DETECTION AND ROBUST ESTIMATION OF CYLINDER FEATURES IN POINT CLOUDS

Yun-Ting Su
James Bethel
Geomatics Engineering
School of Civil Engineering
Purdue University
550 Stadium Mall Drive, West Lafayette, IN, 47907
su0@purdue.edu
bethel@ecn.purdue.edu

ABSTRACT

The objective of this work is to develop new methods for efficient automatic 3D modeling of existing industrial installations from point cloud data. Traditionally, cylinder feature extraction algorithms utilize 5D Hough transforms, resulting in impractically high computational complexity. A more efficient approach uses a 2D Hough transform to estimate orientation followed by a 3D Hough transform to detect position, but still has extensive runtimes and lacks robustness in dense point cloud data. This work endeavors to (1) further decrease the runtime for cylinder feature extraction by implementing a coarse-to-fine approach, and (2) improve the robustness of the algorithm in detecting multiple cylinders by applying a clustering algorithm. In the coarse-to-fine approach, an initial estimate of the cylinder feature is quickly generated by coarsely sampling the Hough space. Subsequently, the search space is iteratively restricted based on the previous estimate while increasing sampling density to generate continually improving feature estimates until a stop criterion is reached. Results show that the implemented coarse-to-fine approach yielded an improvement in orientation estimate accuracy of 20% while reducing runtime by 74%. To improve the robustness of the Hough transform in the presence of multiple cylinders, a clustering technique is implemented on the accumulator. First, cells in the accumulator with small number of accumulations are discarded to facilitate the computation of a hierarchical tree. Clustering is then applied to group the remaining cells into clusters representing different cylinders. This method improves robustness as well as accuracy of feature extraction in point cloud data with diverse cylinders.

INTRODUCTION

Technological advances in light detection and ranging (LIDAR) have enabled acquisition of dense and accurate point clouds at high speeds (Laser scanner survey, 2005). The wealth of detailed point cloud data have necessitated and also facilitated the development of automated 3D reconstruction algorithms. Since planes and cylinders compose up to 85% of all objects in industrial scenes (Petitjean, 2002), research in 3D reconstruction and modeling have largely focused on these two important primitives. Cylinders are especially prevalent in settings such as petrochemical plants, refineries, and nuclear plants – robust automatic methods for the detection and fitting of cylinders in point cloud data are essential for 3D reconstruction of these sites.

The objective of this work is to develop new methods and techniques for efficient semi-automatic or automatic 3D modeling of existing industrial installations from point cloud data. The specific focus of this work is on automatic cylinder feature extraction using extensions of the Hough transform. The tradition approach of cylinder feature extraction involves the direct application of the Hough transform in a 5D space (since cylinders are defined by five features), which becomes impractical due to the extremely high computational complexity in 5D space. To resolve this problem, a modified two-step approach (Rabbani, 2005) was implemented first utilizing 2D and then 3D Hough transforms to reduce computation complexity. The first step estimates cylinder orientation while the second step estimates the remaining three parameters of the cylinder (radius and position) using the estimated orientation from the first step. This document presents results of cylinder feature extraction using synthetically generated point cloud data where the true cylinder parameters are known.

BACKGROUND

Laser scanning provides detailed 3D measurements. In the last ten years, the speed and accuracy of laser scanners have improved dramatically (Laser scanner survey, 2005), providing vast quantities of detailed point cloud data. Much research has been focused on the development of automated reconstruction procedures which can be classified into (1) algorithms requiring a prior segmentation (Lukacs et al., 1998; Marshall et al., 2001) and (2) methods that process raw point cloud data without initial segmentation (Bolles and Fischler, 1981; Fischler and Bolles, 1987). Most of the algorithms requiring prior segmentation are based on non-linear least squares approaches to minimize the orthogonal distance of the points from the fitted cylinder, and consequently requires good initial segmentation as well as good initial parameters for the inherently iterative procedures. The methods in the second processes raw point clouds using robust fitting methods like random sample consensus but are computationally intensive.

Hough transform-based methods are also computationally intensive, but these methods are robust in the presence of outliers and multiple instances (Hough, 1962). An effective way to reduce the complexity of the Hough transform is to use sequential processing by splitting the problem. The approach of Rabbani (2005) divides the problem of cylinder fitting into two separate steps. The first step uses the Gaussian sphere of the point cloud as its input and utilizes the 2D Hough transform to finds strong hypothesis for the direction of the cylinder axis, followed by a second step that computes a 3D Hough transform to estimate the position and radius of the cylinder. This sequential processing reduces the complexity of the Hough transform-based algorithm.

SEQUENTIAL HOUGH TRANSFORM APPROACH

An effective way to reduce the complexity of the Hough transform is to use sequential processing by splitting the problem into a set of manageable sub-problems. The approach of Rabbani (2005) divides the problem of cylinder fitting into two separate steps. The first step uses the Gaussian sphere of the point cloud as its input and utilizes the 2D Hough transform to finds estimate for the direction of the cylinder axis, followed by a second step that computes a 3D Hough transform to estimate the position and radius of the cylinder. This sequential processing drastically reduces the complexity of the Hough transform-based algorithm.

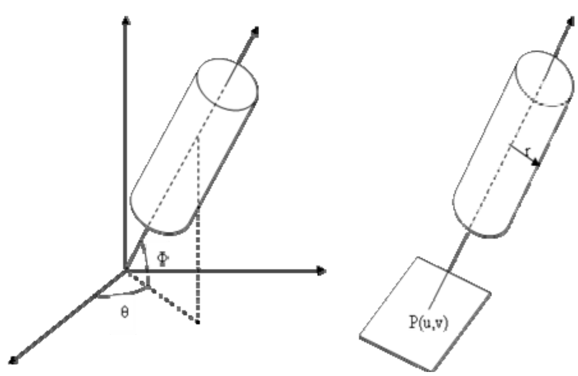

Figure 1. Five parameters of the cylinders.

Estimate the Orientation of the Cylinder

The first step of the sequential Hough transform procedure (Rabbani, 2005) is to estimate the orientation of the cylinder axis given a set of point cloud data. Carmo et al. (1976) observed that the normal of each point on the cylinder makes a great circle in the Gaussian sphere. This great circle is a result of the intersection of the unit sphere with a plane passing through the origin that is perpendicular to the normal of a given point in the point cloud. Each point of the cylinder point cloud therefore creates a separate great circle on the unit sphere. The intersection of these great circles represents the orientation of the cylinder axis.

The Hough transform is essentially a voting procedure. The previous paragraph was written in terms of continuous space, but to implement orientation estimation using Hough transforms on a computer, the algorithm must operate in discrete space. Therefore instead of using the continuous space unit sphere, the unit sphere is sampled and represented as consisting of a finite number of adjoining cells in discrete space. The i^{th} cell is defined by its center position (θ_i,φ_i) in polar coordinates. Similarly, instead of being continuous, each great circle is uniformly sampled to consist of a finite number of points.

The voting procedure can consequently be visualized as an accumulation process. Given a point on the cylinder point cloud data, a great circle with a finite number of samples is generated that is perpendicular to the normal of that point. The discrete great circle passes through a finite number of cells on the discrete unit sphere. Each point on the great circle votes (ex. increments the accumulator) for the cell it is located in. This process is repeated for all the great circles, and the cell on the unit sphere with the highest accumulations is chosen as the estimate for cylinder orientation. The details for the first sequential step (Rabbani, 2005) are summarized as follows:

(1) Generate a discrete Hough space (i.e. discrete unit sphere).
(2) Compute the normal of a point in the point cloud data using k-nearest neighbors and plane fitting, and then generate a discrete great circle with orientation perpendicular to computed normal.
(3) Increment the accumulator of each cell that the discrete great circle passes through.
(4) Repeat steps 2-3 for each point on the cylinder point clouds, revealing a region of intersection of the great circles.
(5) Select the cell with the highest accumulator value as axis orientation (i.e. with greatest number of intersections)

Estimate the Position and Radius of the Cylinder

Following cylinder orientation estimation as in the first step, the next step is to estimate the position and radius of the cylinder (Rabbani, 2005). For the second sequential step, all the points are first projected to the plane perpendicular to the cylinder axis estimated in first step. Then the position and radius of the cylinder are calculated using circle fitting on the projected points. The details for this second sequential step are summarized below:

(1) Apply single value decomposition to project each point in 3D dataset onto a 2D plane perpendicular to estimated orientation.
(2) Form Cartesian voting grid $(u\ v)$ for range.
(3) For each projected point $(u_0\ v_0)$, generate sampled circle with radius r_i centered at $(u_0\ v_0)$.
(4) Sampled circle votes for cells it passes through.
(5) Repeat steps 3-4 for each projected point.
(6) Repeat steps 2-5 for a range of radius values.

ACCURATE AND FAST DETERMINATION OF ORIENTAION

In the original approach by Rabbani (2005), the unit sphere was sampled at a given number of points to generate the Hough space, in which voting subsequently occurs. The number of cells on the unit sphere impacts the accuracy of the orientation estimate. A larger number of cells would yield a more accurate estimate of the cylinder orientation, but comes at the cost of speed in terms of computational runtime. A smaller number of cells, on the other hand, would result in a less accurate estimate, but would reduce the runtime. An approach that would yield a better orientation estimate while reducing runtimes is the coarse-to-fine procedure.

Coarse-to-fine Approach

The idea of the coarse-to-fine approach is as follows. Initially, the whole unit sphere is coarsely sampled and an initial estimate of orientation made. The initial estimate gives an indication of the orientation, but due to the coarse sampling of the unit sphere, does not yield an accurate estimate. But the information of the initial estimate can be used to restrict the search region of the unit sphere while increasing the sampling density of the unit sphere in that restricted region. The second iteration uses the restricted region with the denser sampling to generate a more accurate estimate. Subsequent iterations continue to restrict the search region while increasing the sampling density, improving the accuracy of the orientation estimate until the stop criterion is reached.

For performance assessment, the coarse-to-fine approach was implemented and illustrated as follows. Initially, 98 cells are used covering the whole unit sphere (90°) to generate the initial orientation estimation for a cylinder point cloud containing 588 points. Once the estimate is made, subsequent iterations (Figure 2) reduce the search region by 32° each time around the orientation estimate of the previous iteration. During each iteration, the number of cells in the restricted search region is continually increased for denser sampling of the unit sphere where the cylinder is likely to be orientated.

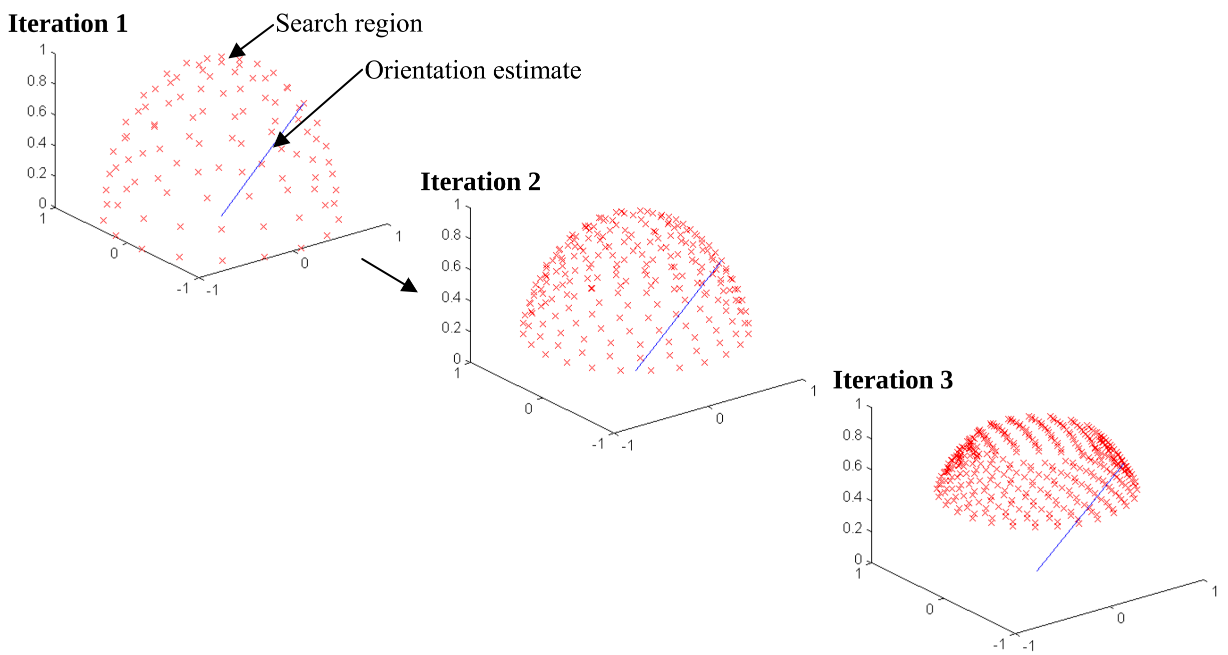

Figure 2. Illustration of coarse-to-fine approach showing initial orientation estimate using a sparsely populated unit sphere. Each subsequent iteration focuses on the previous estimate by restricting the search region as well as increasing the number of cells in the restricted region, improving accuracy as well as decreasing the overall runtime.

The number of cells as well as the estimated orientation for each iteration are tabulated in Table 1. The actual orientation of the cylinder is (0.7687,-0.1848,0.6124). Note that the orientation estimate continually improves (in terms of mean absolute error) from the initial iteration to the third iteration. The computational runtime for this course-to-fine approach is 3.2 seconds. Using the original approach with 1366 cells covering the whole unit sphere, the orientation estimate is (0.7799,-0.1922,0.5957). The corresponding mean absolute error and runtime are 0.1879 and 12.1 seconds, respectively. Therefore the coarse-to-fine approach yielded an improvement in orientation estimate accuracy of 20% while reducing runtime by 74%.

Table 1. Results of coarse-to-fine iterative procedure

	Number of Cells	Orientation Estimate (x,y,z)	Mean Absolute Error
Iteration 1	98	(0.6124,-0.3536,0.7071)	0.6479
Iteration 2	223	(0.7228,-0.2481,0.6450)	0.3767
Iteration 3	398	(0.7784,-0.1845,0.6001)	0.1494

ROBUST ESTIMATION OF ORIENTATIONS FOR MULTIPLE CYLINDERS

The sequential Hough transform approach described above estimates the cylinder orientation in the case of a single cylinder. If multiple cylinders are present in the point cloud data or region of interest within the data, the Hough space would exhibit multiple peaks in terms of number of accumulations. Simply choosing the cells with the maximum number of accumulations would not yield accurate estimates of multiple orientations. In the case of multiple cylinders, robust estimation of multiple orientations can be achieved using clustering, as detailed in the next section.

Clustering

Clustering is useful for detecting multiple peaks in Hough space, given an input dataset containing multiple cylinders. For example for a dataset containing two cylinders (Figure 3), if the two highest values of the accumulator are simply chosen, these may not correspond to the two separate cylinders, due to noise and differences in the size of the cylinders. Clustering, on the other hand, is a more robust method that can be used to separate the two clusters that correspond to the two cylinders.

The basic idea for getting a uniform sampling of the Hough Gaussian sphere is to sample Φ uniformly and change the sampling density along θ adaptively. Because of uniform sampling, there are points almost everywhere on the Hough sphere. Clustering cannot be performed in the situation where points are uniformly distributed on the Gaussian sphere. But given the availability of the number of accumulations at each point, certain points can be discarded prior to clustering.

Since there are a lot of points with small accumulations while very few of them have large accumilations, points with small accumulations could be discarded to facilitate clustering. First, we treat the accumulator as a random variable, using its histogram to estimate its cumulative distribution function. Let a be the value the accumulator can realize, and A be the accumulator as a random variable. The objective is to determine a threshold a_0 so that $P(A<a_0)=p_{thr}$, where p_{thr} is user defined (typically 0.95 works well for this application). All points in the accumulator with less than a_0 accumulations are discarded.

Then the remaining points are used to compute spherical distance from each point to every other point. Given points u and v, the spherical distance is defined as follows:

$$d_{sph} = \frac{\cos^{-1}(u \cdot v) \times 2\pi}{360}$$

Following histogram-based thresholding of the accumulator cells, the next step is to create a hierarchical tree from computed distances followed by clustering to group the remaining data points into clusters. After finding the separate clusters, the cell within each cluster that has the highest number of accumulations is chosen as an estimate for cylinder orientation. Figure 4 illustrated the detailed clustering steps.

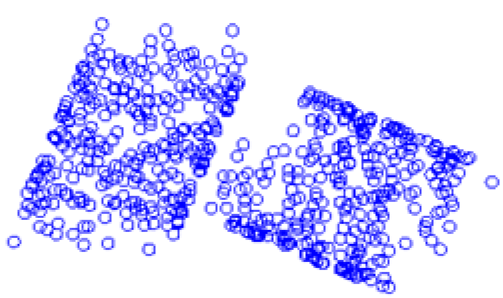

Figure 3. Example two cylinder point clouds dataset.

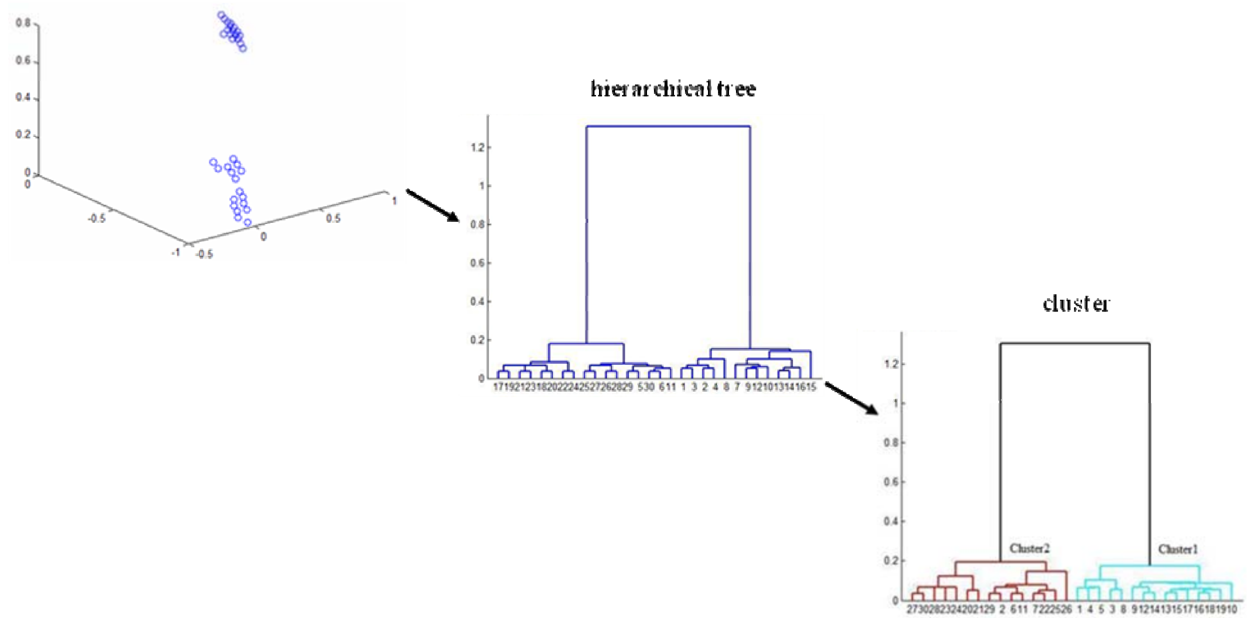

Figure 4. The steps involved in the clustering procedure are illustrated above. First, cells with few accumulations are discarded using a histogram-based approach. Spherical distance of each cell to every other cell is computed, and the results are used to produce a hierarchical tree and subsequently clustered to represent the different cylinders in the point cloud data.

The results of the clustering method as applied to point cloud data with two cylinders are tabulated below in Table 2. As can be observed, the clustering method can be used to robustly estimate the orientation of multiple datasets within the cylinder.

Table 2. Results of robust detection method for two cylinders in point cloud data

	x orientation	y orientation	z orientation
Actual Orientation 1	0.7687	-0.1848	0.6124
Actual Orientation 2	0.0875	0.9706	-0.2241
Estimated Orientation 1	0.7697	-0.1923	0.6088
Estimated Orientation 2	0.0975	0.9690	-0.2271

CONCLUSION

This work extends upon the sequential Hough transform procedure of Rabbani (2005) for automatic cylinder extraction in point cloud data, decreasing runtimes as well as improving the robustness of orientation estimation. The implemented course-to-fine approach not only reduces runtime, but improves the accuracy of the orientation estimates; the clustering approach furthermore improves the robust of cylinder feature extraction in the presence of multiple cylinders within dense point cloud data.

REFERENCES

R. C. a. F. Bolles, M. A., "A ransac-based approach to model fitting and its application to finding cylinders in range data." pp. 637-643.

M. P. D. Carmo, Differential Geometry of Curves and Surfaces: Prentice-Hall, 1976.

M. A. a. B. Fischler, R. C., “Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography,” Communications of ACM vol. 24, no. 6, pp. 381, 1981.

M. A. a. B. Fischler, R. C., "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography." pp. 726-740.

P. V. C. Hough, Method and means for recognizing complex patterns, U. S. P. 3069654, 1962.

C. Kimme, Ballard, D. and Sklansky, J., “Finding circles by an array of accumulators,” Communications of ACM vol. 18, no. 2, pp. 120-122, 1975.

G. Luk'acs, Martin, R. and Marshall, D., "Faithful least-squares fitting of spheres, cylinders, cones and tori for reliable segmentation." pp. 671-686.

A. D. Marshall, Luk'acs, G. and Martin, R. R., “Robust segmentation of primitives from range data in the presence of geometric degeneracy,” PAMI, vol. 23, no. 3, pp. 304-314, 2001.

S. Petitjean, “A survey of methods for recovering quadrics in triangle meshes,” ACM Computing Surveys, vol. 34, no. 2, pp. 211-262, 2002.

T. a. H. Rabbani, F. A. v. d., "3d industrial reconstruction by fitting csg models to a combination of images and point clouds." pp. 7-12.

T. a. H. Rabbani, F. A. v. d., "Methods for fitting csg models to point clouds and their comparison.." pp. 279-284.

L. s. survey. "Laser Scanner Survey," www.pobonline.com/POB/FILES/HTML/PDF/POB0105-LaserScanningSurvey.pdf <http://www.pobonline.com/POB/FILES/HTML/PDF/POB0105-LaserScanningSurvey.pdf>.