Simple Surface Segmentation

Dr Alan M. McIvor, David W. Penman and Peter T. Waltenberg

Industrial Research Limited
PO Box 2225, Auckland
New Zealand
E-mail: {a.mcivor,d.penman,p.waltenberg}@irl.cri.nz

Abstract

Methods are presented for the extraction of surface patches from simple geometric shapes from 3D range data. These provide useful cues for recognising the objects visible in the scene. The geometric shapes considered are planes, spheres, cylinders, and ruled surfaces. Test results based on both Monte Carlo simulation and real range images are presented.

Keywords: *3D, feature extraction, object recognition.*

1. Introduction

This paper is concerned with the recognition of simple curved surface patches from dense 3D range data, such as that provided by a structured light system. Patches from planes, spheres, cylinders, and ruled surfaces are considered. The approach can be summarised as follows. The first step is to estimate the local surface geometry (the principal quadric) at each visible surface point [1, 2, 3, 4]. Then points at which the signs of the Gaussian and mean curvatures are inconsistent with those of a particular surface type are rejected from further consideration. Each remaining point is mapped to a point in the parameter space of the surface type. By using an unsupervised Bayesian classification or region growing algorithm, the clusters in parameter space that correspond to surface patches are identified, and the parameters of that surface can be determined.

Many methods exist for segmenting range data into patches based on approximating the data by low-order polynomials [5, 6, 7], deformable surfaces [8], etc. The work presented here differs in two ways. Firstly, the approach is based on using principal quadric estimates as a fundamental tool. Secondly, the aim is to identify surface patches of known form, for use as features in object recognition, instead of a complete surface segmentation. Further details on the segmentation methods are given below. This is followed by the results from an experimental evaluation of their performance.

2. Planar Surfaces

A planar surface can be characterised as a connected set of 3D surface points at which the two principal curvatures (alternatively, the Gaussian and mean curvatures) are zero. Thus, planar regions can be extracted by thresholding the principal curvatures (i.e., rejecting all points at which one or both principal curvatures have a magnitude above some threshold), and dividing the remaining points into connected regions, using a standard (binary) region labelling algorithm. The performance of this method is illustrated in Figure 1.

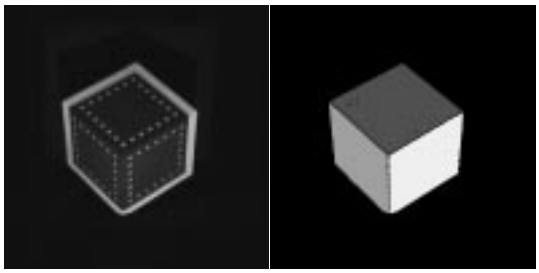

Figure 1: The image on the right shows the connected planar regions identified in 3D data taken from the scene on the left.

3. Spherical Surfaces

A spherical surface can be characterised as a connected set of 3D surface points at which the two principal curvatures have the same value, and this value is the same at all points (i.e., the principal curvature is spatially constant). They represent a more complex class of surfaces than planar regions in that the value of the principal curvatures is unknown. To extract such regions, all points at which the principal curvatures are not the same are rejected. The remaining points lie on a sphere of radius

$$r = \frac{1}{|H|} \tag{1}$$

where H is the mean curvature, with centre at

$$\mathbf{c} = \mathbf{x} + \frac{1}{H}\mathbf{n} \tag{2}$$

where $\mathbf{x}$ is the point location and $\mathbf{n}$ the surface normal estimate there. Thus each point can be mapped into the 4D parameter space for spheres (with parameters r and $\mathbf{c}$). Region growing or unsupervised Bayesian classification can be used to detect clusters in this parameter space. Each of the clusters corresponds to a set of connected spherical surface patches from the same sphere.

4. Cylindrical Surfaces

A right cylindrical surface can be characterised as a connected set of parabolic 3D surface points (i.e., points at which one of the principal curvatures is 0). Furthermore, the associated principal directions must all be aligned. The non-zero principal curvature must also be spatially constant. Given a parabolic point, the parameters of the cylinder it lies on (if indeed it lies on one) can be calculated as follows:

1. Radius estimate r

$$r = \frac{1}{2|H|} \tag{3}$$

2. Axis direction estimate $\mathbf{d}$

$$\mathbf{d} = \mathbf{r} \tag{4}$$

   where $\mathbf{r}$ is the principal direction in which the principal curvature is 0.

3. Unique point on axis $\mathbf{p}$

$$\mathbf{p} = (\mathbf{I} - \mathbf{d}\mathbf{d}^T)(\mathbf{x} + \frac{1}{2H}\mathbf{n}) \tag{5}$$

   where $\mathbf{n}$ is the surface normal. $\mathbf{p}$ is the point on the axis closest to the origin.

Thus, a cylinder has a 5D parameterisation. A region growing or unsupervised Bayesian classification based segmentation method is used to detect clusters in this parameter space. Each of the clusters corresponds to a distinct cylindrical surface patch. The performance of these methods is illustrated in Figure 2.

5. Ruled Surfaces

Ruled surfaces are a common type of manufactured surface and are a natural progression, in terms of complexity, from those considered in the previous sections. At all points on a ruled surface, there is a direction in which a line can be drawn in 3-space through the point such that the line is always in the surface. Many surfaces folded out of sheet material without stretching it are ruled surfaces.

5.1 Definitions and Properties

A ruled surface is a surface parameterised by

$$\mathbf{x}(u, v) = \alpha(u) + v\omega(u) \tag{6}$$

where $|\omega| = 1$ and α, ω are smooth functions of u [9, pp 188–197]. The unit vector $\omega(u)$

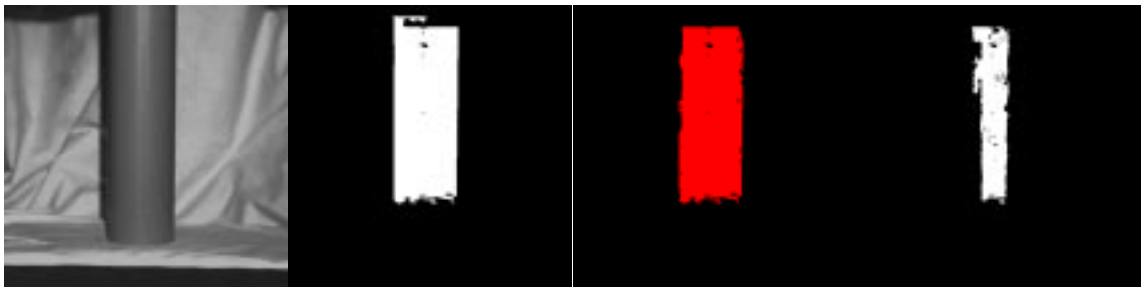

Figure 2: Example of cylindrical region extraction. From left to right: An image of a scene containing a piece of pipe; A mask indicating the largest connected region of parabolic points; The largest cylindrical region extracted by region growing; The largest cylindrical region extracted by Bayesian classification.

is the direction of the *ruling* at a point parameterised by (u, v). The curve α is called a *directrix* of the ruled surface.

If $\omega'(u) = 0, \forall u$ then $\mathbf{x}$ is a *cylindrical* ruled surface. The rulings at all points on such a surface are all in the same direction.

5.2 Cylindrical Ruled Surfaces

It is possible to devise a method for segmenting a ruled surface by identifying points that share a common ruling, and then identifying a directrix. But such an approach fails because the finite sampling means that each data point probably sits on a unique ruling. Thus, it is necessary to introduce some assumptions about the local behaviour of α and ω, which results in algorithms that differ little from a general deformable surface approach.

As noted above, a cylindrical ruled surface is a special ruled surface in which all points have a ruling in the same direction. A right cylinder is a special case of a cylindrical ruled surface where the directrix is a planar circle. Cylindrical ruled surfaces can be extracted by the following approach:

1. Reject points which have positive Gaussian curvature.
2. Estimate the ruling direction ω at each remaining data point.
3. Cluster constant ω regions.
4. For a constant ω region, choose a plane perpendicular to ω.
5. Project the data points onto this plane.
6. Thin the data points to identify a directrix.

The normal curvature in the ruling direction ω is 0. Hence points on a ruled surface must have a non-positive Gaussian curvature. If directions in the tangent plane are parameterised by the angle θ from the maximum principal direction to the given direction, then the normal curvature is

$$\kappa_n(\theta) = \kappa_1 \cos^2\theta + \kappa_2 \sin^2\theta \quad (7)$$

$$= H + \sqrt{H^2 - K}\cos 2\theta \quad (8)$$

where κ_1, κ_2 are the principal curvatures, and K is the Gaussian curvature. Hence, the direction of the ruled surface, i.e., where the normal curvature is zero, is given by

$$\cos 2\theta = \frac{-H}{\sqrt{H^2 - K}} \quad (9)$$

This can be converted into a vector in the data coordinate frame by

$$\omega = \mathbf{A}^T \begin{bmatrix} \cos\theta \\ \sin\theta \\ 0 \end{bmatrix} \quad (10)$$

where $\mathbf{A}$ is the attitude matrix.

As before, the segmentation of constant ω regions can be done using region growing or unsupervised Bayesian classification techniques.

6. Experimental Results

The algorithms presented in the previous sections were evaluated in two ways. The first test was the application of an algorithm to synthetic 3D data generated from models of a surface of the given type and a structured light system. The structured light system model incorporates system noise, and the performance of the algorithm was measured as a function of the noise variance. This can be used to evaluate the sensitivity of a segmentation algorithm to system noise. The size of the extracted regions for the various types of surface are presented in Figure 3. Data on the parameters of the extracted regions is presented in [10]. As the noise standard deviation increases, the performance drops as expected. The methods all fail before the noise standard deviation reaches 0.1. This is due to the principal quadric estimator failing at this noise standard deviation [1]. The large region sizes at high noise levels are artifacts due to the validation gate growing unacceptably large. Thus, for these surface extraction methods to be useful, the noise level of a 3D scanner must be less than 0.1.

The algorithms were also evaluated by forming two segmenters from them, and applying these to a standard set of curved surface images [11]. One segmenter used the planar surface extraction algorithm, along with region growing methods for spherical and cylindrical surfaces. The other used unsupervised Bayesian classification for the spherical and cylindrical surfaces along with the same planar surface extraction algorithm. This experiment measures not only the ability of the algorithms to correctly segment regions of the given type but also their ability to reject regions of different type. Results are presented in Table 1. The poor performance in comparison to the other two methods is due to a number of reasons:

1. The test data set contains very noisy data, generated by a K2T GRF-2 scanner which used no substripe interpolation (see [11, Figure 12, page 30]). This was made worse by the compression applied to the data. This resulted in poor principal quadric estimates.

2. Torodial surfaces in the data set are classified as "missed" since the segmenters don't contain this type. Corresponding extracted regions become "noise".

3. Segmentation boundary detection is poor. This is because our approach is surface-based not edge based, and dropping data close to boundaries does not affect the intended use (feature detection). This results in small regions being missed and causes disproportionate oversegmentation (where a single surface in the true segmentation has been reported as multiple surfaces by the segmenter, but most of the true surface is contained in just 1 reported segment).

4. The simple cluster number selection process used in the unsupervised Bayesian classification rejects outliers as wanted but tends to fragment large clusters into small ones (i.e., over-segment the data).

The algorithms need either a region growing or an unsupervised Bayesian classification sub-algorithm. There are two aspects to a region growing approach: selecting neighbour of the current region to consider, and deciding whether to add a neighbour. The neighbour selection was based on the "Basic Connected Graph Traversal Algorithm" [12, pg 102], which is an 8-connected connected set extraction algorithm. Neighbours were only added if their surface parameter measurements fitted inside a validation gate centred on the current region's surface parameter estimates. The validation gate size was selected on the basis of the variance of the current region's surface parameter estimates and measurement error variances. The current region's surface parameter statistics were calculated using uniformly weighted estimators, written in a recursive manner. They were updated after each neighbour was added.

The unsupervised Bayesian classification approach also has two components: assigning data to clusters; and selecting the number of clusters. The "Basic Isodata Procedure" [13, pg 201] for assigning data to clusters, and this was repeated as the number of clusters was increased, until the sum-of-squared-error criterion didn't decrease significantly. Further details are given in [10].

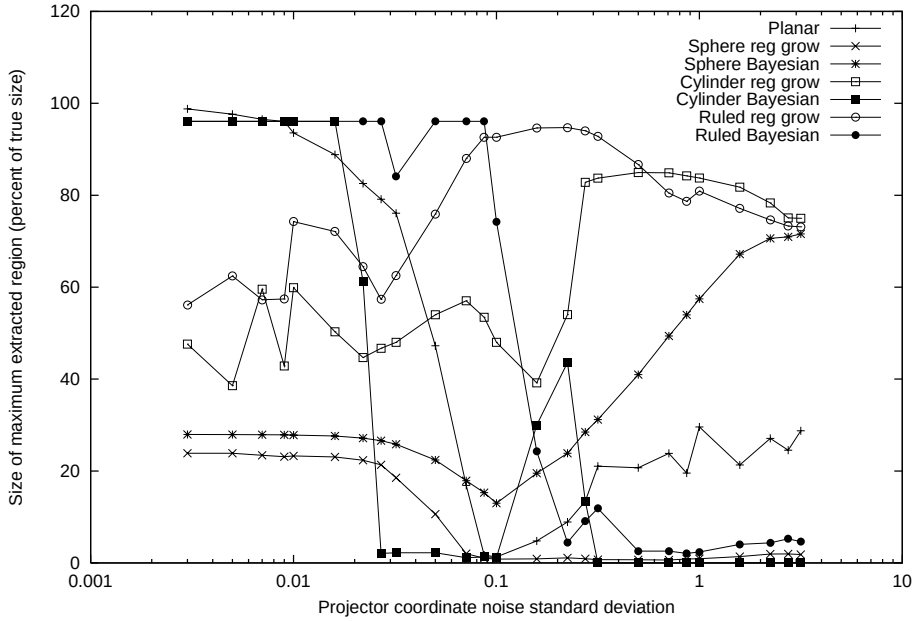

Figure 3: The size of the extracted regions, as a percentage of the true size, for the Monte-Carlo simulations of the methods.

7. Conclusion

This paper has presented methods for extracting from 3D range data surface patches which are from planar, spherical, (right) cylindrical, or cylindrical ruled surfaces. The methods are based on using the principal quadric estimated at each data point. Points which do not have the correct Gaussian and mean curvature signs are eliminated from further consideration. The principal quadric parameters at the other points are used to map them to the parameters of the surface (of a given type) that they lie on. Clusters are then detected in this parameter space to extract the surface patches.

The algorithms were tested using both region growing and unsupervised Bayesian learning for the cluster detection. First, a Monte Carlo simulation was used to evaluate the sensitivity to noise in the 3D measurements. This highlighted the need for accurate data and good principal quadric estimates, an issue that has already been addressed elsewhere [1]. Secondly, a generally available curved surface range data test set was used to compare the performance with other segmenters, even though the goal of this work is feature extraction not segmentation. It was observed from this that the unsupervised Bayesian classification method results in highly over-segmented surface patches. Correcting this deficiency will be a major focus of future work.

Acknowledgement

This work was funded by the New Zealand Foundation for Research, Science, and Technology under Contract CO8410 Objective 2 “3D Vision”.

Measure	Region Growing	Bayesian Class.	U. Bern	Bessel & Jain
Correctly Segmented (% no GT)	23.79	6.72	68.85	15.33
Over Segmented (% no GT)	23.60	40.34	2.29	62.94
Under Segmented (% no GT)	34.30	34.95	26.89	11.60
Missed Regions (% no GT)	18.31	17.99	1.97	10.13
Proportionate Over Seg. (% no GT)	1.72	2.46	0.22	22.65
Disproportionate Over Seg. (% no GT)	21.88	37.88	2.07	40.29
Correct or Disproportionate (% no GT)	45.67	44.60	70.92	55.62
Noise Regions (% no MS)	19.24	25.81	1.45	18.79

Table 1: Results for two segmenters, one constructed from the region growing methods and one constructed from the unsupervised Bayesian classification methods, applied to the USF curved range data test set, along with the results of the University of Bern and Bessel and Jain range image segmentation routines. The results are presented for a comparison threshold of 51% (the minimum overlap of a machine segmented (MS) region with a ground-truth (GT) region for the former to be counted as part of the segmentation of the latter).

References

[1] A. M. McIvor and R. J. Valkenburg, "A comparison of local surface geometry estimation methods," *Machine Vision and Applications*, vol. 10, no. 1, pp. 17–26, 1997.

[2] F. P. Ferrie, J. Lagarde, and P. Whaite, "Darboux frames, snakes, and super-quadrics: Geometry from the bottom up," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 15, pp. 771–784, Aug. 1993.

[3] P. J. Flynn and A. K. Jain, "On reliable curvature estimation," in *CVPR*, pp. 110–116, 1989.

[4] E. Trucco and R. B. Fisher, "Experiments in curvature-based segmentation of range data," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 17, pp. 177–182, Feb. 1995.

[5] K. L. Bowyer, M. J. Mirza, and G. Ganguly, "The robust sequential estimator: A general approach and its application to surface organization in range data," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 16, pp. 987–1001, Oct. 1994.

[6] N. S. Raja and A. K. Jain, "Obtaining generic parts from range images using a multiview representation," *CVGIP: Image Understanding*, vol. 60, pp. 44–64, July 1994.

[7] B. Parvin and G. Medioni, "B-rep object description from mulitple range views," *International Journal of Computer Vision*, vol. 20, pp. 81–112, Oct. 1996.

[8] D. Terzopoulos, J. Platt, A. Barr, and K. Fleischer, "Elastically deformable models," *Computer Graphics*, vol. 21, no. 4, pp. 205–214, 1987.

[9] M. P. do Carmo, *Differential Geometry of Curves and Surfaces*. Englewood Cliffs: Prentice-Hall, 1976.

[10] A. M. McIvor and P. T. Waltenberg, "Recognition of simple curved surfaces from 3D surface data," IRL Report 690, Industrial Research Limited, Mar. 1997. URL ftp://ftp.vision.irl.cri.nz/pub/doc/1997/irl-report-690.ps.

[11] M. Powell, "Comparing curved-surface range image segmenters," Master's thesis, Department of Computer Science and Engineering, University of Southern Florida, Apr. 1997.

[12] T. Pavlidis, *Algorithms for graphics and image processing*. Springer-Verlag, Berlin – Heidelberg: Computer Science Press, 1982.

[13] R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis*. New York: Wiley, 1973.