

Performance Evaluation of Consumer-Grade 3D Sensors for Static 6DOF Pose Estimation Systems

J. A. Marvel^{*a,b}, M. Franaszek^a, J. L. Wilson^c, and T. Hong^a

^aNational Institute of Standards and Technology, Intelligent Systems Division, 100 Bureau Drive, MS-8230, Gaithersburg, MD, USA 20899; ^bInstitute for Research in Electronics and Applied Physics, University of Maryland, Bldg. #223, Paint Branch Drive, College Park, MD USA 20742; ^cDepartment of Computer Science, Princeton University, 35 Olden Street Princeton, NJ, USA 08544

ABSTRACT

Low-cost 3D depth and range sensors are steadily becoming more widely available and affordable, and thus popular for robotics enthusiasts. As basic research tools, however, their accuracy and performance are relatively unknown. In this paper, we describe a framework for performance evaluation and measurement error analysis for 6 degrees of freedom pose estimation systems using traceable ground truth instruments. Characterizing sensor drift and variance, and quantifying range, spatial and angular accuracy, our framework focuses on artifact surface fitting and static pose analysis, reporting testing and environmental conditions in compliance with the upcoming ASTM E57.02 standard.

Keywords: Computer vision for robotics and automation, range sensing, directional surface fitting

1. INTRODUCTION

The market has witnessed an increase in low-cost, consumer-grade sensors coinciding with the steady growth of the hobbyist and “maker” communities, and the evolution of video game technologies. Roboticists can rapidly transition from concept to prototype with little detailed knowledge of the internal mechanisms driving sensors in part to the establishment of common interfaces and ready availability of plug-and-play solutions. New 3D technologies thus see fast, widespread utility in robot applications including world sensing [1-2], haptics [3], interactive robot control [4], and 6 degree of freedom (DOF) object localization [5].

Despite their pervasive use in robotics, however, relatively little significant analysis is conducted to validate these sensor systems’ performances. This becomes particularly problematic as their utility in robot testbeds increases. There is thus a real need for methods for the systematic evaluation of the accuracy and performance of these sensor systems targeted toward robot applications.

In this report, we describe a framework for evaluating consumer-grade sensor systems for 3DOF object pose estimation, and apply it to the analysis of a low-cost 3D sensor using traceable instruments. A 3D sensor is given to be a logical black-box combination of software and hardware that generates collections of coordinate points in Cartesian space. Section 2 provides the procedures for fitting geometric primitives for performance analysis, while Section 3 outlines the framework for evaluating consumer-grade 3D sensors. Section 4 provides a sample analysis using the sensor under test.

2. ARTIFACT FITTING AND ASSESSMENT

A central component of the performance evaluation of any 3D sensor is the necessary ability to assess the fit of the depth data to any artifacts employed for pose analysis. Here we briefly describe the algorithms used to fit depth data to the geometric primitives used in the performance tests.

*jeremy.marvel@nist.gov; phone 1-301-795-4592; fax 1-301-990-9688

2.1 Sphere Fitting

Point clouds are segmented to remove non-artifact data points for each sample repetition. The remaining N measured data points, $\mathbf{p}_j \in P_{\{N\}}$ $\mathbf{p}_j = [x_j, y_j, z_j]$, were fitted to points on the surface of a sphere of known radius r , and center $\hat{\mathbf{s}} = [\hat{x}, \hat{y}, \hat{z}]$. Rather than orthogonally fitting points to the sphere, which matches a measured point to the closest point on the surface, a directional fitting algorithm is used for evaluation. The details of this algorithm are described in [6], but the basic fit error equations are summarized here.

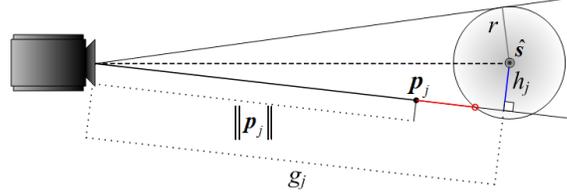


Figure 1. The error function used in the sphere fitting is based on a distance between a measured point (black dot) and its projection onto a sphere surface (red dot) in the scanning direction.

The directional sphere fit error function is based on the distance between \mathbf{p}_j and its projection onto the surface closest to the sensor along the direction of \mathbf{p}_j from the sensor (Figure 1). To aid in the computation of the sphere fitting error, the following notations are defined:

$$\mathbf{s}_j = \frac{\mathbf{p}_j}{\|\mathbf{p}_j\|_2} \quad \mathbf{g}_j = \mathbf{s}_j \bullet \hat{\mathbf{s}}$$

$$\mathbf{h}_j = \mathbf{s}_j \times \hat{\mathbf{s}}$$

Where \bullet and \times are the dot product and cross product, respectively, and $\|\cdot\|_2$ is the 2-norm length. The distance from the sphere center to the line defined by the unit vector \mathbf{s}_j is solved as the $h_j = \|\mathbf{h}_j\|_2$, and g_j and h_j are expressed in Cartesian coordinates as

$$g_j = \hat{x}x_j + \hat{y}y_j + \hat{z}z_j \quad (1)$$

$$h_j = \sqrt{(\hat{z}y_j - \hat{y}z_j)^2 + (\hat{x}z_j - \hat{z}x_j)^2 + (\hat{y}x_j - \hat{x}y_j)^2}. \quad (2)$$

The directional sphere fitting error function is expressed as

$$\varepsilon_s = \frac{1}{N} \sum_{j=1}^N E_{S_j}(\hat{\mathbf{s}}) \quad (3)$$

where

$$E_{S_j}(\hat{\mathbf{s}}) = \begin{cases} \left(g_j - \sqrt{r^2 - h_j^2} - \|\mathbf{p}_j\|_2 \right)^2, & \text{if } h_j < r \\ \left(g_j - \|\mathbf{p}_j\|_2 \right)^2 + (h_j - r)^2, & \text{if } h_j \geq r \end{cases}. \quad (4)$$

2.2 Plane Fitting

Depth data segmented to remove non-artifact points. The remaining N measured points, $\mathbf{p}_j \in P_{\{N\}}$, lie on a planar surface with distance d to the sensor when they satisfy

$$\mathbf{p}_j \bullet \mathbf{w} = d \quad (5)$$

where \mathbf{w} is a unit vector, normal to the hypothetical plane, defined by the elevation, ϑ , and azimuth, φ

$$\mathbf{w} = [\cos \vartheta \cos \varphi, \cos \vartheta \sin \varphi, \sin \vartheta] \quad (6)$$

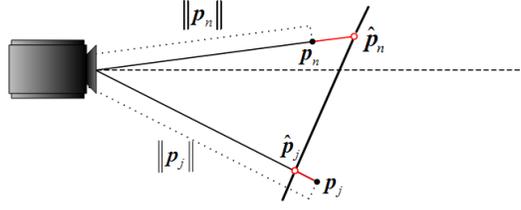


Figure 2. The error function used in the plane fitting is based on a distance between a measured point (black dot) and its projection onto a planar surface (red dot) in the scanning direction.

Unlike orthogonal plane fitting [7], which projects a point along a normal of a theoretical plane, the directional fitting algorithm (Figure 2) projects experimental points along a ray originating from the sensor, passing through \mathbf{p}_j , and intersecting with the plane at point $\hat{\mathbf{p}}_j = [\hat{x}_j, \hat{y}_j, \hat{z}_j]$:

$$\hat{\mathbf{p}}_j = t_j \mathbf{p}_j \quad (7)$$

where $t_j > 0$ is a scalar that approaches 1 as the measurement error approaches 0. The value of t_j can be calculated as

$$t_j = \frac{d}{\|\mathbf{p}_j\|_2 \mathbf{w} \bullet \mathbf{p}_j} \quad (8)$$

where

$$\mathbf{p}_j = [\cos \vartheta_j \cos \varphi_j, \cos \vartheta_j \sin \varphi_j, \sin \vartheta_j]. \quad (9)$$

The directional plane fitting error is then calculated as

$$\varepsilon_p = \frac{1}{N} \sum_{j=1}^N \left(\frac{E_p(\mathbf{w}, P_{\{N\}})}{\mathbf{w} \bullet \mathbf{p}_j} - \|\mathbf{p}_j\|_2 \right)^2 \quad (10)$$

where

$$E_p(\mathbf{w}, P_{\{N\}}) = d = \frac{\sum_{j=1}^N \|\mathbf{p}_j\|_2 (\mathbf{w} \bullet \mathbf{p}_j)^{-1}}{\sum_{j=1}^N (\mathbf{w} \bullet \mathbf{p}_j)^{-2}}. \quad (11)$$

2.3 Relative Transformations

Spatial accuracy between depth readings can be assessed by analyzing the effects of changing artifact poses. By addressing only the relative transformation, the necessity of registering coordinate frames between the ground truth and sensor under test is eliminated provided the scales are congruent. The estimated 4x4 homogeneous transformation matrix, $\hat{\mathbf{H}}$, defines the change in orientation and position of an object from one 6DOF pose to another.

$$\hat{\mathbf{H}}(\hat{\mathbf{R}}, \hat{\mathbf{T}}) = \begin{bmatrix} R_{1,1} & R_{1,2} & R_{1,3} & T_X \\ R_{2,1} & R_{2,2} & R_{2,3} & T_Y \\ R_{3,1} & R_{3,2} & R_{3,3} & T_Z \\ 0 & 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} \hat{\mathbf{R}} & \hat{\mathbf{T}} \\ 0 & 1 \end{bmatrix}. \quad (12)$$

When compared with a ground truth, $\mathbf{H}_{GT}(\mathbf{R}_{GT}, \mathbf{T}_{GT})$, the accuracy of the orientation and translation can be evaluated independently. The angle-axis rotational error is the scalar difference, θ , between the ground truth and estimate:

$$\|\varepsilon_R\|_F^2 = \|\mathbf{R}_{GT} - \hat{\mathbf{R}}\|_F^2 = 6 - 2(1 + 2 \cos \theta) \geq 8, \quad (13)$$

where $\|\cdot\|_F$ is the Frobenius norm. The translation error is the difference in Euclidian distances:

$$\varepsilon_T = \|\mathbf{T}_{GT} - \hat{\mathbf{T}}\|_2. \quad (14)$$

A limitation of the transformation evaluation is the dependence on known transforms from a sensor's coordinate frame to the artifact's coordinate frame. If such a coordinate frame for the artifact has not been explicitly defined, a common reference point must be provided. This step is optional for spheres as rotations about their centers are ubiquitous, but is required for planes, which otherwise require a one-to-one mapping of points between the ground truth and test sensors. Thus a sphere should be placed within the reported ranges of operation of both the ground truth and test sensors. The sphere's center defines the origin of the new coordinate frame.

3. EXPERIMENTAL FRAMEWORK DESIGN

Our objective is the design and assessment of an experimental evaluation of consumer-grade 3D sensors. Five evaluative tests are proposed to register the properties of such sensors that are relevant to robotics:

1. Characterizing measurement drift as a function of time
2. Assessing performance as the artifact distance increases
3. Characterizing spatial accuracy
4. Measuring accuracy as angles of inclination (AOI) vary
5. Capturing variance of sensor data.

Three of the five tests (2, 3, and 4) require a known ground truth in order to assess measurement error, while the remaining two measure the sensor based on the values it reports. As a general rule, a given ground truth sensor must be at least an order of magnitude more accurate than the sensor under test.

The number of samples, Q , for any given test should be large enough that one is confident (95th percentile) in the fidelity of the data within two standard deviations. The sample size is based on the precisions and biases of the sensor vendor's specifications:

$$Q \geq \left(z_\beta \frac{\sigma}{\mu} \right)^2 \quad (15)$$

where z_β is the normal curve value corresponding with a given error rate, β , and σ and μ are the standard deviation and mean, respectively, of the measurement error. The number of repetitions per sample, M , should also be statistically significant in order to capture the deviation in sampling. As a default, a value of $M = 30$ may be used. The uncertainties of angular measurements are smaller than those of range measurements [8], so the uncertainties of the bearings of every p_j are ignored. All measurements are assumed to be statistically independent.

3.1 Sensor Drift

Understanding how a given sensor performs as a function of time characterizes the capacity for short- and long-term sensing accuracy (e.g., as a function of internal temperature). Place a planar artifact d meters from the sensor along the sensor's principal viewing axis, \bar{Z} , with the surface of the plane orthogonal to the direction of \bar{Z} (Figure 3-A), where d is within the vendor's specified operational range. Beginning with a cold sensor (i.e., one that has been unpowered for several hours), power on the sensor and begin recording data, noting any time lapse between power on and the time that recording began. Using a sample size of $N \geq 100$ points at the center of the depth reading for each time step, record the average Z axis depth and the variance of the depth data relative to the nominal sensor distance.

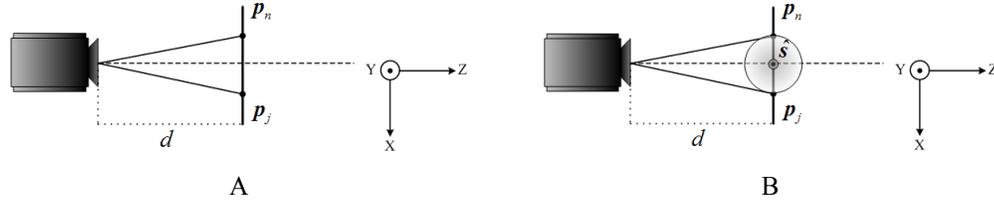


Figure 3. Illustration of the target range setup for planar (A) and spherical (B) targets. For a given nominal distance from the sensor under test, d , along the sensor's principal viewing axis, the sphere should be positioned on its center, and the plane surface should be oriented to be orthogonal to the viewing axis.

3.2 Range Accuracy

A given 3D sensor will report a measured maximum and minimum range of operational efficacy. Evaluating the sensor's performance within and beyond these bounds characterizes the sensor's measurement error as a function of distance from the sensor. For best results, both spherical and planar artifacts should be used for these tests in order to also capture the sensor's ability to capture the full complexity of the effects of distance on artifact fitting.

The artifacts are tested independently, and sampled along \vec{Z} (Figure 3). For each, ground truth is also gathered, and surfaces fit for both sensors. For varying values of d , the Q samples are taken regularly in a range that is inclusive of the vendor's specified range of operation. Compare the relative distance values between the measured spheres and between the spheres and planar targets with those measured by the ground truth and report the variances of the differences.

3.3 Spatial Accuracy

The characterization of the sensor's spatial resolution is based on its performance for relative translations incurred by varying X, Y, and Z coordinates of identical spheres within a scene. For these trials, two sphere targets are placed in known locations within the sensor's range of operation according to the pose schedule shown in Figure 4-A. The planes defined by ABCD and EFGH are parallel to each other, and orthogonal to \vec{Z} . The parallel planes are situated at nominal distances d_1 and d_2 from the sensor (Figure 4-B), with sufficient separation distance to ensure significance. Though it is assumed that the distances from \vec{Z} to the opposing lateral edges of the two depth planes (e.g., \overline{AC} and \overline{BD}) are symmetric, and that the vertices defining the lateral edges are collinear along the X and Y axial planes (Figure 4-A), no such restrictions are imposed.

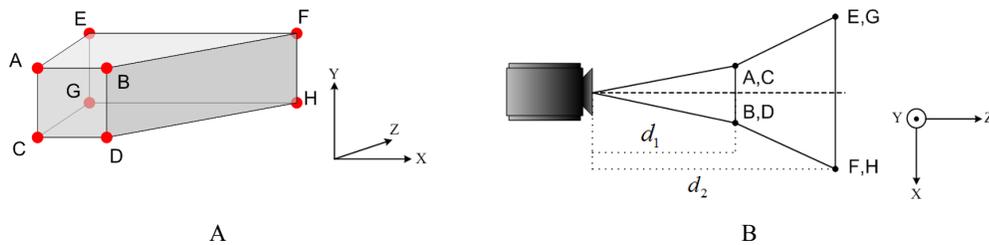


Figure 4. Illustration of the spatial accuracy pose schedule. Sphere targets are evaluated in eight sample positions in a trapezoidal configuration representing X, Y, and Z axis translations (A) configured along two nominal depth planes, d_1 and d_2 (B).

For each pair of vertices, compare the distance values, Δd , between each sphere pair, with those measured for the same diagonals by the ground truth sensor, and report the variances of the differences. Characterize the errors as functions of the adherence to the surface and internal limitations of the evaluation volume.

3.4 Angular Accuracy

A function of any robotic platform requiring localization or identification of objects is the necessary ability to accurately gauge the orientation of surfaces. For this evaluation, a single-axis rotation is sufficient to capture the effects of AOI as compared with a ground truth.

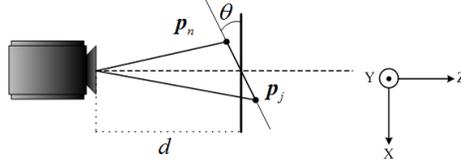


Figure 5. Illustration of the AOI test setup. A planar target is centered at distance d from the sensor under test, and rotated about its Y axis for varying angle values, θ .

Set the target at a nominal distance d from the sensor, with the rotational axis centered along \bar{Z} (Figure 5). The planar target should be at least an order of magnitude larger than the maximum axial measurement uncertainty at distance d . Vary the AOI, θ , between 0° and 90° in nominal 5° increments, and fit a planar surface to the sensor's point cloud data. Compare the relative rotational changes with those calculated from the ground truth's planar surface fits.

3.5 Sensor Variance

An analysis of the variations in the data collected from multiple measurements under identical experimental conditions provides additional insight regarding a sensor's performance as other operational parameters change. Using the sphere and planar artifact data from previous tests, evaluate the variances in the measured fit surfaces to determine sensing consistency by evaluating the average (RMS) standard deviations of the minimized surface fit errors, ε_S (3) or ε_P (10), for varying operational parameters.

4. EMPIRICAL CASE STUDY

To test the framework, we evaluated a widely-available, low-cost ($\leq \$150$) RGB-D triangulation sensor with a 1.3 MP fixed-lens 0.25 in complementary metal-oxide-semiconductor (CMOS) color camera, and an infrared (IR) emitter module paired with a SXGA fixed-lens 0.5 in CMOS depth camera. Erroneous measurements have been reported for this sensor, and some have improved its performance by augmenting depth map generation with temporal filtering [9]. Early validation efforts (e.g., [10]) of the claimed 3.0 mm X/Y spatial and 1.0 cm Z depth resolutions of similar sensor packages [11] focused on the accuracy of the depth sensing and not on the sensor's utility. More thorough investigations of measurement error and variance for specific robot applications (e.g., indoor mapping [12]) have since been published that investigate both theoretical and physical sources of data distortion.

For comparison, we used a 3D laser range finder with a measured accuracy of $15 \mu\text{m}$, $\pm 1.5 \text{ ppm}$ (2 sigma) as ground truth. For sphere artifacts, a 152 mm aluminum spherical target was used. Ground truth was measured by pivoting the artifact to expose a laser tracker target at the sphere's center. A 0.61 cm square target with 50 % visible light reflectivity was used as the planar artifact. A surface was fit to the back of the planar artifact using 20 sample measurements along the perimeter taken via laser tracker targets, and calculated forward using the artifact's known dimensions. For all five tests, $Q \geq 10$, and $M = 30$.

All data was gathered using stock hardware, though the housing of the sensor under test was modified to be rigidly mounted to a frame by means of four posts affixed to the bottom of the sensor. Raw image and depth sensor data was gathered via USB using open source libraries, though only the depth data was evaluated. Neither sensor was moved during the evaluation.

The tests were carried out in the 3D Imaging System Performance Evaluation Facility and the Intelligent Sensing and Automation Testbed at NIST. Although the facilities are not temperature-controlled, the temperatures are constantly monitored by sensors throughout the labs' volumes. Temperatures in both were, on average, 20.03°C , with only minor fluctuations ($\leq 0.0014^\circ\text{C}$) during the evaluation. While the sensor under test uses structured IR light projection, the labs' lighting is provided by fluorescent fixtures 10.24 m and 4.8 m above the ground, respectively, minimizing IR noise. The sensor was tested with the lights both on and off, and no discernable effect of lighting changes was observed.

4.1 Sensor Drift

A planar target was set up according to Section 3.1 at a nominal 0.9 m distance from the sensor under test. An average time of 17.32 s passed between the powering of the sensor and the first sensor reading. Each sample consisted of

approximately 3 600 s of data recording a (30x30) pixel region from the center of the depth image. The average depth reading and variance of the data are plotted in Figure 6.

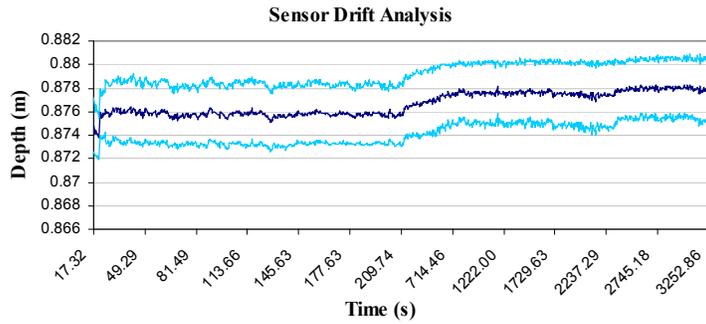


Figure 6. Results of sensor drift testing at 0.9 m nominal distance. The dark line is the average depth reading of a (30x30) pixel area, with the lighter lines being a single standard deviation up and down for each time step.

The variance remained consistent at about ± 2.0 mm, but the depth readings as a whole drifted upward by ≈ 4.0 mm over the duration of the test. Short term drift was well within the purported 1.0 cm Z depth resolution of the sensor. The long-term drift, however, indicates degradation in sensor reliability.

4.2 Range Accuracy

It is reported that the sensor under test performs optimally between 0.8 m and 3.5 m distance from a target [11]. To validate this, the sphere and planar targets were sampled along a path inclusive of this range. The sphere was sampled from 0.7 m to 3.1 m at 0.2 m nominal distances, and a last sample was taken at 4.1 m. The planar target was nominally sampled between 0.6 m and 3.0 m every 0.2 m, and then between 3.0 m and 4.5 m at regular 0.5 m nominal steps. The relative distances between detected artifacts (spheres to spheres, and spheres to planes) were compared with the ground truth; the results are shown in Figures 7 and 8.

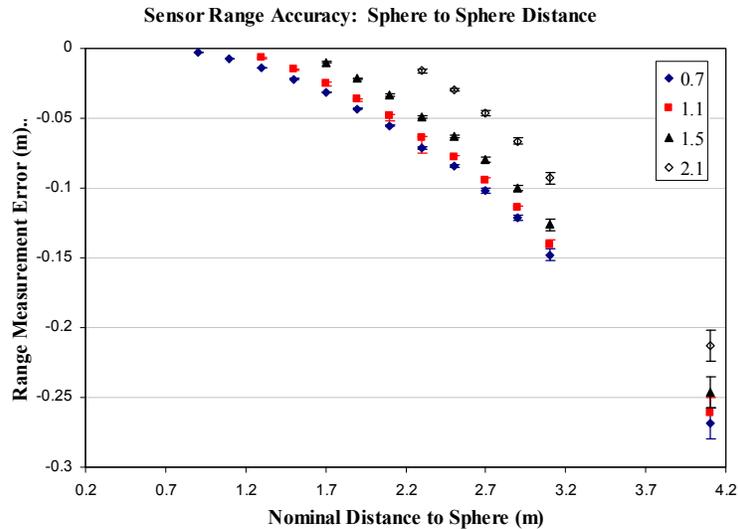


Figure 7. Signed measurement error and variance of spherical targets for the sensor under test compared with the ground truth. Comparisons between spheres at varying ranges (e.g., relative to the sphere at nominal 0.7 m, black diamond, or 1.1 m, red square) result in smaller errors as the relative distances decrease.

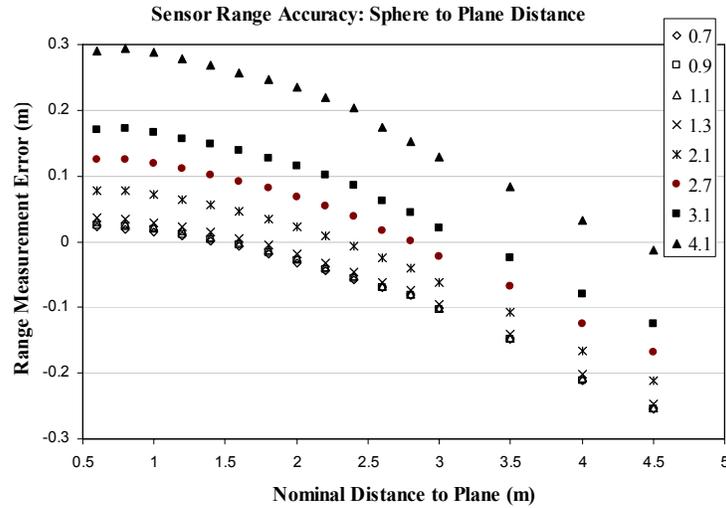


Figure 8. Signed range measurement error of planar targets for the sensor under test compared with the ground truth. Plane surface measurements are made relative to spheres placed at different depths.

While the absolute depth accuracy degrades (which is consistent with [10]), the relative depth accuracy is more stable. For example, with the sphere-to-plane analysis, there is a ≈ 0.25 m error in evaluation for the planar target at the nominal 4.5 m mark when compared with the sphere at the nominal 0.7 m mark. In contrast, there is only a ≈ 1.0 cm error relative to the spherical artifact at nominal 4.1 m.

4.3 Spatial Accuracy

The ground truth of all eight scheduled poses was established using the laser tracker and compared with the fitted spheres using the sensor under test (Table 1). The experimental point cloud data was fitted to two spheres. The translations of the resulting sphere centers were compared with the ground truth; the resulting errors are provided in Table 2.

Table 1. Pose schedule for spatial sensor accuracy.

| | Nominal* | | | Fit Center Point** | | | Ground Truth** | | |
|---|----------|-------|-----|---------------------------|---------------------------|--------------------------|----------------------------|----------------------------|----------------------------|
| | X | Y | Z | X | Y | Z | X | Y | Z |
| A | -0.4 | 0.25 | 1.0 | -0.42184 0.1821 | 0.178105 0.1553 | 0.99551 0.1522 | -8.29264 0.01564 | -0.5691 0.05995 | 0.57102 0.0575 |
| B | 0.4 | 0.25 | 1.0 | 0.47878 0.2924 | 0.16699 0.2319 | 1.04703 0.6361 | -8.24826 0.01763 | -1.39357 0.05753 | 0.57104 0.06162 |
| C | -0.4 | -0.25 | 1.0 | -0.3948 0.3483 | -0.30744 0.2623 | 1.01782 0.5633 | -8.31341 0.01509 | -0.53974 0.06184 | 0.09132 0.06286 |
| D | 0.4 | -0.25 | 1.0 | 0.43813 0.1934 | 0.319 0.2003 | 1.00014 0.1276 | -8.20206 0.01713 | -1.42085 0.05886 | 0.09162 0.06083 |
| E | -1.0 | 0.6 | 2.0 | -0.7918 2.106 | 0.68843 2.258 | 2.10402 1.43 | -7.1027 0.01582 | 0.25518 0.05275 | 0.86964 0.05294 |
| F | 1.0 | 0.6 | 2.0 | 0.89429 1.275 | 0.66674 1.384 | 2.2913 2.373 | -7.20395 0.01921 | -1.58055 0.05494 | 0.90719 0.05408 |
| G | -1.0 | -0.6 | 2.0 | -1.14889 2.239 | -0.59694 1.823 | 2.34835 3.1 | -7.26436 0.0155 | -0.08669 0.05531 | -0.37934 0.0568 |
| H | 1.0 | -0.6 | 2.0 | 0.72118 4.205 | -0.65121 2.668 | 2.04573 6.648 | -6.96993 0.01882 | -1.70204 0.05414 | -0.37944 0.05343 |

* - Nominal coordinates given from perspective of the sensor under test's principal viewing axis, in meters

** - Measured (sensor under test and ground truth) coordinates provided as averages (bold), in meters, and standard deviation (red), in millimeters. Notes that the coordinate systems for the sensor under test and ground truth are disjoint, and that accuracy is correlated via relative transforms between poses.

Table 2. Spatial accuracy mean (bold) and standard deviation (red), in meters, for opposing vertices on the trapezoidal pose schedule.

| | B | C | D | E | F | G | H |
|---|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|---------------------------|
| A | -0.0075 0.0004 | -0.0058 0.0003 | -0.0129 0.0002 | -0.0785 0.0016 | -0.0321 0.0013 | -0.0452 0.003 | -0.0733 -0.0076 |
| B | | -0.0118 0.0003 | -0.0076 0.0003 | -0.0711 0.0021 | -0.0201 0.0021 | -0.0437 0.0031 | -0.0647 0.0072 |
| C | | | -0.014 0.0005 | -0.0761 0.0017 | -0.0337 0.0015 | -0.0394 0.003 | -0.0727 0.0083 |
| D | | | | -0.0694 0.002 | -0.0148 0.002 | -0.0357 0.0029 | -0.0551 0.0078 |
| E | | | | | -0.0563 0.0026 | -0.0513 0.0028 | 0.0775 0.0049 |
| F | | | | | | -0.0493 0.0028 | -0.0384 0.0045 |
| G | | | | | | | -0.0547 0.0059 |

As was seen in [10], the spatial accuracy declines as a function of distance. Although, from the results, it can be shown that the sensor’s Y axis accuracy is marginally better than that of its X axis. In all cases, however, the fitted surfaces from the sensor’s point cloud data consistently placed the sphere’s center too far laterally from \bar{Z} .

4.4 Angular Accuracy

The target was positioned nominal 1.5m from the sensor, with the central pivot on the plane’s Z axis situated directly on \bar{Z} . For each angular sample, ground truth was gathered using the laser tracker and fitted to a planar target. The point clouds were also fit to planar surfaces, and the resulting angular estimations were compared with the ground truth. Although AOIs varied between 0° and 90°, beyond 75° not enough of the target could be seen by the sensor under test to accurately fit a plane. So the results of the evaluation do not include measurements for the values 80° through 90°. The results of comparing angular readings relative to the nominal 0° plane fitting are shown in Figure 9.

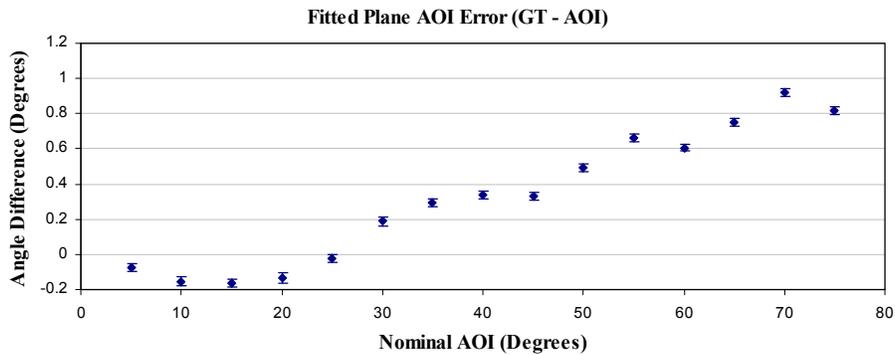


Figure 9. Angles of inclination error (solid line) and measurement variance (error bars). As the AOI increases, so too does the magnitude of error. Beyond $\theta = 75^\circ$ a planar surface could not be accurately fit to the data.

The angular error increased as a function of AOI, though variances in estimation error remained small. Further, at the nominal distance of 1.5 m, the magnitude of all fit errors for 5° to 75° remains less than 1°.

4.5 Sensor Variance

Using the planar and spherical artifact data from the depth and AOI tests, the variances in the measured fit surfaces were evaluated to determine sensing accuracy. The results of these evaluations are illustrated in Figure 10 to Figure 12. In the case of the range variance data (Figure 10 and Figure 11), we saw the fitting error increase exponentially as a function of distance. With the spherical artifact data, we also witness an increase in the spread of the fit error, while the spread of the planar depth data remained confined. In contrast, the AOI test data’s fit error and spread remained small

and consistent throughout the test at 4mm, but increased sharply as the visible areas of the planar artifact's surface decreased.

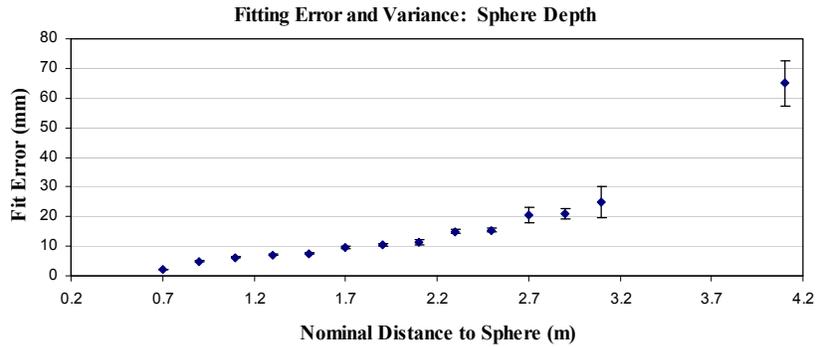


Figure 10. Average and one standard deviation of fit error for the range accuracy tests using the sphere artifacts. As the distance from the sensor to the sphere artifact increased, so too did the fit error and spread of the data.

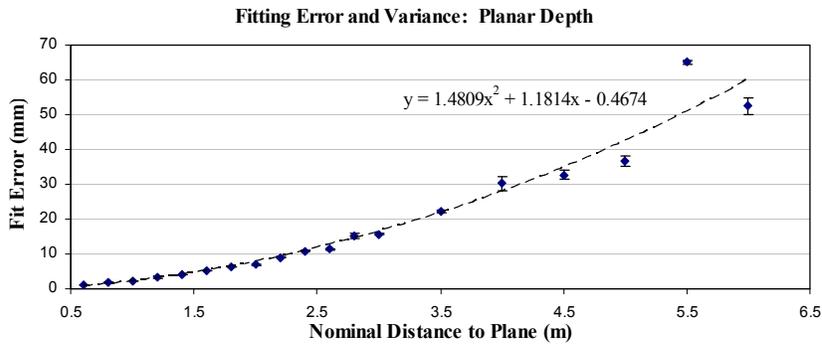


Figure 11. Average and one standard deviation of fit error for the range accuracy tests using the planar artifacts. The fit error for the planar artifact increased as a function of distance, though the variance remains small.

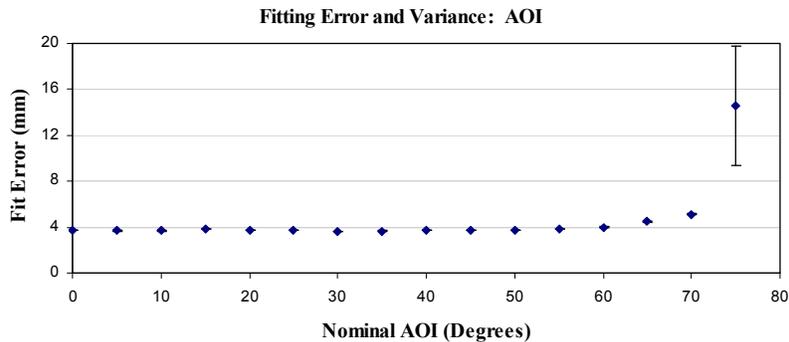


Figure 12. Average and one standard deviation of fit error for the AOI tests using the planar artifacts. The fit error remained consistent through the test, but increased significantly as the target surface became unobservable.

In both cases, the variance was a function of both sensor error and exponential drops in observable data points. For instance, with the range test, the number of data points fit to a sphere drops from $\approx 13\ 000$ points at 0.7 m to $\approx 6\ 500$ points at 0.8 m, and all the way down to ≈ 500 data points at 4.1 m. Similarly, the AOI tests drop from $\approx 52\ 600$ points at 0° down to $\approx 11\ 000$ points at 75° .

5. DISCUSSION

In this report we described a 3D sensor measurement evaluation framework for robotics applications, and applied it to a low-cost, consumer-grade 3D sensor. As the availability and utilization of low-cost sensors increases, their performance and reliability must be systematically assessed. Understanding the limitations of these sensors helps us to improve the reliability of the data gathered from them. The framework provided herein is conceptually comprehensive, and provides a user with information sufficient to classify the accuracy and precision of sensors for potential application in their research.

REFERENCES

- [1] Khandelwal, P, and Stone, P., "A Low Cost Ground Truth Detection System for RoboCup Using the Kinect," Proc. RoboCup Int. Symp., (2011).
- [2] Fiala, M., and Ufkes, A., "Visual Odometry Using 3-Dimensional Video Input," Proc. Can. Conf. Comput. & Rob. Vision., 86-93 (2011).
- [3] Frati, V., and Prattichizzo, D., "Using Kinect for Hand Tracking and Rendering in Wearable Haptics," Proc. IEEE World Haptic. Conf., 317-321 (2011).
- [4] Ramey, A., González-Pacheco, V., and Salichs, M. A., "Integration of a Low-Cost RGB-D Sensor in a Social Robot for Gesture Recognition," Proc. 6th Int. Conf. Human-Rob. Interact., 229-230 (2011).
- [5] Newman, M., and Balakirsky, S., "Contests in China Put Next-Gen Robot Technologies to the Test," RA Mag., 18(3), 131-133 (2011).
- [6] Franaszek, M., Cheok, G. S., Saidi, K. S., and Witzgall, C., "Fitting Spheres to Range Data from 3-D Imaging Systems." IEEE Trans. Instrum. Meas., 58(10), 3544-3553 (2009).
- [7] Franaszek, M., Cheok, G. S., Saidi, K. S., "Experimental Verification of Formulas for Variances of Plane Parameters Fitted to Three-Dimensional Imaging Data," IEEE Trans. Instrum. Meas., 61(1), 103-110 (2012).
- [8] Stone, W. C., Juberts, M., Dagalakis, N., Stone, J., and Gorman, J., "Performance Analysis of Next-Generation LADAR for Manufacturing, Construction, and Mobility," NISTIR, 7117, 4.7-4.12 (2004).
- [9] Mayunin, S., Vatolin, D., Berdnikov, Y., and Smirnov, M., "Temporal Filtering for Depth Maps Generated by Kinect Depth Camera," Proc. 3DTV Conf.: The True Vision – Capture, Transm., and Disp. 3D Video, 1-4 (2011).
- [10] Viager, M., "Analysis of Kinect for Mobile Robots," Thesis project report: Technical University of Denmark, Department of Electrical Engineering, (2011).
- [11] Primesense Ltd., The PrimeSensor™ Reference Design 1.08 Data Sheet, (2010).
- [12] Khoshelham, K, and Elberink, S. O., "Accuracy and Resolution of Kinect Depth Data for Indoor Mapping Applications," Sensors, 1437-1454 (2012).