A Laser-based Volumetric Measurement Approach for Industrial Settings

Jérôme Rutinowski[1*], Niklas Schrötler[1*], Simon Klüttermann[2], Moritz Roidl[1], and Christian Janiesch[3]

***Abstract*— This work introduces a novel approach to the currently available automatic volumetric measurement systems used in the industry. The proposed system uses a ceiling-mounted laser in tandem with an RGB camera to generate a point cloud of the object that is to be measured, utilizing an adaptation of a structured light scanning approach. The resulting point clouds are post-processed, removing detected structural anomalies to improve the yielded results. The resulting system is then tested and validated on a set of seven objects of various sizes commonly encountered in industrial environments. For these experiments, our approach yields outlier-corrected volumes which are accurate at up to 1% volume divergence. The results differ vastly when going beyond the system's optimal object size range. The limits of the system are tested for smaller and larger objects, showing a notably higher inaccuracy when tested on relatively small or large objects. Further experiments, exploring optimal laser line distance, lighting conditions, and (in the case of containers) filling degrees are also conducted.**

I. Introduction

In industrial environments, the effective tracking and management of goods and other objects are pivotal for streamlined business operations [1], [2]. This entails not only the identification of objects but also gathering crucial data on them, such as their dimensions, their origin, or their class. While visual identification [3] and classification of objects [4] are well-established research fields, determining further properties, especially an object's volume, lacks comparable research depth [5].

Currently, to estimate an object's volume, contemporary commercial systems typically adopt one of two approaches. The first utilizes a one-dimensional, static laser projection to analyze the object's height and width within a set projection area [6]. Multiple measurements are taken as the object moves beneath the sensor at a known and static speed, often enabled by the use of conveyor belts, allowing for volume inference but also limiting implementation flexibility. Alternatively, the second approach employs one or more fixed-point depth sensors to generate a point cloud, enabling rapid capture and analysis of the object's attributes [7]. While the second method offers more adaptability by not requiring stringent guiding infrastructure like conveyors, it still requires the installation and acquisition of additional hardware (e.g., LIDAR), which in practice is often not feasible.

[1] Chair of Material Handling and Warehousing, TU Dortmund University: jerome.rutinowski@tu-dortmund.de
[2] Chair of Data Science and Data Engineering, TU Dortmund University
[3] Chair of Enterprise Computing, TU Dortmund University
[*] Equal contribution

Given these limitations, we propose an alternative solution for volumetric measurements in industrial settings. Our approach is capable of measuring the volume of common industrial objects while offering flexibility in the required necessary infrastructure, thereby negating the drawbacks of currently utilized solutions. As such, the proposed approach has the potential to be an additional asset in the realm of volumetric measurement methods in the industry and beyond.

The paper is organized as follows: the subsequent section will provide a brief overview of the relevant literature concerning methods utilized in volume estimation and the post-processing of the resulting volumes. Following that, the methodology employed for the approach proposed in this study will be presented. Section IV will then discuss the results obtained from testing the proposed system on various objects and explore the influence of real-world effects on the accuracy of our measurements. Finally, the last section will provide a final overview of this work and outline future research directions that would allow us to enhance our system's efficacy further.

II. Related Work

In this section, we will briefly introduce the three main parts that make up a volume measurement setup based on the use of point clouds. As such, the concept of point cloud measurement will be addressed. Subsequently, we will introduce approaches for anomaly detection for point cloud correction. Finally, we will offer a brief overview of volume estimation strategies.

A. Point Cloud Measurement

The term point cloud [8] is used to describe sets of points in three-dimensional space, which can be acquired using various methods, one of which involves the utilization of depth cameras. Depth cameras [9] employ different technologies to capture depth information, resulting in point clouds with distinct characteristics depending on the underlying depth-capturing technology. Among these technologies are stereo systems, time-of-flight systems, and structured light systems [10].

In stereo systems [11], two cameras are strategically positioned to capture offset views of the same scene. By leveraging the perspectives provided by these cameras and the known spatial relationship between them, individual points' positions can be determined. However, accurately identifying corresponding points in both camera images can be challenging, particularly on smooth surfaces where distinguishing between points on the surface becomes increasingly

difficult. Natural environmental infrared light noise can aid in this process by introducing variations.

In contrast, time-of-flight systems [12] utilize a light source, typically a laser emitter, to illuminate the respective scene gradually. A camera, or a similar sensor, measures the time it takes for the emitted light to return to its origin after striking the scene. By analyzing the time difference between light emission and reception, the system can calculate the distance to each point, leveraging the (known) speed of light in the process. An example of this are LIDAR systems [13], [14].

Similar to time-of-flight cameras, only one camera is used for structured-light 3D scanning [15] (also known as strip projection, or structured-light scanning). The simplest example of structured-light scanning with a static light source are so-called 1D laser displacement sensors [16]. These use a static light source similar to a laser pointer that is directed straight downwards and thus hits the point to be detected. A one-dimensional light sensor, which is positioned to the side of the emitter, can then detect the angle at which the light falls back on it. In this way, the distance to the sensor can be estimated for a given point in space. A more elaborate concept in structured-light scanning is the projection of a pattern onto the scene. The distortion of the pattern based on the objects it hits can then be used to determine their dimensions. This volumetric derivation, which is also called reconstruction, works by triangulation, i.e., the construction of a triangle between the point, the camera, and the laser for each point in space using cosine theorems.

B. Anomaly Detection

Anomaly detection [17] refers to the process of identifying samples that deviate from an expected pattern, behavior, or attribute [18]. Initially employed for data cleaning purposes (i.e., for the removal of erroneous samples) [19], [20], anomaly detection now finds applications in various domains, including the detection of fraud and faults [21], [22], [23] as well as diverse fields of research [24]. Given the propensity of image data to contain outliers that can significantly impact their accurate analysis (i.e., the accurate estimation of an object's volume), anomaly detection plays a crucial role in image processing tasks, such as the one presented in this work.

To address this challenge, a diverse array of algorithms is available, each with its strengths and limitations [18]. In our scenario, which is characterized by relatively low-dimensional data and which needs an efficient and reliable solution, we are particularly interested in the use of shallow algorithms. Such algorithms offer fast execution and yet significant reliability. In doing so, they align well with the requirements of our task, in which we would like to remove outliers in the point cloud data to provide more accurate volumetric measurements. However, even within shallow algorithms, variations exist in their approaches. Some algorithms, such as DBSCAN and KNN [25], [26], specialize in identifying local anomalies, while others, like Isolation Forest (iForest) and One-Class Support Vector Machines (OC-SVM) [27], [28], excel at detecting more global and structural outliers. We will employ a combination of both types of algorithms to ensure comprehensive anomaly detection.

C. Volumetric Estimation

When presented with a point cloud, the next step commonly is the inference of a volume [5]. For this purpose, various approaches do exist, each suited to different use cases.

A rather simple approach involves determining the width, length, and height of the object represented by the point cloud and calculating the product of these dimensions, thus obtaining an estimate of the object's volume. One approach is to utilize the minimum and maximum values along each axis. Accordingly, the object's width is calculated as $x_{\max} - x_{\min}$, its depth as $y_{\max} - y_{\min}$, and its height as $z_{\max} - z_{\min}$. However, this method is limited in its accuracy, as this approach assumes perfectly cuboid objects.

Alternatively, the volume can be directly inferred from the point cloud. Numerous implementations of this approach exist and are frequently employed in volume estimation tasks involving point clouds [29], [30], [31]. Typically, these methods involve constructing a convex hull from the given point cloud and subsequently computing the volume enclosed by the points within this new configuration. Such techniques have been extensively utilized and explored in literature concerning volume estimation from point clouds [5], [32].

III. Methodology

In this chapter, we will first introduce our general point cloud capturing and correction method and describe how an object's volume can be calculated from this point cloud. Subsequently, the practical implementation and the experimental setup of our approach is presented. For this, the calibration of the sensors as well as the remaining used hardware and test subjects, will be discussed.

A. Point Cloud Generation

As previously mentioned, creating a point cloud is an intermediate step in the desired volume estimation. It provides a representation of the captured volume (i.e., the object in it, the ground and other elements that are located in the captured area). Similarly, this process can be extended to a line of points using a two-dimensional light sensor (a camera) and a line projector (a laser).

As described in the preceding section, we use a structured-light approach to create these lines of points. For this, the distance between the camera and the laser must be known through calibration. Furthermore, it follows from the calibration of the camera that a vector can be constructed which leads from the camera to the observed point. We can calculate the angle β (as shown in Fig. 1) between this vector and the vector pointing from the camera to the laser. Similarly, we can calculate α, shown in Fig. 1, from the position of the camera relative to the laser and the direction the laser is pointing in.

To combine the correct vectors from the projector with those from the camera in the next step, the pattern must be observed. As described above, this encodes the individual coordinates for each point in the projection. This is done both vertically and horizontally, with each pattern encoding whether the coordinate's grey code representation is a one (by displaying a line) or a zero (by not displaying a line) at the position in the bitstring represented by the pattern. This results in the pattern series, which can then be can then be projected onto the scene. The pattern used can be seen in Fig. 2.

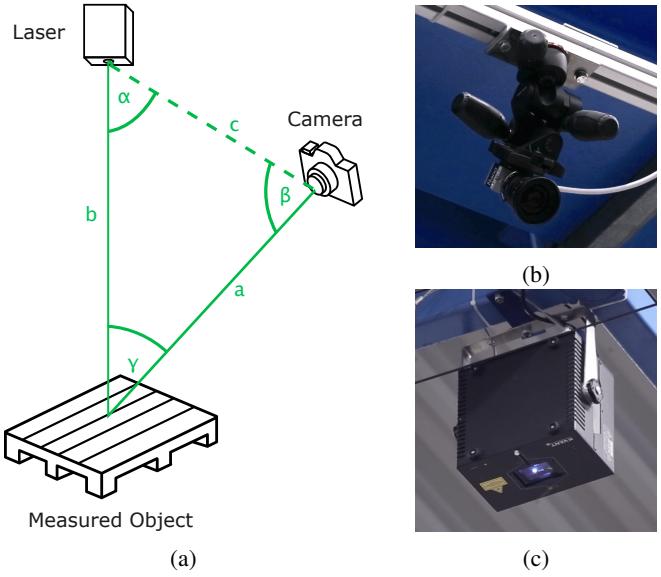

Fig. 1: Illustration of the constructed triangle (a) between the RGB camera (b), the laser (c), and the object that is to be measured.

Once the image series has been recorded, it must then be recognized in which images which pixel was illuminated or not. To this end, a minimum and maximum brightness is first determined for each pixel of the series. The resulting brightness range is then separated using a predefined factor. Brightness values below the pixel-individual threshold determined in this way are then marked as "no light" (0), values above are interpreted as "illuminated" (1), as inspired by [33]. This produces the required pairs of coordinates in the camera image coordinate system and the projection coordinate system.

To find the real-world coordinates of an observed point, the previously constructed vector from the camera in the direction of the observed point can be normalized into a unit vector and then multiplied by the distance found. Adding the support vector of the camera (i.e. the position of the camera) then yields the real-world coordinates of the point to be determined. This procedure is repeated for each pair of points found until all pairs have been processed. The set of real-world coordinates found for the points finally forms the point cloud.

In addition, an image of the scene in its initial state can be created before the scan. This has no direct use in the system presented here. However, it is possible to use this image to colorize the resulting point cloud. In this way, an RGB point cloud can be made available to subsequent analysis processes.

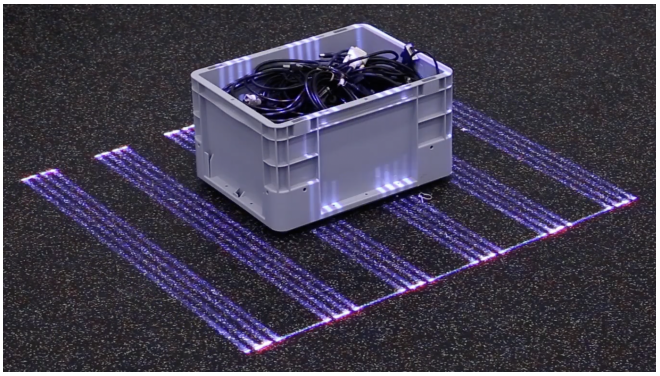

Fig. 2: Pattern projected by the laser on object (b) when filled.

B. Point Cloud Post-Processing

Regardless of the method employed for generating a point cloud, artifacts and noise are an inevitability in the resulting data. As we calculate a convex hull around all points of a point cloud, including its anomalous points, singular outliers can significantly affect the resulting estimated volume. Due to this phenomenon, refining the resulting point cloud becomes essential. With this aim in mind, we choose to utilize anomaly detection methods to filter out point cloud outliers. It is worth noting that erroneously removing correct samples does not impact the convex hull as significantly as not removing erroneous ones.

We opt to employ a two-step post-processing approach, leveraging two distinct algorithms in sequence. First, DBSCAN [25] is utilized to eliminate local anomalies, thereby reducing the overall noise level of our data. Subsequently, an isolation forest algorithm (iForest) [27] is applied to identify and filter out the remaining structural anomalies.

DBSCAN clustering uses hyperparameters to characterize the forming of clusters, these are set to $\epsilon = 0.1$ and $min\_samples = 10$, while iForest is built from 100 trees and removes a fixed fraction of 10% of the remaining samples.

C. Context-Adapted Volume Estimation

Due to the fact that the point clouds generated by the presented system may only show an incomplete, one-perspective view of the object, regular 3D point cloud volume estimation approaches, for which multiple perspectives would be needed, cannot be employed. Instead, a modified approach inspired by the ones described in [29], [34] will be used.

For this purpose, all points of the point cloud are first classified into groups with similar z-axis coordinates. Two points are considered similar if they have the same z-axis value rounded to the nearest step, given a predetermined step size h, which controls precision. The groups are then processed one after another, starting with the group with the

largest z-axis value. The points forming a two-dimensional convex hull for the corresponding slice and the area enclosed by these points are then determined. This area, multiplied by the step size h, is then added up to form the total volume. This is then repeated per group, whereby the points that formed the convex hull of the previous group are added to the set of points of the group, thus only allowing slices to either grow or maintain their size.

Due to this approach, it is sufficient for the upper surface of an object to be recorded. This, however, also implies that inward facing 'bulges' at the side of an object cannot be accurately measured. Nevertheless, such object shapes are seldom relevant in the industry. Consequently, the final volume estimation equals the area of the recorded upper surface multiplied by its height.

D. Testing Environment

The experiments in this work were conducted in a defunct warehouse, now used as a research facility providing a highly realistic industrial environment. As a laser, we employ a Kvant ClubMax 3000 FB4 RGB S1 laser with a Pangolin ScannerMax Saturn 1 $90kpps$ scanner. The laser is mounted on a steel beam under the warehouse's ceiling, pointing downwards at a $90°$ angle. In the laser's setup, it is important to ensure that the projection completely covers the upper surface of the objects to be scanned and that it is placed orthogonally to the floor.

For the RGB camera, we use a Teledyne DALSA Nano-C2590 camera. The camera is also mounted on a steal beam under the warehouse's ceiling. It is mounted in such a way that it looks down onto the laser projection field from above at an angle such that it can record the upper surface of the objects to be captured particularly well.

The laser is controlled using the Pangolin BEYOND Software Development Kit (BEYOND SDK) from the Unity game engine. The laser line distance of our projected pattern is, by default, set to $0.02m$. The lighting in the warehouse is dimmable and is measured to have a brightness of $784lx$ (measured with a UNI-T UT383 lux-meter), representing an ordinary lighting condition in an industrial setting. The images captured with the camera are collected using a Python script, which also reports the pattern to be projected to Unity via MQTT. A series of calibration images were taken and processed with OpenCV to calibrate the intrinsic parameters of the camera. The position of the laser was set to $x = 0m, y = 0m, z = 7.541m, rx = 0°, ry = 0°, rz = 180°$, with the height z having been measured manually using a Leica DISTO S910 laser distance meter. The position of the camera ($x = -0.245m, y = 3.58m, z = 4.803m, rx = 162.333°, ry = 0.159°, rz = 179.533°$) was determined manually by adjusting its position while observing the resulting point clouds and is thus approximate in nature. The height of the camera was set to $4.803m$, also measured with the Leica DISTO S910 laser distance meter.

E. Test Subject Selection

To evaluate the system's performance, a selection of commonly used industrial objects were chosen. These small to large sized objects reflect both the materials (polymers, metal, wood) as well as common size categories (small, medium, large) found in industrial contexts. The objects used can be seen in Fig. 3.

Object (b), a small, gray load carrier, was selected as the default object that we perform our experiments with, due to its widespread use and its flexibility in terms of alignment and loading. By default, the objects used are placed centrally under the laser. If applicable, the objects are empty and, with their opening pointing upwards, aligned in parallel to the projected pattern. Furthermore, the color white was chosen for laser projection, to generate sufficient contrast with the object to be scanned, even if certain wavelength ranges are absorbed (e.g., with colored objects such as the small, blue load carrier).

Fig. 3: Objects used for the evaluation of our method: (a) open-fronted bin, (b) small load carrier (gray), (c) small load carrier (blue), (d) cardboard box, (e) large load carrier, (f) barrel, (g) wrapped and loaded Euro pallet.

F. Experimental Design

To validate the viability of our approach, a series of experiments are conducted. These experiments aim to analyze the robustness of our approach, especially its accuracy when real-world impreciseness is emulated.

- Experiment I: In the first experiment, we aim to measure the volume of the objects shown in Fig. 3, repeating the measurement for each object ten times. The results are then compared with their respective ground truth (*GT*). Thus, we compare both the variance in our resulting volumes as well as the potential limitations of our approach when encountering particularly small or large objects.
- Experiment II: In the second experiment, we compare multiple laser line distances *LLD* ($0.01m$, $0.02m$, $0.05m$) with one another, exploring the impact of such a change on the system's accuracy. For this experiment, as well as the subsequent ones, the default object (the small, gray load carrier) is used.
- Experiment III: In the third experiment, we investigate the impact of the object being filled or not, on the

measurement's accuracy. For this purpose, the load carrier is filled to the brim with cables.

- Experiment IV: In the fourth and final experiment, we compare the impact of different degrees of lighting on the accuracy of our measurements. For this, we use the default lighting brightness of $784lx$ and a dimmed light of $114lx$, measuring the volume of our default object (b) and the cardboard box (d).

For the experiments we conduct, especially Experiment I, our measurement will be repeated multiple times. In such cases, the resulting volumes may vary, which provides a more realistic picture of the precision of a method but also introduces measurement uncertainties. The representation of a combined resulting accuracy in the form of a simple arithmetic mean decreases the given uncertainty significantly [35], but is susceptible to outliers. These are especially important in this case since a slight overestimation in each dimension can translate into a significant difference in the final volume (i.e., a divergence of 26% per dimension would lead to an overall divergence of the total volume by 100%). To mitigate this effect, the median could be used, as the median is drastically less affected by outliers. However, the median does not decrease the yielded uncertainty as much as the mean would. We therefore propose to use a combination of the two, which we call the midpoint mean or *MPM*:

$$MPM(x) = \frac{Mean(x) + Median(x)}{2} \tag{1}$$

Given this context, we also list the minimum and maximum volume measured to highlight the outliers occurring during our measurements.

IV. Results

This section will present the reader with the results of our work. The section is subdivided per experiment, as presented in the preceding section. The code used for these experiments and the resulting data can be found at https://github.com/psorus/laser.

A. Experiment I: Overall Volume Estimation Accuracy

In this subsection, we compare the results for our first experimental setup, probing the robustness of our measurements. According to the base configuration described above, all seven objects were recorded ten times.

TABLE I: Volumetric measurements for all objects shown in Fig. 3. The results are based on ten measurements each. GT represents the Ground Truth volume.

Obj.	GT [cm^3]	Min. [cm^3]	Max. [cm^3]	MPM [cm^3]	Deviation [%]
(a)	3,510	15,396	19,854	17,735	405.28
(b)	26,400	21,739	30,805	24,924	-5.59
(c)	26,400	23,973	30,464	26,840	1.67
(d)	34,810	18,023	181,974	31,225	-10.3
(e)	52,800	24,646	280,001	51,974	-1.56
(f)	233,824	51,875	1,564,753	232,170	0.71
(g)	748,800	182,211	927,067	438,988	-41.37

Looking at the data from Tab. I, it can be noted that the accuracy of our measurements sharply decreases for the smaller and larger objects in our list. Across all raw measurements recorded during this work, none came out with a volume of under $15,000cm^3$. We consider this value to be the lower bound of volumes that are measurable with our recording setup. For larger objects, it was not possible to determine an upper bound of measurable volume from the data. However, objects that have a volume matching or exceeding that of a loaded Euro pallet can be assumed to have larger volumetric deviations. This is also visually evident since the laser pattern that was projected could be seen flickering, due to the scanner not being fast enough, likely leading to inaccurate point cloud generation.

Objects that have a volume between the estimated bounds, however, exhibit low deviations relative to their ground truth. Especially objects (c), (e), and (f) have been measured with deviations of below 2%. Except for the objects representing the lower and upper bounds (objects (a) and (g)), all MPM values are relatively close to the respective ground truth and thus represent promising results. The measurement (i.e., laser pattern projection) time as well as the computation time were recorded as approximately $10s$ each. The specific times per experiment can be taken from the link provided above.

B. Experiment II: Laser Line Distance

In this experiment, different laser line distances and, by extension, different laser patterns were tested. Given the results of this experiment, which can be seen in Tab. II, it is apparent that the laser line distance of $0.02m$ yields the by far smallest deviations from the ground truth.

TABLE II: Volumetric measurements for object (b) with different laser line distances. The results are based on ten measurements each.

LLD [m]	GT [cm^3]	Min. [cm^3]	Max. [cm^3]	MPM [cm^3]	Deviation [%]
0.01	26,400	28,764	1,482,985	144,186	446.16
0.02	26,400	21,739	30,805	24,924	-5.59
0.05	26,400	36,847	860,766	96,655	266.12

Increasing the laser line distance to $0.05m$ produces a worse results, possibly because the lack of granularity ensures that the object is estimated to be larger than it is. The results obtained at a laser line distance of $0.01m$ are again considerably worse than the ones obtained using the default laser line distance. This might be understood as counter-intuitive, as a higher granularity would be expected (apart from the side effects described above). However, the increasing pattern complexity could lead to a poorer pattern representation of the laser. This, in turn, can lead to significant deviations during the reconstruction, e.g., when the pattern was not projected in its entirety.

Furthermore, the increased overall amount of laser light in the scene may amplify reflections. For example, in the images captures by the RGB camera, surrounding pixels could be interpreted as illuminated by the reflection, which would falsify the determination of their coordinates.

C. Experiment III: Filling Degree

When exploring the effect of whether a load carrier is filled or not (i.e., is equipped with a payload or not), it can be seen that our proposed method is somewhat insensitive regarding the filling level of a container.

Looking at Tab. III, the relative deviations differ by 14.7%. Additionally, the MPM of our obtained results deviate negatively. The increased deviation from the ground truth that can be observed for the filled object might be due to the increased complexity of the geometry at the top of the object (see Fig. 2).

TABLE III: Volumetric measurements for object (b), both filled and empty. The results are based on ten measurements each.

Payload	GT [cm^3]	Min. [cm^3]	Max. [cm^3]	MPM [cm^3]	Deviation [%]
Empty	26,400	21,739	30,805	24,924	-5.59
Filled	26,400	17,977	25,730	21,043	-20.29

D. Experiment IV: Lighting Conditions

For this experiment, objects (b) and (d) were chosen, since they should enable us to recognize the effects of material reflectivity in darker environments, being made of smooth, reflective polymer and non-reflective cardboard, respectively.

TABLE IV: Volumetric measurements for objects (b) and (d), measured under two different lighting conditions. The results are based on ten measurements each.

Obj.	Bright. [lx]	GT [cm^3]	Min. [cm^3]	Max. [cm^3]	MPM [cm^3]	Dev. [%]
(b)	114	26,400	22,383	35,320	25,113	-4.88
(b)	784	26,400	21,739	30,805	24,924	-5.59
(d)	114	34,810	17,484	136,384	32,849	-5.63
(d)	784	34,810	18,023	181,974	31,225	-10.3

Analyzing the resulting values in Tab. IV, it is apparent that a change in lighting conditions only has a marginal impact. The effect is so minimal in fact, that it is well within the expected range of error across identical measurements.

V. Conclusion, Limitations, and Outlook

It is conceivable that, for cubic objects, only a select few surfaces are analyzable by the system. In this work, we presented a novel, laser-based volumetric measurement approach. This approach employs a laser that projects a pattern onto storage containers, or objects in general, to be measured. This pattern is recorded by an RGB camera, permitting us to generate a point cloud of the object to be measured. We post-process the yielded point clouds with anomaly detection methods, enabling us to increase the accuracy of our results. In doing so, our approach enabled us to measure volumes with deviations of as low as 1% compared to their ground truth, when applied on a set of seven objects commonly found in the industry. Still, more significant deviations could be observed for particularly small and large objects.

In addition, certain factors, such as laser line distance changes, the degree to which the objects are filled as well as lighting conditions, were explored. In doing so, further insight into the strengths and limitations of the approach could be gained: While most common situational deviations do not affect the performance significantly, the laser parameters require careful tuning. However, further research into the impact of the objects' spacial orientation as well as the use of differently, e.g., cubic-shaped objects, would be beneficial.

Our system was implemented in a realistic industrial setting, a defunct warehouse, in order to assess its viability for industrial use. However, there are currently two main limitations concerning the introduction of our approach, which could be understood as potential avenues for future research: the system's calibration and the objects' rotation. Especially a careful calibration is required since small errors can spiral into large deviations for each measurement, as was found while manually calibrating the camera. The obtained point clouds currently tend to be tilted with respect to the floor, increasing inaccuracies. This effect worsens when objects are not perfectly centered below the laser. As such, an automated calibration method would greatly increase the reproducibility as well as the accuracy of our approach.

Another relevant factor, the scanning speed of the system, could also be increased further. With the current setup, the duration of the measurement is largely determined by the image acquisition script waiting for a fixed time after changing the pattern to be projected. Closer integration into the laser system would solve this challenge and could significantly reduce the current acquisition time. Additionally, the use of multiple cameras and lasers would be a potential way of increasing the system's functionality. This would further increase the spatial flexibility of the volumetric measurement process. In addition, the currently hidden surfaces of an object could be captured from other perspectives and be taken into account when calculating the object's volume.

Acknowledgements

This work is part of the project "Silicon Economy Logistics Ecosystem" which is funded by the German Federal Ministry of Transport and Digital Infrastructure.

This research has been funded by the Federal Ministry of Education and Research of Germany and the state of North-Rhine Westphalia as part of the Lamarr-Institute for Machine Learning and Artificial Intelligence.

The Linux HPC cluster at TU Dortmund University, a project of the German Research Foundation, provided the computing power.

This paper is based on the Bachelor's Thesis "Laser-based Volume Estimation of Industrial Entities" by Niklas Schrötler.

References

[1] Z. Soleimanitaleb and M. A. Keyvanrad, "Single object tracking: A survey of methods, datasets, and evaluation metrics," *arXiv preprint arXiv:2201.13066*, 2022.

[2] J. Cai, M. Xu, W. Li, Y. Xiong, W. Xia, Z. Tu, and S. Soatto, "Memot: Multi-object tracking with memory," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 8080–8090.

[3] S. Klüttermann, J. Rutinowski, C. Reining, M. Roidl, and E. Müller, "Towards graph representation based re-identification of chipwood pallet blocks," in *IEEE International Conference on Machine Learning and Applications (ICMLA)*, 2022, pp. 1543–1550.

[4] X. Shen, "A survey of object classification and detection based on 2D/3D data," *arXiv preprint arXiv:1905.12683*, 2019.

[5] W.-C. Chang, C.-H. Wu, Y.-H. Tsai, and W.-Y. Chiu, "Object volume estimation based on 3d point cloud," *International Automatic Control Conference (CACS)*, pp. 1–5, 2017.

[6] F. Zeng, Q. Wu, X. Chu, and Z. Yue, "Measurement of bulk material flow based on laser scanning technology for the energy efficiency improvement of belt conveyors," *Measurement*, vol. 75, pp. 230–243, 2015.

[7] W. Y. Amaglo, "Volume calculation based on lidar data," Ph.D. dissertation, 2021. [Online]. Available: https://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-299594

[8] Y. Guo, H. Wang, Q. Hu, H. Liu, L. Liu, and M. Bennamoun, "Deep learning for 3d point clouds: A survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 12, pp. 4338–4364, 2020.

[9] A. Haider and H. Hel-Or, "What can we learn from depth camera sensor noise?" *MDPI Sensors*, vol. 22, no. 14, 2022.

[10] M. Aboali, N. Abd Manap, A. Darsono, and Z. Yusof, "Review on three dimensional (3-d) acquisition and range imaging techniques," *International Journal of Applied Engineering Research*, vol. 12, pp. 2409–2421, 2017.

[11] J. Wang, D. Scharstein, A. Bapat, K. Blackburn-Matzen, M. Yu, J. Lehman, S. Alsisan, Y. Wang, S. Tsai, J. Frahm, Z. He, P. Vajda, M. F. Cohen, and M. Uyttendaele, "A practical stereo depth system for smart glasses," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023, pp. 21 498–21 507.

[12] R. Horaud, M. Hansard, G. Evangelidis, and C. Ménier, "An overview of depth cameras and range scanners based on time-of-flight technologies," *Machine vision and applications*, vol. 27, no. 7, pp. 1005–1020, 2016.

[13] L. A. Wasser, "The basics of lidar: Light detection and ranging: Remote sensing," 2023. [Online]. Available: https://www.neonscience.org/resources/learning-hub/tutorials/lidar-basics

[14] J. Kernhof, J. Leuckfeld, and G. Tavano, "LiDAR-Sensorsystem für automatisiertes und autonomes Fahren," in *Automobil-Sensorik 2: Systeme, Technologien und Applikationen*. Springer, 2018, pp. 29–54.

[15] M. Gupta, A. Agrawal, A. Veeraraghavan, and S. G. Narasimhan, "Structured light 3d scanning in the presence of global illumination," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011, pp. 713–720.

[16] Y. Suh, "Laser sensors for displacement, distance and position," *MDPI Sensors*, vol. 19, no. 8, pp. 19–24, 2019.

[17] S. Han, X. Hu, H. Huang, M. Jiang, and Y. Zhao, "Adbench: Anomaly detection benchmark," *Advances in Neural Information Processing Systems*, vol. 35, pp. 32 142–32 159, 2022.

[18] L. Ruff, J. R. Kauffmann, R. A. Vandermeulen, G. Montavon, W. Samek, M. Kloft, T. G. Dietterich, and K.-R. Müller, "A unifying review of deep and shallow anomaly detection," *Proceedings of the IEEE*, vol. 109, no. 5, pp. 756–795, 2021.

[19] A. Zhang, S. Song, J. Wang, and P. S. Yu, "Time series data cleaning: From anomaly detection to anomaly repairing," *Proceedings of the VLDB Endowment*, vol. 10, no. 10, pp. 1046–1057, 2017.

[20] C. C. Aggarwal, *Outlier Analysis*. Springer, 2013. [Online]. Available: http://dx.doi.org/10.1007/978-1-4614-6396-2

[21] W. Hilal, S. A. Gadsden, and J. Yawney, "Financial fraud: a review of anomaly detection techniques and recent advances," *Expert Systems with Applications*, vol. 193, 2022.

[22] L. Dong, L. Shulin, and H. Zhang, "A method of anomaly detection and fault diagnosis with online adaptive learning under small training samples," *Pattern Recognition*, vol. 64, 2017.

[23] D. Miljković, "Fault detection methods: A literature survey," in *34th International Convention MIPRO*, 2011, pp. 750–755.

[24] V. Mikuni, B. Nachman, and D. Shih, "Online-compatible unsupervised nonresonant anomaly detection," *Physical Review D*, vol. 105, 2022.

[25] N. S. Chauhan, "Dbscan clustering algorithm in machine learning: An introduction to the dbscan algorithm and its implementation in python." 2022. [Online]. Available: https://www.kdnuggets.com/2020/04/dbscan-clustering-algorithm-machine-learning.html

[26] X. Gu, L. Akoglu, and A. Rinaldo, "Statistical analysis of nearest neighbor methods for anomaly detection," *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[27] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation forest," in *8th IEEE International Conference on Data Mining (ICDM)*, 2008, pp. 413–422.

[28] A. Bounsiar and M. G. Madden, "One-class support vector machines revisited," in *International Conference on Information Science Applications (ICISA)*, 2014, pp. 1–4.

[29] W.-C. Chang, C.-H. Wu, Y.-H. Tsai, and W.-Y. Chiu, "Object volume estimation based on 3d point cloud," in *International Automatic Control Conference (CACS)*, 2017, pp. 1–5.

[30] D. Moreno and G. Taubin, "Simple, accurate, and robust projector-camera calibration," in *2nd IEEE International Conference on 3D Imaging, Modeling, Processing, Visualization & Transmission*, 2012, pp. 464–471.

[31] H. Huadong, C. Xianlei, S. Haolei, L. Xuemin, and Y. Pengju, "The automatic measurement system of large vertical storage tank volume based on 3d laser scanning principle," in *13th IEEE International Conference on Electronic Measurement & Instruments (ICEMI)*, 2017, pp. 211–216.

[32] Y. Zhi, Y. Zhang, H. Chen, K. Yang, and H. Xia, "A method of 3d point cloud volume calculation based on slice method," in *International Conference on Intelligent Control and Computer Application (ICCA)*, 2016, pp. 155–158.

[33] S. Armstrong, "gray-code-structured-light." [Online]. Available: https://github.com/sjnarmstrong/gray-code-structured-light

[34] W. Hongyan, Y. Ning, C. Hui, L. Weibin, and B. Ahmad, "Volume calculation for power equipment point cloud based on concave hull slice method," in *6th IEEE International Conference on Power and Renewable Energy (ICPRE)*, 2021, pp. 302–306.

[35] S. G. Kwak and J. H. Kim, "Central limit theorem: the cornerstone of modern statistics," *Korean Journal of Anesthesiology*, vol. 70, no. 2, p. 144, 2017.