2023 IEEE International Symposium on Robotic and Sensors Environments (ROSE) | 979-8-3503-0804-4/23/$31.00 ©2023 IEEE | DOI: 10.1109/ROSE60297.2023.10410732

Improving Fitting CAD to 3D Point Cloud Acquired with Line-of-Sight Sensor

Marek Franaszek*, Prem Rachakonda†, Kamel S. Saidi‡,
National Institute of Standards and Technology
Gaithersburg, MD 20899-8239, USA
* marek@nist.gov † prem.rachakonda@nist.gov ‡ kamel.saidi@nist.gov

***Abstract*—Due to self-occlusions in 3D point clouds acquired with line-of-sight sensors, an incomplete representation of an object's surface is used in fitting a CAD (computer-aided design) model to the data. For certain types of object geometry, CAD pose fitted using the Iterative Closest Point (ICP) procedure is systematically misaligned, even when the selected initial pose is very close to the expected, correct pose. We demonstrate on experimental data obtained from manufacturing-relevant parts that the final CAD alignment can be greatly improved if only a part of the CAD surface is used in the ICP registration. The resulting residual ICP error is then reduced by three to four times.**

***Index Terms*—3D point cloud, CAD fitting, ICP registration, line-of-sight sensor, bin picking**

I. Introduction

Accurate determination of six degrees of freedom (6DOF) pose is a critical prerequisite for successful deployment of vision systems in many manufacturing applications. When a CAD model of a part is known, fitting a model to a 3D point cloud of the part requires a minimization of a properly defined error function. For a CAD model represented by a mesh, the error function can be calculated as a root mean square (rms) of point-to-point or point-to-surface distances. It is therefore a common practice to gauge the quality of the fit by reporting a residual value (in millimeters) of the error function, where a small value is usually interpreted as an indication of a good fit.

A metric used in CAD fitting may be misleading since the residual error is affected by two different and independent factors: 1) accuracy of the model selected for fitting, and 2) the amount of noise and bias in the experimental data. In industrial applications such as bin picking, the proper CAD model of the part is known a priori. However, segmentation of a single part in a point cloud representing an unstructured pile of parts is a challenging task which is prone to the inclusion of outliers. The presence of outliers, in turn, may lead to an inflated value of the residual error, which may be interpreted as an indicator of bad alignment. This may be at odds with a subjective impression following a visual inspection. Such a situation is illustrated in Fig. 1 which shows the same CAD model fitted to two different point clouds acquired from the same part using two different sensors. In spite of much cleaner 3D data in Fig. 1b and a clearly smaller value of the residual error, visual observation provides ambiguous impression. While none of the CAD alignments in Fig. 1 is accurate, one may well

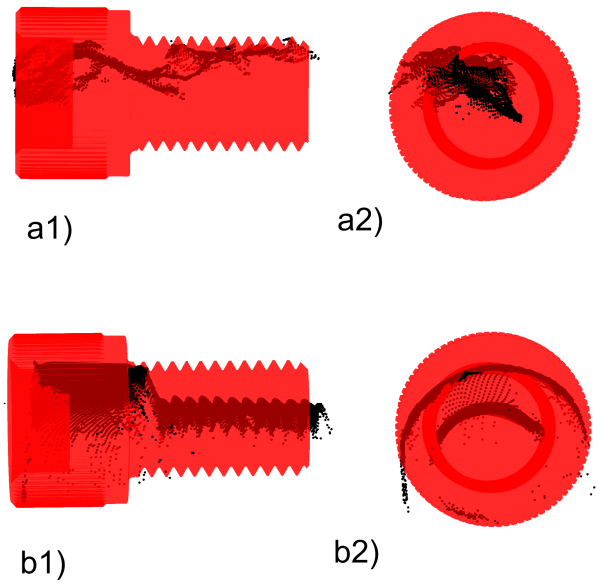

Fig. 1. Outcome of fitting the same CAD model to two different point clouds (rows a and b), shown from two viewpoints (columns 1 and 2). Residual error: a) 4.93 mm; b) 3.9 mm.

argue that the better fit is that of Fig. 1a. A similar ambiguity was reported earlier: two different registration techniques of the same datasets yielded very different residual error values and yet, better alignment corresponded to the larger error, see Fig. 7 in [1].

Data segmentation is the first step in pose determination and it is natural to try to relate the outcome of the fitting to the accuracy of the segmentation. In the research domain, when simulated data are used, the quality of the segmentation and the resulting object's pose can be directly compared with ground truth. For example, Intersection over Union (IoU), which provides a percentage of segmented data overlapping with the ground truth, was used to train Deep Learning (DL) networks, such as PointNet [2], [3], [4] or SMA-Net, that are capable of both part identification and CAD fitting [5]. This approach was also adopted in Benchmark for 6D Object Pose Estimation (BOP) Challenge [6]. The biggest problem with such approaches is the realism of the simulated noise. For

Authorized licensed use limited to: NIST Virtual Library (NVL). Downloaded on March 04,2024 at 18:56:08 UTC from IEEE Xplore. Restrictions apply.

example, 3D point clouds collected by structured light systems are notoriously contaminated by phantom points (data points floating in empty small regions with no correspondence to the physical objects in a scanned scene). Such points cannot be easily simulated as traced rays, in a synthetic scene, that are perturbed using some form of Gaussian noise [7].

When a ground truth is not available, an entropy-based method could be used to evaluate a quality of segmentation. Originally developed for Red Green Blue (RGB) images, it could be applied to 3D data derived from depth images [9], but the metric only provides a relative evaluation of two different segmentations.

In another approach to gauge accuracy of fitting, simple geometric primitives were extracted from more complex CAD models and fitting accuracy was expressed as a difference between the nominal, ground truth CAD parameters and the fitted parameters (e.g., sphere radius, cone height, and apex angle, etc.) [8]. However, this metric is limited to CAD models in which such geometric primitives can be identified and it gauges only a selected aspect of the fitting process. A similar approach was used when an acquired point cloud and *a priori* a CAD model of industrial pipelines were used to adjust CAD parameters [10].

When it comes to registering two 3D point clouds, the Iterative Closest Point (ICP) algorithm is the most common procedure [11], [12]. It alternates between finding the corresponding points and finding a rigid body transformation that minimizes the distances between them. ICP requires a good starting pose to avoid being trapped in incorrect local minima [13]. This restriction can be mitigated by redesigning the minimization process to ensure that the global minimum is reached [14], but these solutions suffer from long run times. The residual value of the ICP error is then used to gauge the quality of the final alignment.

To address the fact that a relation between the residual error value and the error of the fitted 6DOF pose is not generally known, a rigorous concept of certifiable registration was introduced in [15] and tested in [17], [18]. Originally developed for two datasets with established correspondences, it can also handle a high ratio of outliers (99%) by decoupling scale, translation, and orientation searches. Taking input parameters characterizing noise, such as signal to noise or noise bounds, the implemented TEASER++ procedure [16] could output estimated pose together with some error bounds, without explicit knowledge of the ground truth pose. In cases where a simultaneous pose and correspondences must be determined, the procedure may fail if there are not enough inliers to identify a unique registration. Existence of symmetries in the scanned scene may also be challenging as they make the registration non-unique.

In this paper, we investigate a fitting of CAD model to 3D point cloud acquired with the line-of-sight sensor. We show that ICP can provide systematically inaccurate poses even as the starting pose is very close to the expected, final pose. A remedy for this deficiency is proposed which decreased the residual error three to four times and shortened execution time four to nine times.

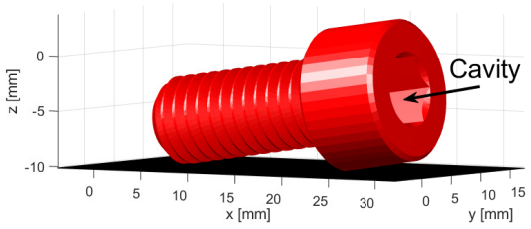

Fig. 2. Example of a part with a void cavity. Such parts are prone to erroneous CAD fitting when scanned with a line-of-sight sensor.

II. Perils of Fitting Some Types of Self-Occluded Surfaces

In many industrial applications, line-of-sight sensors are placed in a fixed pose within world coordinates. When 3D point clouds are acquired with these types of sensors, the resulting points are sampled only from the visible parts of the object's surface in the sensor's field of view. For many surfaces this limitation does not degrade the final output of the CAD fitted to the point cloud. However, for some more complex surfaces, portions of the surface may be convex and others may be concave. If both portions are close to each other and are located in such a manner that one can occlude the other, then this may lead to incorrect alignment of the CAD model to the surface when the point cloud is acquired from only one direction. An example of such an object is shown in Fig. 2. As can be seen, the head of the screw has a cavity and a part of the inner surface defining its boundary is shadowed by the external surface of the head. The use of a line-of-sight sensor to acquire a point cloud for such object can have negative impact on accuracy of fitted CAD pose. In Fig. 3, the initial starting CAD pose (in red) together with the acquired 3D point cloud (in black) are shown, and the resulting fitted end pose is displayed in Fig. 1b. In spite of a relatively good initial alignment, the ICP procedure was trapped in an incorrect minimum. To demonstrate this problem, a series of experiments described in the next section was conducted.

III. Experiment

Four different sized black-oxide socket head screws were used, see Fig. 4. For each selected size, $M = 16$ screws were placed in a random orientations on a flat table and scanned with Zivid One+Small sensor[1]. An example screw placement on a table is shown in Fig. 5. All adjustable parameters of the sensor were set to their default values. Approximated distance between sensor center and the center of table was 0.5 m. For comparison, a bin filled with a pile of randomly placed screws was also scanned with the same sensor. Additionally, a

[1]Certain commercial equipment, instruments, or software are identified in this paper to foster understanding. Such identification does not imply recommendation or endorsement by NIST, nor does it imply that the equipment or software identified are necessarily the best available for the purpose.

Authorized licensed use limited to: NIST Virtual Library (NVL). Downloaded on March 04,2024 at 18:56:08 UTC from IEEE Xplore. Restrictions apply.

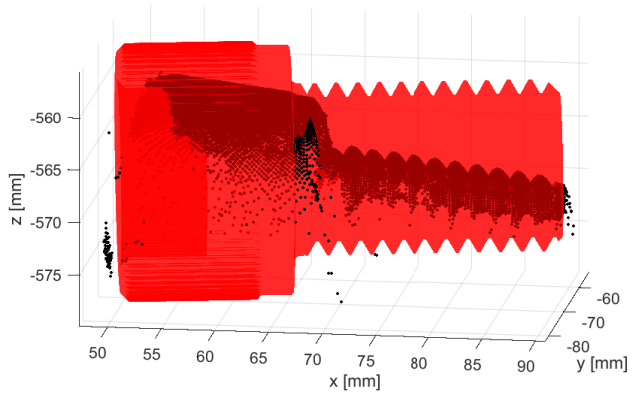

Fig. 3. Example of a 3D point cloud acquired with a line-of-sight sensor and an initially aligned CAD pose.

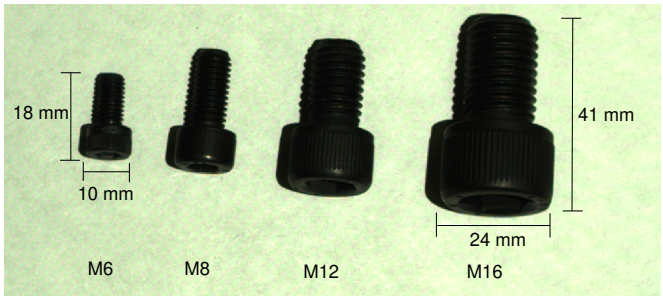

Fig. 4. Four sizes of screws used in experiment. Labels M indicate the bolt diameter in millimeters.

single screw was placed on a turntable and scanned with Faro Quantum S sensor mounted on 7DOF arm (Articulating Arm Coordinate Measuring Machine - AACMM) which automatically registered scans from different directions and output a 3D point cloud covering almost all of the part's entire surface.

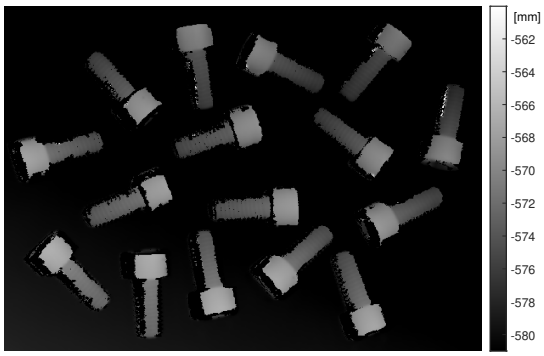

Fig. 5. Example of a depth image of a random configuration of M8 screws used in the experiment.

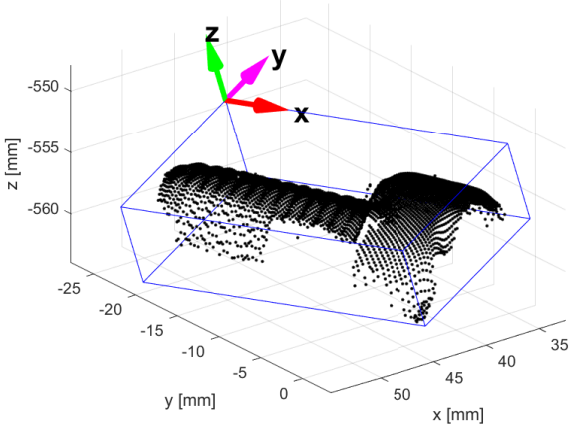

Fig. 6. Example of oriented 3D bounding box and properly oriented local coordinate frame shown with color arrows.

IV. Data Processing

All calculations and graphic visualizations were performed in Matlab using its Image Processing Toolbox. For each acquired 3D point cloud of the screws on the table, such as in Fig. 5, individual screws were manually segmented and the 3D oriented bounding box was calculated using the $minboundbox()$ procedure [19] for each screw. In order to initially align a CAD model with its segmented point cloud, a properly defined orientation matrix $\mathbf{R}_{init}$ was derived from the corners of the 3D bounding box and some features characterizing the CAD model. This step was necessary since eight corners of the bounding box output by $minboundbox()$ do not resolve ambiguity caused by the definition of the local CAD coordinate frame (for example, if axis of symmetry is parallel to $\hat{\mathbf{z}}$, the head of the screw may be oriented up or down). An example of a calculated oriented 3D bounding box with properly defined local unit vectors is shown in Fig. 6. CAD vertices were rotated using $\mathbf{R}_{init}$ and the point-to-point version of the ICP procedure $pcregistericp()$ was used to finalize a CAD alignment with the segmented part of the 3D point cloud. The residual error of the fitting was recorded. This procedure of using all CAD vertices in the ICP registration will hereafter be referred to as Method A.

A. Improving CAD alignment

Due to the peculiar geometrical feature of the screw described in Section II, the final CAD alignment with the 3D point cloud is not very good, as can be seen in Fig. 1b. Since line-of-sight sensors acquire data from only a portion of an object's surface, it is reasonable to use only the subset of CAD vertices that are visible in the initial orientation $\mathbf{R}_{init}$ in the ICP minimization. For good initial alignment, such a strategy eliminates the risk that the ICP minimization will try to match the occluded part of the CAD surface with experimental 3D points. Such risks exists when a ray originating from the sensor's center passes through an object's surface in two

Authorized licensed use limited to: NIST Virtual Library (NVL). Downloaded on March 04,2024 at 18:56:08 UTC from IEEE Xplore. Restrictions apply.

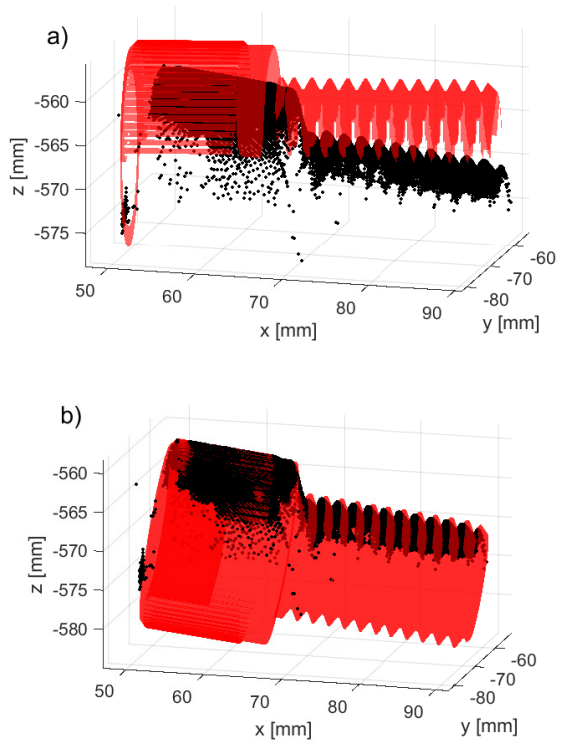

Fig. 7. 3D point cloud and: a) visible part of CAD surface from Fig. 3; b) full surface of CAD fitted with Method B.

close points which have local normals pointing in opposite directions. Visible CAD vertices are obtained using the Hidden Point Removal (HPR) procedure [20] with a viewpoint set at $(0, 0, 0)$ (the origin of sensor's coordinate frame). In Fig. 7a, the same 3D point cloud and the visible part of the CAD surface, in the same orientation as in Fig. 3, are shown. The visible portion of the CAD vertices is then fed into the *pcregistericp()* procedure and the resulting CAD alignment with the 3D point cloud is shown in Fig. 7b. This should be compared with Fig. 1b, which shows the outcome of the registration when all CAD vertices were used. This procedure of using only a subset of the CAD vertices visible in the initial orientation $\mathbf{R}_{init}$ for the ICP registration is hereafter referred to as Method B.

V. Results

More examples of fitting to three sizes of screws are plotted in Fig. 8: top row is for M6 screws, the middle row for M12 screws, and the bottom row is for M16 screws. The left column in Fig. 8 shows the 3D point clouds and the visible CAD surface in the pose determined from the 3D oriented bounding box, as described in Section IV. This pose serves as an initial, starting pose for the ICP registration. The middle column in Fig. 8b shows the outcomes of Method A and the right column shows the outcomes of Method B. In Fig. 9 and Fig. 10 the residual ICP errors are shown for all $M = 16$ poses and two screw sizes: M16 and M12, respectively. For each part index, n, the CAD starting pose was the same for both methods A and B.

The longest triangle side in the mesh representing the CAD model was restricted to be shorter than 0.2 mm for screw sizes M16, M12, and M8, and to 0.1 mm for the M6 screws.

For all four screw sizes, each individual screw was manually segmented from the 3D point cloud. This careful and labor-intensive operation resulted in very clean subsets of 3D points. The average number of points N in the segmented subsets depends on the screw size and varies between 2,000 and 14,000 points.

Central Processing Unit (CPU) times, T_A and T_B, for running the $pcregistericp()$ procedure were recorded for both methods A and B, respectively. Recorded times vary greatly for different screw sizes as they depend on the number of segmented points N and the number of vertices in the CAD model. A smaller variation in CPU times was also observed for all $M = 16$ screws of the same size. On average, the ratio T_A/T_B varied between 4 and 9. The extra time needed to execute the HPR procedure took less than 1% of T_A.

Fig. 11 shows the example of fitting a CAD model to a 3D point cloud that covers the majority of a part's surface (data acquired with AACMM).

Fig. 12 shows a fragment of an unstructured pile of screws in a bin with the CAD fitted to one of them using Method A and B.

VI. Discussion

The initial CAD pose for the ICP registration derived from the 3D oriented bounding box, such as that shown in Fig. 6, resulted in a good alignment with 3D point cloud. The average angle of relative rotations between the initial and the final CAD orientations, such as shown in the left and the right columns of Fig. 8, was about $3.7°$. However, in spite of a relatively good initial alignment, the final CAD pose fitted to an experimental point cloud using all CAD vertices (Method A) is clearly incorrect, as can be seen in Fig. 1b and Fig. 8b. On the other hand, using only the subset of CAD vertices visible in the initial pose for the ICP registration (Method B), such as that shown in Fig. 7 and Fig. 8a, greatly improved the final CAD alignment. The improvement can be quantified by a reduction in the ICP residual error for Method B vs. Method A, as shown in Fig. 9 and Fig. 10. The outlier in Fig. 9b indicates the only case where a use of only the visible portion of the CAD vertices failed to reduce the ICP residual error. For all other cases, the reduction was between three to four times.

Use of visible only CAD vertices in Method B eliminates existence of undesired, local minimum caused by closeness of convex and concave regions on object's surface.

Residual error for Method A in Fig. 12a is 2.00 mm and 0.57 mm for Method B in Fig. 12b. This indicates that Method B is superior over Method A not only for the particular configuration, such as in Fig. 5, where all screws are placed

Authorized licensed use limited to: NIST Virtual Library (NVL). Downloaded on March 04,2024 at 18:56:08 UTC from IEEE Xplore. Restrictions apply.

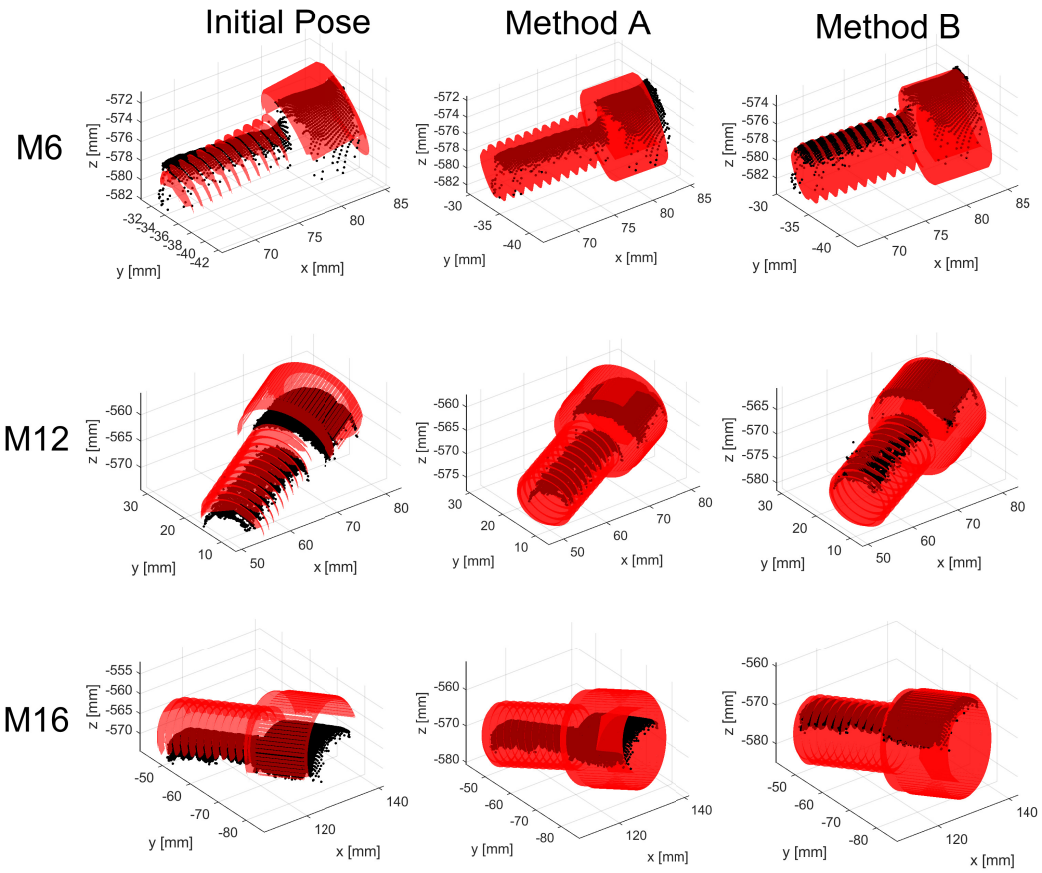

Fig. 8. Examples of CAD (red) and segmented 3D point clouds (black) for three screw sizes: M6, M12 and M16 (top to bottom). Left to right column: visible only part of CAD in the initial pose used by ICP; final CAD alignment using Method A; final CAD alignment using Method B. The right column shows clearly better alignment of CAD with 3D data than the results in the middle column.

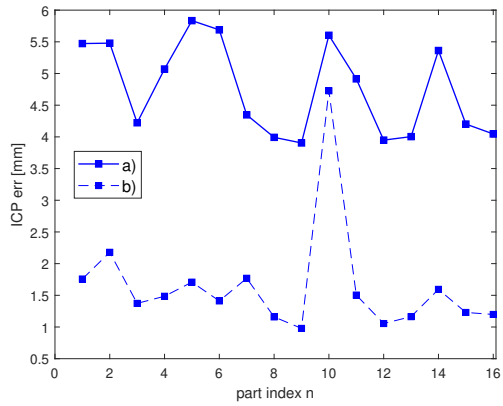

Fig. 9. Residual ICP error for: a) Method A; b) Method B. Plotted results are for M16 screws.

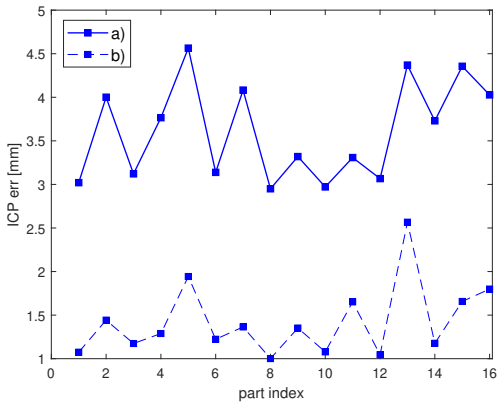

Fig. 10. Residual ICP error for: a) Method A; b) Method B. Plotted results are for M12 screws.

Authorized licensed use limited to: NIST Virtual Library (NVL). Downloaded on March 04,2024 at 18:56:08 UTC from IEEE Xplore. Restrictions apply.

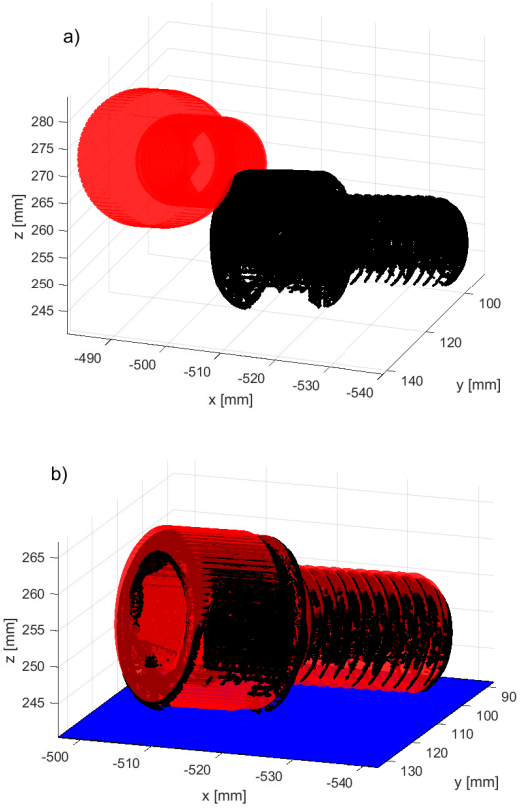

Fig. 11. Fitting CAD with Method A to 3D data covering most of part's surface (black points): a) CAD in starting pose; b) aligned final CAD pose, plane fitted to the table plotted in blue. Plotted results are for M16 screws.

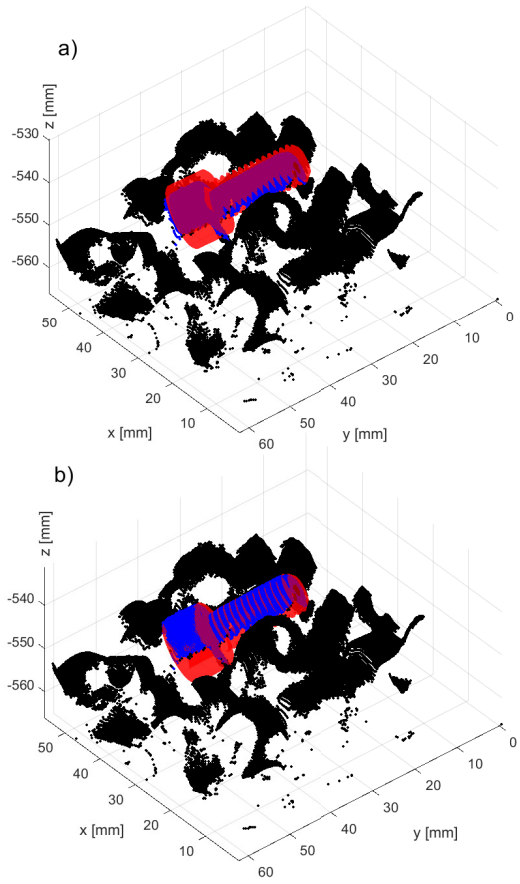

Fig. 12. Fitting CAD (shown in red) to the segmented points (blue) selected from a 3D point cloud (black) acquired from an unstructured pile of parts and using: a) Method A; b) Method B. Plotted results are for M8 screws.

on a flat surface and are clearly separated from each other. Automated segmentation of a single part from an unstructured pile of parts is much more challenging and may include more outliers than manually segmented parts distributed in the special configuration shown in Fig. 5. Thus, one may expect a persistent problem with 3D data acquired by line-of-sight sensors from certain objects and Method B to generally yield smaller values of the residual errors for such cases. While the described phenomenon was demonstrated only on screws, other parts with convex and concave surfaces (such as bushings, spacers, shaft collars, etc.) are expected to reveal the same pattern when scanned with line-of-sight sensors.

The problem shown in Fig. 8b did not occur when the CAD models were fitted to the point clouds covering the majority of the part's surface and acquired with the AACMM sensor. Even when the starting CAD model pose for the ICP registration was very poorly aligned with the 3D data, as in Fig. 11a, Method A yielded a very good final result. As discussed earlier, a proximity of convex and concave regions on screw's surface caused ICP to be trapped in incorrect local minimum when line-of-sight sensor was used. While there are no many points from the inner wall of the cavity in the point cloud acquired by AACMM, majority of part's external surface was sampled with AACMM and this was sufficient for ICP to converge to the right pose.

VII. Conclusions

In vision guided robotic applications such as bin picking, a decision to accept or reject a part's pose is based on the residual error of the CAD fitting. Use of only the visible subset of the CAD model's vertices in the ICP registration (Method B) can greatly improve the alignment between certain types of CAD models with point clouds acquired by line-of-sight sensors. Even for solid objects with geometries not affected by the problem described in this paper, it may be beneficial to use Method B as it substantially reduces execution time of the ICP procedure.

References

[1] Alvaro Parra Bustos, Tat-Jun Chin, and David Suter, "Fast rotation search with stereographic projections for 3d registration," Proceedings of the IEEE conference on computer vision and pattern recognition, pp.3930–3937, 2014.

[2] Charles R. Qi, Hao Su, Kaichun Mo, Leonidas J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," Proceedings of the IEEE conference on computer vision and pattern recognition, pp.652–660, 2017.

Authorized licensed use limited to: NIST Virtual Library (NVL). Downloaded on March 04,2024 at 18:56:08 UTC from IEEE Xplore. Restrictions apply.

[3] Michael Danielczuk, Matthew Matl, Saurabh Gupta, Andrew Li, Andrew Lee, Jeffrey Mahler, and Ken Goldberg, "Segmenting unknown 3d objects from real depth images using mask r-cnn trained on synthetic data," Int. Conf. on Robotics and Automation (ICRA), pp.7283–7290, 2019.
[4] Yajun Xu, Shogo Arai, Fuyuki Tokuda, and Kazuhiro Kosuge, "A convolutional neural network for point cloud instance segmentation in cluttered scene trained by synthetic data without color," IEEE Access, vol.2, pp.70262–70269, 2020.
[5] Sijie Hu, Arnaud Polette, and Jean-Philippe Pernot, "SMA-Net: Deep learning-based identification and fitting of CAD models from point clouds," Engineering with Computers, pp.1–22, 2022.
[6] Tomáš Hodaň, Martin Sundermeyer, Bertram Drost, Yann Labbé, Eric Brachmann, Frank Michel, Carsten Rother, and Jiří Matas, "BOP challenge 2020 on 6D object localization," Computer Vision–ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16, pp.577–594, 2020.
[7] Jérémy Montlahuc, Ghazanfar Ali Shah, Arnaud Polette, and Jean-Philippe Pernot, "As-scanned point clouds generation for virtual reverse engineering of cad assembly models," Computer-Aided Design and Applications, vol. 16, pp.1171–1182, 2019.
[8] Chiara Romanengo, Andrea Raffo, Yifan Qie, Nabil Anwer, and Bianca Falcidieno, "Fit4cad: A point cloud benchmark for fitting simple geometric primitives in CAD models," Computers and Graphics, vol. 102, pp.133–143, 2021.
[9] Hui Zhang, Jason E. Fritts and Sally A. Goldman, "An Entropy-based Objective Evaluation Method for Image Segmentation," Proc. SPIE Storage and Retrieval Methods and Applications for Multimedia, vol.5307, pp,38–49, 2004.
[10] Aurelien Bey, Raphaelle Chaine, Raphael Marc, Guillaume Thibault, Samir Akkouche, "Reconstruction of consistent 3D CAD models from point cloud data using a priori CAD models," ISPRS workshop on laser scanning, vol. 1, 2011.
[11] P. Besland N. McKay, "A Method for Registration of 3-D Shapes," IEEE Trans. PAMI, vol. 14, pp.239–256, 1992.
[12] Xiaoshui Huang, Guofeng Mei, Jian Zhang, and Rana Abbas, "A comprehensive survey on point cloud registration," arXiv:2103.02690, 2021.
[13] Adam Leon Kleppe, Lars Tingelstad, and Olav Egeland, "Coarse alignment for model fitting of point clouds using a curvature-based descriptor," IEEE Trans. on Automation Science and Engineering, vol. 16, pp.811–824, 2018.
[14] Jiaolong Yang, Hongdong Li, Dylan Campbell, and Yunde Jia, "Go-ICP: A globally optimal solution to 3D ICP point-set registration," IEEE transactions on pattern analysis and machine intelligence, vol. 38, pp.2241–2254, 2015.
[15] Heng Yang, Jingnan Shi, and Luca Carlone, "TEASER: Fast and Certifiable Point Cloud Registration," IEEE Transactions on Robotics, vol. 37, pp.314-333, 2021.
[16] Heng Yang, Jingnan Shi, and Luca Carlone, "TEASER: Fast and Certifiable Point Cloud Registration," arXiv: 2001.07715v2, 2020.
[17] Simone Fontana, Daniele Cattaneo, Augusto L. Ballardini, Matteo Vaghi and Domenico G. Sorrenti, "A benchmark for point clouds registration algorithms," Robotics and Autonomous Systems, vol.140, p.103734, 2021.
[18] Yu-Kai Lin, Wen-Chieh Lin, and Chieh-Chih Wang, "K-Closest Points and Maximum Clique Pruning for Efficient and Effective 3-D Laser Scan Matching," IEEE Robotics and Automation Letters, vol. 7, pp. 1471-1477, 2022.
[19] Johannes Korsawe, "Minimal Bounding Box," https://www.mathworks.com/matlabcentral/fileexchange/18264-minimal-bounding-box, 2015.
[20] Sagi Katz, Ayellet Tal, and Ronen Basri, "Direct Visibility of Point Sets," ACM Trans. on Graphics, vol. 26, pp.1–11, 2007.

Authorized licensed use limited to: NIST Virtual Library (NVL). Downloaded on March 04,2024 at 18:56:08 UTC from IEEE Xplore. Restrictions apply.