

Evaluation of point cloud filtering techniques in object dimensioning via vision-based geometric measurement systems

Karol Kuziora^{1*}, Robert Sereda^{1*}, Marcin Badurowicz¹, Konrad Smoliński²

¹ Department of Computer Science, Lublin University of Technology, Nadbystrzycka 36B, Lublin, Poland

² Department of Computational Intelligence, Lublin University of Technology, Nadbystrzycka 38, Lublin, Poland

* Corresponding author's e-mail: kuziora.karolx@gmail.com

ABSTRACT

In this study, we evaluate the influence of different cloud point filtering algorithms on the process of accurately dimensioning objects. This is critical in vision-based measurement systems, particularly for logistics and packaging applications. We assess three smoothing algorithms: bilateral filtering, statistical outlier removal, shadow filtering algorithm alongside baseline unfiltered data. We extract object dimensions by fitting a convex hull applied to the processed point cloud, and evaluate across different positions, parcel types, and edge lengths. We employ various statistical metrics to evaluate algorithm performance. Our research utilizes point clouds of cardboard boxes for evaluation, collected with the ToF Kinect v2 depth camera. Study includes both cuboidal objects and distortion-simulated shapes. We assessed a dataset of 639-point cloud samples. The data was collected under controlled lighting with top-down camera orientation and processed using the PCL library. Our findings show that shadow filtering consistently and significantly outperforms the other methods on standard cuboid geometries. However, in the presence of shape distortions, it occasionally introduces large-magnitude outliers, reflecting overly aggressive filtering behaviour. Additionally, we observe scale-dependent error pattern across all object types, with dimensional accuracy decreasing as object size increases.

Keywords: point cloud processing, time of flight depth camera, filtering algorithms.

INTRODUCTION

Over the past two decades, depth-sensing technologies such as light detection and ranging (LiDAR), time-of-flight (ToF) cameras, structured-light RGB-D sensors, and passive stereo vision systems have evolved from specialized laboratory tools to widely adopted instruments across numerous applied domains. Despite significant advances in acquiring detailed three-dimensional (3D) data, point cloud processing continues to face persistent challenges that remain evident in photogrammetric workflows [1]. These include noise, occlusions, non-uniform point densities, and difficulty preserving accurate geometric features. These challenges have become more prominent as depth-sensing system transition from specialized metrology laboratories to a broader

range of practical applications in uncontrolled or dynamic environments. This research addresses these challenges within the specific context of object dimensioning using consumer-grade depth sensors, specifically the Kinect sensor. It focuses on improving the interpretation of sensor data and enhancing volumetric measurement accuracy under real-world conditions, where cost-effective and robust solutions are critical for logistics, packaging, and industrial automation.

Applications in industrial environments

Depth sensing technologies have found widespread use in industrial environments, enabling accurate 3D reconstruction, quality control, and robotic automation. Hoegg et al. developed a multi-camera Time-of-Flight system for real-time

3D vehicle reconstruction, improving automated car-wash optimization [2]. In a separate automotive context, Józwick and Dziedzic employed optical 3D scanning to analyse the geometric accuracy of polymer fuel tanks, enabling dimensional verification of complex components directly on the production line [3]. In rubber processing, Huang et al. used an RGB-D sensor and point-cloud analysis to measure tread profiles with millimetre-level accuracy, providing a reliable inline alternative to manual callipers [4]. More recently, some approaches have explored the fusion of RGB image features with point cloud data to aid object boundary refinement in cluttered or reflective environments [5], highlighting ongoing interest in improving point cloud interpretation through complementary sensing modalities. An illustrative example of close-range human-sensor interaction is presented in the context of human-robot collaboration [6]: Kinect sensors were deployed in a shared robotized work-cell, with a focus on sensor calibration and integration within ROS-based control environments. This demonstrated the practical viability of low-cost ToF sensing for interactive, safe, and dynamically reconfigurable industrial setups.

Precise 3D object maps are crucial for domestic robots in cluttered environments. Rusu et al. used a laser scanner on a manipulator arm to create semantically annotated kitchen models, linking geometric perception with symbolic action planning [7]. Finally in the context of transport logistics, dimension detection of irregularly shaped items such as airline baggage has also been addressed using purely geometric methods. The minimum elastic bounding box (MEBB) algorithm adapts bounding surfaces to point cloud structure, effectively reducing over-estimation errors from appendages and noise in cluttered settings [8].

Consumer-grade depth-sensing hardware

Depth cameras form the backbone of many 3D workflows, but their performance ultimately hinges on sensor quality. In the consumer and prototyping space, Microsoft and Intel remain the dominant vendors, with a variety of third-party and industrial systems available for more specialized needs.

Scientific studies have shown that new advancement Azure Kinect, as the successor to Kinect v2, significantly improves depth precision (≤ 17 mm SD) and distance accuracy (< 11 mm + 0.1%) thanks to its time-of-flight sensing [9]. It

demonstrates about half the random depth error of Kinect v2 at close range and maintains under 2 mm spatial accuracy from 1 to 2 meters. These improvements make it more reliable in controlled indoor settings, though limitations remain under reflective or bright outdoor conditions [10].

Comparative studies confirm the Kinect v2's past value, while emphasizing ToF systems' edge over stereo-based ones. Azure Kinect and Blaze 101 outperformed Intel RealSense D435 and OAK-D in depth accuracy, residuals, and performance under varied lighting and textures [11]. In medical use, Kinect v2 also showed better 3D reconstruction and repeatability than the D435, despite the latter's faster frame rates and customization options [12].

Stereo-based systems like the RealSense D415, which uses active stereoscopy, have proven effective in challenging settings with transparency and translucency. Its IR-enhanced stereo method is less impacted by light scattering than the ToF-based L515 or Structured Light-based SR305 [13]. Metrological studies also show the D415 excels at close range, the L515 is best on flat surfaces, and the D455 offers the strongest overall accuracy by ISO standards [14].

Recent LiDAR advances, particularly in photon-counting LiDAR, show strong potential for high-precision use in extreme conditions. These systems perform well in bright sunlight, offering over 24 km range and high-resolution 3D imaging – far beyond what consumer-grade sensors can achieve [15].

In this context, Kinect v2 continues to serve as an important reference benchmark for evaluating newer systems. While it has been surpassed in several technical dimensions, its role in shaping expectations and methodologies for depth accuracy and performance across real-world scenarios remains vital.

Challenges in LIDAR devices and algorithms

As 3D sensors produce more detailed and dense data, algorithmic solutions play a key role in handling common point cloud issues. Challenges like noise, outliers, occlusions, uneven non-uniform densities, and geometric feature loss can seriously affect tasks such as segmentation, classification, and measurement if not properly addressed.

Early methods like the bilateral mesh filter (BMF) [16] improved on standard bilateral filtering by better preserving edges and reducing noise

in mesh data. It also showed higher PSNR (peak signal-to-noise ratio) scores in grayscale contexts, highlighting its effectiveness.

A detailed review [17] found L_0 minimization and Edge Aware Resampling (EAR) to be highly effective for noise removal and structure preservation in 3D point clouds, though they are computationally intensive. In contrast, simpler methods like the voxel grid (VG) filter offer faster performance at the cost of geometric accuracy for precision-critical tasks.

Recent methods like Elliptic Gabriel Taubin (EGT) smoothing [18] improved efficiency by using only point coordinates. EGT outperformed techniques like moving least squares (MLS), Jet Smoothing, and Weighted Locally Optimal Projection (WLOP) in preserving features, while using less memory and running faster by avoiding repeated neighbour searches.

Segmentation and classification are key for object-level understanding in point clouds. RANSAC remains a fast, reliable choice for simple geometric forms, outperforming the Hough Transform in noisy conditions. However, it lacks automation and struggles with complex geometries, highlighting the need for adaptive or learning-based methods [19]. New denoising techniques based on normal voting tensors and constrained quadratic error metrics – better preserve sharp features and structure, with lower angular deviation and minimal shrinkage compared to earlier methods [20].

The field has increasingly moved toward deep learning methods. Architectures such as graph convolutional networks and autoencoders show strong performance in handling noise and preserving fine details [21]. While they often outperform classical techniques in noise reduction and feature retention, they require substantial training data and high computational resources which can hinder real-time application in lightweight systems.

Object dimensioning approaches

Recent advancements in object dimensioning integrate algorithmic improvements with enhanced depth-sensing hardware. Ladplee et al. (2022) [22] showed real-time parcel dimensioning with a single LiDAR camera using flood fill-based top-plane detection. Rodriguez et al. (2023) [23] improved accuracy on noisy ToF data through bounding-based superquadric fitting for both cuboid and cylindrical objects.

This study extends prior approaches by systematically evaluating the impact of point cloud filtering techniques on volumetric estimation accuracy in depth-sensing systems. It specifically addresses practical challenges related to noisy data interpretation and geometric models fitting under varying object shapes and data acquisition conditions. This work contributes to the development of robust, accurate, and computationally efficient, object dimensioning pipelines suitable for practical deployment in real-world applications.

MATERIAL AND METHODS

This study focuses on the acquisition, processing, and geometric analysis of point cloud data to enable accurate dimension extraction of rigid objects. The experimental methodology integrates consumer-grade depth sensing using Kinect v2 device, data acquisition protocols, and a customized point cloud processing pipeline designed to assess the influence of filtering techniques on volumetric measurement accuracy.

The process is detailed in the further parts of the following study. This approach ensures reliable and repeatable dimensional measurements based on point cloud analysis.

Sensor configuration and experimental setup

Depth Data were captured using a Microsoft Kinect v2 ToF sensor mounted in a nadir (top-down) orientation, clearly visualizing the surface of a dedicated workbench. The experimental setup featured a laboratory workbench with a horizontal surface elevated 60 cm from the floor. The Kinect sensor was attached to an overhead arm positioned 140 cm above the workbench surface. The workspace was illuminated uniformly by fluorescent lighting, provided by six lamps, each rated at approximately 5000 lumens, distributed across a 16 m² controlled laboratory environment to minimize shadowing and specular interference.

Evaluated object dataset

The dataset preparation involved capturing point clouds and corresponding colour, infrared (IR) images of 14 types of objects, each placed at the center/edges of the table in varying orientations. While these objects are referred to as boxes in the technical setup, we sometimes use the term

parcels in the context of our application-oriented approach, as it better reflects practical, real-world scenarios. Parcels were systematically positioned on each of three unique side surfaces (considering distinct form-factor features) and rotated along the Z-axis. In total, the final dataset encompassed 639 captures of the table's upper surfaces alongside the parcels, incorporating relevant calibration data for each measurement scenario.

The dataset includes both standard cuboidal geometries and distorted variants with non-orthogonal edges intended to simulate real-world packaging irregularities. The test objects varied in edge lengths size from as small as 3.3 cm up to 80 cm. They were grouped into the following types:

- Standard cuboidal objects (8 variants): rigid, cuboidal-shaped items with well-defined orthogonal edges and planar surfaces.
- Distorted objects (6 variants): items with low height profiles, slanted edges, non-parallel faces or irregular shapes.

This package set was deliberately chosen to test the robustness and adaptability of dimensioning methods across a range of challenging, practical scenarios.

Data processing pipeline

Point cloud data were processed using a custom-built pipeline designed to isolate and measure individual packages from unfiltered sensor output. The pipeline integrated calibration, filtering, segmentation, and geometric analysis, with particular focus on evaluating the impact of smoothing techniques on dimensional accuracy.

Calibration and ROI definition

We performed initial calibration using physical markers placed in the scene, enabling the identification of table boundaries and defining the region of interest (ROI). The table surface was subsequently aligned with the XY-plane. This alignment allowed us to extract object height directly from Z-coordinates relative to the tabletop.

Noise reduction

We filtered out non-numeric or undefined values, points beyond the sensor's operational range, and spurious points such as those arising from shadowing effects at depth discontinuities.

Smoothing and surface refinement

Three algorithms - Bilateral filtering, statistical outlier removal (SOR) and custom developed Shadow filtering – were tested to assess their impact on local surface consistency and the final dimension estimates (parameters for those steps were carried manually to ensure proper parameters for general use). This step was crucial given the known variability in point density and noise across the sensor's field of view.

Segmentation and object isolation

We segmented each object in two stages. First, we identified and removed the dominant plane corresponding to the table surface using SAC-based model fitting. Next, we applied Euclidean Cluster Extraction to isolate individual objects. We refined cluster boundaries using Progressive Morphological Filtering with adaptive thresholds.

Dimension estimation

We extracted object dimensions using two steps approach. We estimated height by averaging the Z-values of the topmost 1.5% points relative to the table plane in each cluster. To estimate width and length, we computed minimal surface in the XY-plane of the segmented point cloud and extracted the axis-aligned bounding box. We then compared the estimated dimensions against manually recorded ground-truth measurements to evaluate the accuracy of the pipeline under different smoothing configurations.

Accuracy metrics

We assess accuracy for each of the three dimensions independently, reflecting their different computational origins: length and width are derived from a 3D bounding box, while height is taken from the highest points in the 3D point cloud. We used mean absolute error (MAE) to quantify the average magnitude of absolute error between estimated and ground-truth measurements.

To investigate spatial consistency in sensor performance, we introduced Corner-to-Center Error Ratio. This metric compares dimensioning accuracy between two spatial placements of the object on the workbench: the geometric centre and the table corners. It highlights potential edge-field performance degradation due to sensor perspective and occlusion.

EVALUATION AND RESULTS

We analysed how each filtering technique affected both dimensional accuracy and computational cost. By evaluating under uniform hardware and software conditions, we ensured a fair comparison of algorithmic efficiency and effectiveness and therefore practical relevance.

Overall dimensioning accuracy

The use of various point cloud filtering techniques has demonstrated a considerable impact on the accuracy of parcel measurement. The analysis of dimensional measurement accuracy extends beyond individual spatial parameters to include derived metrics of significant commercial relevance. We evaluated the volumetric weight (DIM weight), an industry-standard metric used in shipping cost calculations. This metric serves as a billing determinant in the logistics and transport industry where volume rather than actual weight is the limiting factor.

$$DIM_{Weight} = \frac{L \cdot W \cdot H}{SF} \quad (1)$$

where: L – length of the object, W – width of the object, H – height of the object, SF – shipping factor (e.g. 5000).

Figure 1 illustrates the mean absolute error in volumetric weight estimation for each point cloud filtering technique, accompanied by the standard deviation bars. The results demonstrate that unfiltered baseline data exhibited the highest error (0.199 kg), whereas the Shadow Filtering algorithm consistently achieved the lowest (0.079 kg), corresponding to a 60.3% relative improvement over unfiltered data. Statistical Outlier Removal achieved the second-best accuracy (0.106 kg), followed by Bilateral Filtering (0.148 kg) showing moderate improvements compared to unfiltered data.

To further analyse the results, the overall measurement accuracy was decomposed into individual dimensional components. By evaluating the Mean Absolute Error, calculated according to Eq. (2), for each principal dimension (height, width, and length) we gain deeper insights into the specific strengths and limitations of each filtering method.

$$MAE_d = \frac{1}{\sum_{m=1}^M n_m} \sum_{m=1}^M \sum_{i=1}^{n_m} |x_{im}^{(d)} - \widehat{x}_{im}^{(d)}| \quad (2)$$

where: d – index of the measured dimension, with $d \in \{1,2,3\}$, corresponding to height, width and length, respectively, M – total number of objects, n_m – number of samples for the m -th object, $x_{im}^{(d)}$ – ground truth value of dimension d for the m -th sample in package m , $\widehat{x}_{im}^{(d)}$ – estimated value of dimension d of the m -th object in the i -th sample.

Dimensional error analysis, as visualized in Figure 2, reveals that measurement accuracy varies significantly across the three principal axes. Height measurements consistently exhibited the lowest error magnitudes (0.329–0.343 cm) regardless of filtering approach, likely due to their orthogonal orientation relative to the sensor. In contrast, width (0.482–0.849 cm) and particularly length measurements (0.463–1.581 cm) showed much greater sensitivity to the filtering method employed.

We evaluated the distributional properties of dimensional and volumetric weight measurement errors using the Shapiro-Wilk normality test for each filtering method (Table 1). Across all filtering techniques and measurement types, the resulting p-values were significantly below the conventional threshold of 0.05, indicating strong evidence against normality. This outcome suggests that the errors in both dimensional and volumetric weight measurements do not follow a normal distribution, regardless of whether point cloud filtering was applied. This justifies the use of non-parametric statistical methods in subsequent analysis.

Given the consistent non-normality observed in the Shapiro-Wilk tests, we employed the Kruskal-Wallis non-parametric test to assess whether the differences in measurement errors across filtering methods were statistically significant (Table 2). The results show a clear divergence in filtering impact depending on the measured dimension. Results showed no significant difference in height error ($p = 0.83$), aligning with its overall measurement stability. In contrast, width, length, and volumetric weight




Figure 1. Comparative analysis of volumetric weight estimation precision across filtering techniques, presented as mean absolute error (kg) with corresponding standard deviation error indicators




Figure 2. Quantitative assessment of dimensional measurement accuracy across point cloud filtering methods, illustrating Mean Absolute Error values with 95% confidence intervals for principal spatial dimensions. For each dimension, the optimal results were annotated

errors differed significantly ($p < 0.05$), confirming that filtering has a statistically significant effect – particularly on planar dimensions and derived metrics like volumetric weight.

Following the identification of statistically significant differences through the Kruskal-Wallis test (Table 2), we performed a comprehensive post hoc pairwise comparisons analysis

using Dunn’s test to determine specific differences between filtering methodologies. Figure 3 presents the results of this analysis through heat maps of p-values. For width, length and volumetric weight error demonstrated significant differences in the initial analysis.

Post hoc comparisons reveal distinct performance clusters among the filtering approaches.

Table 1. Shapiro-Wilk normality test for dimensional and volumetric weight errors across different point cloud filtering methods, with p-values indicating the probability of data following normal distribution

Dimension error	Unfiltered	Bilateral filter	Shadow filter	SOR filter
Width (cm)	3.67E-26	1.33E-19	3.36E-20	2.09E-13
Length (cm)	3.77E-18	4.77E-15	4.90E-21	4.70E-11
Height (cm)	3.45E-14	8.26E-14	3.45E-14	3.56E-14
Volumetric weight (kg)	2.29E-33	7.28E-36	1.05E-28	4.64E-28

For width measurements, both the statistical outlier removal (SOR) and Shadow filters demonstrated statistically significant improvements ($p < 0.0001$) compared to the Bilateral filtering method and the unfiltered baseline (i.e., unfiltered data without any point cloud manipulation). Similarly, for length error analysis showed significant superiority of SOR and Shadow filters to two other methods, though a slight statistical difference was also detected between the SOR and Shadow methods ($p = 0.0003$).

The volumetric weight error analysis presents a more complex pattern of relationships. Statistical Outlier Filter demonstrated significant improvements over all other methodologies ($p \leq 0.0001$), while the Bilateral filter showed significant advantages over both unfiltered and Shadow approaches ($p < 0.0001$). Interestingly, no significant difference was observed between unfiltered and Shadow filtering methods for volumetric weight calculations ($p = 1.0000$), suggesting that the Shadow filter's improvements in dimensional accuracy may not translate proportionally to volumetric weight estimation.

While the Shadow filter achieved the lowest mean absolute error in volumetric weight estimation, the difference was not statistically significant (as seen in Figure 3) compared to the unfiltered method. In contrast, the SOR filter demonstrated statistically significant improvements over all other filters, indicating greater robustness and consistency. Therefore, while Shadow may offer the highest observed precision, SOR emerges as

the most statistically reliable filtering approach for practical deployment.

Parcel type

Standard objects, characterised by uniform cardboard faces and consistently sharp edges, present a geometrically regular cuboid structure. Contrastingly, distorted boxes exhibit pronounced morphological distortion. Most notably, these objects present fundamentally different edge characteristics across dimensions:

- Length measurement (longer dimension) involves faces with sharp edges, though the perpendicular cardboard face is inset by several centimetres.
- Width measurement (shorter dimension) involves faces with non-sharp, bulging edges

This inherent structural dichotomy establishes fundamentally different measurement challenges for each dimension.

Figure 4 presents a comprehensive comparison of width and length measurement errors across different filtering methods for both parcel types. Particularly striking is the disparity observed in width measurements for distorted parcels with the SOR filter, which exhibits dramatically elevated error levels (1.24 cm) compared to length measurements (0.43 cm). This pronounced discrepancy corresponds directly to the morphological differences, as SOR filtering struggles with the bulging edges characteristic of width measurement in distorted parcels.

Table 2. Summary of Kruskal-Wallis non-parametric test results examining statistical differences in measurement errors across filtering methodologies. Test statistics (H) and corresponding p-values indicate statistical significance for the given sample

Metric	Test Statistic (H)	p-value	Significant
Height error (cm)	0.87	0.83	False
Width error (cm)	153.42	4.82E-33	True
Length error (cm)	526.77	7.53E-114	True
Volumetric weight error (kg)	91.48	1.05E-19	True




Figure 3. Post Hoc analysis of filter performance using Dunn's Test. Heat maps displaying p-values from Dunn's test for pairwise comparisons of filtering methods across width error, length error, and volumetric weight error dimensions

Conversely, the unfiltered and Bilateral filtered datasets exhibit the opposite pattern for distorted parcels, with substantially higher errors for length measurements (2.56 cm and 2.43 cm respectively) than width measurements (0.87 cm and 0.88 cm). This suggests these methods handle the bulging edges of width measurements more effectively than the sharp-but-inset edges of length measurements.

Each filtering method demonstrates distinctive responses to the geometric challenges presented by different parcel types. The bilateral filter, despite offering modest improvements over unfiltered processing for standard parcels, struggles substantially with distorted parcel length determination (2.43 cm error), offering negligible improvement over unfiltered processing (2.56 cm), despite the presence of sharp edges.

Shadow filtering demonstrates more balanced performance, achieving substantial error reductions across both dimensions and parcel types. Particularly noteworthy is its relatively consistent performance for both width and length measurements of distorted parcels (0.43 cm and 1.44 cm respectively), suggesting effective handling of both bulging and sharp-but-inset edges.

Most intriguing is the SOR filter's performance pattern. It demonstrates exceptional length accuracy for distorted parcels (0.43 cm, representing an extraordinary 83.2% improvement over unfiltered processing), indicating remarkable efficacy in handling the sharp-but-inset edges. However, this comes at considerable cost to width accuracy for the same parcel type, where error increases to 1.24 cm – 42.5% worse than unfiltered processing. This suggests SOR filtering




Figure 4. Mean absolute error comparison for width and length measurements across different filtering methods, segregated by parcel type. Bars are also marked with standard error of the mean

particularly struggles with the bulging edges used for width determination in distorted parcels.

Figure 5 illustrates the impact of parcel type on volumetric weight determination, a crucial metric for logistics operations. Unfiltered processing shows nearly identical errors for both parcel types (0.200 kg for standard, 0.195 kg for distorted), suggesting minimal inherent bias in unfiltered measurement.

The application of filtering methods produces distinctive effects based on parcel geometry. The bilateral filter demonstrates stronger improvement for distorted parcels (43.1% reduction to 0.111 kg) than for standard parcels (18.0% reduction to 0.164 kg). This contrasts with its dimensional measurement performance, highlighting the complex relationship between individual dimension accuracy and volumetric calculation.

Shadow filtering achieves the lowest absolute errors for both parcel types (0.077 kg for standard, 0.085 kg for distorted), with more substantial relative improvement for standard parcels (61.5% versus 56.4%). SOR filtering produces intermediate results, with greater accuracy for standard parcels (0.098 kg) than distorted parcels (0.126 kg).

The algorithm-specific dimensional weaknesses observed in Figure 4 reflect fundamental mismatches between filtering mechanisms and geometric measurement requirements. The shadow filter's poor performance in length estimation for skewed parcels occurs because noise characteristics near edge features fall outside its

detection parameters, allowing spurious points to persist and thereby compromise edge localization accuracy. Conversely, the SOR filter poor performance in width estimation likely results from the width measurements in skewed plots dominated by curved edges, where point-to-neighbour distances exhibit maximum variance. This variance causes SOR to systematically remove the exact boundary-defining points that are critical for accurate width determination. Such dimensional selectivity reflects each algorithm's fundamental design limitations.

Scale dependent error

Given the large size and heterogeneity of the dataset used in this study, we were able to investigate how dimensional measurement errors evolve with increasing object size. This allowed for a deeper analysis of scale-dependent phenomena that are often hidden in smaller or more homogeneous datasets. Most importantly, since each spatial dimension is inferred using a slightly different methodological approach, it is necessary to break down the analysis based not only on the measurement axis (length, width, height) but also on the dimensional type – planar versus vertical.

This distinction stems from operational differences in how dimensions are estimated: vertical measurements (i.e., height) are typically derived from the total height of the point cloud relative to a stable reference plane, such as a




Figure 5. Volumetric weight measurement error by filtering method and parcel type. Between package types marked percent difference

scanning table. In contrast, horizontal dimensions (length and width) are calculated from the distribution of points on the upper surface, making them more susceptible to point cloud noise, edge irregularities, and algorithmic interpretation of object boundaries. As a result, the scale-dependent nature of the error is not a general sensor artifact, but rather a direct consequence of the geometry-based estimation methods used for different types of dimensions

To better understand this behaviour, we conducted a detailed analysis of scale dependence. This included comparing absolute measurement errors for each basic dimension (height, length, width) as well as for the derived volume mass. By examining how error magnitudes change with increasing object size, we were able to reveal systematic trends unique to each filtering method. These insights are illustrated in Figure 6 and 7, which shows the error distributions across dimensional scales for all tested algorithms.

Figure 6 presents the dimension-specific regression analysis for height, width, and length measurements. For height, all filters demonstrate equivalent performance (slope = 0.004) with coefficient of determination ($R^2 \approx 0.04$ -0.05), suggesting inherent stability in height measurements across all filtering methods.

Width measurements reveal more nuanced behaviour. The bilateral filter and unfiltered processing exhibit identical sensitivity (slope = 0.028) but with higher coefficient of determination ($R^2 = 0.18$ and 0.10, respectively). The shadow filter's negative slope (-0.007, $R^2 = 0.03$) represents inversion of the typical size-error relationship, whilst the

SOR filter effectively neutralises width-dependent errors (slope = 0.002, $R^2 = 0.00$).

Length measurements display the most pronounced differentiation between filtering methods. The unfiltered data shows strong size dependence (slope = 0.051, $R^2 = 0.39$), with the bilateral filter offering modest improvement (slope = 0.038, $R^2 = 0.38$). The shadow filter substantially reduces size sensitivity (slope = 0.014, $R^2 = 0.13$), whilst the SOR filter achieves near-complete elimination of length-dependent errors (slope ≈ 0.000 , $R^2 = 0.00$).

Figure 7 illustrates the relationship between expected volumetric weight and absolute error for each filtering methodology. The regression analysis reveals compelling differences in filter performance. The unfiltered data exhibits the steepest slope (0.1008, $R^2 = 0.473$), indicating that for each kilogram increase in volumetric weight, measurement error increases by approximately 0.1 kg – a 10% error growth rate. The bilateral filter offers modest improvement (slope = 0.0845, $R^2 = 0.462$), reducing the error growth rate by approximately 16% compared to unfiltered data.

More substantial improvement is observed with the SOR filter (slope = 0.0409, $R^2 = 0.461$), which reduces the error growth rate by nearly 60%. Most impressive, however, is the shadow filter's performance, with a remarkably low slope of 0.0192 ($R^2 = 0.184$). This represents an 81% reduction in error growth rate compared to the unfiltered data – a significant achievement in measurement stability. The shadow filter's substantially lower R^2 value suggests its error




Figure 6. Dimension-specific regression analysis showing the relationship between expected dimension (height, width, length) and absolute error for each filtering method




Figure 7. Relationship between expected volumetric weight and absolute error across four filtering methods, showing regression lines and statistical parameters

behaviour is less determined by size and more influenced by other factors.

The impact of the object's position

The position of a parcel within the measurement field of view represents a critical variable in dimensional measurement systems. This study specifically investigated how parcel placement -comparing corner versus center positioning on the measurement surface – affects the accuracy of dimensional measurements across different filtering methods.

The influence of parcel position on measurement accuracy was quantified by comparing error metrics at center versus corner placements. The ratio of corner to center errors was calculated for each filter and dimension:

$$p_{c/c} = \frac{MAE_{corner}}{MAE_{center}} \quad (3)$$

where: $p_{c/c}$ – denotes the corner-to-center Mean Absolute Error ratio, MAE_{corner} – MAE measured for objects placed in corner of the setup, MAE_{center} – MAE measured for objects placed in center of the setup.

Figure 8 presents the corner-to-center error ratio across different filtering methods, providing an immediate visualization of position sensitivity. A ratio exceeding 1.0 indicates higher measurement errors in corner positions compared to

center positions, while ratios below 1.0 suggest superior performance in corner positions.

In order to examine the impact of the position of the package in the field of view on measurement accuracy, a statistical analysis of measurement errors was performed for the following filters: unfiltered, Bilateral, Shadow, and SOR, comparing the results for the center and corner positions. The Kruskal-Wallis test was used to assess the significance of the differences. As shown in Table 3, in the vast majority of cases, regardless of the filter type and the dimension analysed, no statistically significant differences were found, which indicates high stability of the filtration methods with respect to the position of the object in the measurement space.

However, the result obtained for the SOR filter in the case of length error is particularly noteworthy, where the difference between the corner and center positions proved to be statistically significant ($p = 0.0267$). This value was associated with a clear reduction in error in corner positions – the average difference was -0.05 cm, which corresponded to an approximately 10.14% relative improvement in accuracy. The test statistic ($H = 4.91$) further confirms that this effect is not accidental. This result suggests that when using the SOR filter, corner positioning of the object can positively affect the precision of length measurement, probably due to the way the filter handles the distribution of points within the more exposed edges of the package.




Figure 8. Corner-to-center error ratio by filter and dimension

Table 3. Summary of Kruskal-Wallis test results examining statistical differences in measurement dimension errors across filtering methods

Filter	Metric	Mean difference	Percent difference (%)	H statistic	p-value
Unfiltered	Height error (cm)	0.03	9.18	2.23	0.136
	Width error (cm)	0.02	2.10	0.04	0.849
	Length error (cm)	0.03	1.68	0.38	0.536
	Volumetric weight error (kg)	0.01	3.66	0.02	0.896
Bilateral filter	Height error (cm)	0.02	5.37	0.27	0.602
	Width error (cm)	0.07	10.24	1.79	0.180
	Length error (cm)	0.10	7.33	3.17	0.075
	Volumetric weight error (kg)	0.01	8.38	0.95	0.330
Shadow filter	Height error (cm)	0.03	9.18	2.23	0.136
	Width error (cm)	-0.04	8.29	2.20	0.138
	Length error (cm)	-0.06	9.19	2.26	0.133
	Volumetric weight error (kg)	0.00	0.76	1.70	0.193
SOR filter	Height error (cm)	0.03	9.82	2.34	0.126
	Width error (cm)	-0.02	2.52	0.25	0.619
	Length error (cm)	-0.05	10.14	4.91	0.027
	Volumetric weight error (kg)	0.00	0.08	2.03	0.154

The other metrics and filters did not show a significant impact of the package position on measurement accuracy, confirming that all filter proposals are resistant to changes in the position of the tested object.

Time performance

While accuracy of measurements is essential, as examined in the previous section, the time required to process each point cloud and extract

dimensional information determines the practical applicability of the system in real-world setups. The growing interest in edge processing reflects a critical paradigm shift in modern automation, emphasizing on-site data processing where hardware expansion is often impractical.

To better understand the computational characteristics of each filtering approach, we decomposed the total processing time into its constituent components: filtering, segmentation, dimensioning, cropping, and alignment. Figure 9 illustrates

this breakdown. Our analysis reveals a critical insight: the filtering stage is the only component that varies significantly across methods, while all other processing stages maintain remarkably consistent timing regardless of the filtering approach used. This consistency in downstream processing demonstrates that the choice of filter affects only the initial point cloud processing stage without introducing cascading computational impacts on subsequent operations.

When comparing the three actual filtering methods (excluding the unfiltered approach which serves as our baseline), the differences in filtering time are substantial, as shown in Figure 9. The Bilateral filter is the most computationally efficient at 7.42 ms (7.52% of its total processing time), while the Shadow filter requires 42.66 ms (33.70% of total time) and the SOR filter demands 119.83 ms (60.01% of total time). This means the Shadow filter's filtering stage takes 5.75 times longer than the Bilateral filter, while the SOR filter requires 16.15 times the processing time of the Bilateral approach and 2.81 times that of the Shadow filter.

The Bilateral filter's localized approach to noise reduction requires minimal processing, the Shadow filter's moderate complexity results in intermediate processing demands, and the SOR filter's statistical analysis of point neighbourhoods creates the heaviest computational load. It is also worth noting that the relatively long processing time for the SOR filter is due not only to its computational complexity, but also to

implementation limitations. This algorithm, in the available version of the PCL library, has been implemented sequentially and does not allow for the effective use of multithreading. The lack of support for parallel processing significantly increases the total filtering time, which may limit its usefulness in practical applications.

DISCUSSION

The comparative analysis of the Shadow, Bilateral, and SOR smoothing algorithms reveals dimension-specific strengths and trade-offs in object dimensioning accuracy using depth data. While all three filters outperformed unfiltered data processing in terms of dimensional accuracy (mainly width and length), their effectiveness varied depending on the measured dimension and the geometry of the object.

The Shadow filter consistently demonstrated strong results, particularly in width measurements, and exhibited the lowest mean absolute error (MAE) in volumetric weight estimation. Its key advantage lies in striking a favourable balance between precision and computational efficiency, making it suitable for time-sensitive applications.

The SOR filter, on the other hand, achieved the highest accuracy for length measurements, performing particularly well on objects with sharp, inset edges. This outcome suggests that its statistical treatment of outliers makes it well-suited for handling well-defined geometric




Figure 9. Processing time components across filtering methods per one sample dimensioning cycle

boundaries. However, its performance was less stable for width in distorted packages with bulging edges, where its error exceeded that of unfiltered data. This trade-off points to a context-dependent suitability of the SOR filter, where geometric characteristics of the object can significantly affect its reliability.

Interestingly, the position of objects within the field of view did not lead to statistically significant accuracy variations in most cases. However, minor trends emerged – most notably with the Shadow filter showing improved lateral performance and the SOR filter yielding a statistically significant improvement in length measurements for objects placed in corner positions. This latter finding may relate to the SOR's sensitivity to edge distribution, which is more pronounced in these spatial placements.

An important observation pertains to scale-dependent error. Across all algorithms, larger objects tended to yield higher dimensional errors, especially for planar dimensions derived from convex hull estimates. The Shadow filter exhibited the error decrease with object width size, indicating strong generalizability, but in the case of length it indicated positive slope. In contrast the SOR filter showed greater sensitivity to object width size, but for the length it exhibited near-zero slope indicating its resistance to growth in size in this dimension. This indicates that further evaluation is needed to study the impact of the geometry on the algorithm performance and accuracy.

In terms of processing efficiency, the Bilateral filter proved to be the fastest, followed by the Shadow and SOR filters. Despite its speed, the Bilateral filter offered only moderate improvements in accuracy, especially for distorted shapes. The Shadow filter strikes a practical balance – moderate processing demands and consistent accuracy – making it suitable for near real-time applications. Conversely, the high computational cost of the SOR filter, exacerbated by the lack of multithreading in its current implementation, may limit its deployment in systems requiring rapid throughput.

Finally, comparative analysis with related works reveals that the proposed approach using Shadow and SOR filtering achieves superior or comparable accuracy with significantly faster processing times. The results obtained in this study are comparable to the work of Ladplee et al. [26], in which a single LiDAR camera was

used together with a flood fill algorithm to measure package dimensions. There, the average error was less than 5%, with a processing time of approximately 1.01 seconds. Compared to this approach, our solution using the Shadow filter achieves a lower average absolute error for volumetric weight (~ 0.079 kg) and a significantly shorter processing time (~ 126.6 ms), making it potentially more effective in the context of real-time applications.

In the study by Rodriguez et al. [27], a superquadric fitting method was applied to noisy and partially occluded ToF data, achieving average errors of 4–13% depending on the orientation of the object and the visibility of the surface. In comparison, our approach based on processing the full point cloud using filters such as Shadow and SOR showed lower absolute errors (MAE for lengths up to ~ 0.463 cm), without the need for special geometric models.

CONCLUSIONS

This study demonstrates that point cloud smoothing techniques significantly improve object dimensioning accuracy when compared to unfiltered data. Among the evaluated methods, the Shadow algorithm yielded the lowest MAE in volumetric weight estimation (~ 0.079 kg), followed by the SOR (~ 0.106 kg) and Bilateral (~ 0.148 kg) filters, whereas unfiltered data resulted in the highest error (~ 0.199 kg). For individual dimensions, the Shadow filter was most effective in width estimation (~ 0.482 cm), while the SOR filter performed best for length (~ 0.463 cm), highlighting the dimension-specific strengths of each method.

While package placement in the field of view had generally no statistically significant impact on measurement accuracy, specific cases – such as improved length accuracy with SOR filtering in corner positions – did show localized significance. Furthermore, the Shadow filter exhibited high resistance to scale-dependent error, reflected in the smallest regression slope (~ 0.041) with increasing object size.

In terms of efficiency, the Bilateral filter offered the shortest processing time (~ 98.6 ms), followed by the Shadow (~ 126.6 ms), and SOR (~ 199.7 ms). The Shadow filter's favourable trade-off between accuracy and computational

demand makes it particularly suited for real-time applications.

The complementary performance characteristics observed between Shadow and SOR filters present an opportunity to develop adaptive filtering pipelines that dynamically select or combine methods based on preliminary object geometry assessment, while simultaneously implementing custom parallelized versions of these algorithms to overcome existing computational limitations in libraries such as PCL.

Future research directions should pursue a comprehensive advancement of object dimensioning accuracy through multi-modal sensor fusion while simultaneously optimizing the complementary filtering approaches identified in this study. The primary focus involves investigating YOLO (you only look once) object detection integration with LiDAR depth data to create a hybrid dimensioning system, where LiDAR point clouds provide precise height measurements - leveraging the demonstrated stability across all filtering methods (MAE $\approx$ 0.329–0.343 cm) – while YOLO bounding boxes determine width and length through calibrated pixel-to-centimetre conversions. This sensor fusion approach could address the scale-dependent errors observed in planar measurements and enable direct accuracy comparison between traditional point cloud filtering and hybrid computer vision methodologies.

REFERENCES

1. Kwieciński K, Paluch M, Lisowska A. Analysis of the dimensional accuracy of point clouds created by photographic scanning. 2024; 18(6): 121–132. <https://doi.org/10.12913/22998624/191430>
2. Hoegg T, Lefloch D, Kolb A. Time-of-flight camera-based 3-D point-cloud reconstruction of a car. *Comput Ind.* 2013; 64(9): 1099–1114. <https://doi.org/10.1016/j.compind.2013.06.002>
3. Józwik J, Dziedzic K. Digital shape and geometric dimension analysis of polymer fuel tanks. *Adv Sci Technol Res J.* 2021;15(4): 38–48. <https://doi.org/10.12913/22998624/141649>
4. Huang L, Chen M, Peng Z. Point cloud measurement of rubber tread dimension based on RGB-depth camera. *Appl Sci (Basel).* 2024; 14(15): 6625. <https://doi.org/10.3390/app14156625>
5. Raaj Y., Nair S., Knoll A. Precise measurement of cargo boxes for gantry robot palletization in large scale workspaces using low-cost RGB-D sensors. In: *Lecture Notes in Computer Science.* 2017; 10114: 472–486. Presented at: 13th Asian Conference on Computer Vision; 2016 Nov 20–24; Taipei, Taiwan. https://doi.org/10.1007/978-3-319-54190-7_29
6. Kumičáková D., Rengevič A., Císar M., Tlach V. Utilisation of kinect sensors for the design of a human-robot collaborative workcell. *Adv Sci Technol Res J.* 2017; 11(4): 270–278. <https://doi.org/10.12913/22998624/80937>
7. Rusu RB, Marton ZC, Blodow N, Dolha M, Beetz M. Towards 3-D point-cloud-based object maps for household environments. *Robot Auton Syst.* 2008; 56(11): 927–941. <https://doi.org/10.1016/j.robot.2008.08.005>
8. Park H.-M., Van Messem A., De Neve W. Box-Scan: An efficient and effective algorithm for box dimension measurement in conveyor systems using a single RGB-D camera. In: *Proceedings of the 7th IIAE International Conference on Industrial Application Engineering*; 2019 Apr; Kitakyushu, Japan. <https://doi.org/10.12792/iciae2019.032>
9. Tölgyessy M, Dekan M, Chovanec L, Hubinský P. Evaluation of the Azure Kinect and its comparison to Kinect V1 and Kinect V2. *Sensors (Basel).* 2021; 21(2): 413. <https://doi.org/10.3390/s21020413>
10. Kurillo G, Hemingway E, Cheng ML, Cheng L. Evaluating the accuracy of the Azure Kinect and Kinect v2. *Sensors (Basel).* 2022; 22(7): 2469. <https://doi.org/10.3390/s22072469>
11. Neupane C, Koirala A, Wang Z, Walsh KB. Evaluation of depth cameras for use in fruit localization and sizing: finding a successor to Kinect v2. *Agronomy.* 2021; 11(9): 1780. <https://doi.org/10.3390/agronomy11091780>
12. Chiu CY, Thelwell M, Senior T, Choppin S, Hart J, Wheat J. Comparison of depth cameras for three-dimensional reconstruction in medicine. *Proc Inst Mech Eng H.* 2019; 233(9): 938–947. <https://doi.org/10.1177/0954411919859922>
13. Curto E, Araujo H. An experimental assessment of depth estimation in transparent and translucent scenes for Intel RealSense D415, SR305 and L515. *Sensors (Basel).* 2022; 22(19): 7378. <https://doi.org/10.3390/s22197378>
14. Servi M, Mussi E, Profili A, Furferi R, Volpe Y, Governi L, et al. Metrological characterization and comparison of D415, D455, L515 RealSense devices in the close range. *Sensors (Basel).* 2021; 21(22): 7770. <https://doi.org/10.3390/s21227770>
15. Jiang Y, Liu B, Wang R, Li Z, Chen Z, Zhao B, et al. Photon counting LiDAR working in daylight. *Opt Laser Technol.* 2023; 163: 109374. <https://doi.org/10.1016/j.optlastec.2023.109374>
16. Noel G, Djouani K, Van Wyk B, Hamam Y. Bilateral mesh filtering. *Pattern Recognit Lett.*

- 2012; 33(9): 1101–7. <https://doi.org/10.1016/j.patrec.2012.02.008>
17. Han XF, Jin JS, Wang MJ, Jiang W, Gao L, Xiao L. A review of algorithms for filtering the 3D point cloud. *Signal Process Image Commun.* 2017; 57: 103–12. <https://doi.org/10.1016/j.image.2017.05.009>
18. Agathos A, Azariadis P, Kyratzi S. Elliptic Gabriel Taubin smoothing of point clouds. *Comput Graph.* 2022; 106: 20–32. <https://doi.org/10.1016/j.cag.2022.05.009>
19. Grilli E, Menna F, Remondino F. A review of point clouds segmentation and classification algorithms. *Int Arch Photogramm Remote Sens Spatial Inf Sci.* 2017; XLII-2/W3:339–44. <https://doi.org/10.5194/isprs-archives-XLII-2-W3-339-2017>
20. Yadav SK, Reitebuch U, Skrodzki M, Zimmermann E, Polthier K. Constraint-based point set denoising using normal voting tensor and restricted quadratic error metrics. *Comput Graph.* 2018; 74: 234–43. <https://doi.org/10.1016/j.cag.2018.05.014>
21. Zhou L, Sun G, Li Y, Li W, Su Z. Point cloud denoising review: from classical to deep learning-based approaches. *Graph Models.* 2022; 121: 101140. <https://doi.org/10.1016/j.gmod.2022.101140>
22. Ladplee N, Pimpin A, Srituravanich W, Damrongplasit N. Volumetric measurement of rectangular parcel box using LiDAR depth camera for dimensioning and 3D bin packing applications. In: *Proceedings of the IEEE International Conference on Consumer Electronics-Asia (ICCE-Asia)*; 2022. <https://doi.org/10.1109/ICCE-Asia57006.2022.9954650>
23. Rodriguez B, Rangarajan P, Zhang X, Raj D. Dimensioning cuboid and cylindrical objects using only noisy and partially observed time-of-flight data. *Sensors (Basel).* 2023; 23(21): 8673. <https://doi.org/10.3390/s23218673>