Automated Removal of Planar Clutter from 3D Point Clouds for Improving Industrial Object Recognition

T. Czerniawski [a], M. Nahangi [a], S. Walbridge [a], C. Haas [a]

[a] Ralph Haas Civil Infrastructure Sensing Laboratory, University of Waterloo, Waterloo, Canada.
E-mail: tczernia@uwaterloo.ca, mnahangi@uwaterloo.ca, swalbridge@uwaterloo.ca, chaas@uwaterloo.ca

Abstract

The industrial construction industry makes use of prefabrication, preassembly, modularization and off-site fabrication (PPMOF) for project execution because they offer a superior level of control as compared to on-site operations. This control is enabled by systematic and thorough performance feedback loops. Improvement of the feedback systems within these facilities will require a transition away from suboptimal manual data collection to more reliable automated data collection and processing. Laser scanners are an effective tool for automatically gathering dimensional data but extraction of useful information from point clouds remains a challenge. The speed of 3D object recognition methods depends on the size of the search space. Methods for reducing this search space are needed in order to improve the performance of 3D object recognition and subsequent information extraction. Large planar objects (e.g. floors and walls) constitute a large portion of the search space in fabrication facilities, yet are rarely the objects of interest for analysis. In this paper, an automated framework for detecting and removing large planes in point clouds is presented to speed up object recognition. The raw point cloud is first Guassian mapped to normal vector space by calculating normal vectors at each point. The Gaussian sphere is clustered using a density-based clustering algorithm and major parallel planes are segmented from the rest of the point cloud. The major planes are removed and the remaining objects in the scene continue on to 3D object recognition. Results show the algorithm for automatic plane removal can reduce the search space for object recognition by as much as 60% or 70%.

Keywords:

Fabrication and process control, industrial fabrication, pipe spool, 3D laser scanning, 3D point cloud, 3D object recognition, clutter removal.

1 Introduction

Laser scanning has been profoundly affecting project surveying in the construction industry. These 3D imaging sensors capture existing structural and terrestrial conditions accurately, objectively, and with greater continuity than any manual metrology methods. Despite these benefits, the potential of 3D imaging for applications like automated progress tracking [1] and automated dimensional compliance control [2-9] remains limited because isolating the object of interest from the collected data remains a manual process. Manually extracting information from the raw 3D images and running analysis is painstaking, requires many man-hours and specialized personnel training, and therefore discourages adoption by industry. Industrial object recognition is the fundamental enabler for further developments in automated spatial analysis and information flow in construction applications.

The search space in many industrial scans is largely comprised of massive planer objects (i.e. walls, floor, and ceiling). For cases where these planes are not the focus of the analysis, they clutter the search space and substantially slow the recognition of the object of interest. Therefore, quick removal of these planes before the object recognition process begins is desirable.

2 Background

2.1 Industrial Object Recognition

3D object recognition is the process of detecting the presence of an object in a 3D image with similar characteristics to a reference image or model and mapping the 3D coordinates of the reference to the 3D coordinates (or world coordinates) of the detected object in 3D space [10,11]. For industrial specific object recognition, some of the popular methods found in the literature are based on voxel connectivity [12-14], 2D primitive fitting [15-18], classifier training [19], and feature and shape descriptor matching [20-23]. Although a variety of approaches exist, all are negatively impacted by unnecessarily cluttered search spaces.

2.2 Plane Recognition

A few plane removal algorithms exist in the literature, each with varying degrees of performance.

2.2.1 RANSAC

Basic RANSAC [24] is comprised of two repeating steps: (1) minimal set selection and (2) minimal set evaluation. The minimal set for plane removal is a single point along with its normal vector, as this provides a complete description for a plane. RANSAC randomly samples minimal sets from the scan data, fits a plane using their description, and counts the number of points in the scan that are consistent with the fitted plane. After a given number of trials, a plane is considered to be recognized at the locations defined by the minimal set which achieved a score higher than a predefined threshold. Although basic RANSAC is conceptually simple, a direct application to plane recognition is computationally intensive. Methods for speeding up RANSAC have been explored [25].

2.2.2 Hough Transform

The general Hough transform [26] can be used to recognize planes within noisy data. It is comprised of three steps, (1) repeated transform mapping, (2) application of a voting rule and (3) finding the shape parameters within the accumulated array of votes. Use of the 3D Hough transform for extraction of planar faces from point clouds was investigated by Vosselman and Dijkman. Randomized Hough Transform is a variant of the 3D Hough transform that has proven to be especially effective for plane detection in point clouds [27,28].

2.2.3 Gaussian Mapping

An elegant solution for identifying major planes within point cloud data includes mapping normal vectors to a Gaussian sphere [29,30]. Each cluster on the Gaussian sphere represents a direction that is perpendicular to major sets of parallel planes. Gaussian mapping will be used by the framework presented in this paper because of its superior speed of execution on large point sets.

3 Research Methodology

The framework for automatic plane removal is a pre-processing step for 3D object recognition. The method for planar clutter removal has five steps as outlined in Figure 1. Following plane removal, 3D object recognition is then performed using a novel method based on local data level curvature estimation, clustering, and bag-of-features matching [31].

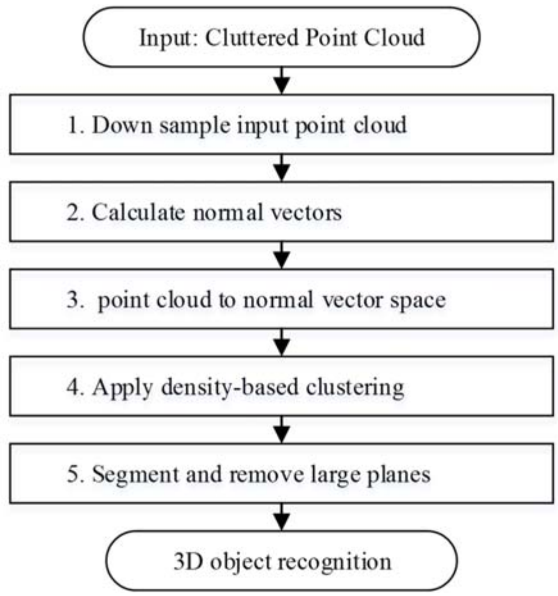

Figure 1: Framework for automatic removal of planar clutter objects.

3.1 Down Sampling

The point cloud is first down sampled because the original point set is usually massive in size with a density of points exceeding the amount necessary for accurate spatial analysis of any objects of interest. For down sampling, a random filter is applied on the original point cloud to reasonably reduce its size. The down sampled point cloud still adequately represents the objects and components without containing an excessive amount of points. Down sampling resolves the problem of redundancy that causes expensive computation with no additional benefit.

3.2 Calculating Normal Vectors

In order to detect planar clutter and filter it out from the point cloud, normal vectors are calculated and provide the primary feature for clustering and segmentation. A three step procedure is applied to each point to calculate its normal vector:

1. Identify the nearest neighbours of a point using K-Nearest Neighbours (KNN) with K=12. The 12 closest points to each point constitute the nearest neighbourhood for that point and are the set of points used for plane fitting.
2. Plane fitting is performed on the point and corresponding neighbourhood identified in the previous step. For fitting a plane to a dataset, least-square adjustment is employed.
3. The normal vector of the plane, which also represents the normal vector of the point being investigated, is extracted from the plane fitted to the investigated point.

3.3 Transforming the Point Cloud into the Normal Vector Space

Once a normal vector has been calculated for each point in the down sampled point cloud, the point cloud is again down sampled and mapped to vector space in a process called Gaussian mapping. So rather than each point being represented using (x,y,z) in Cartesian space, each point is represented by their normal vector (n_x, n_y, n_z) on a Gaussian sphere (Figure 2). This coordinate space transformation will result in identification of major planar regions existing in the point cloud. A similar transformation is used for clustering parallel branches of a pipe spool in a 3D point cloud [32]. The transformation will results in a sphere with scattered points.

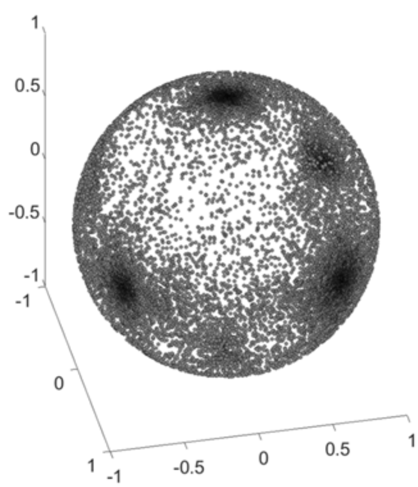

Figure 2: A typical Gaussian sphere of a point cloud. Dense point clusters on the Gaussian sphere correspond with directions perpendicular to major planar objects.

3.4 Clustering

For extracting the planar regions in the point cloud, a density-based clustering algorithm is used. In the normal vector space, planar regions are all concentrated around a point, which is the normal vector of the plane. In this paper, density-based clustering of applications with noise (DBSCAN) is employed [33]. The advantage of using density-based methods rather than centroid-based methods is that the noise and outliers are removed simultaneously using DBSCAN. While in centroid-based methods (e.g. k-means clustering), outliers and noise is also clustered with the inliers, which causes errors in the plane detection phase [34]. In addition, DBSCAN does not require the number of clusters to be known as a-priori, while in k-means clustering, the value of k is to be known to initiate the clustering process. : (a) Clustered points in the point cloud based on DBSCAN clustering. (b) Corresponding points are shown with similar colours in the down sampled point cloud.

Figure 3 shows typical results of major parallel plane set detection using DBSCAN on a point cloud.

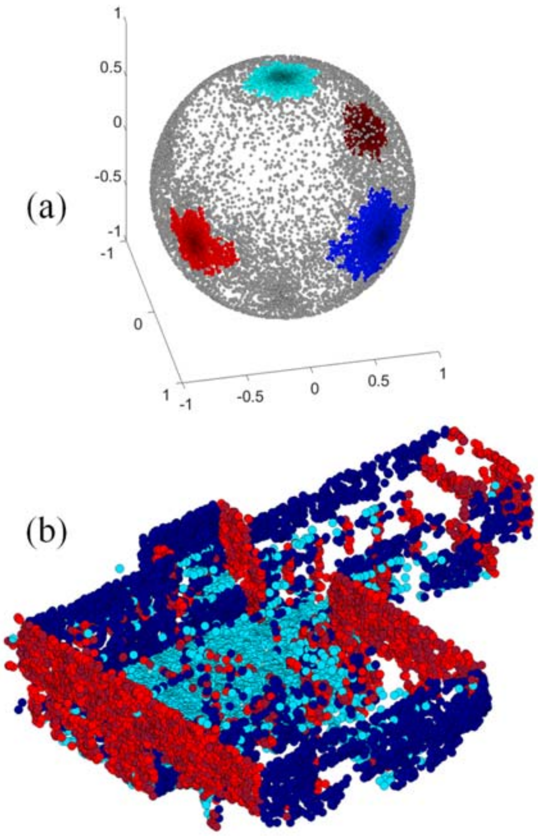

Figure 3: (a) Clustered points in the point cloud based on DBSCAN clustering. (b) Corresponding points are shown with similar colours in the down sampled point cloud.
As seen in

Figure 3, major parallel plane sets are clustered, because parallel planes have the same normal vectors. The major parallel plane sets can then be segmented based on Euclidean space separation using an appropriate approach such as hierarchical clustering (Figure 4). The same approach is used to separate parallel branches of a pipe spool that are clustered similarly [32].

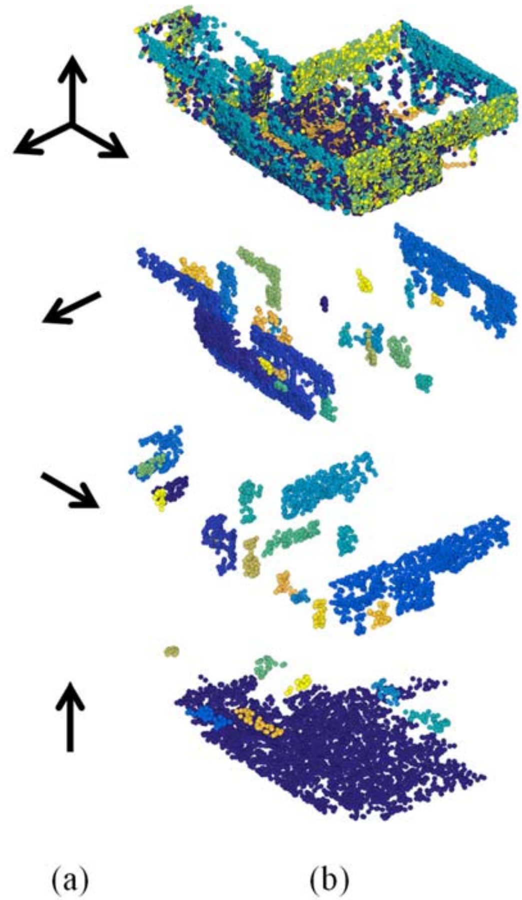

Figure 4: Segmenting and removing planar objects. (a) Normal vectors to major parallel plane sets (b) segmented major parallel plane sets

3.5 Filtering out the Planar Clutter

The segmented planes or clusters, which represent the major planes in the point cloud are then removed from the original point cloud along with all other points within their bounding box. In other words, once the planes are detected as the major clusters, the outliers which represent the points that are not belonging to the planes are kept in the point cloud.

4 Experimental Validation

The automated plane removal framework proposed is tested on two case studies: a laser scan of the Civil Infrastructure Sensing (CIS) laboratory at the University of Waterloo (Figure 5) and a scan of an industrial pipe spool fabrication facility (Figure 6).

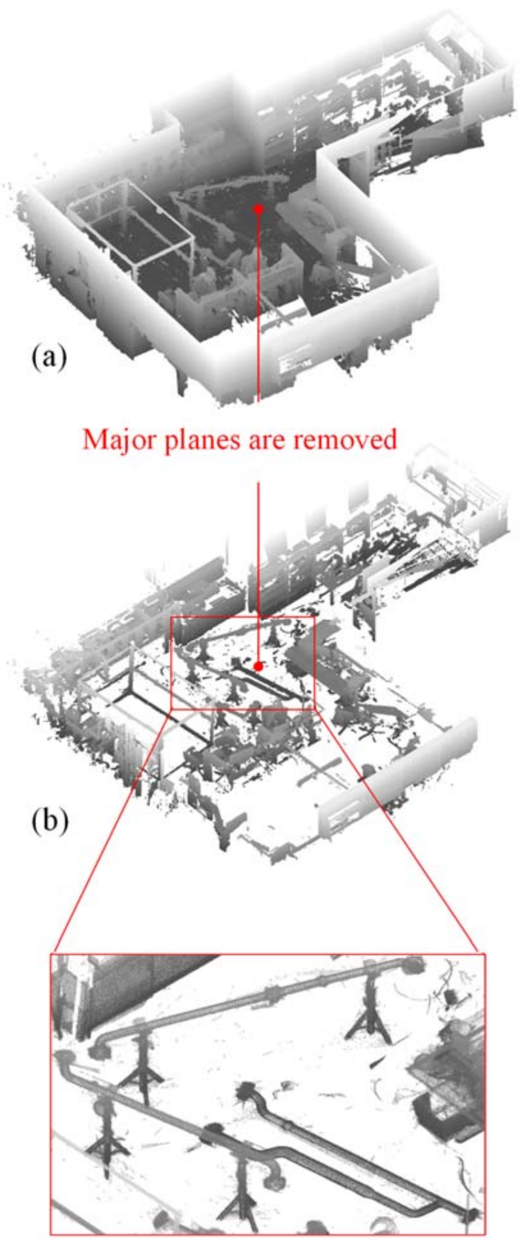

Figure 5: : Results of automated plane removal case 1, (a) original scan with 8.2 million points, (b) scan with major planar clutter removed with3.3 million points maintains rich information on the pipe spools of interest.

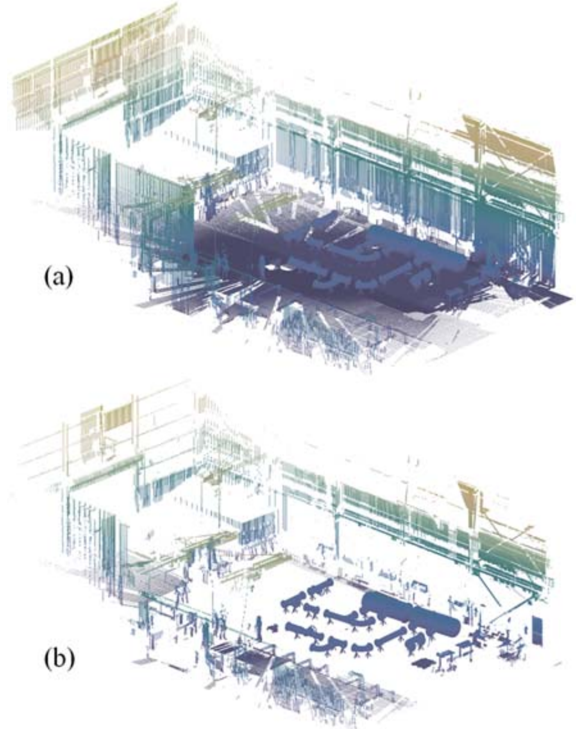

Figure 6: Results of automated plane removal case 2, (a) original scan with 7 million points, (b) scan with major planar clutter removed with 2.1 million points.

In both case studies, a set of pipe spool objects must remain in the scans for 3D object recognition and must not be affected by the plane removal. During the data acquisition, walls and floors are captured by the indiscriminant laser scanner and clutter the point cloud search space. Such clutter and unwanted objects are to be automatically removed, in order to improve the 3D object recognition step.

The object recognition framework used by the authors [31] relies on curvature calculation. Knowing mechanical/electrical/piping (MEP) components are usually non-planar, removing planar regions will reduce the processing time for curvature calculation.

As seen in Figures 5 and 6, major planes captured during laser scanning were successfully removed automatically using the framework for planar regions removal. This removal will result in a significant reduction of the point cloud search space. In case study 1, the original size of the point cloud was approximately 8.2 million. The plane removal framework reduced the size of the point cloud to approximately 3.3 million points, a 60% reduction. In case study 2, the original size of the point cloud was approximately 7 million points. The plane removal framework reduced the size of the point cloud to approximately 2.1 million points, a 70% reduction Applying the automated plane removal framework will down sample the point cloud in a supervised way. In other words, down sampling is performed without losing the required details in the regions of interest (**Error! Reference source not found.**). The clutter and unwanted regions are removed automatically. Applying indiscriminant random down sampling methods on the original point cloud will also reduce the size; however, such algorithms have no control on keeping the regions of interest. This will impact the accuracy of object recognition and subsequent object analysis. Using this plane removal framework will resolve this issue.

The plane removal execution time for both the CIS laboratory and industrial fabrication facility is approximately 3 minutes. The analyses are benchmarked on a computer with a 64 GB RAM and a 2×1.90 GHz×12 cores. This 3 minute execution time consequently reduces the execution time of object recognition or object search. The execution time for object recognition is proportional to the point cloud size. Therefore, the execution time for object recognition in case 1 and case 2 were reduced by 60% and 70% respectively.

5 Conclusions and Recommendations

Unsupervised down sampling of massive point clouds resolves the problem of processing time to some extent; However, it removes critical information about the object(s)-of-interest, and will therefore impact the accuracy of extraction of semantic information. A supervised down sampling algorithm was proposed. The proposed algorithm removes major planar regions from the massive laser scans acquired on fabrication shops and construction sites. Knowing that the majority of the industrial objects (focused in this paper), are non-planar, removing planar regions will reduce the point cloud size without losing detail and point density in the regions that contain the objects to be recognized and analyzed. The algorithm requires calculating normal vectors on a sub-sampled point cloud. Mapping the point cloud into the normal vector space will result in concentration of the planar regions around a point, which is normal vector to the estimated planar region. Applying a density-based clustering step will automatically segment the points that are belonging to a major plane in the point cloud. The inliers of the fitted plane are then removed to reduce the point cloud search space. The key observations and results are summarized as follows:

- Density-based clustering method is suitable for planar regions segmentation without knowing the number of planes a-priori. Specifically, DBSCAN approach is employed to automatically segment the planar regions on the normal vector space.
- The object recognition accuracy is not reduced because the down sampling framework only removes the planar regions.
- Experimental tests showed that the proposed framework can typically reduce the processing time for cylindrical object recognition down by as much as 60% or 70%.

Evaluating effective variants on the validation metrics can be a potential extension for future research. In this paper, effective variants such as the search region for plane fitting, and the thresholds for DBSCAN clustering are fixed. Obviously, these parameters must be calibrated to generalize the framework. Integrating the developed framework with some other processing steps, such as automated discrepancy analysis and automatic realignment of defective assemblies are other potential applications that the authors are currently investigating.

Acknowledgement

The authors would like to thank Natural Science and Engineering Research Council (NSERC) and Aecon for the financial support.

References

[1] Y. Turkan, F. Bosche, C.T. Haas, R. Haas, Automated progress tracking using 4D schedule and 3D sensing technologies, Autom. Constr. 22 (2012) 414-421.

[2] M. Nahangi, C.T. Haas, Automated 3D compliance checking in pipe spool fabrication, Advanced Engineering Informatics. (2014).

[3] F. Bosché, A. Guillemet, Y. Turkan, C. Haas, R. Haas, Tracking the Built Status of MEP Works: Assessing the Value of a Scan-vs-BIM System, J. Comput. Civ. Eng. 28 (2014) 05014004.

[4] F. Bosché, M. Ahmed, Y. Turkan, C.T. Haas, R. Haas, The value of integrating Scan-to-BIM and Scan-vs-BIM techniques for construction monitoring using laser scanning and BIM: The case of cylindrical MEP components, Autom. Constr. 49 (2015) 201-213.

[5] P. Tang, E. Anil, B. Akinci, D. Huber, Efficient and effective quality assessment of as-is building information models and 3D laser-scanned data, (2011) 19-22.

[6] F. Bosche, C. Haas, Automated retrieval of 3D CAD model objects in construction range images,

Autom. Constr. 17 (2008) 499-512.

[7] F. Bosche, C. Haas, B. Akinci, Automated Recognition of 3D CAD Objects in Site Laser Scans for Project 3D Status Visualization and Performance Control, J. Comput. Civ. Eng. 23 (2009) 311-318.

[8] F. Bosché, Automated recognition of 3D CAD model objects in laser scans and calculation of as-built dimensions for dimensional compliance control in construction, Advanced Engineering Informatics. 24 (2010) 107-118.

[9] M. Nahangi, T. Czerniawski, T.C. Haas, S. Walbridge, J. West, Parallel systems and structural frame realignment planning and actuation strategy, Journal of Computing in Civil Engineering. (2015).

[10] D. Hoiem, S. Savarese, Representations and Techniques for 3D Object Recognition and Scene Interpretation, San Rafael, Calif. (1537 Fourth Street, San Rafael, CA 94901 USA) : Morgan & Claypool, c2011, 2011.

[11] L. Brown, A survey of image registration techniques, ACM Comput. Surv. 24 (1992) 325.

[12] D. Gilsinn, G.S. Cheok, C. Witzgall, A. Lytle, Construction Object Identification from LADAR Scans: An Experimental Study using I-Beams, US Department of Commerce, Technology Administration, National Institute of Standards and Technology, 2005.

[13] J. Teizer, C.H. Caldas, C.T. Haas, Real-time three-dimensional occupancy grid modeling for the detection and tracking of construction resources, J. Constr. Eng. Manage. 133 (2007) 880-888.

[14] G. Erdos, T. Nakano, G. Horváth, Y. Nonaka, J. Váncza, Recognition of complex engineering objects from large-scale point clouds, CIRP Ann. Manuf. Technol. 64 (2015) 165-168.

[15] M. Ahmed, C.T. Haas, R. Haas, Autonomous modeling of pipes within point clouds, Proceedings of the 30th ISARC, Montréal, Canada. (2013) 1093-1100.

[16] M. Ahmed, C. Haas, R. Haas, Automatic Detection of Cylindrical Objects in Built Facilities, J. Comput. Civ. Eng. 28 (2014) 04014009.

[17] R. Qiu, Q. Zhou, U. Neumann, Pipe-run extraction and reconstruction from point clouds, in: Anonymous Computer Vision–ECCV 2014, Springer, 2014, pp. 17-30.

[18] Y. Liu, J. Zhang, J. Hou, J. Ren, W. Tang, Cylinder detection in large-scale point cloud of pipeline plant, Visualization and Computer Graphics, IEEE Transactions on. 19 (2013) 1700-1707.

[19] G. Pang, U. Neumann, Training-based object recognition in cluttered 3d point clouds, (2013) 87-94.

[20] B. Akinci, F. Boukamp, C. Gordon, D. Huber, C. Lyons, K. Park, A formalism for utilization of sensor systems and integrated project models for active construction quality control, Autom. Constr. 15 (2006) 124-138.

[21] A.E. Johnson, M. Hebert, Using spin images for efficient object recognition in cluttered 3D scenes, Pattern Analysis and Machine Intelligence, IEEE Transactions on. 21 (1999) 433-449.

[22] H. Son, C. Kim, C. Kim, Fully automated as-built 3D pipeline extraction method from laser-scanned data based on curvature computation, J. Comput. Civ. Eng. (2014) B4014003.

[23] H. Date, Y. Kaneta, A. Hatsukaiwa, M. Onosato, S. Kanai, Object Recognition in Terrestrial Laser Scan Data using Spin Images, Computer-Aided Design and Applications. 9 (2012) 187-197.

[24] M.A. Fischler, R.C. Bolles, Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography, Commun ACM. 24 (1981) 381-395.

[25] R. Schnabel, R. Wahl, R. Klein, Efficient RANSAC for Point-Cloud Shape Detection, Comput. Graphics Forum. 26 (2007) 214-226.

[26] D.H. Ballard, Generalizing the Hough transform to detect arbitrary shapes, Pattern Recognit. 13 (1981) 111-122.

[27] D. Borrmann, J. Elseberg, K. Lingemann, A. Nüchter, The 3D Hough Transform for plane detection in point clouds: A review and a new accumulator design, 3D Research. 2 (2011) 1-13.

[28] L. Xu, E. Oja, P. Kultanen, A new curve detection method: randomized Hough transform (RHT), Pattern Recog. Lett. 11 (1990) 331-338.

[29] Y. Liu, Y. Xiong, Automatic segmentation of unorganized noisy point clouds based on the Gaussian map, Comput. -Aided Des. 40 (2008) 576-594.

[30] Y. Wang, W. Hao, X. Ning, M. Zhao, J. Zhang, Z. Shi, X. Zhang, Automatic Segmentation of Urban Point Clouds Based on the Gaussian Map, The Photogrammetric Record. 28 (2013) 342-361.

[31] T. Czerniawski, M. Nahangi, C. Haas, S. Walbridge, Pipe spool recognition in cluttered point clouds using a curvature-based shape descriptor, (Submitted (2016)).

[32] M. Nahangi, C.T. Haas, Skeleton-based discrepancy feedback for automated realignment of industrial assemblies, Autom. Constr. 61 (2016) 147-161.

[33] M. Ester, H. Kriegel, J. Sander, X. Xu, A density-based algorithm for discovering clusters in large spatial databases with noise. 96 (1996) 226-231.

[34] A.K. Jain, M.N. Murty, P.J. Flynn, Data clustering: a review, ACM computing surveys (CSUR). 31

(1999).