprints and aerial lidar survey data. A high fidelity survey database of the NASA Langley Research Center has been built and maintained using a variety of techniques including traditional tripod-based measurements, aerial and ground lidar scans, and airborne photogrammetry via survey-grade UAVs. Ground structure footprints are periodically entered and revised as graphical elements (circles, ellipses and Bezier curves) and stored with corresponding structure elevation estimates using the ArcGIS geographical information systems toolset.

| 2.5D bounding boxes extruded to maximum elevation | 2.5D ground footprints extruded to highest enclosed point | 3D bounding cylinders and boxes |

Figure 1. Polyhedra construction methods. Left to right: 2.5D Maximum Elevation Box, 2.5D Ground Map Extrusion, 3D Bounding Cylinder, and 3D Bounding Box. The 2.5D Maximum Elevation Box method is simplest and produces the greatest data compression at the cost of feature fidelity and 'dead space'. The 2.5D Ground Map Extrusion method is less compact and follows the lidar point cloud more closely, but requires hand-drawn ground footprints and fails to enclose lidar points that do not overly the footprints. Feature fidelity is highest for the 3D Bounding Cylinder and 3D Bounding Box methods, but these methods yield the lowest data compression. © *Map data: Google & DigitalGlobe*

This method attempts a compromise between the inherently simplified graphical footprints and the exact but enormous raw lidar survey data by combining them to constructing 2.5D right prisms with complex, multi-vertex base polygons. Aerial scan lidar is converted from its initial LAS format to an ArcGIS 'multipoint' format in the same local state plane coordinate system as the footprints. Python code running in the ArcGIS tool then performs an ArcGIS 'spatial join' operation, which finds all (x,y,z) triplets of the point cloud which are enclosed by the footprint polygons, and an 'extrude' operation which finds the maximum altitude of the enclosed points. Finally, custom Python code exports each footprint curve from ArcGIS, discretizes it to a polygon, constructs a 2.5D prism with the height found in the extrude operation, and converts it to the WGS84 (latitude/longitude/altitude) coordinate system. Footprint curve discretization is conducted such that there is no greater than a two centimeter deviation from the original hand-drawn curve.

The 2.5D geometries (Fig. 2, left) are geospatially verified by visual inspection after conversion to KML (Keyhole Markup Language) and importing to Google Earth (Fig. 2, top right) and also after conversion to XML (Extensible Markup Language) and importing to the MavProxy [30] ground station platform (Fig. 2, bottom right). Additionally, the XML is imported to the NASA Polycarp [31] geometry engine to verify that the polyhedra are free of problematic geometries such as repeated vertices and nonoriented ("bow-tie") topologies. This geometry engine is used by the Safeguard and ICAROUS safety technologies (see Discussion).
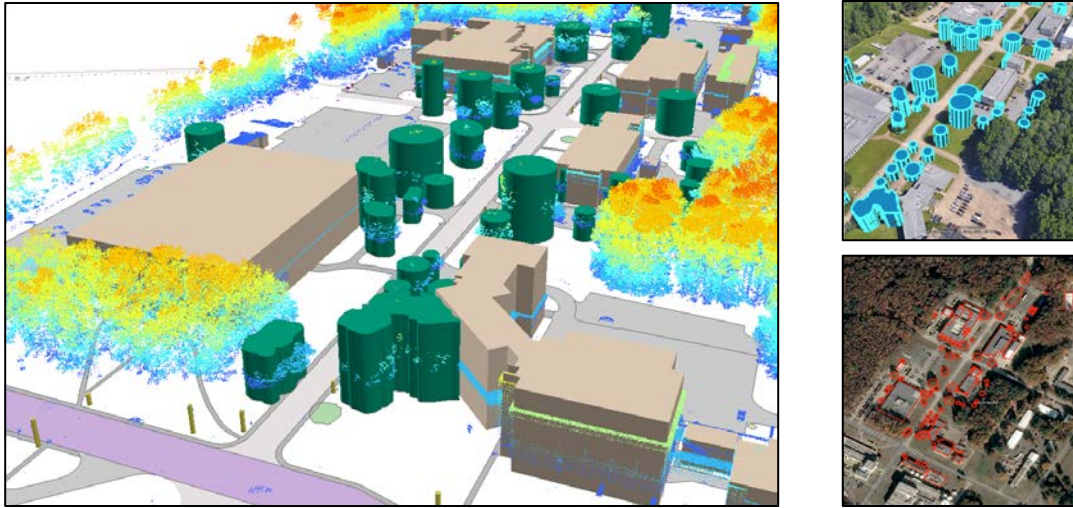
Figure 2. 2.5D Ground Map Extrusion method. Left: Hand-drawn footprint curves for vegetation are digitized and extruded to the highest elevation of the lidar points enclosed. Building footprints are extruded to their surveyed height. Right: The constructed polyhedra for vegetation are verified as KML in Google Earth (top right) and, along with the building footprints, as XML in the MavProxy [30] ground station map (bottom right). © *Map data: Google & DigitalGlobe*

3) 3D Bounding Cylinder and Box Methods

The 3D bounding methods (Fig. 1, right) cluster an input lidar point cloud via the steps described in the following paragraphs. Spatial constants (e.g., d_{min}) used in the steps were optimized for datasets with an average point spacing of 0.3 m or less.

Overview. After an initial grouping of lidar points into clusters based on proximity, each lidar point may be a member of multiple, spatially overlapping neighborhoods. The 3D methods are designed to find a set of neighborhoods just sufficient to enclose the point cloud. This is accomplished by repeatedly growing existing neighborhoods and eliminating neighborhoods made redundant by that growth. Initially the neighborhood count is on the order of the lidar point count, and as the growth/elimination iterations proceed, the neighborhood count drops rapidly.

The growth/elimination sequence and the metrics used to determine which neighborhoods to grow/eliminate are:

Create base neighborhoods. Aerial lidar survey data is preprocessed by removing points in the ground plane and converting from LAS format in a survey coordinate system (e.g., State Plane Georgia) to a point array $P_{1...n}$ of *(x,y,z)* triplets in the earth-centered-earth-fixed (ECEF) coordinate system. For each point P_n a local neighborhood N_n is found which contains all points no more than d_{min}=1m distant from P_n. For each neighborhood, an eigenaxis-aligned oriented cylinder or oriented bounding box is found using the Point Cloud Library (PCL) moment of inertia estimator, which computes the major, middle and minor eigenvalues and eigenvectors. The eigenvectors form an orthogonal basis $(\hat{a}, \hat{b}, \hat{c})$, and the oriented cylinder or box is aligned with a major axis in the $\hat{a}$ direction and dimensions sufficient to contain all of the points in the neighborhood.

Remove outliers. The distribution of a neighborhood's point distances from the neighborhood centroid is computed, and those points with a distance greater than 1.5 times the interquartile range are considered as outlier points and removed from the neighborhood.

Sort neighborhoods (cylindrical method). Neighborhoods are sorted from largest to smallest according to the number of points enclosed in the neighborhood. If two or more neighborhoods enclose the same number of points, they are sorted according to a measure of ideal cylindricality (the ratio of major eigenvalue to middle eigenvalue).

Sort neighborhoods (box method). A tally *neighborhood_agreement* is initialized to zero for each base neighborhood. Duplicate neighborhoods are eliminated by finding neighborhoods which have the exact same points and removing all except one from further consideration. The *neighborhood_agreement* tally is incremented for each duplicate of a particular neighborhood that is found and eliminated. The primary, secondary, and tertiary eigenaxes $(\hat{a}, \hat{b}, \hat{c})$ of each remaining oriented bounding box is recalculated and all points
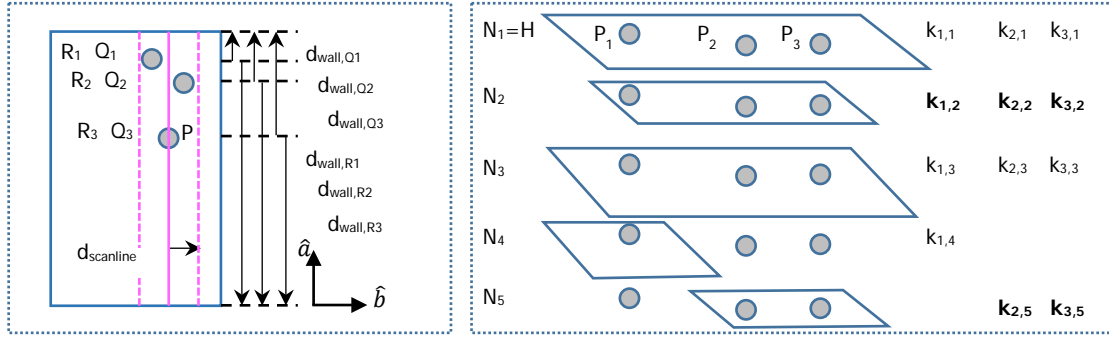
Figure 3. Spatial and membership metrics used to sort and eliminate neighborhoods in the 3D Bounding Box method. The spatial fitness scanline construction and wall distance measurements are shown at left. At right is a representation of the membership decisions that arise from considering points P_1, P_2, and P_3 in a seed neighborhood H which has points that are contained in four other, overlapping neighborhoods (rendered in two dimensions and exploded in the vertical dimension to remove overlap). The neighborhood designations N_1 to N_5 are shown to the left of the diagram and the corresponding fitness metrics $k_{point,neighborhood}$ are enumerated at the right of the diagram. In this example, neighborhoods N2 and N5 fit their constituent points most tightly (boldface k values). The remaining three neighborhoods will be removed in this or later rounds of fitness ranking.

that fall within its volume are added to its neighborhood grouping.

Next, the fitness of the bounding box is calculated using a spatial metric which attempts to quantify how well the points fill the bounding box, by finding the distance of the points to each wall of the bounding box for all six possible permutations of axes $(\hat{a}', \hat{b}', \hat{c}') \in \{ (\hat{a}, \hat{b}, \hat{c}), (\hat{a}, \hat{c}, \hat{b}), (\hat{b}, \hat{a}, \hat{c})...\}$. This calculation is depicted at left in Fig. 3. First, the coordinates of each point in the neighborhood are transformed from the input axes to the current local eigenaxes $(\hat{a}', \hat{b}', \hat{c}')$. Then in the plane described by primary and secondary eigenaxes $(\hat{a}', \hat{b}')$, each point P enclosed in the box is selected in turn, and points within a fixed buffer distance $d_{scanline}$ from P along the secondary axis $\hat{b}'$ are considered. The point Q with the lowest distance d_{wall} to the bounding box wall intersected by the primary axis in the positive direction is found, and the point R with the lowest distance to wall in the opposite (negative) direction is found. The identities and distances of Q and R are 'tagged' to avoid double counting and, the distances $d_{wall,Q,min}$ and $d_{wall,R,min}$ are added to a running sum for $(\hat{a}', \hat{b}')$. The next point in the box is chosen as P, a new scan line is re-centered on it, new Q and R points within the scan line are selected, and $d_{wall,min}$ values are found and tallied, etc. After the scanline is swept across the box along the $\hat{b}'$ axis, the next axis combination $(\hat{a}', \hat{b}', \hat{c}')$ is considered, and these summations are calculated, and so on until sums for all axis combinations are found.

The sums are combined for pairs of the six axis permutations with the same primary axis (e.g., for $\{\hat{a}, \hat{b}, \hat{c}\}$ and $\{\hat{a}, \hat{c}, \hat{b}\}$), resulting in three unnormalized eigenaxis fitness values and three counts of how many points were used for each eigenaxis fitness value. The three counts of points used are summed to find a total count of points used, and eigenaxis fitness values are normalized by dividing the three fitness values by this total point count. This yields the fitness value for each possible primary eigenaxis for the bounding box.

At this point in the computation, individual lidar points may be members of more than one box group (neighborhood). The eigenaxis fitness values are used to sort the set of boxes and remove boxes with poor fit and low point membership as depicted at the right in Fig. 3. An arbitrary 'seed' neighborhood H is chosen, and each point $P_i = P_1, P_2, ..., P_h$ inside H is considered in turn. P_i commonly belongs to $k_i > 1$ neighborhoods; the k_i neighborhoods which contains P_i are sorted from best to worst fitness, where the best has the lowest sum of

the three normalized eigenaxis fitness values. There is great spatial overlap amongst the resultant set of neighborhoods $k_1 \ldots k_h$. For example, many neighborhoods containing P_1 will also contain P_2, and therefore sets k_1 and k_2 will generally have common members. The neighborhoods contained in $k_1 \ldots k_h$ are ranked according to the number of times (from 0 to a maximum possible count of h) that they have the best fitness. All neighborhoods with a rank of 0 are eliminated.

Grow/eliminate neighborhoods (cylindrical method). Unenclosed points are added to existing neighborhoods in order of the previous sorting, i.e., from largest to smallest fitness. For each point in a neighborhood N, any point within a 0.75m radius of the major axis is found and added to N if it meets the following criteria:

a) The point falls within an infinite cylinder with 0.4m radius, aligned to the major eigenvector of N.
b) The point is 'sufficiently parallel', as determined by creating a two-point neighborhood N_{test} from the point under consideration and the point in N nearest to it. If the angular difference between the major eigenvector of N and the major eigenvector of neighborhood N_{test} is less than ~26 degrees, the point is added to N.

The eigenvalues and eigenvectors are recalculated using PCL's moment of inertia estimator and then neighborhoods are 'greedily' regrown as follows. For each neighborhood, starting with its origin point the neighborhood is grown (more points added) using the criteria a) and b), but using the recalculated eigenvalues and eigenvectors. Once a point is added to a neighborhood in this step, it is not added to any other neighborhood.

If a resulting neighborhood has less than four points, the neighborhood is eliminated. The points which were contained in an eliminated neighborhood are now unenclosed 'orphans' which will be added to surviving neighborhoods in subsequent growth steps.

Grow/eliminate neighborhoods (box method). Unenclosed (orphan) points are added to surviving neighborhoods in a way that approximates incremental, greedy Manhattan box growth on a 3D grid defined by a neighborhood's eigenaxes. For each neighborhood, 'orphan' points are tentatively added if they are within a distance $d_{orphan}=1$m of points in the neighborhood, and fall within the bounding box if it were projected infinitely along its primary eigenaxis. Second, the bounding box and corresponding fitness metrics are recalculated, and if the fitness values for each eigenaxis of the new bounding box is less than a fixed value $F_{min}=0.3$, the added 'orphan' points in the neighborhood are retained and removed from the pool of 'orphan' points. Third, this process is repeated until no new points are added to the neighborhood in the direction of its primary eigenaxis. If any of the fitness values exceed F_{min}, then the temporarily added points are removed.

These three steps are repeated for the secondary and tertiary eigenaxes, including or rejecting 'orphan' points depending on their fitness. If at least one point is added to the neighborhood along any axis, the three steps are restarted, beginning again with the primary axis. After this axis-wise Manhattan growth, the neighborhood's oriented bounding box is recalculated and all points that fall within its (likely expanded) 3D box are added to the neighborhood, except for Tukey outliers as determined by a 1.5 interquartile range threshold along each eigenaxis.

After this growth step, neighborhoods are identified which are duplicates, i.e., contain the exact same lidar points, and all but one duplicate is eliminated.

Merge neighborhoods and create bounding polyhedra. At this writing, merging the final set of boxes is not implemented for the 3D Bounding Box method, so that unmerged box neighborhoods, converted to WGS84 coordinates, serve as the bounding polyhedra.

For the 3D Bounding Cylinder method, as a final data reduction step, surviving cylindrical neighborhoods are merged iteratively based on adjacency and overlap. Neighborhood N and neighborhood M are considered adjacent if there exists a pair of points, one from N and one from M, which are within a distance $d_{adjacent}=1$m of each other. This value is set to prevent merging of well-separated parallel powerline conductors while enforcing merging of long lines of lidar points from single conductors which may be incidentally separated by sampling offsets. To handle the special case of conductor pairs (see Fig. 6 for an example), two candidate neighborhoods N and M with points closer than $d_{adjacent}=1$m are considered "super-adjacent" and merged only if for every point in N there exists a point in M that is within $d_{adjacent}$ of it. Merging proceeds for all neighborhoods and is then repeated until no merging occurs.

Finally, using PCL's Moment of Inertia Estimator, an oriented bounding box is computed for each cylindrical neighborhood. The set of oriented bounding boxes found in this way is converted to the WGS

latitude/longitude/altitude coordinate space and serve as the bounding polyhedra for linearly oriented lidar points such as powerline conductors.

III.   Results

A.   *Enclosure Style and Fidelity for 2.5D Maximum Elevation Box and 3D Bounding Cylinder Methods*

Two areas from dataset A were chosen to compare the 2.5D Maximum Elevation Box method and 3D Bounding Cylinder method (top of Fig. 4). At left is a section of a 500kV transmission line with a steel truss tower and two conductors (referred to herein as the 500kV Section). About 8000 lidar points remain after removal of ground points and the termination structures and associated guy wires at either end of the structure. At right is a nearby heterogeneous area in the same facility (referred to herein as the North Quarter) with a mixture of trees, buildings, poles, towers, electrical overhead conductors, and ground bus bar structures. About 2.3M lidar points remain after removal of ground points from survey data of this area.

The 2.5D Maximum Elevation Box method is adequate for the lidar point cloud of the 500kV Section (middle left of Fig. 4). Though the method produces bounding polyhedra that are overly conservative in the z direction, bounding all of the flight volume below the conductors, it is efficient in the lateral plane. The bounding polyhedra created by the 3D Bounding Cylinder method (bottom left of Fig. 4) tightly envelope the conductors, though the steel truss tower is inefficiently represented.

The more diverse North Quarter obstacle field a much richer basis for comparison between the 2.5D Maximum Elevation Box method and 3D Bounding Cylinder method. The bounding polyhedra created by the 2.5D Maximum Elevation Box method (middle right of Fig. 4) delineate the bus bars in the foreground efficiently, but the trees and powerlines in the background are grouped together into loosely bounding polyhedra. In contrast, the bounding polyhedra created by the 3D Bounding Cylinder method (bottom left of Fig. 4) tightly envelope the bus bars and power lines. The computed boundary of the irregular foliage geometry, while tightly enveloping, is a jumble of polyhedra.

These visual metrics are reflected in the relative data compression of the two methods (Table 1). The 2.5D Maximum Elevation Box method reduces the data size by three to four orders of magnitude for the areas tested [4]. That is, the 2.5D Maximum Elevation Box method polyhedral vertex count is 1000-5000 times lower than the number of raw lidar points (after ground removal) as tested with these two data sets. This reduction is largely due to two geometric simplifications:

a)   Clustering by z value encloses space underneath an object even if there are no underlying lidar points.

b)   Elevation along a line of latitude is simplified to a single (max) elevation for long runs in longitude, so that open air over an object may be enclosed.

In contrast, the 3D Bounding Cylinder method reduces the data volume approximately by 1.5 to 2 orders of magnitude (a ratio of 50-100) with the same data set. While the polyhedral fit is superb, compression is poor for foliage and complex structures.

A close-up view of the North Quarter results for the two methods is shown in Figure 5. Large swaths of empty area within the array of bus bar structures at right are enclosed by the 2.5 Maximum Elevation Box method, while the 3D Bounding Cylinder method bounds them tightly. At left, the mixture of foliage and powerlines is clumped together by the 2.5 Maximum Elevation Box method. The 3D Bounding Cylinder method handily segregates power lines from foliage, but the foliage is enclosed by many polyhedra with diverse orientation.

As expected based on their relative complexity, the run time of the 2.5 Maximum Elevation Box method is much shorter than the run time for the 3D Bounding Cylinder method. Using a single thread on a 3.1 GHz Xeon processor, the North Quarter point cloud requires about two minutes of run time for the 2.5D method and two hours for the 3D cylindrical method.

**Table 1. Data reduction of two areas using 2.5D Maximum Elevation and 3D Bounding Cylinder Methods**

|	Input lidar	2.5D Maximum Elevation		3D Bounding Cylinder	
	Point count (without ground)	Polyhedra count	Ratio	Polyhedra count	Ratio
500 kV Section	7,965	6	1328	146	55
North Quarter	2,388,823	406	5883	22,380	107
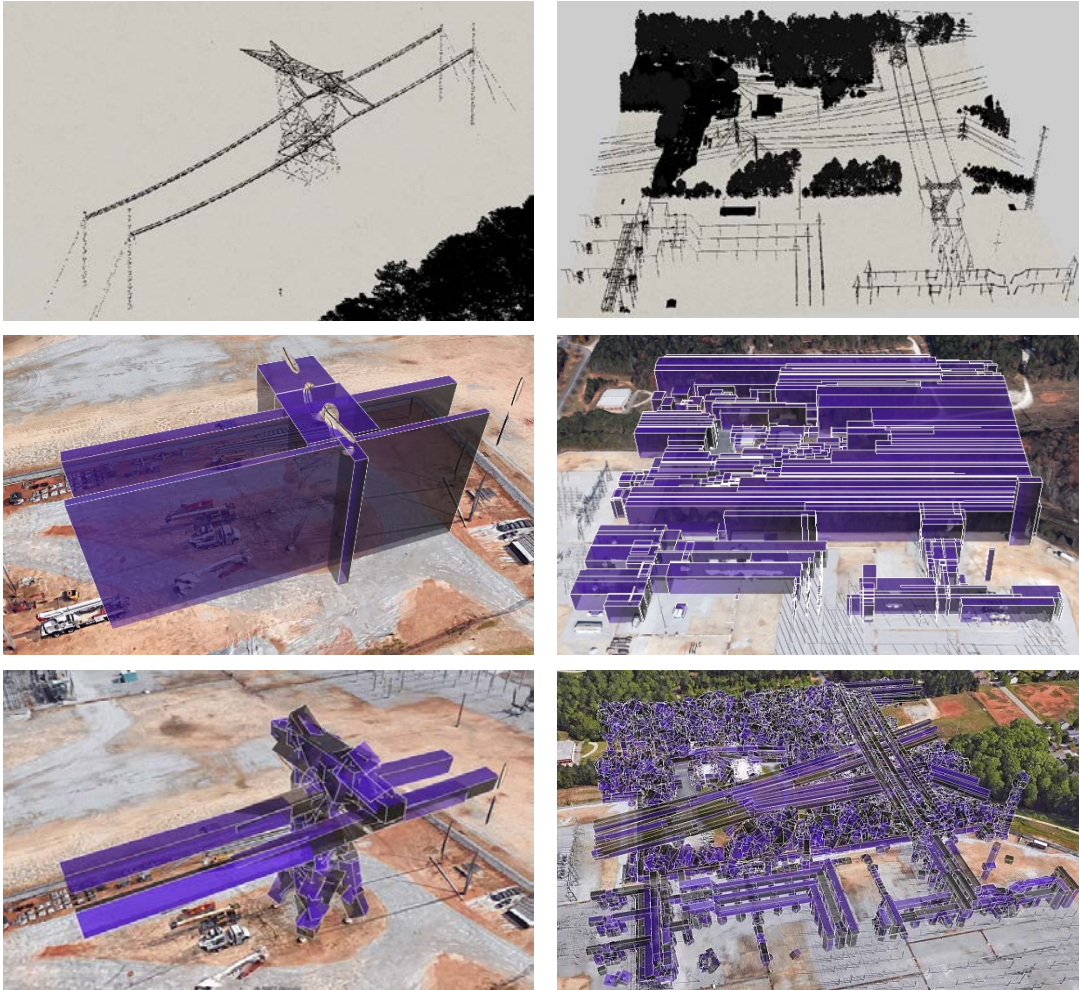
Figure 4. Comparison of 2.5D Maximum Elevation Box and 3D Bounding Cylinder methods on two areas from dataset A. Left: 500 kV Section. Right: North Quarter. Top row: input lidar after ground point removal. Middle row: 2.5D Maximum Elevation Box representation. Bottom row: 3D Bounding Cylinder representation. © *Lidar data: Southern Company; Map data: Google & DigitalGlobe*

The internal representation of the 3D Bounding Cylinder method is illustrated in Figure 6 for a section of a 360kV Dominion Energy transmission line structure from dataset B. Three conductor pairs are suspended by two steel truss towers in this structure. At top, the individual lidar points are colored according to the membership in a neighborhood as described in Section II.B.3. The bounding polyhedra constructed with a one meter safety buffer are shown at bottom. The constraints employed to enforce cylindrical clustering (0.4m radial proximity) cause the catenary curve of the suspended cables to be segmented. The data compression afforded by the method is considerable (~ 100:1) for the linear conductors. Most of the 370 polyhedra that enclose the approximately 15000 lidar points are located on the truss elements of the steel towers.

While the 3D Bounding Cylinder method is adequate for most linear structures (poles, conductors, bus bars), very thin objects such as guy wires are not possible to enclose completely due to lidar sampling errors. Figure 7 shows an attempt to bound a lidar scan of a switching tower at the Dominion Energy Chester facility. The beams and poles of the tower resolve well in the aerial scan and were bounded adequately by the method, but there are dropouts in the aerial lidar survey scan data for the thin (~3/8 in) guy wires which mechanically support the tower. The one meter maximum gap d_{min} rule used by the method to avoid bridging across spatially separated objects breaks the bounding geometry of the guy wires into multiple polyhedra (Figure 7, right).
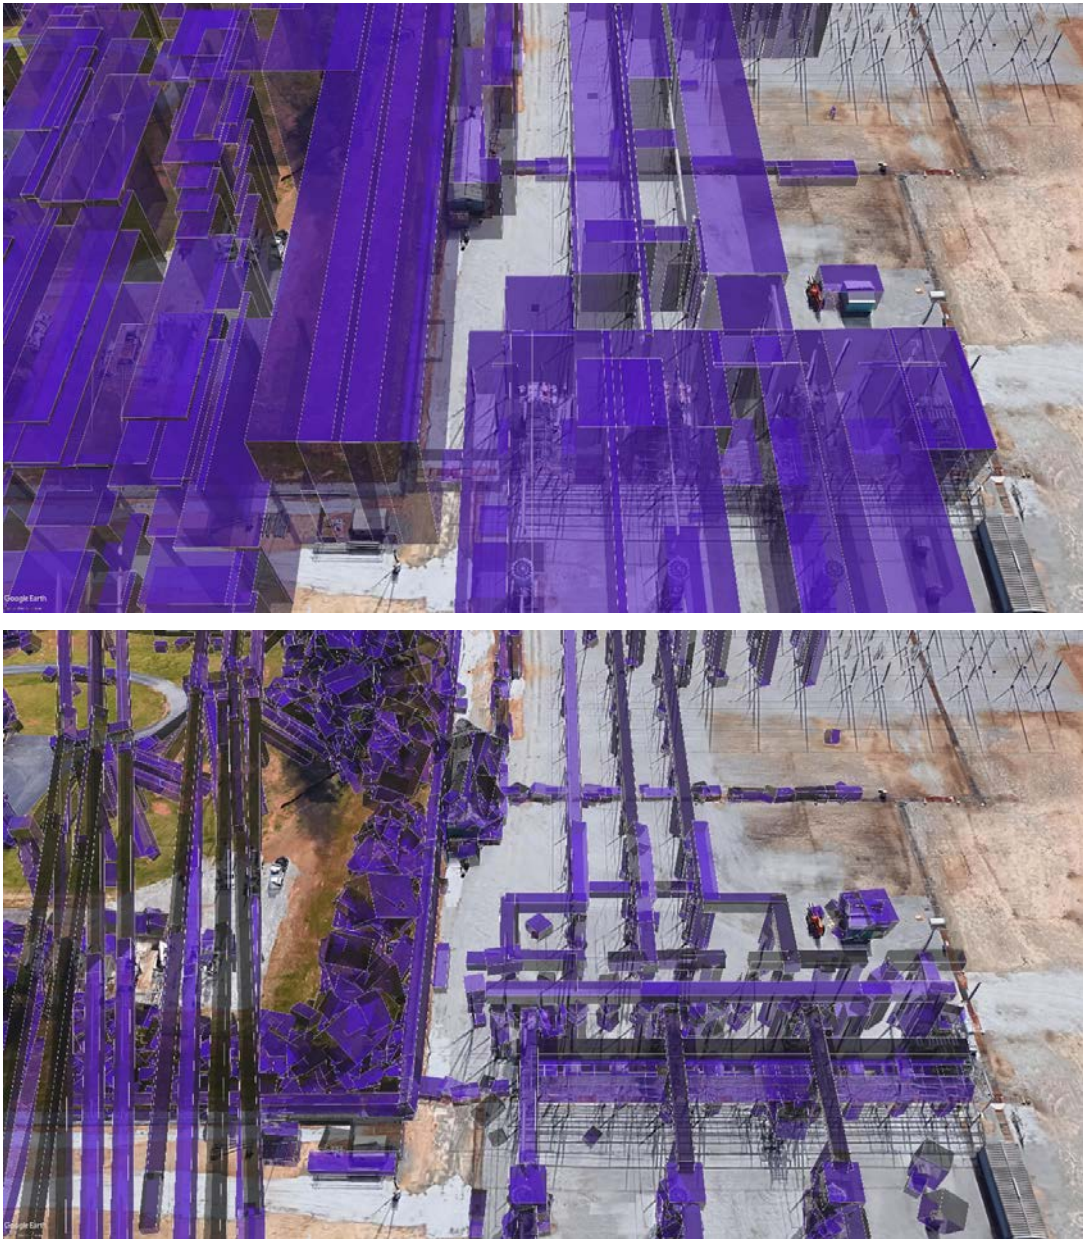
Figure 5. Close up of bounding polyhedra for the heterogeneous North Quarter section of the Southern Company Klondike test facility. Top: 2.5D Maximum Elevation Box method; bottom: 3D Bounding Cylinder method. © *Lidar data: Southern Company; Map data: Google & DigitalGlobe*
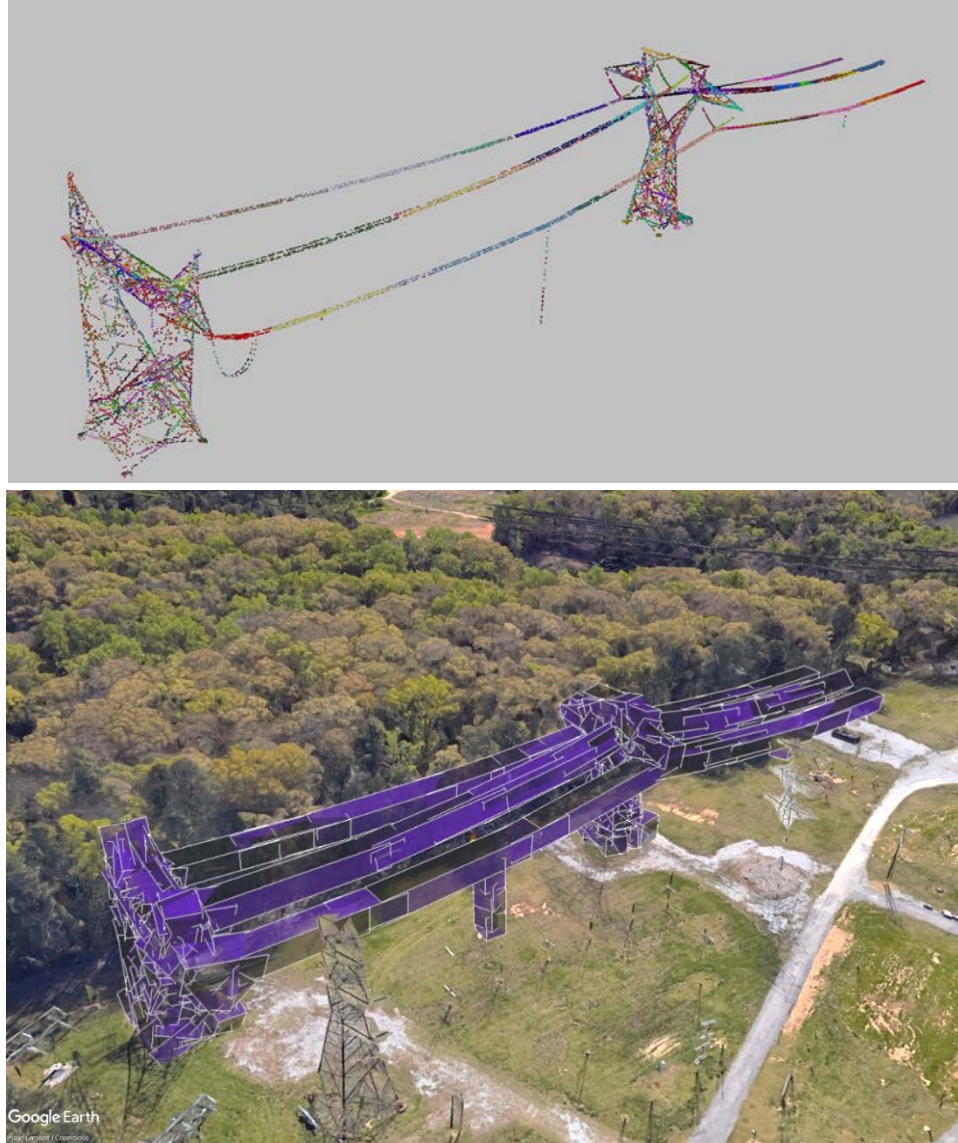
Figure 6. The internal representation of the 3D Bounding Cylinder method (top) to enforce collinearity results in segmenting of the catenary arc of conductors at bottom. The approximately 15000 lidar points from dataset B are enclosed by 370 bounding polyhedra. © *Lidar data: Dominion Energy; Map data: Google & DigitalGlobe*

Figure 7. Aerial lidar scans fail to adequately sample fine structures such as guy wires. The 3D Bounding Cylinder method attempts to bridge gaps in sampling by projecting 1 meter along a given cylindrical eigenaxis to search for points in neighborhood construction. If there are no collinear points within a meter, a guy wire with drops in scan sampling is represented as more than one neighborhood (left), yielding multiple bounding polyhedra for the span (right). © *Lidar data: Dominion Energy; Map data: Google & DigitalGlobe*
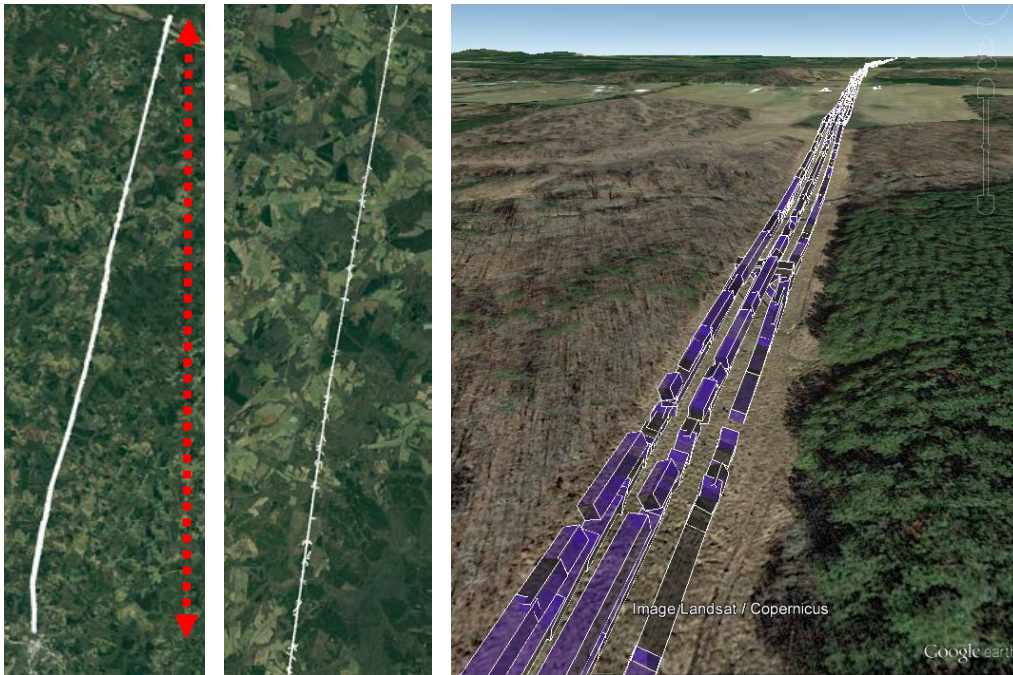

Figure 8. Enclosure of a 29 mile transmission line corridor from dataset C with the 3D Bounding Cylinder method. Left: entire set of polyhedra; red arrow is 27 miles long. Middle: detail of 11 mile flight region. Right: low altitude view. The cylindrical model parameters used to created bounding polyhedra in this report were optimized for datasets with finer sampling (~0.3m). At the relatively low point spacing (0.6m average) of this dataset, lidar sampling dropouts result in gaps in the enclosure geometry, but fidelity is otherwise excellent. 2.2M lidar points are compressed to 71K polyhedra in this example. © *Lidar data: Dominion Energy; Map data: Google, Landsat/Copernicus & DigitalGlobe*

Dataset C has a sampling (point spacing) twice as coarse as the other transmission line survey datasets used in this study, but spans the entire length of a 29 mile transmission line. Extensive manual inspection and data preparation are needed to produce survey sets such as this one, which are used monitor foliage [32], since lidar

points along the transmission corridor must be relabeled by object type (i.e., foliage, conductor, ground). An eleven mile section of this corridor selected for a pilot UAV inspection overflight [33] was used to test the data capacity and enclosure fidelity of the 3D Bounding Cylinder method. The results (Figure 8) exhibit the features of the method described above: high data compression (2.2M lidar points enclosed by 71K polyhedra), efficient conductor rendering, and scattered dropouts from lidar sampling errors.

B. *Enclosure Style and Fidelity for 2.5D Ground Extrusion Method*

The 2.5D Ground Map Extrusion method takes advantage of existing geographic information system (GIS) data to simplify vegetation in lidar survey scans. A dedicated GIS group at NASA Langley Research Center maintains a database of key geographical features at the Center, and continually updates its breadth, accuracy, and resolution as new survey technologies become available. Major permanent structures such as buildings and roads are represented with sub-meter accuracy sets of ground boundary (footprint) coordinates and associated elevations in commercial GIS software. Trees, signage and other minor structures are entered on an *ad hoc* basis as graphical footprint drawings and a single elevation.

To build an obstacle map of a flight range on North Dryden Boulevard, Hampton, Virginia, building footprints and elevations were combined with polyhedra generated by the 2.5D Ground Map Extrusion method on foliage in two sets of aerial lidar scans. Results from the dataset D (with 0.2m average point spacing) are shown at left in Figure 9, with buildings colored blue, light poles colored yellow, and foliage (bounded with a best-fit rectangle) colored green. After software to digitize the ground footprints was developed, the method was applied to a second, newer lidar survey (dataset E) with 0.1m average point spacing. As shown at the right in Figure 9, the method bounds most of lidar points at locations where footprints exist.

Limitations of this method are also evident in Figure 9. The foliage footprint outlines are quite close to the actual footprint, but since they are created manually, and since structures (especially trees) change over time, they are not exact. Trees without a footprint entry in the GIS database are not captured at all (yellow arrow). Thin, vertically oriented structures such as street signs and light posts are poorly sampled by an aerial scan, so that the computed polyhedron height is often inaccurate (red arrow). The data compression produced by the 2.5D Ground Map Extrusion method cannot be computed readily, as it incorporates a mixture of building geometry which is already simplified and misses the foliage areas that have no drawn ground footprint. Qualitatively, compression appears to be much higher than the 3D Bounding Cylinder method, due to the reduction of foliage data volume, and somewhat lower that the 2.5D Maximum Elevation Box method, since its base polyhedra have many more vertices than the rectangular bases of the 2.5D Maximum Elevation Box method.
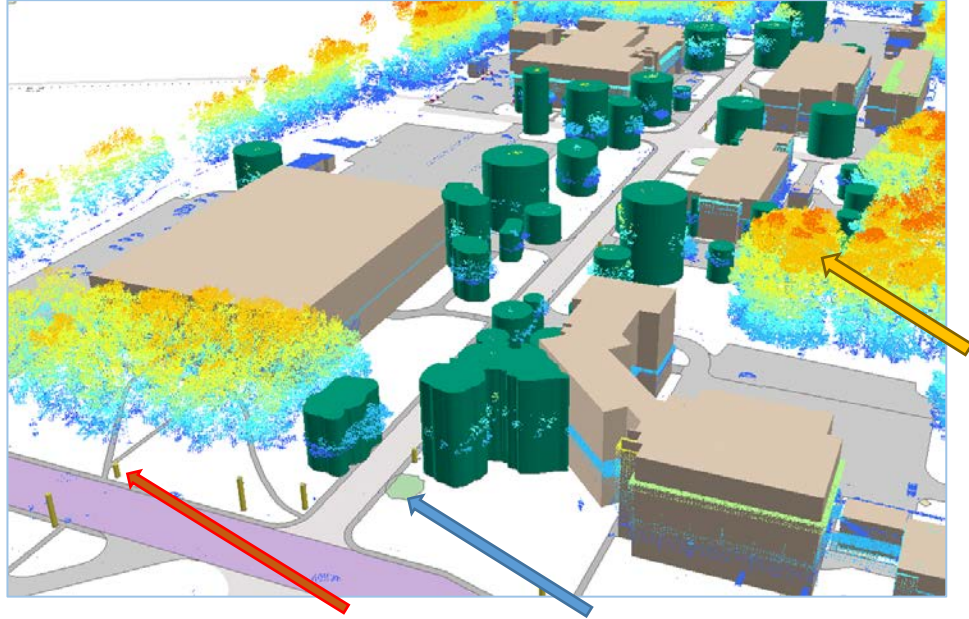
Figure 9. Obstacle field created with the 2.5D Ground Map Extrusion method maps for a flight corridor at NASA Langley Research Center. Building footprints with sub-meter accuracy are taken from the LaRC GIS database and extruded to the corresponding elevation. Extruded buildings (brown) are combined with computed boundary results from lidar dataset E. Raw lidar, after ground point removal, is rendered with a color gradient according to elevation (blue-to-red dots). The bounding polyhedra for trees (green) and light poles (yellow) are computed with the 2.5D Ground Map Extrusion method. Arrows point to exemplars of the method's limitations: footprints are not drawn for some groups of trees (yellow arrow at right), some trees have been cut down since the footprints were created (blue arrow at bottom), and lidar sampling of thin light pole structures produces height errors (red arrow at bottom).

C.   *Preliminary Results for 3D Bounding Box Method*

The 3D Bounding Box Filter method was developed last in this research project. Merging of box neighborhoods is not yet implemented, so compression statistics are not available; the results for this method are discussed qualitatively in this section.

A portion of the heterogeneous North Quarter section of the Southern Company Klondike test facility from dataset A with trees, poles and buildings is shown in Figure 10 to compare the 3D Bounding Cylinder method and 3D Bounding Box method. At top left is the lidar point cloud with ground points removed; at top right, the lidar is shown again with annotated regions of interest; at bottom left, the internal representation of the 3D Bounding Cylinder neighborhoods in the Point Cloud Library tool is shown; and at bottom right, the internal representation of the 3D Bounding Box neighborhoods in the Point Cloud Library tool is shown.

The 3D Bounding Cylinder method excels at clustering lidar points of power lines and poles. This can be seen at the upper right of the two images at the bottom of Figure 10. Most of the conductor spans are rendered with a single color in the 3D Bounding Cylinder method results, representing a single neighborhood cluster. In contrast, the spans are comprised of multiple neighborhood clusters in the 3D Bounding Box method results.

The box filter method excels at clustering planar lidar points of rooftops. This can be seen in the foreground of the two images at the bottom of Figure 10. The 3D Bounding Cylinder method 'paints stripes' across the rooftops, while the 3D Bounding Box method tiles the roofs with large rectangles. With the addition of a merge step, this method appears capable of representing each roof face with a single neighborhood.

Neither method copes well with the complex geometry of foliage [23][24]. A small area of the same flight range shown in Fig. 9 is used to illustrate this point. As shown in Figure 11, the internal neighborhood representation for foliage for both methods is characterized by topological incoherence and spatial

inhomogeneity. Foliage aside, this example again illustrates the comparative strength of the 3D Bounding Box method for computing compact 3D rooftop boundaries.
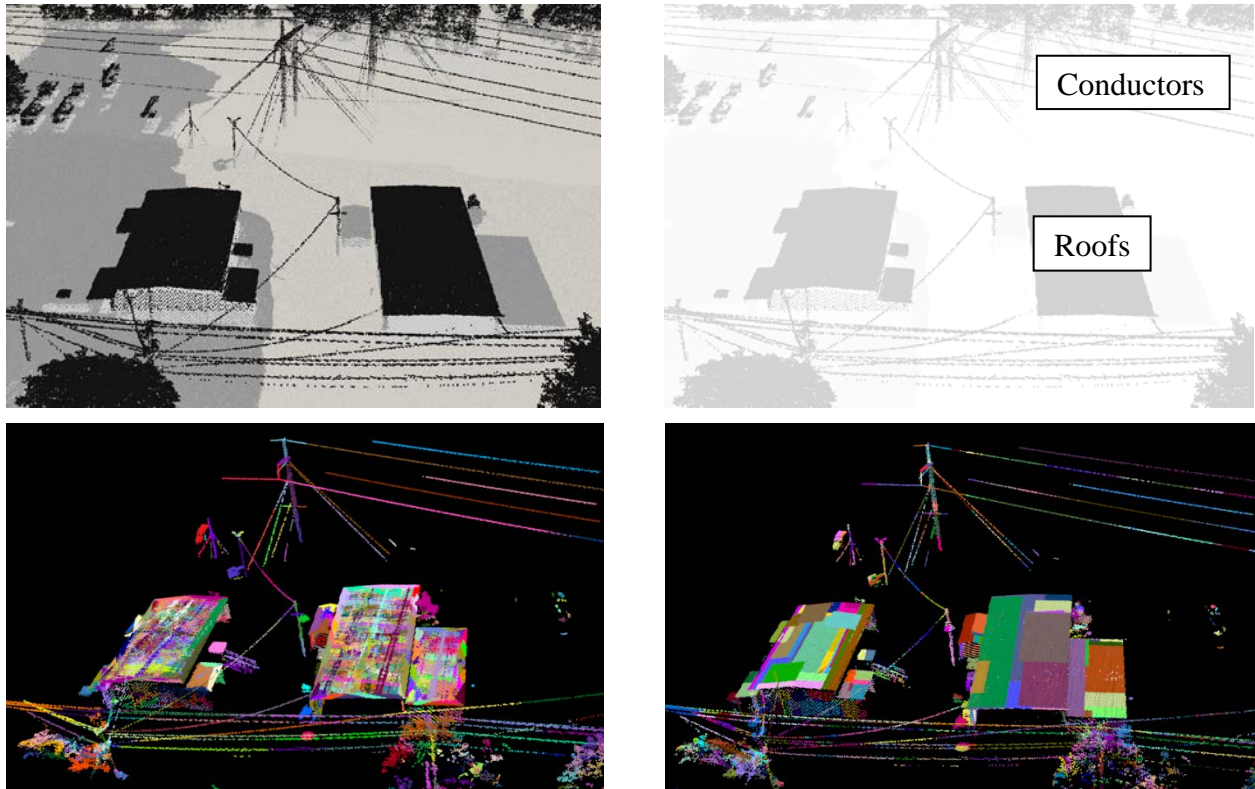

Figure 10. A comparison of the 3D polyhedron bounding methods in an area from dataset A with trees, poles and buildings. Top: input and annotated lidar. Bottom: internal neighborhood representation in PCL for the 3D Bounding Cylinder method (left) and 3D Bounding Box method (right). The 3D Bounding Cylinder method is more effective with linear objects, while the 3D Bounding Box method is more effective with planar objects. *© Lidar data: Southern Company; Map data: Google & DigitalGlobe*
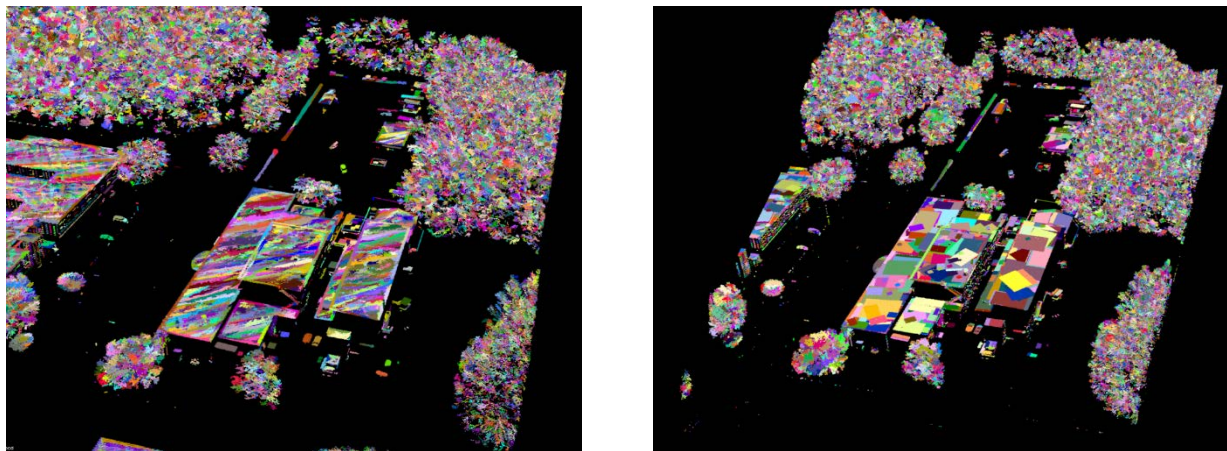

Figure 11. A comparison of the 3D polyhedron bounding methods for a NASA Langley Research Center flight range using the lidar survey from dataset D. The poor performance of both methods for foliage is evident.

IV. Discussion

Summary of results. The results for the four methods primarily show the inverse relationship between spatial fidelity and data compression. The 2.5D methods can construct ground structure boundaries that are compact but spatially approximate, while the boundary geometry constructed by the 3D methods is rather less compact (Table 1) but more tightly encloses the actual boundaries.

Of the 2.5D methods, the Ground Map Extrusion technique strikes the better balance between compactness and fidelity, due to the human judgment applied in drawing the ground footprints used for extrusion. The 3D methods are each superb for one geometric class (linear structures for the 3D Bounding Cylinder method and planar structures for the 3D Bounding Box method). Foliage and complex structures such as lattice towers are handled poorly by both 3D methods.

Figure 4 and 5 summarize well the tradeoff between obstacle boundary fidelity and compactness. The 2.5D Maximum Elevation Box method (middle row of Fig. 4, top of Fig. 5) encloses a great deal of empty space, particularly below elevated structures. The 3D Bounding Cylinder method (bottom row of Fig. 4, bottom of Fig. 5) encloses almost no empty space, at the cost of poor compactness in complex areas such as the trees and the trusswork of the steel lattice towers. Table 1 quantifies this cost for two sections of a survey with 0.3 meter average point spacing: while the 2.5D Maximum Elevation Box method compresses the lidar points by 3-4 orders of magnitude, the 3D Bounding Cylinder method compresses them by 1.5-2 orders of magnitude. The 2.5D Ground Map Extrusion method compression is slightly lower compared to the 2.5D Maximum Elevation Box method, because its base polyhedra generally have many more vertices than the rectangular bases of the 2.5D Maximum Elevation Box method. The 3D Bounding Box method (once a final merge step is implemented) is expected to yield data compression ratios similar to the 3D Bounding Cylinder method, with better compression on planar structures than linear structures.

Optimizing compactness. There is no single bounding method which optimizes the exactness/compactness tradeoff. As mentioned in the Introduction, in some studies an ensemble of methods is applied in series. For
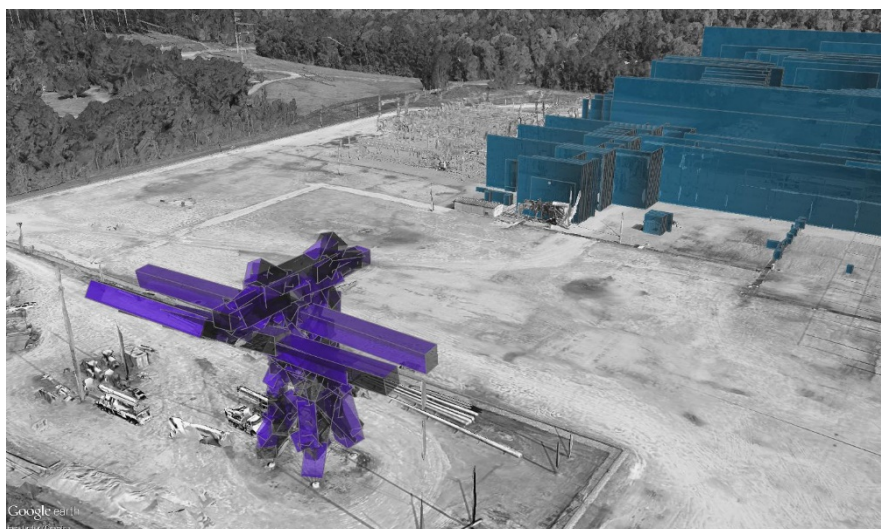

Figure 13. Mixed polyhedra representation for UAV-based inspection of a 500kV transmission line tower and conductors. The 500 kV Section is rendered with the 3D Bounding Cylinder method (foreground, purple) while the North Quarter, of less immediate concern as an obstacle, is rendered with the 2.5D Maximum Elevation Box method (background, blue). © Lidar data: Southern Company; Map data: Google & DigitalGlobe

example, a pipeline could constructed which finds and bounds 3D planar areas (e.g., buildings) and then removes the bounded planar lidar points first, finds and bounds 3D linear structures (e.g., power lines) and then removes the bounded linear lidar points second, and encloses the remaining lidar points (presumably foliage) in a convex hull or surface spline. Another approach is suggested in a study which addressed compactness for compute-bound onboard aerial navigation (Figure 13, from [4]). The capacity and throughput of onboard processors, while

improving, is lower than desktop or laptop processors by at least an order of magnitude. The two portions of dataset A compared in Table 1 were combined to produce a reference obstacle field in a detailed compute capacity analysis for a version of the ICAROUS (Independent Configurable Architecture for Reliable Operations of Unmanned Systems) flight path conformance software [31][35][36] running on a single-threaded processor [4]. That analysis found that for conservative UAV velocities, the obstacle field geometry should not exceed about 200 polyhedra. While more powerful onboard processors and more optimized software have become available since that study, it is safe to say that the volume of raw lidar data (at over 2.3 million points for this reference obstacle field) still exceeds the capability of onboard obstacle avoidance autonomy. At the time of the analysis, only the 2.5D Maximum Elevation Box method reduced the data volume sufficiently to avoid overwhelming the onboard processor. A hybrid geometric representation was proposed (Figure 13) to manage the data volume without loss of fidelity. In this hybrid approach, 3D bounding polyhedra are computed for structures near the current UAV position (foreground, purple) and 2.5D bounding polyhedra for more distant structures (background, blue).

Applications to UAV autonomy. Widely available geo-conformance monitors, such as the one embedded in the ArduCopter [37] controller for the Pixhawk [38] autopilot, and more sophisticated monitors such as Safeguard [39][40] enforce UAV containment within a prescribed flight corridor that is specified as a 2.5D geofence. (See [10] for research flights that used both of these monitors.) The bounding polyhedra produced by the 2.5D Maximum Elevation Box method can serve as the geofence source for both ArduCopter and Safeguard. Geofence geometries created by the 2.5D Ground Map Extrusion method such as those shown in Figure 2 will exceed ArduCopter's vertex capacity, however, while the high capacity geometry processing in the current Safeguard technology can easily cope with its higher vertex count.

The 3D bounding methods are particularly useful for onboard autonomy such as ICAROUS, which continually computes flight paths between planned waypoints that avoid obstacles, such as geofences and 'safety bubbles' around other vehicles in the airspace. Figure 14 shows an example of the internal representation that ICAROUS computes (adopted from [36]) and a typical flight path [10] with ICAROUS constrained to plan 2D avoidance maneuvers. With this or similar autonomy onboard a UAV, assuming reliable GPS reception [5] and precomputed bounding geometry as described in this report, inspection flights underneath structures such as bridges and powerlines could be conducted with added safety.
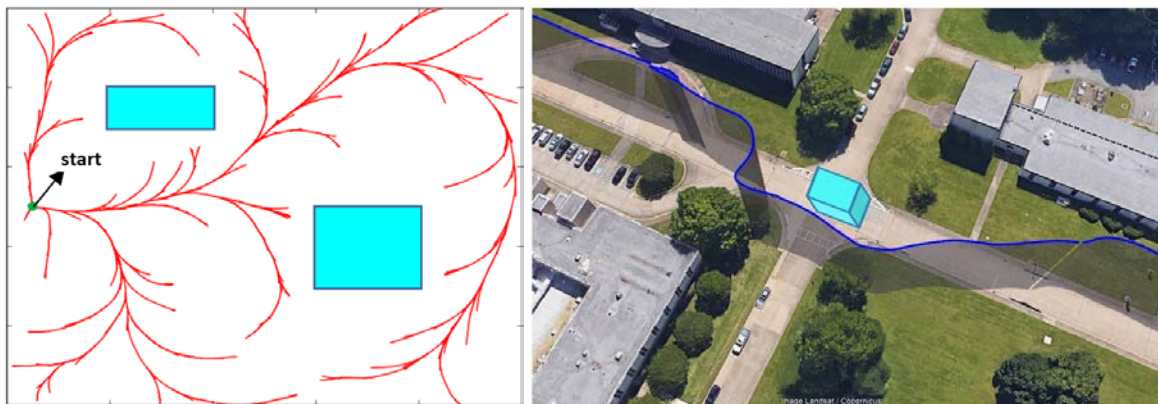
Figure 14. Internal representation (left) and flight results (right) of ICAROUS path conformance autonomy, adopted from [36] and [10], respectively. Geofences are colored cyan in both images. Potential flight paths are shown in red at left. Actual flight path is shown in blue at right. This example illustrates contingency flight paths computed for 2.5D geofences, though the technology can compute and execute 3D paths as well. © *Map data: Google & Landsat/Copernicus*

Emerging satellite survey technology. This report details methods to create bounding polyhedra for survey data gathered using aerial lidar scans. At the time that this research began, surveys using satellite photogrammetry [41] were not widely available with precision (average point spacing) and accuracy [42] comparable to surveys using lidar. Survey data from satellite photogrammetry is now readily obtainable at low cost, in formats that the methods described herein can process. The common figure of merit for satellite

photogrammetry, LE90, is a statistical estimate of errors in accuracy at the ninety percent level [43]; for commodity satellite-based surveys the LE90 distance is 3m or more [44], about ten times higher than current lidar surveys. The spatial sampling, a measure of precision, is coarser than that of lidar (about 50cm for satellite photogrammetry *vs.* 10cm for high fidelity lidar data such as in dataset E), so that gaps in bounding polyhedra (Figure 7, Figure 8) that arise from data dropouts will be more severe than for lidar.

The research described in this report demonstrated, on a diverse set of lidar survey data, that both sub-meter sampling and sub-meter accuracy are required to resolve and map power lines, guy wires and small structures such as road signs. The current precision and accuracy of satellite photogrammetry are therefore suitable only for large obstacles such as buildings and trees. To ensure low-altitude UAV collision avoidance, satellite-based survey data must be supplemented with lidar survey data to safely map all collision hazards.

V.   Conclusion

This report described in detail four methods to compute bounding geometries of ground obstacles from lidar point clouds. The four methods are: 1) 2.5D Maximum Elevation Box, 2) 2.5D Ground Map Extrusion, 3) 3D Bounding Cylinder, and 4) 3D Bounding Box. Custom software was used for the 2.5D methods. For the 3D methods, the open source Point Cloud Library software was used to find clusters of point cloud data, and custom software tested clusters according to tunable geometric criteria, purged outliers, merged clusters, and constructed polyhedra from the resultant point clouds.

The methods were evaluated with heterogeneous survey data including buildings, vegetation, power lines, and small structures such as road signs and guy wires. The level of detail of the test data ranged from 0.1m to 0.6m.

The performance of each method was described using areas from the lidar survey that highlight strengths and weakness of each, and the methods were compared on the bases of boundary fidelity and data compression. The 2.5D Maximum Elevation Box method reduces the data volume the most, as measured by the ratio of the input lidar point count and the output polyhedral vertex count. However, the bounding polyhedra it produces enclose a relatively high volume of empty space. The 2.5D Ground Map Extrusion method bounding fidelity is higher and its compression is slightly lower, because the 2D polygons which are extruded generally have more vertices than the four rectangle corners characteristic of the 2.5D Maximum Elevation Box method. The extruded geometry is human drawn, and so this is not a fully automated method.

Both 3D methods have high spatial fidelity, and enclose very little empty space. Each 3D method is most effective on a distinct geometry type. The 3D Bounding Cylinder method excels at enclosing suspended cables such as power lines and linear structures such as bus bars and fences. The 3D Bounding Box method excels at enclosing planar structures such as buildings. The high fidelity comes at the cost of a compression ratio that is ~30 times lower than the 2.5D Maximum Elevation Box method. Another consequence of their geometric selectivity is that the geometry count is high on vegetation and structures with a high degree of angular variation such as lattice towers. The 3D Bounding Cylinder method is particularly susceptible to sampling dropouts for structures much thinner than the average point spacing, which produce gaps in the output polyhedra.

Both the 2.5D and 3D methods are useful for low altitude UAV navigation. Geo-containment autonomy is typically based on 2.5D keep-in geofences. Flight path conformance autonomy is capable of planning avoidance paths both in 2D and 3D. The compute power of onboard processors is relatively limited, however, and so a mix of 2.5D and 3D obstacle representation may be needed to stay within the limits of airborne processor throughput and capacity.

Satellite-based photogrammetry data was not processed in this report, but its spatial accuracy and resolution was evaluated in the light of these results. While suitable for low resolution applications such as forestry surveys and terrain mapping, it is not currently capable of resolving all of the obstacles needed for safe UAV navigation at low altitudes.

VI.   References

1. Wehr, A. and Lohr, U., 1999. Airborne laser scanning—an introduction and overview. ISPRS Journal of photogrammetry and remote sensing, 54(2-3), pp.68-82.
2. Chida, Akisato, and Hiroshi Masuda. "Reconstruction of polygonal prisms from point-clouds of engineering facilities." Journal of Computational Design and Engineering 3.4 (2016): 322-329.
3. Doyle, David Raymond. Development of the National Spatial Reference System. US Department of Commerce, National Oceanic and Atmospheric Administration, National Ocean Service, 1995.
4. Moore, Andrew J., Matthew Schubert, Nicholas Rymer, Swee Balachandran, Maria Consiglio, Cesar Munoz, Joshua Smith, Dexter Lewis, and Paul Schneider. "UAV Inspection of Electrical Transmission Infrastructure with Path Conformance Autonomy and Lidar-based Geofences NASA Report on UTM Reference Mission Flights at Southern Company Flights November 2016." NASA Technical Memo 2017-219673 (2017). – 2017b
5. Moore, A. J., Schubert, M. and Rymer, N., 2017. Autonomous Inspection of Electrical Transmission Structures with Airborne UV Sensors-NASA Report on Dominion Virginia Power Flights of November 2016. NASA Technical Memo 2017-219611 – 2017a
6. Rymer, N., Moore, A .J. and Schubert, M., 2018, June. Inexpensive, lightweight method of detecting coronas with UAVs. In 2018 International Conference on Unmanned Aircraft Systems (ICUAS) (pp. 452-457). IEEE.
7. Moore, A., Schubert, M. and Rymer, N., 2018. Autonomous Inspection of Electrical Transmission Structures with Airborne UV Sensors and Automated Air Traffic Management. In 2018 AIAA Information Systems-AIAA Infotech@ Aerospace (p. 1628). – 2018a
8. Moore, A. J., Schubert, M. and Rymer, N., 2018, August. Technologies and Operations for High Voltage Corona Detection with UAVs. In 2018 IEEE Power & Energy Society General Meeting (PESGM) (pp. 1-5). IEEE. – 2018b
9. Moore, A. J., Schubert, M., Rymer, N., Balachandran, S., Consiglio, M., Munoz, C., Smith, J., Lewis, D. and Schneider, P., 2018b, February. Inspection of electrical transmission structures with UAV path conformance and lidar-based geofences. In 2018 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT) (pp. 1-5). IEEE. – 2018c
10. Moore, A., Balachandran, S., Young, S.D., Dill, E.T., Logan, M.J., Glaab, L.J., Munoz, C. and Consiglio, M., 2018b. Testing enabling technologies for safe UAS urban operations. In 2018 AIAA Aviation Technology, Integration, and Operations Conference (p. 3200). – 2018d
11. Lefsky, M.A., Cohen, W.B., Parker, G.G. and Harding, D.J., 2002. Lidar remote sensing for ecosystem studies: Lidar, an emerging remote sensing technology that directly measures the three-dimensional distribution of plant canopies, can accurately estimate vegetation structural attributes and should be of particular interest to forest, landscape, and global ecologists. BioScience, 52(1), pp.19-30.
12. Vosselman, George, et al. "Recognising structure in laser scanner point clouds." International archives of photogrammetry, remote sensing and spatial information sciences 46.8 (2004): 33-38.
13. Kwon, Soon-Wook, et al. "Fitting range data to primitives for rapid local 3D modeling using sparse range point clouds." Automation in construction 13.1 (2004): 67-81.
14. Franklin, J., 1995. Predictive vegetation mapping: geographic modelling of biospatial patterns in relation to environmental gradients. Progress in physical geography, 19(4), pp.474-499.
15. Molenaar, Martien. "A formal data structure for three-dimensional vector maps." Proc. EGIS'90 Amsterdam Vol. 2 (1990) 770-781. Ook: Proc. Commission III ISPRS, Wuhan, PR China (1990) 535-550. Ook: Proc. 4th Int. Symp. Spatial data handling, Zürich, Switzerland Vol. 2. 1990.
16. Demantke, J., Mallet, C., David, N. and Vallet, B., 2011. Dimensionality based scale selection in 3D lidar point clouds. Int. Arch. Photogramm. Remote Sens. Spat. Inf. Sci, 38(5), p.W12.
17. Blomley, R., Jutzi, B. and Weinmann, M., 2016. 3D semantic labeling of ALS point clouds by exploiting multi-scale, multi-type neighborhoods for feature extraction. In: Proceedings of the International Conference on Geographic Object-Based Image Analysis, Enschede, The Netherlands, pp. 1–8.
18. Bentley, L.P., Stegen, J.C., Savage, V.M., Smith, D.D., Sperry, J.S., Reich, P.B. and Enquist, B.J., 2013. An empirical assessment of tree branching networks and implications for plant allometric scaling models. Ecology letters, 16(8), pp.1069-1078.
19. Filin, S. and Pfeifer, N., 2005. Neighborhood systems for airborne laser data. Photogrammetric Engineering & Remote Sensing, 71(6), pp. 743–755.

20. Melzer, Thomas, and Christian Briese. Extraction and modeling of power lines from ALS point clouds. 2004. AAPR. Proceedings of the 28thWorkshop of the Austrian Association for Pattern Recognition, Hagenberg, Austria, June17r18, 47, p.r54.
21. Ippolito, Corey, Kalmanje Krishnakumar, and Sebastian Hening. "Preliminary results of powerline reconstruction from airborne LiDAR for safe autonomous low-altitude urban operations of small UAS." SENSORS, 2016 IEEE. IEEE, 2016.
22. Matikainen, Leena, et al. "Remote sensing methods for power line corridor surveys." ISPRS Journal of Photogrammetry and Remote Sensing 119 (2016): 10-31.
23. Côté, Jean-François, Jean-Luc Widlowski, Richard A. Fournier, and Michel M. Verstraete. "The structural and radiative consistency of three-dimensional tree reconstructions from terrestrial lidar." Remote Sensing of Environment 113, no. 5 (2009): 1067-1081.
24. Prusinkiewicz, Przemyslaw, and Aristid Lindenmayer. The algorithmic beauty of plants. Springer Science & Business Media, 2012.
25. Xu, S., Vosselman, G. and Elberink, S.O., 2014. Multiple-entity based classification of airborne laser scanning data in urban areas. ISPRS Journal of photogrammetry and remote sensing, 88, pp.1-15.
26. Tymkow, P., Borkowski, A. and Norwida, C.K., 2010. Vegetation modelling based on TLS data for roughness coefficient estimation in river Valley. International Archives of the Photogrammetry, Remote Sensing and Spatial Information Science, 38(8), pp.309-313.
27. Mitasova, H., Mitas, L. and Harmon, R.S., 2005. Simultaneous spline approximation and topographic analysis for lidar elevation data in open-source GIS. IEEE Geoscience and Remote Sensing Letters, 2(4), pp.375-379.
28. Pearse, G.D., Dash, J.P., Persson, H.J. and Watt, M.S., 2018. Comparison of high-density LiDAR and satellite photogrammetry for forest inventory. ISPRS journal of photogrammetry and remote sensing, 142, pp.257-267.
29. Nex, Francesco, and Fabio Remondino. "UAV for 3D mapping applications: a review." *Applied geomatics* 6, no. 1 (2014): 1-15.
30. Tridgell, A. and Dade, S., 2016. Mavproxy: a UAV ground station software package for MAVLink based systems. Ardupilot Project.
31. Narkawicz, Anthony, and George E. Hagen. "Algorithms for collision detection between a point and a moving polygon, with applications to aircraft weather avoidance." In *16th AIAA Aviation Technology, Integration, and Operations Conference*, p. 3598. 2016.
32. Wolf, Gene. "LiDAR Meets NERC Alert." Transmission & Distribution World 63.8 (2011): 9-13.
33. Lillian, Betsy. "Virginia Team Testing out BVLOS Drones for Power Line Inspections." (2017, April 10). Retrieved from https://unmanned-aerial.com/virginia-team-testing-bvlos-drones-power-line-inspections
34. Skiljan, I., 2012. IrfanView. URL http://irfanview. tuwien. ac. at/. Retrieved June.
35. Consiglio, María, et al. "ICAROUS: Integrated configurable algorithms for reliable operations of unmanned systems." Digital Avionics Systems Conference (DASC), 2016 IEEE/AIAA 35th. IEEE, 2016.
36. Balachandran, S., Narkawicz, A., Muñoz, C. and Consiglio, M., 2017. A path planning algorithm to enable well-clear low altitude UAS operation beyond visual line of sight. In Twelfth USA/Europe Air Traffic Management Research and Development Seminar (ATM2017).
37. ArduPilot Open Source Autopilot. http://ardupilot.org/. Accessed:2018-02-25
38. Meier, L., Tanskanen, P., Fraundorfer, F. and Pollefeys, M., 2011, May. Pixhawk: A system for autonomous flight using onboard computer vision. In 2011 IEEE International Conference on Robotics and Automation (pp. 2992-2997). IEEE.
39. E. T. Dill, S. D. Young and K. J. Hayhurst, "SAFEGUARD: An assured safety net technology for UAS," 2016 IEEE/AIAA 35th Digital Avionics Systems Conference (DASC), Sacramento, CA, 2016, pp. 1-10.
40. R. V. Gilabert, E. T. Dill, K. J. Hayhurst and S. D. Young, "SAFEGUARD: Progress and test results for a reliable independent on-board safety net for UAS," 2017 IEEE/AIAA 36th Digital Avionics Systems Conference (DASC), St. Petersburg, FL, 2017, pp. 1-9.
41. Aguilar, M.A., Aguilar, F.J., Agüera, F. and Sánchez, J.A., 2007. Geometric accuracy assessment of QuickBird basic imagery using different operational approaches. Photogrammetric Engineering & Remote Sensing, 73(12), pp.1321-1332.
42. Aguilar, M.A., del Mar Saldana, M. and Aguilar, F.J., 2013. Assessing geometric accuracy of the orthorectification process from GeoEye-1 and WorldView-2 panchromatic images. International Journal of Applied Earth Observation and Geoinformation, 21, pp.427-435.
43. ASPRS, 2014. "ASPRS Positional Accuracy Standards for Digital Geospatial Data." Photogrammetric Engineering & Remote Sensing, 81, pp. A1–A26.

44. Vricon point cloud. https://www.vricon.com. Accessed 2019-7-10

REPORT DOCUMENTATION PAGE

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE		3. DATES COVERED *(From - To)*
1-09-2019	Technical Memorandum		

4. TITLE AND SUBTITLE

Bounding Methods for Heterogeneous Lidar-derived Navigational Geofences

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Moore, Andrew J.; Schubert, Matthew; Fang, Terry; Smith, Joshua; Rymer, Nicholas

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER
334005.04.10.07.01

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

NASA Langley Research Center
Hampton, VA 23681-2199

8. PERFORMING ORGANIZATION REPORT NUMBER

L-21053

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

National Aeronautics and Space Administration
Washington, DC 20546-0001

10. SPONSOR/MONITOR'S ACRONYM(S)

NASA

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

NASA-TM-2019-220399

12. DISTRIBUTION/AVAILABILITY STATEMENT

Unclassified-
Subject Category 01
Availability: NASA STI Program (757) 864-9658

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Safe Unmanned Aerial Vehicle (UAV) operations near the ground require navigation methods that avoid fixed obstacles such as buildings, power lines and trees. Aerial lidar surveys of ground structures are available with the precision and accuracy to geolocate obstacles, but the high volume of raw survey data can exceed the compute power of onboard processors and the rendering ability of ground-based flight planning maps. Representing ground structures with bounding polyhedra instead of point clouds greatly reduces the data size and can enable effective obstacle avoidance, as long as the bounding geometry envelopes the structures with high spatial fidelity. This report describes in detail four methods to compute bounding geometries of ground obstacles from lidar point clouds. The four methods are: 1) 2.5D Maximum Elevation Box, 2) 2.5D Ground Map Extrusion, 3) 3D Bounding Cylinder, and 4) 3D Bounding Box.

15. SUBJECT TERMS

UAV; Geofence; Lidar

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			STI Help Desk (email: help@sti.nasa.gov)
U	U	U	UU	26	19b. TELEPHONE NUMBER *(Include area code)* (757) 864-9658

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18