Automated Building Corner Detection for Validating 3D Point Cloud Data

Amy Neuenschwander, Jeff Perry, and Lori Magruder
University of Texas at Austin
Austin, Texas
amy@csr.utexas.edu, jsp@utexas.edu, lori.magruder@austin.utexas.edu

ABSTRACT

The utilization of high-resolution 3D point cloud data is becoming more common to a variety of Department of Defense (DoD) applications and many relevant data are collected via small UAVs to satisfy the growing need. There are a number of factors that can affect the quality of geospatial data, including the accuracy of the measurements, the precision of the instrument/sensor used to collect the data, and the methodologies used to process and analyze the data. Warfighters must be aware of these factors and take steps to ensure that the data they are using is of the highest quality. Although UAV data are common, they do not always contain information regarding their geolocation uncertainties. Determination of the geolocation accuracy of any 3D data set typically involves the labor-intensive process of manually extracting the coordinates of recognizable features in the point cloud as a reference for relative data quality validation. These features are often in the form of building corners. Here, we present two automated 3D building corner finding techniques and demonstrate the capability to determine the geolocation offsets (translation and rotations) of high-resolution point clouds to quantify the quality of the data. The first method develops a likelihood ratio between point density and a theoretical model. The second method fits a primitive roof model to the raw data. Both methods provide a faster, efficient way to isolate building corner points as a means to calculate geolocation accuracies or 3D registration quality. What is unique about this capability is that it works directly on the 3D point cloud rather than being image based. The importance of geospatial data quality with accurate geolocation knowledge is clear in all realms of GEOINT. By ensuring that the data uncertainties are well known, warfighters can make better decisions and improve their chances of success.

ABOUT THE AUTHORS

Amy Neuenschwander is a Senior Research Scientist at the University of Texas at Austin Center for Space Research. She is a past NASA Fellow, and has over 29 years of experience in EO/IR imagery, signal processing and LiDAR analysis. Recently, her efforts are primarily associated with point cloud manipulation and she serves as the leader of the vegetation science group for the NASA ICESat-2 mission. Dr. Neuenschwander holds a B.S. and a Master's degree in Aerospace Engineering, and a doctorate degree in Geography and the Environment from the University of Texas at Austin.

Jeff Perry is a Senior Engineering Scientist at the University of Texas at Austin Center for Space Research where he is the lead software developer. He has expertise in image analysis, machine learning, 3D data metrics, automated feature extraction, software development, and software integration. Mr. Perry holds a B.S. in Computer Science from Texas A&M University.

Lori Magruder is the Director of the University of Texas at Austin Center for Space Research and an Associate Professor in the Aerospace Engineering and Engineering Mechanics Department. She has expertise with laser diagnostics, laser ranging techniques and LiDAR experimentation efforts and she has developed many algorithmic approaches for enhanced data exploitation for dedicated scientific and operational applications from other 3D data sources. Dr. Magruder served as the science team leader for the NASA ICESat-2 mission for over 10 years and continues to develop capabilities for future space missions. Dr. Magruder received her B.S. in Aerospace Engineering from the University of Southern California, her Master's degree from Princeton University and her doctorate degree from the University of Texas at Austin.

Automated Building Corner Detection for Validating 3D Point Cloud Data

Amy Neuenschwander, Jeff Perry, and Lori Magruder
University of Texas at Austin
Austin, Texas
amy@csr.utexas.edu, jsp@utexas.edu, lori.magruder@austin.utexas.edu

INTRODUCTION

Light detection and ranging, or LiDAR, technology, a standard for high resolution 3-dimensional geospatial data has recently been joined by electro-optical (EO) sensors that provide stereo imagery to derive a 3D data equivalent. In either case, 3D data is often the critical component along the path of strategic environmental analysis and satisfies the knowledge gap present in many tactical applications for specific geospatial (geospecific) representation. Since 3D data accessibility is growing and expanding, the focus of much research and development is centered on development of automated exploitation algorithms in support of content feature extraction, object recognition and reconstruction of geospecific structures, vegetation and terrain.

Although space-borne observations often are the primary resource for 3D geospatial applications and algorithm development, the utilization of high-resolution 3D point cloud data is becoming more common to a variety of DoD applications. These data are most effectively collected using photogrammetric sensors via small Unmanned Aerial Vehicles (UAVs). These small UAVs (e.g. Anafi Parrot or SkyDio) utilize the less precise GPS frequency signal (similar to automotive location capacity) which impacts the quality of the individual images collected. Once collected the images are input to a Structure for Motion (SfM) software package (e.g. PhotoMesh or ContextCapture) and similar points among multiple images are matched and used to derive the final 3D point cloud or mesh product. Depending upon which software package is utilized, the bundled solution for matching the photos may include local errors during the reconstruction process, but these errors are all calculated with respect to the photos and they have no true geolocation assigned to them. Often, the final point clouds from a photogrammetric UAV might conform well internally, but the entire set could be tilted, translated, or rotated from an absolute position. Thus, one use case for the approach we present here is to establish an automated way to determine the translation offsets and rotation angles necessary to align the point cloud with a reference source to extract a geolocation in a known reference frame and poised for conflation with other data sources.

Agnostic to the dimensionality of the geospatial data (2D, 3D) there is a significant importance on the quality of the intelligence. Certainly, the quality of the measurements or data products dictates the level at which the warfighter can utilize the information. Several special forces scenarios emphasize the need for high level understanding of the errors and error sources, as an example. Mission objectives that require insertion into regions of interest (ROI) at specific locations. These types of operations dictate accurate knowledge of building locations and areas of vegetation but also include building rooftop height, building rooftop extent, vegetation height and vegetation structure when considering the third dimension. Further, geometric details of vertical obstructions are critical to successful missions associated with airborne ingress or egress. Geospecific and accurate data is also crucial to correct determination of line of sight (LOS) spatial assessment, as a means of planning communications, or conceal and cover operations.

Beyond the operational needs in strategic or tactical planning with accurate 3D data is the aspect of situational awareness, change detection and safety of navigation. This is particularly applicable to needs for geospatial model updates with newly acquired data from forward deployed autonomous assets. Although the newest data provides temporally relevant assessments of the terrain, it also could dilute the quality of the overall positioning. It is important to consider data accuracies during any type of data conflation as to track the propagation of errors especially with location-based services and actionable intelligence. These data as a component of a larger ecosystem of multi-INT data rely on sufficient capture of error sources or it will become a larger issue with data usability, data transformations and data registration.

In many situations there is a distinct separation between simulation training and actual training. In the past the separation was a reasonable transition from synthetic to real-life scenarios for enhanced operational experiences. In the past, simulation leveraged existing 'geotypical' models of ROIs that gave the warfighter a sense of the landscape, terrain dynamics, and architecture but not specific dimensionality or content of the region. Regardless, geotypical training assisted in both tactical and situational response practice and acted as a proxy for approaches using multi-modal, location-based services in general. Now in the era of accurate, geospecific data availability the simulation provides a realistic environment that exactly mimics the ROI such that the warfighter can experience unlimited operational scenarios in preparation for the mission. The luxury of comprehensive and complete training in a simulated environment is often not even matched by actual training. The advantage of accurate 3D data simulations is also related to the ability to change the environment if the training needs to transition to a different season (removing leaves on deciduous trees) for accurate conceal and cover geometries or implementing synthetic damage to lines of communication (rivers, roadways, bridges, etc.) to explore possible adaptations in a given ROI.

The importance of 3D point cloud data validation

Although 3D data are common, they do not always contain information regarding their geolocation uncertainties. Determination of the geolocation accuracy of any 3D data set typically involves the labor-intensive process of an analyst manually identifying and extracting the coordinates for building corners. The method to automatically identify building corner points outlined in this paper will improve the process for data validation.

The tool that is presented here is unique in that it works directly on the 3D point cloud. Most other methods dedicated to finding corners or edges utilize 2D or raster imagery (Wang and Brady, 1995) . In those methods they often utilize a corner response function and identify pixels whose threshold exceeds that response (Park et al., 2017). Another benefit of this method is that determining the corners can be solved purely mathematically and without the need for supervised learning and training data (Huang et al., 2021; Jiang et al., 2023). Currently, the authors know of no other software package that automatically identifies corners for direct mensuration of geolocation errors and uncertainties.

The National Geospatial Intelligence Agency (NGA) has developed the Generic Point-cloud Model (GPM) that provides a standard for error uncertainties and error propagation for 3D point cloud data. The GPM defines the errors and uncertainties based on geometry from various sensors to predict the uncertainty of different 3D data products. GPM utilizes two implementations for representing the uncertainty of 3D data: sensor-space and ground-space. The sensor-space element is a model of predicted errors based on the sensor geometry and detection capabilities. The underlying sensor-space model is based on a LiDAR collection where the errors are a function of the navigation solution (GPS) and hardware uncertainties (i.e. ranging and scanning). The ground-space model is the determination of the uncertainties with respect to known reference points and is used when the sensor-space model is impractical (Rodarmel et al, 2015). The methodology outlined within this paper provides a path to determining the GPM ground-space model.

METHODOLOGY

This paper presents two automated 3D building corner finding methods that can be used to determine the XYZ coordinates of building corners as a means to quantify the geolocation offsets of high-resolution 3D point clouds using a reference point cloud with known uncertainties (e.g. reference point cloud). Although two paths to identifying building corner points are presented, both methods require the same initial data preparation processing shown in Figure 1 prior to producing the multi-dimensional offsets of the UAV 3D product later in the workflow.

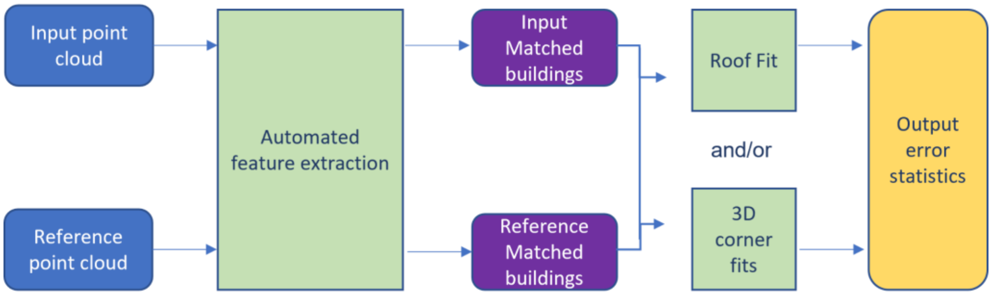

Figure 1. Flowchart of general steps used in this approach.

The first step of data preparation is building segmentation from the 3D point cloud data. Buildings and other built structures are relatively permanent on the landscape and are typically visible in both imagery and 3D point clouds making a natural choice for permanent geolocated features/objects. To segment buildings from other feature classes, building points in the 3D point cloud are identified using an automated feature extraction capability developed internally by the authors. Once identified, building points are consolidated as unique building objects based upon a proximity threshold. The proximity threshold utilized consists of flattened ellipsoidal discs (1 – 2 m semi-major axis in XY plane, 0.5 m axis in Z plane) which is used as the local search threshold criterion. With points segmented into individual buildings, utilizing a disc as a proximity threshold provides additional segmentation on the roof by separation of the actual roof surface points from structural anomalies associated with things such as HVAC units and parapets, as examples. To proceed with the workflow, further filtering of extracted buildings (as objects) occurs by removing those buildings that have too few point measurements or those that do not meet a minimum size in area (e.g.. 60 m2). An example of a 3D point cloud over a building and the segmented regions within the roof structure is shown in Figure 2. Notice in Figure 2 (right panel) that only roof points are considered in the segments and small structures are removed. With the roof points segments into unique heights, the data are ready for the next step in the building corner selection workflow.

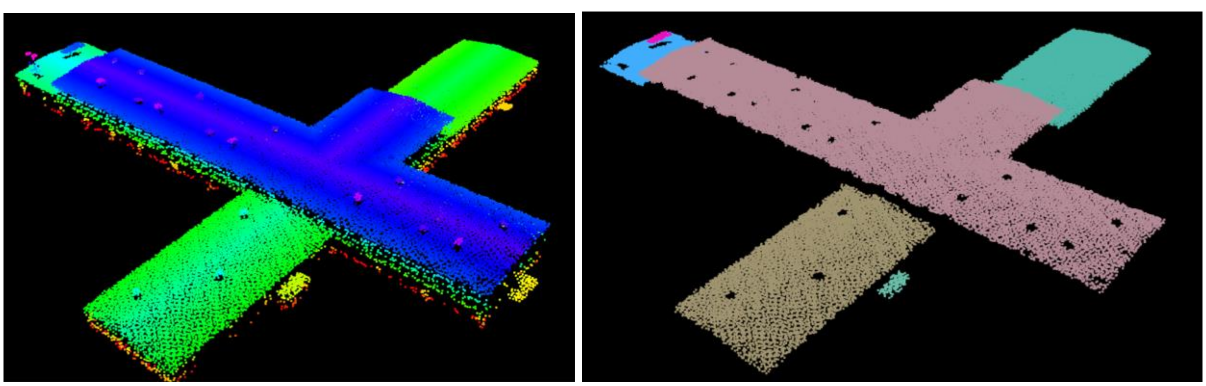

Figure 2. (Left) Raw Point Cloud of a building colored by height and (Right) segmented roof structure based on proximity threshold.

Method 1: 3D corner likelihood

The first method presented here is an innovative approach for isolating roof corner points using a mean shift algorithm. Typically implemented in 2D, the mean shift algorithm searches for the highest density (or mode) of a set of points. For the 3D case, the search radius is expanded to 3 dimensions producing points with an associated density. Under this construct, the lowest density points occur at the perimeter of a building. In Figure 3 (right panel), these low density points are shown as light blue points around the perimeter of this building roofline. Next, for each low density point, a probability distribution function (PDF) is generated from the angles formed by the point and all combinations of neighboring edge point pairs within a certain radius. This PDF is then compared to two idealized PDFs, one formed from points along a ninety-degree corner, and one formed from points along a straight edge. The response (a dot-product) of each point's angle distribution can then be used to compute a likelihood ratio by placing the response to a corner PDF in the numerator and the response to a straight edge PDF in the denominator. This likelihood ratio is greater than 1.0 when points are more likely to come from a corner and less than 1.0 otherwise, shown as red points in Figure 3b. As highlighted in Figure 3 (right panel), the edge points and corner points do not necessarily fall along a straight line due to the non-uniform sampling of the sensor. For this reason, a second methodology for identifying roof corners on buildings is also presented.

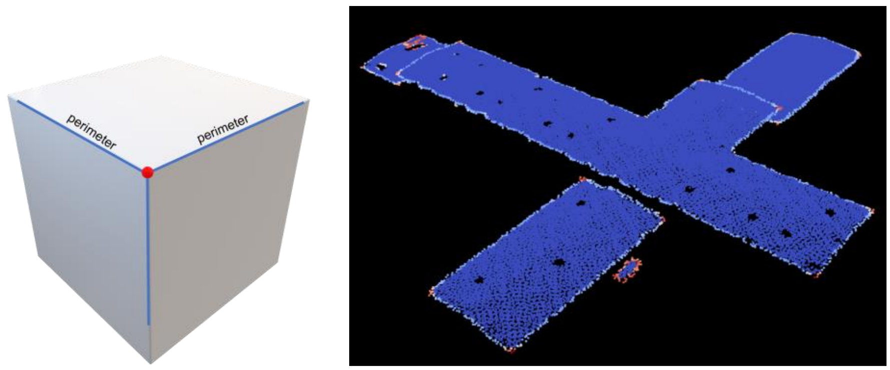

Figure 3. (Left) 3D corner structure used to develop the likelihood ratio and (Right) roof edge density likelihoods with light blue values indicating a likely roof edge and red indicating a corner point.

Method 2: Roof Fitting

A second approach to the automation of identifying building corners (and potentially roof ridge lines) in 3D point clouds is through a comparative roof fitting process. Since sensor systems often collect data as samples, statistically speaking, the true corner of a building will not likely be represented in the data as illustrated in Figure 3 (left). This means that the data collected by the sensor is likely not spatially consistent with the exact corner point of the building, but instead is represented by neighboring measurements that infer the true point location. As a first step in this method, it is assumed that building points have already been identified and grouped into individual buildings and for complex building shapes and the roof is decomposed into distinct roof segments. Next, a rectangle is formed in the XY plane from the smallest extent that encloses all points within the roof segment. The height of each corner point along the rectangle is pulled from the roof raw data heights as the 5th and 95th percentile heights representing a gutter height and ridge height, respectively. A rectangle is then sampled in a regular pattern at a specific sampling rate (e.g. 20 cm) along both the X, Y, and Z axes to create the roof top modeled data. The fit of the modeled roof points is compared to established, basic roof types shown in Figure 4. The height of each point in the roof model is then matched with the closest point in 3D space from the roof structure to calculate a residual error relative to the model, that is each primitive roof type is fit to the data and a residual error is calculated between the input points and the roof model. If the fit error of each building primitive exceeds a certain threshold, the entire roof structure is rejected as it likely represents a complex roof structure that is not representative to the basic types modeled. The roof model with the lowest residual error to the input data points are used to

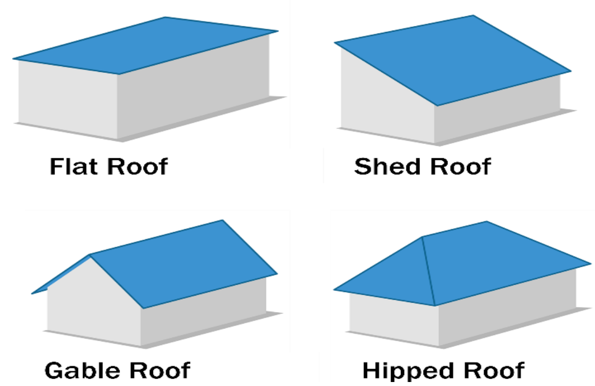

Figure 4. Primitive roof structures that are estimated in the roof fitting method

designate a roof-type label. Using this philosophy potentially allows for a more accurate estimate of the corner location based upon the entire collection of building points. Figure 5 illustrates the mean square error (MSE) between each sampled building segment and the best fitting roof model. In this particular example, the largest roof segment is a "T" shape. The primitive roof fitting procedure, however, only fits rectangles; thus the MSE is 1.56 m because of the disparity in shape comparison. Compared to the other roof segments with MSE scores less than 0.05 m, that "T" shape segment is an outlier and would be rejected. The residuals of the raw building points to the valid roof fit model in the Z direction provide a metric on quality of the data and this term is often referred to as the vertical point spread.

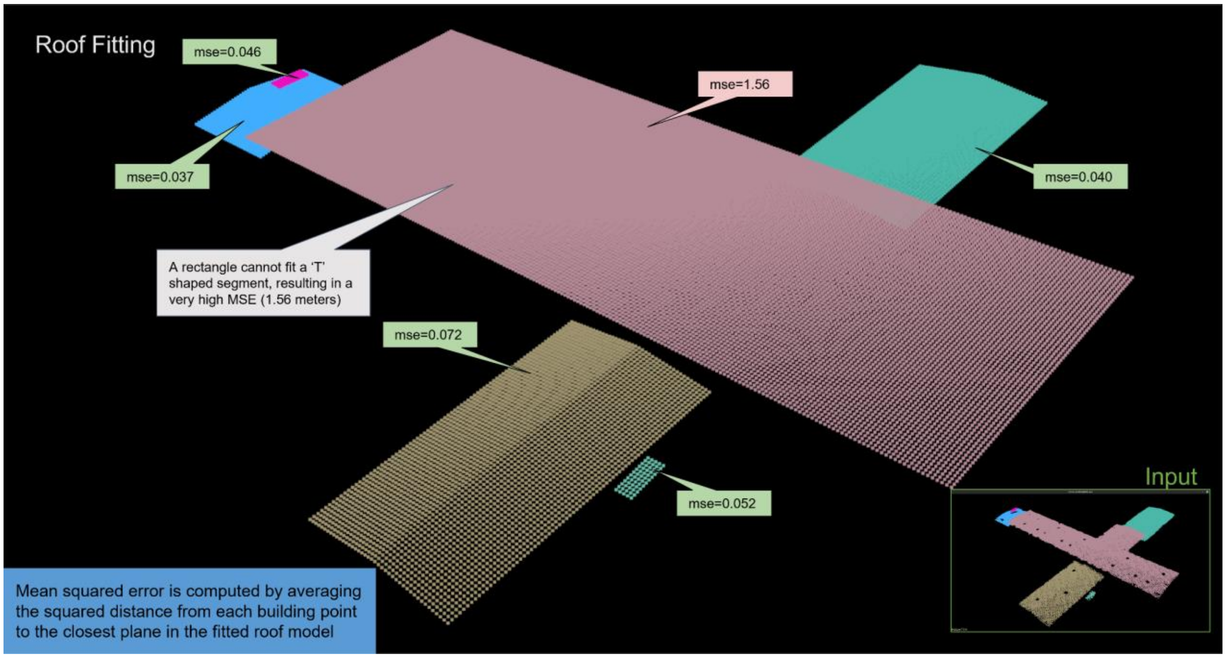

Figure 5. Example of Roof Fitting on input point cloud. The Mean Square Errors (MSE) is calculated between the closest roof models point and 3D point cloud input point.

RESULTS

Using either the 3D Corner Finder or Roof Fitting methodology described above, it becomes feasible to calculate the geolocation accuracy of a 3D point cloud relative to a reference 3D point cloud of the same ROI using a fully automated approach. By determining the corner points of buildings in both the input and reference data sets that have valid roof models (i.e. MSE scores below an acceptable threshold), the translation X, Y, and Z offsets as well as the rotations about the X, Y, and Z axis (see Figure 6) can be calculated. First, the corner points from the input data are matched to the reference data to determine the linear offsets for X, Y, and Z. Next, the translation values are applied to the input data and a singular value decomposition is utilized to determine a principal vector which can be decomposed into the rotations about each axis. The origin of rotation was assumed to exist at the center of the ROI.

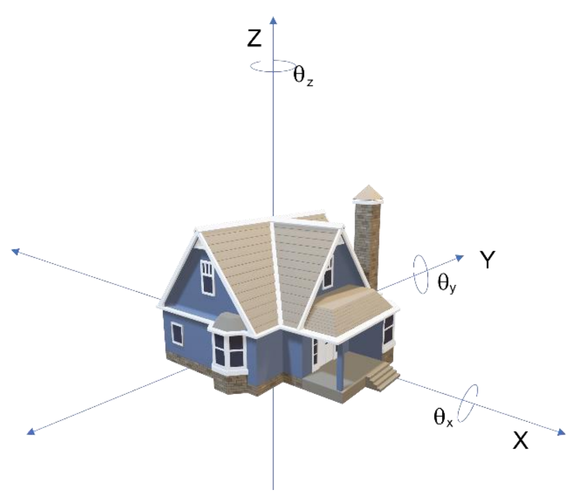

Figure 6. Illustration of errors; three translations X, Y, and Z and three rotations about the X, Y, Z axis that should be accounted for when calculating the geolocation accuracy of a 3D point cloud.

This approach was implemented on a photogrammetric point cloud collected from a Parrot-Anafi drone to calculate the 3D conformal translation parameters and a global geolocation uncertainty. Many UAV platforms (such as the Parrot-Anafi) do not utilize GPS Real Time Kinematic or Precise Point Kinematic in the navigation solutions and users rely upon the Structure for Motion software to compensate for positioning blunders. While these software packages work well, there may likely be translation and rotations relative to known reference positions. The photogrammetric input point cloud used in this study is shown in Figure 7 (left panel) and with individual buildings identified through Automated Feature Extraction (right panel).

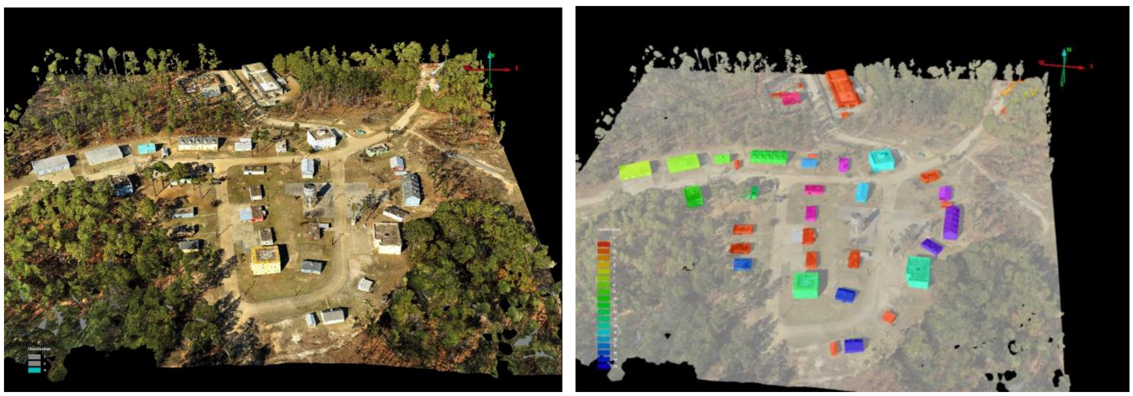

Figure 7. (LEFT) UAV Photogrammetric point cloud collected in 2023 using an Anafi Parrot system and (RIGHT) automated building extraction segmented into unique buildings (color coded by arbitrary building number).

For this example, we acquired airborne LiDAR data collected from the USGS in 2018 to serve as our reference source, or truth dataset. Figure 8 illustrates the difference between the roof fit derived from the EO (photogrammetry) UAV point cloud and the reference LiDAR 3D point cloud (USGS airborne LiDAR survey from 2018) for one of the buildings in the ROI. Using this approach outlined previously, the XYZ differences can be computed on a per building basis or combined with other buildings to determine a per ROI global translation assessment. Similarly, a global rotation can be determined from all corner points within a ROI about a center rotation axis.

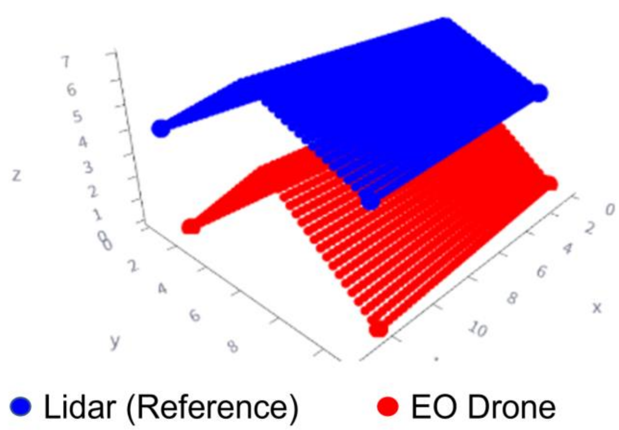

Figure 8. Example of roof model fit between EO drone roof model and LiDAR roof model.

For our example case, the individual buildings within the ROI were found to have a vertical error ranging from 2 to 4.5 m as shown in Figure 9. Each sample in Figure 9 (blue dot) represents the vertical error of a building within the ROI and is plotted as a function of the easting (left panel) or northing (right panel) direction. Using the identified roof corner points, parameters for a 3D Conformal Transformation (translation and rotation) values are calculated and reported in Table 1. The translation and rotation values were calculated using a least squares minimization and would be reported within a GPM error model. Table 2 reports the baseline geolocation uncertainty (RMSE) of the input point cloud as being 6.02 m. After application of the calculated translation parameters, the geolocation RMSE is reduced to 4.85 m. After applying both the translation and rotation parameters, the final geolocation uncertainty is reduced to 2.36 m. The uncertainty values reported in Table 2 include all corner points, including those that might not be a true fit. In some instances, a building might have a notch removed from the roof as an architectural detail which results in a potential mismatch in an automated corner point. In these instances, outliers can be statistically rejected which reduces the overall uncertainty. For the test site chosen here, three buildings have a large error and can be rejected from consideration. Removal of the outlier results in a final RMSE value of 1.62 m.

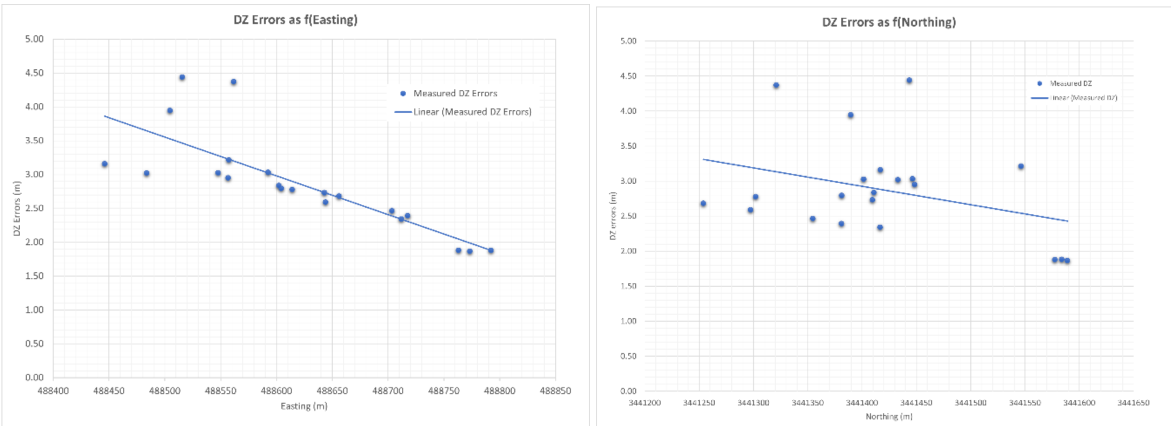

Figure 9. (LEFT) Results of vertical geolocation error calculation as a function of the Easting (X) direction and (RIGHT) Results of vertical geolocation error calculation as a function of the Northing (Y) direction.

Table 1. Calculated 3D Conformal Parameters for minimizing residuals between input and reference data

Axis	Translation (m)	Rotation (degrees)
X	0.34	-0.019
Y	-1.37	-0.032
Z	3.27	0.094

Table 2. Overall geolocation errors of input point cloud

Baseline RMSE (m)	Translation RMSE (m)	Translation and Rotation RMSE
6.02	4.85	2.36

The method presented here provides a capability to calculate the geolocation offsets between two different data sources (i.e. LiDAR and EO point clouds) through automated roof-fitting. Multi-source data often have data at different spatial resolutions and point densities, however, by utilizing the technique of fitting a model to each roof and then calculating the errors, those differences are minimized. The full process of the automated feature extraction, building segmentation, roof fitting, and geolocation accuracy calculation required approximately 30 minutes of computing time on a desktop PC. The same analysis if done by hand from an analyst would require the better part of a day.

CONCLUSIONS

High-resolution 3D point cloud data is becoming increasingly important for a variety of DoD applications and continues to gain momentum in reaching geospatial intelligence objectives and operational capabilities. However, the geolocation accuracy of this data can be difficult to determine as UAV data often does not contain information about its geolocation uncertainties and is more difficult to determine the applicability of the data for a given operational scenario. Here, two automated 3D building corner finding techniques that can be used to determine the geolocation offsets (translation and rotations) of high-resolution point cloud data are presented. The first method develops a likelihood ratio between the calculation of point density and a theoretical model. The second method fits a primitive roof model to the raw data. Both methods provide a faster, more efficient way to isolate building corner points, leading to automated calculation of geolocation accuracies. The roof-fit model approach was applied to calculate the geolocation accuracies of a UAV photogrammetric point cloud with respect to a LiDAR 3D point cloud serving as the reference. It was found that application of the calculated translation and rotation errors reduced the overall geolocation uncertainty from 6.02 m to 1.62 m. The methodologies presented in this paper demonstrate a capability to automatically and efficiently extract building corner locations from 3D point clouds, regardless of sensor type.

ACKNOWLEDGMENTS

The authors would like to thank our Leidos collaborators (Tu Lam and Scot Shiflett) for their contributions to this effort. We would also like to thank our I/ITSEC bird dog James Paul Rowlett for his guidance and suggestions. Funding for this project was made possible by Clayton Burford at the U.S. Army Simulation and Training Technology Center.

REFERENCES

Huang, X., Mei, G., Zhang, J., Abbas, R., 2021. A comprehensive survey on point cloud registration.

Jiang, T., Wang, Y., Zhang, Z., Liu, S., Dai, L., Yang, Y., Jin, X., Zeng, W., 2023. Extracting 3-D Structural Lines of Building From ALS Point Clouds Using Graph Neural Network Embedded With Corner Information. IEEE Trans. Geosci. Remote Sens. 61, 1–28. https://doi.org/10.1109/TGRS.2023.3278589

Park, J., Zhou, Q.-Y., Koltun, V., 2017. Colored Point Cloud Registration Revisited, in: 2017 IEEE International Conference on Computer Vision (ICCV). Presented at the 2017 IEEE International Conference on Computer Vision (ICCV), IEEE, Venice, pp. 143–152. https://doi.org/10.1109/ICCV.2017.25

Rodarmel, C., Lee, M., Gilbert, J., Wilkinson, B., Theiss, H., Dolloff, J., and O'Neill, C. 2015. The Universal Lidar Error Model. Photogrammetric Engineering & Remote Sensing, 81:7. https://doi.org/10.14358/PERS.81.7.543

Wang, H., Brady, M., 1995. Real-time corner detection algorithm for motion estimation. Image Vis. Comput. 13, 695–703. https://doi.org/10.1016/0262-8856(95)98864-P