EFFICIENT HOUGH TRANSFORM FOR AUTOMATIC DETECTION OF CYLINDERS IN POINT CLOUDS

Tahir Rabbani and Frank van den Heuvel
Section of Photogrammetry and Remote Sensing,
Faculty of Aerospace Engineering,
TU Delft, The Netherlands
t.rabbani@lr.tudelft.nl, F.A.vandenHeuvel@lr.tudelft.nl

Working Group V/I

KEY WORDS: Cylinder detection, Object recognition, Hough Transform, Reverse Engineering, Industrial Reconstruction, Automatic 3D processing.

ABSTRACT

We present an efficient Hough transform for automatic detection of cylinders in point clouds. As cylinders are one of the most frequently used primitives for industrial design, automatic and robust methods for their detection and fitting are essential for reverse engineering from point clouds. The current methods employ automatic segmentation followed by geometric fitting, which requires a lot of manual interaction during modelling. Although Hough transform can be used for automatic detection of cylinders, the required 5D Hough space has a prohibitively high time and space complexity for most practical applications. We address this problem in this paper and present a sequential Hough transform for automatic detection of cylinders in point clouds. Our algorithm consists of two sequential steps of low dimensional Hough transforms. The first step, called *Orientation Estimation*, uses the Gaussian sphere of the input data and performs a 2D Hough Transform for finding strong hypotheses for the direction of cylinder axis. The second step of *Position and Radius Estimation*, consists of a 3D Hough transform for estimating cylinder position and radius. This sequential breakdown reduces the space and time complexity while retaining the advantages of robustness against outliers and multiple instances. The results of applying this algorithm to real data sets from two industrial sites are presented that demonstrate the effectiveness of this procedure for automatic cylinder detection.

1 INTRODUCTION

Recent advances in 3D scanning technologies have made possible high-speed acquisition of dense and accurate point clouds at moderate costs (Laser scanner survey, 2005). The explicit geometric information available from point clouds can be used to automate the 3D reconstruction process, which has been largely manual till now. This is especially true for the reconstruction of industrial sites as due to their man-made origin presence of well-defined CAD primitives can be expected. As-built modelling of industrial sites is required for documentation, and for various emerging technologies that use Virtual and Augmented reality for training and other services (STAR, 2004). A high degree of automation in 3D reconstruction should benefit all these application areas.

Cylinders are one of the most important geometric primitives found on industrial sites. As reported by Nourse et al. 85% of objects found in most industrial scenes can be approximated by planes, spheres, cones and cylinders (Nourse et al., 1980, Petitjean, 2002). This percentage rises to 95% if toroidal surfaces are also included in the set of available primitives (Requicha and Voelcker, 1982). Various methods have been proposed in the literature to fit cylinders to point clouds (Lukács et al., 1998, Marshall et al., 2001, Chaperon and Goulette, 2001).

These methods can be divided into two categories: those requiring a prior segmentation and those processing raw point clouds without segmentation. The methods belonging to the first category fit a cylindrical surface to the segmented point cloud. Most of them use non-linear least squares estimation to minimize the orthogonal distance of the points from the surface of cylinder (Lukács et al., 1998, Marshall et al., 2001). These methods assume that the segmentation algorithm is able to assign correct labels and there are only a few outliers or segmentation errors. As is shown in the comparison of segmentation algorithm for planar surfaces in (Hoover et al., 1996) and for curved surfaces in (Min et al., 2000) these requirements are not met in most of the real-life cases. The sensitivity of least squares based geometric fitting to outliers is well known (Björck, 1996, Press et al., 1988). Furthermore, non-linear least squares is an iterative process, and to avoid local minima it requires good initial values of parameters being estimated. In case of over-segmentation estimated initial values are poor and the algorithm can get trapped in one of the local minima. On the other hand, under-segmentation results in high percentage of outliers, resulting in an unfaithful reconstruction.

The methods belonging to the second category try to avoid these problems by processing raw point clouds using robust fitting methods like RANSAC (Fischler and Bolles, 1987, Chaperon and Goulette, 2001, Bolles and Fischler, 1981). For example in (Chaperon and Goulette, 2001) RANSAC is used on Gaussian sphere to find direction of cylinder axis which is then used as initial value for the least squares fitting. This method assumes a small number of cylinders in the scene. The presence of multiple cylinders of different radii along one orientation can lead to failure. For many real life industrial scenes, as the last section will show, these limitations are too restrictive.

Hough based methods have long been used to tackle prob-

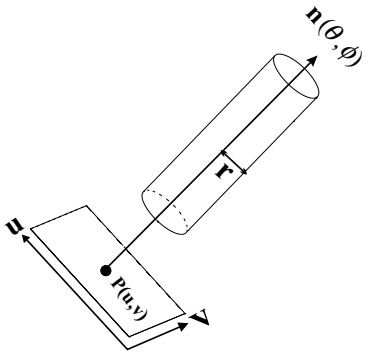

Figure 1: The five parameters for the cylinder. (θ, ϕ) gives the axis direction in spherical coordinates, r is the radius, $P(u, v)$ gives the position in terms of u and v which along with axial direction $n = (\cos\theta \sin\phi, \quad \sin\theta \sin\phi, \quad \cos\phi)$ form the cylinder's local coordinate system.

lems of outliers and multiple instances (Hough, 1962). In noisy and cluttered images they have no parallel in finding lines and curves like circles (Kimme et al., 1975). For range data analysis they have been used to re-construct buildings by finding planar surfaces in $2\frac{1}{2}$D data obtained from an air-borne laser scanner (Vosselman and Dijkman, 2001). Similarly they have been successfully used for characterization of planar fractures from 3D data (Sarti and Tubaro, 2002).

A major drawback of the Hough transform is its time and space complexity. For geometric fitting problems in 3D the space and time complexity can be approximated by $O(s^p)$ and $O(s^{p-1}n)$ respectively, where n is the number of points, p is the number of parameters and s is the number of samples along one Hough dimension. Keeping in view the fact that in most modelling projects employing laser scanning the number of points can be in the order of millions, Hough Transform becomes impractical for fitting of objects having more than three parameters.

Although a cylinder has five degrees of freedom, different parameterizations to represent it have been proposed in the literature. For example seven parameters with two constraints are used in (Lukács et al., 1998). The parameterizations we propose is shown in Figure 1 and is best suited for the Hough transform as it uses a minimum number of free parameters with no constraints.

The five parameters for cylinder make direct use of Hough transform impractical. One effective way to reduce the space and time complexity of Hough transform is to use sequential processing and break the problem into a set of sub-problems of low complexity. That is the approach we employ here and divide the problem of cylinder fitting into two sequential steps. The first step uses Gaussian sphere of the point cloud as its input and consists of a 2D Hough transform to find strong hypothesis for cylinder axes. In step 2 a 3D Hough transform is performed for a few neighboring directions found in step 1, resulting in estimation of radius and position. Thus the sequential processing al-

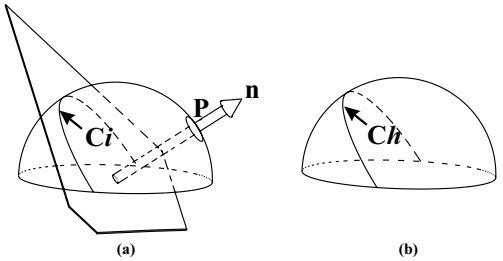

Figure 2: How a point $\mathbf{P}$ in input Gaussian sphere votes for a circle in the Hough space (a) Point $\mathbf{P}$ on Input Gaussian sphere, votes for a great circle on Hough Gaussian Sphere in (b). This great circle $\mathbf{C}$ results from the intersection of Gaussian Sphere with a plane whose normal $\mathbf{n}$ equals $\mathbf{P}$.

lows us to reduce the effective dimension of Hough space to three.

The rest of the paper is organized as follows. The first step, of orientation estimation, is discussed in Section 2. Section 3 contains details about the second step o position and radius estimation. In Section 4 we present the results of applying the proposed method to two data sets captured from real industrial sites. Section 5 contains some conclusions and directions for future work.

2 ORIENTATION ESTIMATION

The first step in the presented sequential Hough transform tries to find strong hypotheses for cylinder orientation. This orientation estimation is based on the observation that for cylinders the normals form a great circle on the Gaussian sphere (Carmo, 1976). This great circle results from the intersection of the unit sphere with a plane passing through the origin. The normal vector of this plane is the same as the cylinder axis. Although, in principle any plane fitting method can be used for detection of this great circle, we use Hough transform to avoid the above-mentioned problems of outliers and multiple-instances.

The standard Hough transform to find planes requires a three dimensional Hough space (Sarti and Tubaro, 2002). There are two parameters corresponding to the direction of the plane normal expressed in spherical coordinates and one for the distance of the plane from the origin. In the current case we have two constraints that will enable us to reduce the dimension of the Hough space from three to two. Furthermore, we will exploit these constraints to formulate a rapid update method for Hough space. Firstly, the plane must pass through the origin, which means that we can take out the third parameter corresponding to the distance from the origin. This leaves us with a 2D Hough space. Secondly, the plane must intersect the unit sphere, meaning each input point votes for one circular region in the Hough space (Figure 2).

Since the Hough transform for this step uses the Gaussian sphere of the given point cloud as its input, we need to estimate the normal for each point. As we are process-

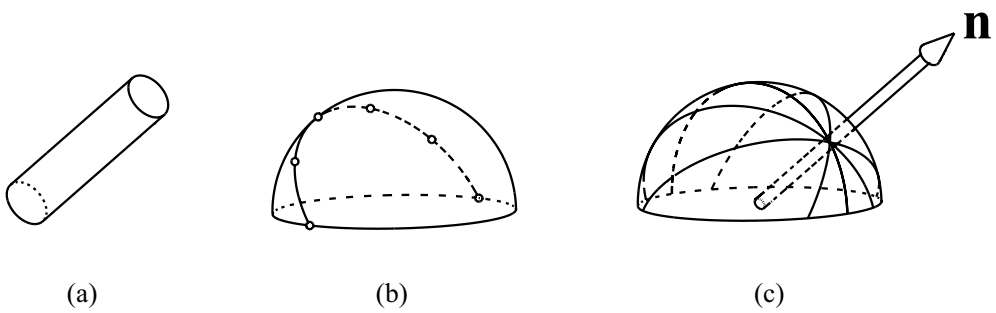

Figure 3: Step 1: *Orientation Estimation.* (a) Input cylinder. (b) Input Gaussian sphere with the great circle corresponding to the input cylinder. (c) The Hough Gaussian Sphere with each point in input Gaussian sphere voting for a great circle. The intersection of all circles gives an estimate for cylinder orientation

ing a raw point cloud without triangulation, the normal estimation methods specific to triangulated data (Petitjean, 2002) cannot be used. Instead, we use the method given by Hoppe in (Hoppe et al., 1992) that is more suited to unstructured point clouds. This method searches for k nearest neighbors for each point, and then estimates the normal by eigen-analysis of its covariance matrix. For more details see (Hoppe et al., 1992).

As we have explained above, the constraints in our problem allowed us to remove the third parameter corresponding to the distance of the plane from the origin from the standard plane-fitting Hough transform. This leaves us with quite a unique situation, because on the input we have the Gaussian sphere of the point cloud, while the Hough space consists of only orientations of the plane normal, which can be interpreted as another Gaussian sphere. To distinguish between these two separate entities we have named the Gaussian sphere resulting from the normals of the point cloud as *Input Gaussian sphere*, while the one in Hough space is named *Hough Gaussian sphere* (Figure 3).

Each point in the input Gaussian sphere votes for all plane orientations on the Hough Gaussian sphere, which are orthogonal to the orientation represented by this point (Figure 2). The set of all points in $\mathbb{R}^3$orthogonal to a given direction forms a plane. But in this case we have another constraint namely the magnitude of each point in this orthogonal set must equal one. This gives us a curve that is the intersection of the plane with the Hough Gaussian sphere i.e., a great circle. Consequently, each point in the Input Gaussian sphere votes for a great circle on the Hough Gaussian sphere. The normal of this great circle equals the normal vector of the current point. This voting scheme for one point is shown in Figure 2, while Figure 3 shows how individual circles arising from points belonging to a cylinder intersect on the Hough Gaussian sphere giving an estimate for the cylinder orientation.

The intersection of these great circles gives a high value in Hough space, and estimates the direction of the cylinder axis. The rapid update method described above requires a parametric equation of the great circle with a given normal expressed in spherical coordinates. Given this parametric equation, each point from the Input Gaussian sphere can be directly mapped to its respective cells in the Hough space. The parametric equation for the great circle in the xy-plane with z-axis as its normal is given by:

$$x = \cos t \quad y = \sin t \quad z = 0 \quad 0 \leq t \leq 2\pi \tag{1}$$

We need to apply a rotation to each point given by Equation 1 to get points for a great circle with a given normal. This rotation matrix must rotate z-axis to the normal of the required circle. As all the points are in one plane, for the same purpose we can use the reflection matrix given by Householder reflection (Golub and Loan, 1991) :

$$z = (0\ 0\ 1)^T \tag{2}$$

$$n = (\cos\theta \sin\phi \quad \sin\theta \sin\phi \quad \cos\phi)^T \tag{3}$$

$$\mathbf{R} = \mathbf{I} - 2\mathbf{b}\mathbf{b}^{\mathbf{T}} \tag{4}$$

Where:

$$\mathbf{b} = \frac{\mathbf{z} - \mathbf{n}}{\|\mathbf{z} - \mathbf{n}\|} \tag{5}$$

Using these expressions for a point $\mathbf{P}(\theta, \phi)$ the resulting rotation matrix $\mathbf{R}$ can be derived.

In the Figure 2 and Figure 3 we see this voting by points on input Gaussian sphere to the cells in the Hough space.

The algorithm for step 1 of our procedure is as follows:

1. Calculate normals for all points in a given point cloud using plane fitting on their k nearest neighbors.
2. Make a sampled Hough space to represent Hough Gaussian Sphere for orientation of cylinder axis using approximate uniform sampling method (Lutton et al., 1994, Rusin, 2004).
3. For each point in the input data, use the spherical coordinates of its normal calculated in step 1 to derive matrix $\mathbf{R}$ using Equation 4.
4. Increment the cells in Hough space given by rotated parametric form of the circle (1 and Equation 4).

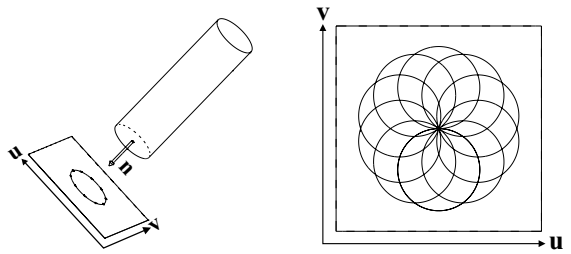

(a) (b)

Figure 4: Step 2. *Position and Radius Estimation* (a) Points projected along the estimated orientation (b) A 2D slice through the 3D Hough space for the correct radius and orientation.

5. Find the points in Hough Gaussian sphere whose accumulator values are greater than a threshold. These are hypotheses for cylinder directions.

6. Select the points who have voted for the highest cell, and proceed with the second step of position and radius estimation.

3 POSITION AND RADIUS ESTIMATION

As explained in Section 2 a cylinder can be represented by five parameters. Step 1 gives the strong hypotheses for direction of the cylinder. As a result we are left with three unknown parameters corresponding to cylinder position and radius, and thus a three-dimensional Hough space is needed for this step. For each orientation found in step 1 we take some of its neighbors on Gaussian sphere and perform Step 2 for each of them. This results in the refinement of the orientation estimate in addition to the estimation of the position and the radius of the cylinder.

We begin Step 2 by projecting all the points to the plane with a normal direction equal to the cylinder axis. For this purpose we need an orthonormal coordinate system, having cylinder orientation as one of its axes. There are two options to calculate such a set of orthonormal bases from a given vector, either we can use Gram-Schmidt Orthogonalization (Press et al., 1988) or alternately singular value decomposition can be employed (Golub and Loan, 1991). By using any of these techniques we get an orthonormal coordinate frame consisting of three bases vectors, $\begin{pmatrix}\mathbf{u} & \mathbf{v} & \mathbf{n}\end{pmatrix}$ where $\mathbf{n}$ equals the cylinder axis. This frame is used for projecting all points to one plane.

Once points have been projected we proceed to calculate the position and the radius of the cylinder using circle fitting. The Hough transform for circle fitting uses formulation given in (Kimme et al., 1975). For a given radius $\mathbf{r}$, each projected point votes for bins in a circular region in the Hough space with the current point as center. If the projected coordinates of a point are given by $(\mathbf{u_p}, \mathbf{v_p})$, it votes for the cells in the Hough space given by:

$$(\mathbf{r}\cos\omega + \mathbf{u_p},\ \mathbf{r}\sin\omega + \mathbf{v_p}) \tag{6}$$

The peak in the Hough space gives the refined orientation and the radius of the cylinder directly. However, the position is still in the projection coordinate system, calculated above, and must be transformed back to the world coordinate system to get a 3D point on the axis of cylinder. This coordinate transformation is given by the following matrix $\mathbf{T}$:

$$\mathbf{T} = \begin{pmatrix} u_x & u_y & u_z \\ v_x & v_y & v_z \\ n_x & n_y & n_z \end{pmatrix} \tag{7}$$

The algorithm for step 2 of our procedure can be summarized as follows:

1. For each orientation found in Step 1, find its N nearest neighbors. For this neighbor search use approximate uniform sampling (Rusin, 2004).

2. For each orientation $\mathbf{n}$ derive an orthonormal coordinate system consisting of three basis vectors given by $\begin{pmatrix}\mathbf{u} & \mathbf{v} & \mathbf{n}\end{pmatrix}$ Project all points along $\mathbf{n}$.

3. For each value of the radius r in a user-specified radius range, increment the Hough cells given by Equation 6 for the projection coordinates of each projected point $(\mathbf{u_p},\ \mathbf{v_p})$.

4. Find the peak in the Hough space. This gives cylinder orientation and radius directly. Transform the position to world coordinate system using Equation 7.

5. Remove the points corresponding to the found cylinder from the data

6. If any points are left proceed with step 1.

Figure 4(a) shows the results of projection, while Figure 4(b) shows a slice through three dimensional Hough space for the correct radius.

4 RESULTS

The method outlined above was applied to three point cloud data sets from industrial sites and the results are shown in Figure 5 and Figure 6. In Figure 5 we show the step-by-step results for a point cloud captured from an L-junction. In 5(a) we show the original point cloud, 5(b) shows the estimated normals as arrows. Note the ambiguity in normal estimation, that justifies the use of half Gaussian sphere. 5(c) shows the input Gaussian sphere with two great circles corresponding to the two cylinders in the the input data. 5(d) shows the Hough Gaussian sphere with two peaks detecting the two strong hypothesis for cylinder orientations. In 5(e) we illustrate the process of neighboring orientation selection on uniformly sampled Gaussian sphere, which is necessary for the refinement of the orientation estimation in the second step. 5(f) shows the Hough space for correct orientation and radius; the peak gives the position of the

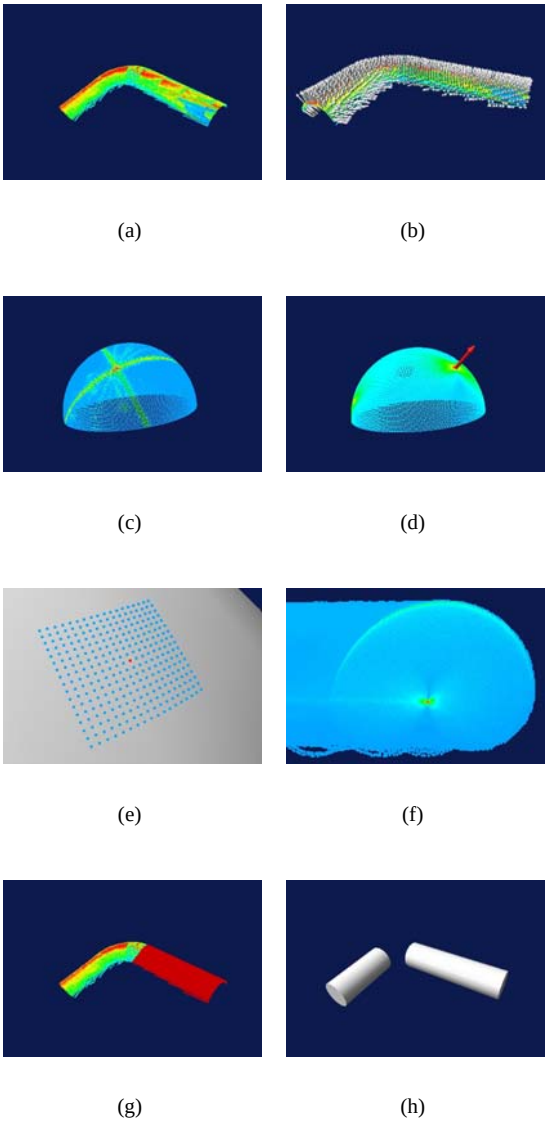

(a) (b) (c) (d) (e) (f) (g) (h)

Figure 5: Step by step processing of an L-junction point cloud (a) Input data (b) Estimated normals (c) Input Gaussian sphere (d) Hough Gaussian Sphere(e) Selection of neighboring orientations (f) Step 2 Position and radius estimation shown for correct orientation and radius. (g) Select points belonging to the cylinder (h) Final result, both cylinders are detected automatically. found.

cylinder. In 5(g) the points belonging to the detected cylinder are selected and removed from the input data based on a distance threshold. The process is repeated for the remaining points resulting in the detection of the second cylinder. 5(h) shows the final result, where both cylinders are successfully detected.

For results shown in Figure 5 the Hough space for Step 1 consisted of $250,000$ cells, whereas Step 2 used a 3D Hough space having $512 \times 512 \times 100 \cong 26 \times 10^6$ cells. In contrast a straight forward 5D Hough space for equivalent results would have required $512 \times 512 \times 100 \times 512 \times 512 \cong 7 \times 10^{12}$ cells in the accumulator, which is impractical, even for today's high computing power. It shows that the presented sequential Hough transform greatly reduces the space and time complexity and makes the problem of automatic cylinder detection manageable.

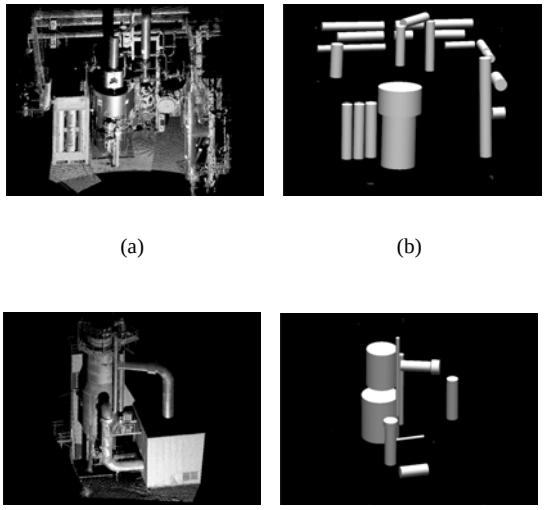

(a) (b) (c) (d)

Figure 6: Results of Hough transform. (a and c) Input point cloud (b and d) Results of cylinder detection

Figure 6 shows the results of our algorithm on two data sets from different industrial sites. During their processing it was found that during the step 1 of orientation estimation the large clusters on Input Gaussian sphere resulting from big planar areas interfered with cylinder axis estimation. To resolve this problem, the data sets were preprocessed with a plane-fitting Hough transform to remove planar areas. Figure 6(a) shows point cloud before processing Figure 6(b) shows the cylinders separately. The sequential Hough transform has been able to detect cylinders in different orientation and multiple radii.

Although the presented algorithm is sequential and errors in step 1 can be expected to affect step 2. To prevent this the second step is carried out for a few neighbors of estimated orientation of step 1. This results in refinement of the orientation estimate along with coupling of both steps.

Figure 6(c) and Figure 6(d) show another data set from an industrial plant along with found cylinders. As for this data set wider radius bounds were specified, cylinders of both small and big radii were successfully detected.

5 CONCLUSIONS

We have presented a sequential Hough transform to automatically detect cylinders in point clouds. A sequential approach enabled us to employ a combination of 2D and 3D Hough transforms instead of a originally required 5D Hough space. As the results showed this sequential breakdown still maintains the advantage of robustness against outliers and multiple instances that necessitated the use of Hough transform in the first place.

As cylinders are one the most commonly found objects, either as pipes or part of other complex objects as cylindrical

patches, their automatic detection can decrease the manual effort required for reverse engineering.

Although we haven't formulated it in this paper, the extension of the second step of *Position and Radius estimation* to elliptical cylinders is quite straightforward. In future we plan to investigate sequential Hough transform for detection of objects arising from planar sweeps of general curves.

REFERENCES

Björck, Å., 1996. Numerical Methods for Least Squares Problems. SIAM, Philadelphia.

Bolles, R. C. and Fischler, M. A., 1981. A ransac-based approach to model fitting and its application to finding cylinders in range data. In: Proc. of the 7th IJCAI, Vancouver, Canada, pp. 637–643.

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

Chaperon, T. and Goulette, F., 2001. Extracting cylinders in full 3d data using a random sampling method and the gaussian image. In: Proc. Vision, Modelling and Visualisation, University of Stuttgart, Germany, pp. 35–42.

Fischler, M. A. and Bolles, R. C., 1987. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. In: M. A. Fischler and O. Firschein (eds), Readings in Computer Vision: Issues, Problems, Principles, and Paradigms, Kaufmann, Los Altos, CA., pp. 726–740.

Golub, G. H. and Loan, C. F. V., 1991. Matrix Computation. John Hopkins University Press. Second Edition.

Hoover, A., Jean-Baptiste, G., Jiang, X., Flynn, P. J., Bunke, H., Goldgof, D. B., Bowyer, K. K., Eggert, D. W., Fitzgibbon, A. W. and Fisher, R. B., 1996. An experimental comparison of range image segmentation algorithms. IEEE Transactions on Pattern Analysis and Machine Intelligence 18(7), pp. 673–689.

Hoppe, H., DeRose, T., Duchamp, T., McDonald, J. and Stuetzle, W., 1992. Surface reconstruction from unorganized points. Computer Graphics 26(2), pp. 71–78.

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

Kimme, C., Ballard, D. and Sklansky, J., 1975. Finding circles by an array of accumulators. Commun. ACM 18(2), pp. 120–122.

Laser scanner survey, 2005. Laser Scanner Survey. http://www.pobonline.com/POB/FILES/HTML/PDF/0205surveyLaser.pdf.

Lukács, G., Martin, R. and Marshall, D., 1998. Faithful least-squares fitting of spheres, cylinders, cones and tori for reliable segmentation. In: ECCV '98: Proceedings of the 5th European Conference on Computer Vision-Volume I, Springer-Verlag, pp. 671–686.

Lutton, E., Maitre, H. and Lopez-Krahe, J., 1994. Contribution to the determination of vanishing points using hough transform. IEEE Trans. Pattern Anal. Mach. Intell. 16(4), pp. 430–438.

Marshall, A. D., Lukács, G. and Martin, R. R., 2001. Robust segmentation of primitives from range data in the presence of geometric degeneracy. IEEE Trans. Pattern Anal. Mach. Intell. 23(3), pp. 304–314.

Min, J., Powell, M. W. and Bowyer, K. W., 2000. Progress in automated evaluation of curved surface range image segmentation. In: ICPR, pp. 1644–1647.

Nourse, B. E., IIakala, D. G., Hillyard, R. C., and Malraison, P. J., 1980. Natural quadrics in mechanical design. In: Proceedings of Autofact West 1, Anaheim, CA., pp. 363–378.

Petitjean, S., 2002. A survey of methods for recovering quadrics in triangle meshes. ACM Comput. Surv. 34(2), pp. 211–262.

Press, W. H., Flannery, B. P., Teukolsky, S. A. and Vetterling, W. T., 1988. Numerical recipes in C: the art of scientific computing. Cambridge University Press.

Requicha, A. A. G. and Voelcker, H. B., 1982. Solid modeling: a historical summary and contemporary assessment. IEEE-CGA 2(2), pp. 9–24.

Rusin, D., 2004. Topics on Sphere Distribution. http://www.math.niu.edu/ rusin/known-math/95/sphere.faq.

Sarti, A. and Tubaro, S., 2002. Detection and characterisation of planar fractures using a 3d hough transform. Signal Process. 82(9), pp. 1269–1282.

STAR, 2004. STAR–Services and Training through Augmented Reality. http://www.realviz.com/STAR/objectives.htm.

Vosselman, G. and Dijkman, S., 2001. 3d building model reconstruction from point clouds and ground plans. In: International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. 34, Annapolis, MA, USA, pp. 37–44.