

# ASSESSING THE ACCURACY AND PRECISION OF IMPERFECT POINT CLOUDS FOR 3D INDOOR MAPPING AND MODELING

J. Chen<sup>1</sup>, O. E. Mora<sup>2</sup>, K. C. Clarke<sup>3</sup>

<sup>1</sup> Department of Geography, University of California, Santa Barbara, USA - jorgechen@geog.ucsb.edu

<sup>2</sup> Department of Civil Engineering, California State Polytechnic University, Pomona, USA - oemora@cpp.edu

<sup>3</sup> Department of Geography, University of California, Santa Barbara, USA - kclarke@geog.ucsb.edu

**KEY WORDS:** Point clouds, remote sensing, indoor mapping

## ABSTRACT:

In recent years, growing public interest in three-dimensional technology has led to the emergence of affordable platforms that can capture 3D scenes for use in a wide range of consumer applications. These platforms are often widely available, inexpensive, and can potentially find dual use in taking measurements of indoor spaces for creating indoor maps. Their affordability, however, usually comes at the cost of reduced accuracy and precision, which becomes more apparent when these instruments are pushed to their limits to scan an entire room. The point cloud measurements they produce often exhibit systematic drift and random noise that can make performing comparisons with accurate data difficult, akin to trying to compare a fuzzy trapezoid to a perfect square with sharp edges. This paper outlines a process for assessing the accuracy and precision of these imperfect point clouds in the context of indoor mapping by integrating techniques such as the extended Gaussian image, iterative closest point registration, and histogram thresholding. A case study is provided at the end to demonstrate use of this process for evaluating the performance of the Scanse Sweep 3D, an ultra-low cost panoramic laser scanner.

## 1. INTRODUCTION

Measurements serve as the basic building blocks of maps and models, and as the relevance of indoor maps continues to grow, so will the demand for fast and affordable techniques for measuring indoor spaces. For many years, professional laser scanning<sup>1</sup> provided the only practical way to measure these spaces on a large scale. However, their high costs limited their use to organizations with large budgets and made them out-of-reach for everyone else.

More recently, a number of low cost alternatives have emerged that make close-range 3D reality capture available to anyone with a modest budget or a digital camera. One popular approach uses photogrammetry based on structure-from-motion with multi-view stereo (SfM-MVS) (Smith et al., 2016), e.g., Agisoft PhotoScan, Autodesk ReCap, and Bentley ContextCapture, while another employs structured light (Khoshelham and Elberink, 2012), e.g., Microsoft Kinect, Google Tango, and Matterport. Other methods include low-cost panoramic LiDAR, calibrated stereo vision, and deriving indoor structure from single images using artificial intelligence (Zou et al., 2018). Their lower prices, however, come at the cost of lower accuracy and precision, which can influence the final quality of a map and its derived information.

All of these approaches present 3D measurements in the form of point clouds, or collections of  $xyz$  coordinate values that can number from thousands to millions of points. Four factors complicate the quality assessment of these point clouds. First, the concepts of *error* and *accuracy* assume a known ground truth (GT) or baseline from which to compare measurements, but establishing this ground truth is often difficult without expensive instruments, especially in non-lab environments. Second, the correspondence between a point in the point cloud and its exact ground truth location is often unknown. Point cloud distortions (such as

from ranging errors and scanner drift) and use of different coordinate reference systems (such as datums and units of measurement) can further add to this uncertainty. Third, errors in point cloud geometry, e.g., global scale, can bias statistical results at a local level when comparisons are directly made using features that do not line up. Finally, non-uniform point cloud densities can distort alignments with GT when some areas have high clustering of points while others have none. Figure 1 provides 2D illustrations of some of these challenges.

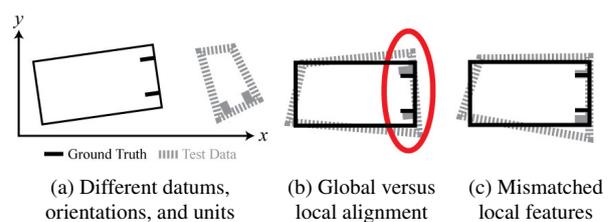


Figure 1. Challenges in assessing low quality point cloud data

This paper presents a robust process for assessing indoor point cloud quality based on two measures: *global accuracy* at the level of a room and *local out-of-plane accuracy and precision* at the level of a flat surface. This process addresses the latter three of the four challenges mentioned earlier (unknown point correspondences with unmatched coordinate reference systems, geometric errors, and non-uniform point densities) and assumes the existence of a ground truth data set. In practice, the ground truth may simply come from a high precision scanner that may or may not have been calibrated to standard distances. A case study is provided at the end of this paper to demonstrate use of this process.

<sup>1</sup>Also referred to as light detection and ranging or LiDAR

## 2. RELATED WORK

### 2.1 Accuracy and precision

While the terms *accuracy* and *precision* are often confused in casual use, they have important and distinctly different meanings when used in a technical context. As used in this paper, accuracy describes the closeness of a measurement from its “true” value and precision as the closeness of repeated measurements of the same object (Zhang and Goodchild, 2002; Mason, 2006).

### 2.2 General approaches

Habib (2008) classified approaches to assessing point cloud quality as being either external or internal, with the former using control data that exist independently of a measured point cloud and the latter checking for relative consistency between different scans. For external assessments, LiDAR point clouds seldom have exact correspondences between measured points and control points due to sampling limitations, which means that control point correspondences will always have some level of uncertainty. On the other hand, SfM-MVS can directly derive control point correspondences from source images, although in practice artifacts from unfavorable scenes (e.g., low texture, reflective, etc.) can render some of these points useless (Bolognesi et al., 2014).

Internal assessments, in theory, measure the consistency between different scans by examining what Habib called the “coincidence of conjugate features” or, simply stated, the level of overlap between identical features in different point clouds. In practice, getting two point clouds to correctly overlap or *register* is difficult due to noise, geometric distortions, and the fact that point clouds in their raw state have no high level semantic information, i.e., points are simply  $xyz$  coordinate tuples with no spatial context. Noting that incorrect registration could be misinterpreted as error, Khoshelham and Elberink (2012) proposed using two variations of the iterative closest point (ICP) approach (Besl and McKay, 1992) for minimizing registration errors, with the first minimizing distances between points and the second between extracted planar patches. Habib (2008) proposed similar ICP and non-ICP approaches using point-to-patch and patch-to-patch registration. In their studies, Khoshelham and Elberink (2012) and Bolognesi et al. (2014) used summary statistics to characterize the accuracy and precision along each of the three Cartesian coordinate axes.

### 2.3 Room-based approaches

Another way to evaluate the accuracy of indoor point clouds involves comparing the derived dimensions of a room. Budroni and Böhm (2009) and Okorn et al. (2010) developed a simple and effective method for deriving the *coordinates* of a room’s boundary surfaces by searching for peaks in axes-aligned histograms of the point clouds, where the salient peaks corresponded to boundary surfaces, e.g., floor, ceiling, and walls. Khoshelham and Díaz-Vilariño (2014) and Díaz-Vilariño et al. (2015) enhanced this method by aligning the point clouds of cuboid-shaped rooms with the  $xyz$  axes using the extended Gaussian image (EGI) (Ikeuchi, 1981), a spherical point plot of the *unit normal vector* for every point, as illustrated in Fig. 2. Notably, they oriented all unit normal vectors in the positive  $xyz$  directions (Fig. 2c), resulting in three clusters that pointed in quasi-orthogonal directions. They then used the *mean* value of each cluster to estimate the normal vectors for the room’s orthogonal walls, i.e., vectors  $uvw$ , and calculated the  $uvw$ -to- $xyz$  transformation using a direct change of basis, i.e., direct projection.

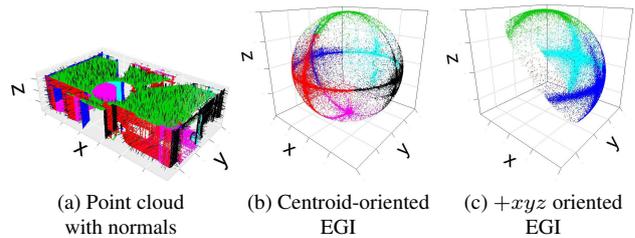


Figure 2. Extended Gaussian image of a cuboid shaped room

## 3. METHODOLOGY

This paper proposes a method for evaluating low cost 3D reality capture systems in measuring the structure of a room, at both the scale of a room and a local flat surface, such as a wall. Since it assumes a Manhattan world scene, it can only work with scans of cuboid-shaped rooms that have solid surfaces. Additionally, it assumes the availability of accurate ground truth data in point cloud form, either previously collected or obtained alongside the test data.

### 3.1 Global room level analysis

Global analysis looks at data quality at the scale of the room and involves aligning the test and ground truth point clouds, deriving the room dimensions, and comparing the results, as illustrated in Fig. 3. However, before starting this analysis, excessively dense data should be subsampled to a more manageable size, noting that subsampling preserves the original coordinates while resampling produces new data points using aggregated data properties, e.g., centroids (Héno and Chandelier, 2014).

**3.1.1 Rough alignment** Since the ground truth and test data can exist anywhere in space and have potentially different units of measurement (Fig. 3a), the first step of the process involves scaling the data to a common unit of measurement and manually aligning them with the  $xyz$  axes (and with one another) using any number of techniques (Fig. 3b). This *rough alignment* sets the initial point cloud orientation and prepares both data sets for fine registration.

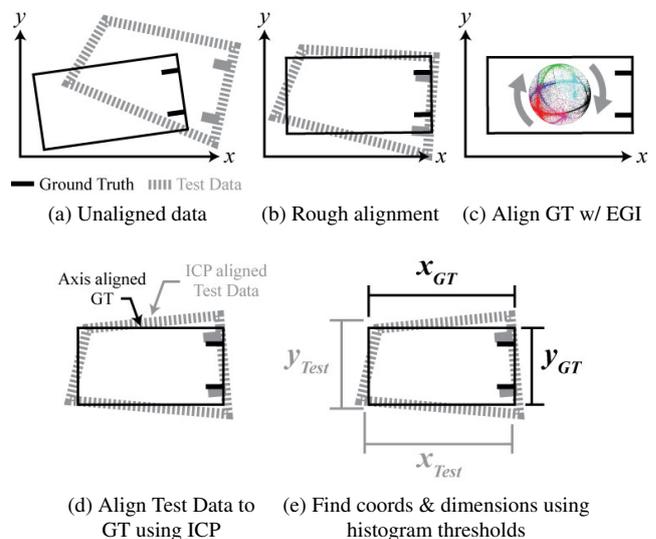


Figure 3. General overview of global analysis

**3.1.2 Axis alignment of GT using the EGI** Next, fine registration of GT with the  $xyz$  axes is performed using the extended Gaussian image. This step involves estimating the *unit* normal vectors<sup>2</sup> for every point in the GT point cloud—making sure to orient them away from the point cloud’s centroid, plotting them from a (0, 0, 0) origin to produce the EGI (as illustrated in Figs. 2 and 3c), finding the center of each EGI cluster to form the local  $uvw$  axes, and finding the optimal rotation to match  $uvw$  with  $xyz$ . The EGI will show six clusters corresponding to each of the cuboid’s six walls, with the *central location* in each cluster corresponding to each wall’s normal vector. A cuboid will therefore have a total of six vectors denoted  $\pm u$ ,  $\pm v$ , and  $\pm w$  that correspond with the  $\pm x$ ,  $\pm y$ , and  $\pm z$  axes.

**Finding the central location** While Díaz-Vilariño et al. (2015) and Khoshelham and Díaz-Vilariño (2014) used the cluster mean for the central location, this paper recommends using the *mode*, which can provide a better measure for clean, dense, and high quality data.<sup>3</sup> Rather than build a complicated spherical histogram to find the mode, a more practical approach involves simply projecting each cluster on to its tangent coordinate plane and finding the coordinates to the location with the highest concentration of points (i.e., the mode) using a 2D histogram, as illustrated in Fig. 4. The third coordinate can then be derived using the Pythagorean theorem.

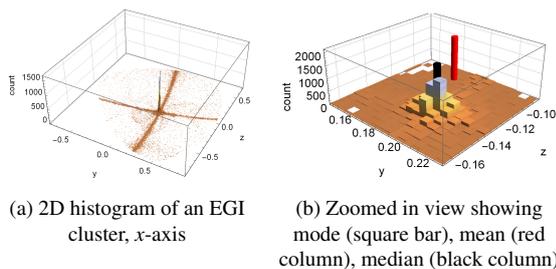


Figure 4. EGI cluster along the  $x$ -axis for high quality data set

**Determining the optimal rotation** The six  $\pm uvw$  vectors that EGI analysis produces will not be truly orthogonal since they come directly from the data. As a result, performing a direct change of basis from  $uvw$  to  $xyz$  can skew the ground truth data. Instead, a best fit approach should be used that preserves the original shape of the point cloud. This paper uses singular value decomposition, with the “scaling” matrix set to identity, to find the pure rotation matrix that provides the best least squares fit of  $uvw$  to  $xyz$  (Umeyama, 1991).

**3.1.3 Aligning the test data with GT** The next step uses the iterative closest point (ICP) technique to optimally register the test data with the axis-aligned GT, which effectively aligns the test data with  $xyz$  as illustrated in Fig. 3d. Any number of ICP implementations can be used, such as the one built into CloudCompare (<https://www.danielgm.net/cc/>). Finding the actual transformation matrix involved selecting three widely spaced points

<sup>2</sup>Unit normals can be computed using any number of methods such as the “Compute normals for point sets” tool in Meshlab (<http://www.meshlab.net/>).

<sup>3</sup>This is due to the fact that the majority of normal vectors on a flat surface will point in the same direction while the (erroneously) angled edge and corner normals will add no value to the data—they can be discarded. Other erroneous normals from furnishings and clutter can be mitigated by cropping out those features prior to plotting the EGI. Nonetheless, the mean and median can perform better than the mode with noisy data.

and finding the transformation required to move those points from their original locations to their new ones after ICP.

To ensure optimal registration, all furnishings, clutter, and minor features should be removed from the point cloud, leaving only large flat surfaces (e.g., floor, ceiling, and bare walls). Additionally, the GT data should have a higher point density than the test data and cover a larger area to prevent many-to-one registrations as shown in Fig. 5.



Figure 5. Point density and coverage

**3.1.4 Deriving and comparing room dimensions** After aligning the GT and test data with the  $xyz$  axes, finding coordinates for the surface boundaries simply involves projecting each point cloud on to the coordinate planes (i.e.,  $xy$ ,  $yx$ ,  $zy$ ), binning the data, and finding the two histogram peaks along each axis/abscissa. For example, Fig. 6 shows two peaks corresponding to walls at both ends of the  $y$ -axis for the given data. Performing this along all three axes results in three min-max pairs that can be used to find the room dimensions.

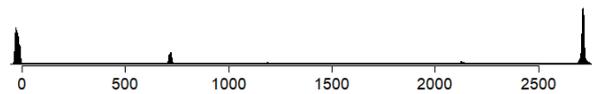


Figure 6. Histogram along  $y$ -axis

**3.2 Local flat surface analysis**

As noted earlier, global alignment does not always result in alignment of local surfaces, making it impossible to directly assess their quality without further processing. For instance, Fig. 7a shows the global alignment of a test data set with ground truth, but the wall of interest, circled in red, does not match its counterpart in the ground truth data. Therefore, preparing local surface for analysis involves two general steps: local alignment via ICP and local adjustments along the plane of the wall.

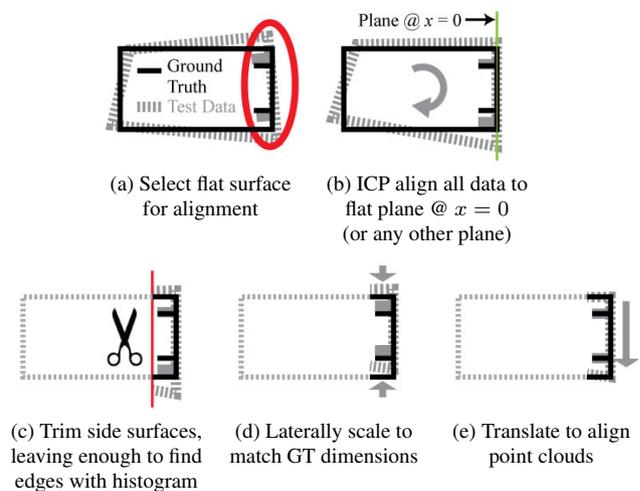


Figure 7. General overview of local flat plane analysis

**3.2.1 Local alignment** Local alignment involves using ICP to register or “snap” the the test and GT point clouds to an artificially generated plane, as illustrated in Fig. 7b. To ensure the best possible fit, the point cloud “surface” must form a flat plane and be free of all furnishings and other protrusions (e.g., white boards, electrical outlets, picture frames, etc.), which can be cropped out. Once registered, flat surfaces in both data sets will rest on the same plane, although they may not line up laterally due to measurement errors or the natural shifting from ICP transformations.

**3.2.2 Lateral matching** Lateral matching of the data will ensure that the analysis compares like items in both point clouds. For instance, directly analyzing the unadjusted walls in Fig. 7c will incorrectly register the two pairs of protrusions (one gray, one black) as measurement errors, when in fact they both correctly measure the same things. After lateral matching, the two pairs will line up as shown in Fig. 7e and the out-of-plane deviations should disappear.

**Trimming and scaling** The first step in performing lateral adjustment involves laterally scaling the test data to match GT based on their side-to-side dimensions, which can be derived using histogram analysis. However, instead of using the entire length of the adjacent surfaces, this process only uses the portion closest to the wall to prevent errors and imperfections further down from biasing the results. Here, histogram analysis produces two pairs of min-max coordinate values that can be used to calculate the width and height (or depth) of the surfaces; these can then be used to scale the test data to match GT, as shown in Fig. 7d.

**Lateral translation** Since scaling shifts all coordinates relative to an arbitrary datum, histogram analysis is performed a second time to re-establish the new boundary coordinates of the test surface. Final alignment involves selecting an arbitrary corner of the GT surface and translating the same point on the test surface to match.

**3.2.3 Assessing out-of-plane performance** Assessing out-of-plane performance involves looking at the accuracy of the test data compared to ground truth and examining its precision. This requires first partitioning the data into equally sized grids and then computing statistics on a cell-by-cell basis. For accuracy, this involves calculating each cell’s mean out-of-plane value for both GT and the test data and then finding the difference between the test data and GT. These cell values can be further aggregated and analyzed in rows, columns, or as a whole, as demonstrated in the later case study.

Precision involves looking at the scanner’s ability to make consistent measurements, which involves comparing the data with itself. The measure of precision used here is the *mean* absolute deviation (MAD), which can be found by first calculating the mean out-of-plane value for each cell and then finding the mean value of the *absolute* differences of all points in the cell. These MAD values can then be aggregated to provide summary statistics for the entire surface.

## 4. CASE STUDY

A case study was performed to evaluate the performance of an ultra-low cost panoramic LiDAR scanner using the quality assessment process proposed in this paper. This study used the US\$750 Scanse Sweep 3D LiDAR scanner (<http://scanse.io/3d-scanning-kit/>) to scan a mid-sized conference room that measured

6.4 m x 2.8 m x 8.5 m and met the process’s requirements, e.g., cuboid shape, solid walls, and relatively clutter free. Ground truth was provided by a high precision Trimble TX8 LiDAR scanner.

### 4.1 Data collection

Each scanner made full 360° spherical scans along the center of the room at 6 m, 4 m, and 2 m distances from the wall that was selected for local assessment. These six scans are denoted as Scanse 1, 2, and 3 and Trimble 1, 2, and 3, with the first station at 6 m and third at 2 m from the wall. Instruments were visually centered over the station points without the use of a sighting tool nor a plumb bob.

All three Scanse scans used the highest quality setting of a 1 Hz vertical rotation rate and a 500 Hz sampling rate, while the Trimble used the highest quality setting for the middle station (Trimble 2) and a slightly lower quality setting for the other two. These produced the six data sets shown in Table 1. The Scanse produced data in units of centimeters while the Trimble delivered data in meters. Trimble 2 was selected as the ground truth scan since it was collected at the highest quality setting. The remaining five point clouds constituted the test data, with Trimbles 1 and 3 providing internal quality checks for Trimble 2.

Station ⇒	Scanse			Trimble		
	1	2	3	1	2	3
# of points	150K	130K	115K	118M	474M	119M
File size	2MB	2MB	2MB	3GB	12GB	3GB
	PLY format			LAS format		

Table 1. Data sets and point counts

### 4.2 Subsampling and rough alignment

The extremely large sizes of the Trimble data made it necessary to reduce them via subsampling, which was performed using CloudCompare to point spacings of 1 mm, 2mm, and 5 mm. However, the Scanse data was unmodified due to their small sizes. After subsampling, all six point clouds were roughly registered together at three common anchor points, converted to units of millimeters, and rotated into an *xzy* orientation, with the *+y*-axis pointing up and the *+z*-axis pointing into the test wall and away from the center of the room.

### 4.3 Global room-level analysis

**4.3.1 EGI alignment** Although the conference room was relatively free of clutter, it still had several large tables at one end of the room. For EGI alignment, the occupied half of the room was completely removed from the ground truth data, leaving five clean and unobstructed surfaces, as shown in Fig. 8 with normal vectors.



Figure 8. Cropped Trimble 2 point cloud used for EGI analysis, with normal vectors displayed as blue lines

Plotting the surface normals resulted in the EGI shown in Fig. 9b, which had five clusters along the  $\pm x$ ,  $\pm y$ , and  $+z$  axes. Extracting the mode from each cluster (e.g., in Fig. 9c) produced the  $\pm u$ ,  $\pm v$ , and  $+w$  vectors, which were then used in SVD to find the rotation matrix needed to align the point cloud with the  $xyz$  axes. The resulting rotation matrix was then applied to the entire ground truth data set, not the cropped version.

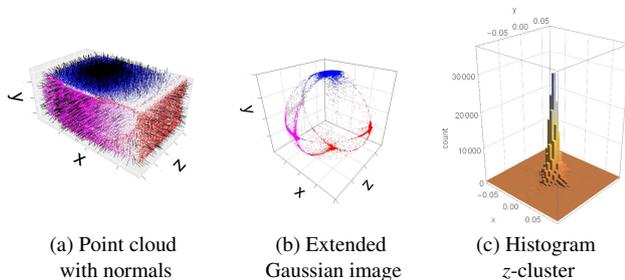


Figure 9. Normals, EGI, and  $z$ -axis histogram for Trimble 2. Removing half of the room resulted in five clusters in the EGI.

**4.3.2 Alignment of test data** After aligning ground truth with the  $xyz$  axes, all test data (i.e., all Scanse data and Trimbles 1 and 3) were aligned with ground truth using a simple brute force ICP process that used an unconstrained nearest neighbor (NN) search and termination criteria of 200 iterations or convergence over 20 iterations at 0.01 mm. Preparing the test data for ICP involved subsampling to a uniform density—the sampling distance varied for each point cloud—and removing all points associated with furnishings as illustrated in Fig. 10. As with EGI, the resulting transformation matrix was applied to the entire point cloud, not the cropped version.

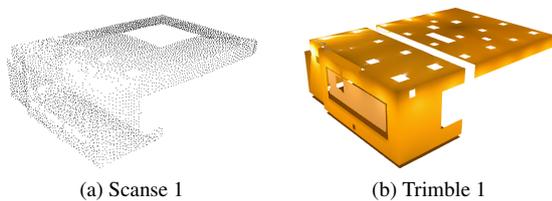


Figure 10. Test data prepared for ICP

**4.3.3 Histogram analysis** Finding the room’s boundary coordinates involved collapsing each point cloud on to two coordinate planes (e.g., Fig. 11), plotting the points in the order  $xy$ ,  $yx$ , and  $zy$ , and performing histogram analysis on those plotted points. Table 2 shows the room dimensions derived from these coordinates.

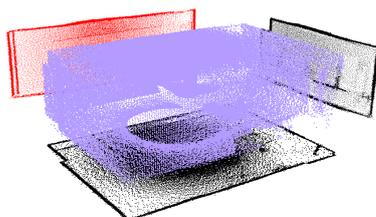


Figure 11. Projections of Scanse 2 point cloud along the three coordinate planes. Only two were used for histogram analysis, shown in black.

	GT	Trimble		Scanse		
		1	3	1	2	3
x	6449	6448	6449	6627	6616	6602
% $\Delta \Rightarrow$		-0.01%	0.00%	2.76%	2.60%	2.38%
y	2745	2738	2744	2852	2883	2891
% $\Delta \Rightarrow$		-0.26%	-0.02%	3.92%	5.05%	5.32%
z	8765	8764	8765	8910	8908	8897
% $\Delta \Rightarrow$		-0.02%	0.00%	1.65%	1.63%	1.51%

Table 2. Comparison of derived room dimensions, in units of millimeter

**4.4 Local flat surface analysis**

**4.4.1 Plane alignment** To perform local analysis, all six point clouds (both GT and test data) had to be aligned with an artificial  $xy$  reference plane that was placed at  $z = 0$ . This required removing all out-of-plane features (e.g., side walls, white board, projection screen holder, fire alarm, electrical outlet, etc.) from the point clouds prior to ICP, as illustrated in Fig. 12. It also required subsampling the points to a uniform density, as shown in Fig. 13, to prevent point clusters from biasing the resulting transformation. As before, the resulting transformation matrices were applied to the entire point cloud, not the cropped version.

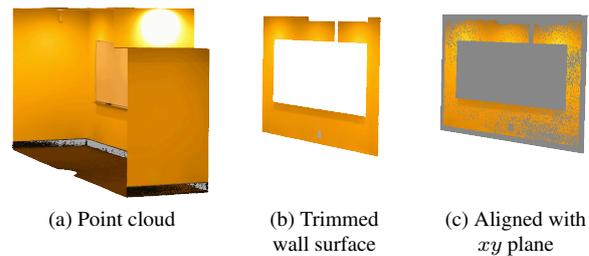


Figure 12. Aligning test wall with  $xy$  plane at  $z = 0$

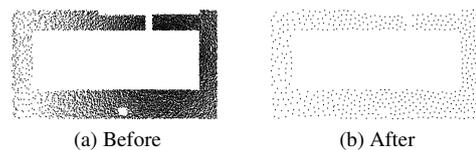


Figure 13. Subsampling of Scanse 3 data

**4.4.2 Lateral scaling and alignment** With all six wall segments resting on the same  $xy$  plane, the next step involved scaling laterally (i.e., only in the  $xy$  direction) to match the dimensions of GT. This involved trimming each adjacent surface to within 1 meter of the test wall (similar to Fig. 12a; also see Fig. 14), performing histogram analysis to find the min-max coordinates of the wall, and using the derived distances for scaling. Scaling was applied to the cropped point clouds, not the entire room.



Figure 14. Profile view of a side wall showing deviations in the Scanse data. Using only the closest 1 m of adjacent surface avoided the influence of errors contained in the remainder of the surface.

Histogram analysis was then performed a second time to re-determine

wall boundary coordinates and all point clouds were then translated to a new datum corresponding to an arbitrarily selected corner points, as shown in Fig. 15.

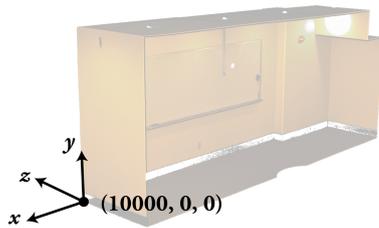


Figure 15. New datum after lateral alignment

**4.4.3 Analyzing out-of-plane performance** After final alignment, the point cloud was cleaned up and segmented into regular grids for statistical analysis. Cleaning up the data involved removing areas that potentially contained artifacts (i.e., random errors), which included the edge of the whiteboard, fire alarm, outlet, overhead projection screen case and cord, and an unscanned area of the whiteboard. The data was then partitioned into 20 cm cells for performing accuracy and precision analysis.

**Accuracy** Accuracy analysis involved calculating the mean  $z$  value for each cell in both the ground truth and test data sets and then performing cell-by-cell comparisons of the test data with ground truth. Table 3 and Fig. 16 show the accuracy results using 200 mm grids.

Units in mm	Trimble		Scanse		
	1	3	1	2	3
mean	-0.07	-0.10	-0.88	4.04	-4.81
median	-0.06	-0.05	0.40	3.87	1.05
std. dev.	0.15	0.36	13.60	10.90	20.32

Table 3. Flat surface accuracy @ 200 mm grids

**Precision** Precision analysis used the same 200 mm cells, but instead of comparing mean cell values to ground truth, it looked at the *mean* absolute deviation (MAD) of individual point values from each cell's mean value. Table 4 and Fig. 17 show the precision results using 200 mm grids.

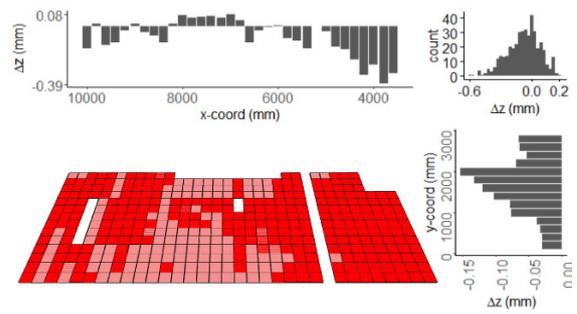
Units in mm	GT	Trimble		Scanse		
		1	3	1	2	3
mean	0.99	0.96	0.98	12.31	16.33	13.28
median	0.96	0.92	0.94	12.47	16.57	13.26
std. dev.	0.14	0.12	0.22	5.49	5.68	3.99

Table 4. Flat surface precision based on mean absolute deviation @ 200 mm grids

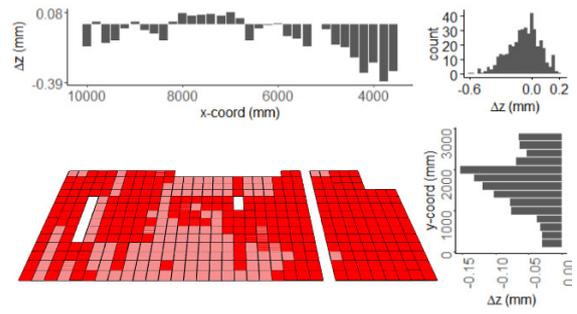
#### 4.5 Summary

This case study demonstrated the effectiveness of the proposed process in assessing the quality of the Scanse data. Internal quality checks using Trimbles 1 and 3 showed 1 mm errors at the level of the room and sub-millimeter errors and precision for the flat surface, validating the ground truth measurements and providing a strong indication that the EGI and ICP registration successfully aligned Trimbles 1 and 3 with ground truth.

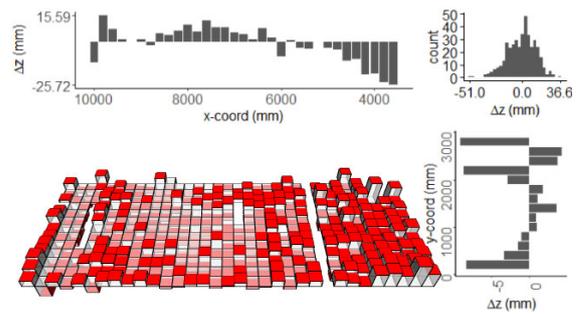
Comparing the Scanse data to ground truth revealed several interesting observations. First, the Scanse consistently overesti-



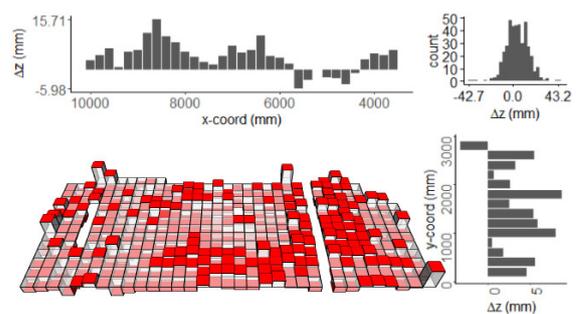
(a) Trimble Station 1



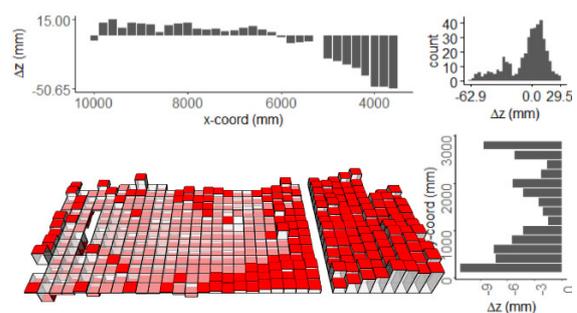
(b) Trimble Station 3



(c) Scanse Station 1

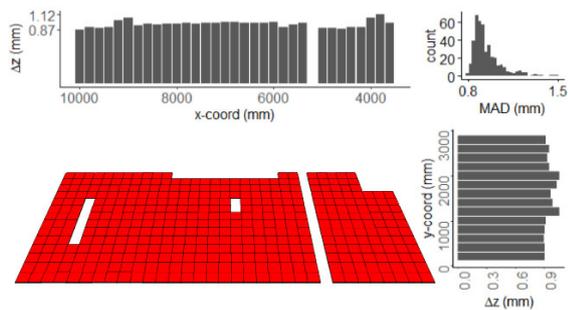


(d) Scanse Station 2

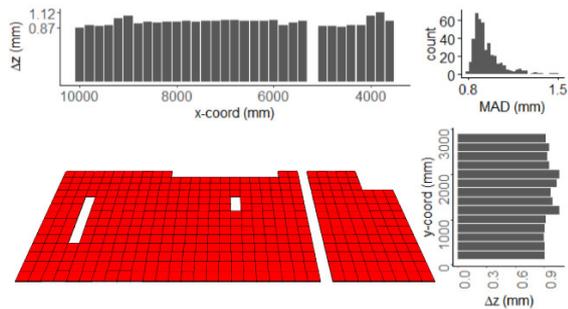


(e) Scanse Station 3

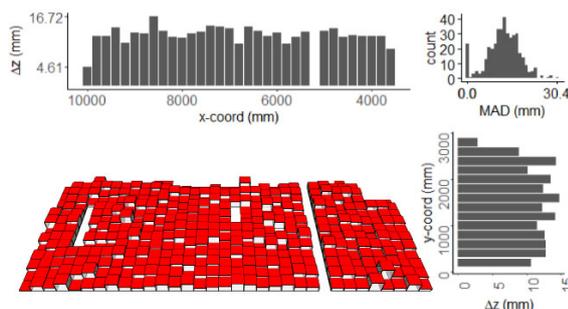
Figure 16. Flat surface accuracy at 200 mm grids



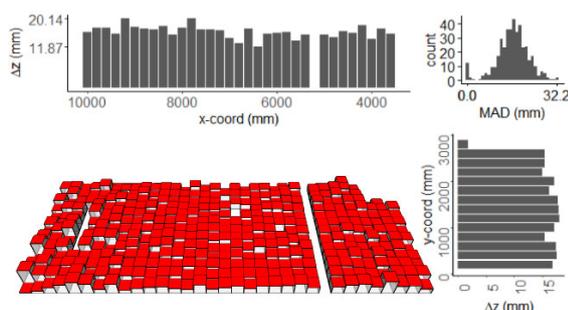
(a) Trimble Station 1



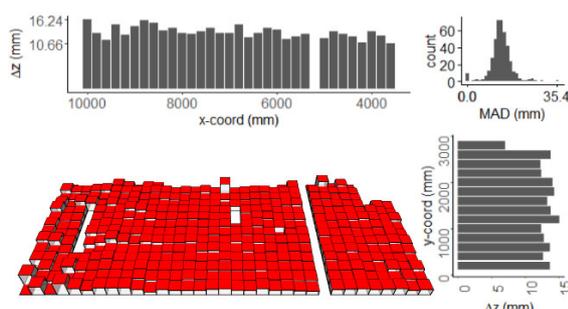
(b) Trimble Station 3



(c) Scans Station 1



(d) Scans Station 2



(e) Scans Station 3

Figure 17. Flat surface precision at 200 mm grids

mated distances, indicating a systematic error that can be corrected through calibration. Second, it appears that the Scans performs better at longer ranges than closer. Evidence for this comes from the increasing values of mean and median errors as the scanner moved closer to the wall from Stations 1 to 3 (Table 3) as well as increasing errors in room measurements from the largest dimension ( $z$  axis) to the smallest ( $y$  axis) (Table 2). This may have been caused by progressive difficulties faced by the LiDAR's timing mechanism to capture very small time differences. Finally, the scanner maintained consistent precision as it swept across the wall in all three scans, although its value bottoms out at the second station. Combining the uniform precision with the high standard deviation in errors seems to indicate a tendency for the scanner to drift, as shown by the undulating profile of the wall in Fig. 14.

## 5. CONCLUSION

A growing demand for indoor maps will not only drive a growing need for indoor measurement capabilities, it will also generate an accompanying desire to know about the quality of the data they produce. The quality assessment process outlined in this paper can help meet this need by providing a way quantify the accuracy and precision of this data and to characterize their measurement devices, especially those of low accuracy and low precision. These quality measures may also be used to select measurement devices that can support creating maps at certain levels of detail and to assign weights to data sources when the crowdsourcing of indoor map data becomes a reality some day in the future.

However, this process has several limitations including the need for a cuboid-shaped room and the need for ground truth data, which usually implies having access to expensive high-precision equipment. Additionally, many aspects require exercising human judgment, such as in the selection of subsampling distances and grid sizes and the removal of features from the point cloud to optimize EGI and ICP alignment. Suggested improvements include expanding the use of this technique beyond the cuboid, providing more rigorous approaches to quantifying quality beyond the simple statistics used here, and accounting spatial variations accuracy and precision.

## ACKNOWLEDGEMENTS

This research was funded by the National Geospatial-Intelligence Agency's Academic Research Program (Award No. HM0476-17-1-2002, Project Title: Automated Level of Detail and Feature Extraction for 3D Indoor Mapping and Modeling). Approved for public release, 18-572. All unattributed graphics and images were produced by the author.

## References

- Besl, P. J. and McKay, N. D., 1992. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14(2), pp. 239–256.
- Bolognesi, M., Furini, A., Russo, V., Pellegrinelli, A. and Russo, P., 2014. Accuracy of cultural heritage 3D models by RPAS and terrestrial photogrammetry. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XL-5, pp. 113–119.

Budroni, A. and Böhm, J., 2009. Toward automatic reconstruction of interiors from laser data. *Proceedings of Virtual Reconstruction and Visualization of Complex Architectures (3D-Arch)* p. 36.

Díaz-Vilariño, L., Khoshelham, K., Martínez-Sánchez, J. and Arias, P., 2015. 3D Modeling of Building Indoor Spaces and Closed Doors from Imagery and Point Clouds. *Sensors* 15(2), pp. 3491–3512.

Habib, A., 2008. Accuracy, Quality Assurance, and Quality Control of LiDAR Data. In: *Topographic Laser Ranging and Scanning*, CRC Press, pp. 269–294.

Héno, R. and Chandelier, L., 2014. *3D Modeling of Buildings: Outstanding Sites*. 1st edn, Wiley-IEEE Press.

Ikeuchi, K., 1981. Recognition of 3-d objects using the extended gaussian image. In: *Proceedings of the 7th International Joint Conference on Artificial Intelligence - Volume 2, IJCAI'81*, San Francisco, CA, USA, pp. 595–600.

Khoshelham, K. and Díaz-Vilariño, L., 2014. 3D Modelling of Interior Spaces: Learning the Language of Indoor Architecture. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XL-5, pp. 321–326.

Khoshelham, K. and Elberink, S. O., 2012. Accuracy and Resolution of Kinect Depth Data for Indoor Mapping Applications. *Sensors* 12(2), pp. 1437–1454.

Mason, F., 2006. Inaccurate? Precisely! *Quality Digest*.

Okorn, B., Xiong, X., Akinci, B. and Huber, D., 2010. Toward Automated Modeling of Floor Plans. In: *Proceedings of the Fifth International Symposium on 3D Data Processing, Visualization and Transmission*, Paris, France.

Smith, M., Carrivick, J. and Quincey, D., 2016. Structure from motion photogrammetry in physical geography. *Progress in Physical Geography* 40(2), pp. 247–275.

Umeyama, S., 1991. Least-squares estimation of transformation parameters between two point patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 13(4), pp. 376–380.

Zhang, J. and Goodchild, M. F., 2002. *Uncertainty in geographical information*. CRC press.

Zou, C., Colburn, A., Shan, Q. and Hoiem, D., 2018. Layoutnet: Reconstructing the 3d room layout from a single rgb image. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2051–2059.