Recognition of simple curved surfaces from 3D surface data

Dr Alan M. McIvor and Peter T. Waltenberg

Industrial Research Limited
P. O. Box 2225, Auckland
New Zealand

Abstract. This paper is concerned with the robust identification of patches from simple geometric surfaces such as planes, spheres, and cylinders. This is done from 3D data acquired using a range scanner, along with principal quadric estimates at each point. Properties of the Gaussian and mean curvature (e.g., their sign) at points on surfaces of the target class are used to reject all points that could not be on such a surface. Each remaining point is then mapped to the parameters of the unique surface in its class that contains it, if it were on such a surface. Region growing and clustering techniques are used to identify those points actually on a surface of the target class and to extract the surface parameters.

1 Introduction

This paper is concerned with the segmentation of smooth, connected surface patches identified in range data into surfaces of simple shape classes (planes, spheres, cylinders), and the estimation of the parameters of the identified surface instances. It assumes that the principal quadric has been estimated at each surface point [1, 2], and this information is used to map each point to the parameters of the surface it lies on, if it is indeed of a given class. Clusters in this parameter space are identified, thereby segmenting the surfaces. The use of principal quadric data as part of the shape recovery process differentiates this method from others such as deformable surfaces [3], scan line grouping [4], tripod operators [5], and least-squares fitting [6, 7].

2 Planar Surface Patches

At all points on a planar surface, both principal curvatures (and hence the Gaussian and mean curvatures) are zero. Although the principal frame is not well-defined, the principal quadric is. Points at which both principal curvatures are zero are called *planar points*. A smooth surface patch containing only planar points is a planar surface. Thus planar surface patches can be extracted by detecting points at which both principal curvatures are zero and grouping them into connected regions.

The operation of such a planar surface detection scheme is illustrated in Figure 1, which shows the planar patches (corresponding to the cube's faces) extracted from 3D data off a cube.

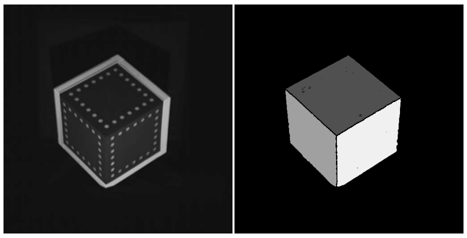

Fig. 1. The image on the right shows the connected planar regions identified in 3D data taken from the scene on the left.

The performance of the planar surface extraction scheme was evaluated using a Monte Carlo simulation [8]. Results are presented in Figure 2 as a function of the standard deviation of the noise added to the stripe data. For all values of standard deviation, the largest extracted planar region was used for calculating the statistics. The first graph shows the number of data points in the region. The other graph plots the RMS standard deviation of the surface normal perturbations [9, pg 281], where the RMS standard deviation is defined as the square root of the average of the perturbations' two eigenvalues. The bias in the estimator was measured by the angle between the true surface normal and the estimated one. In the worst case it was 12°, but in general it was insignificant.

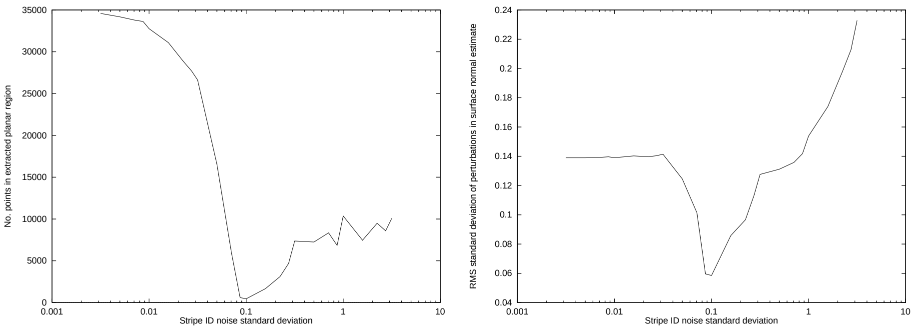

Fig. 2. Performance of planar region extraction method, as measured by the number of data points in the extracted planar region (left), and the RMS standard deviation of the normal perturbations.

The results in Figure 2 suggest very stable performance at low levels of stripe noise, with a rapid fall off in performance as the stripe id noise standard deviation increases from 0.01 to 0.1. There is some improvement beyond this.

3 Spherical Surface Patches

Points on a spherical surface patch are *umbilic*. Patches consisting of umbilic points, all having the same principal curvatures, are spherical surface patches. As such, it can be seen that planar surface patches are special cases of spherical surface patches. Spherical surface patches can be segmented by first identifying points which satisfy $|\kappa_1 - \kappa_2| < \epsilon$. These are then segmented into regions of spatially constant principal curvature, which correspond to spherical surface patches. This can be done with region growing or Bayesian classification techniques as described below.

3.1 Region Growing

The "Basic Connected Graph Traversal Algorithm" [10, pg 102] was used as a basis for a region growing method. Points on the boundary of the already extracted spherical region are added to it if they are "similar" enough to the underlying sphere, based on the current estimate of the sphere's parameters (radius and centre). Similarity is measured as follows, where the ith data point is being considered and there are n points in the already extracted spherical region. The parameters of the sphere on which the data point lies, given that it is an umbilic point, are estimated by

$$r_i = \frac{1}{|H_i|} \tag{1}$$

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

where H_i is the mean curvature at the point, $\mathbf{x}_i$ its location in space, $\mathbf{n}_i$ the surface normal there, r_i the radius of the sphere that the point lies on, and $\mathbf{c}_i$ the sphere's centre. The point is only accepted if it satisfies both

$$(r_i - \bar{r}_n)^2 \leq \alpha^2(\check{r}_n + \acute{r}) \tag{3}$$

$$\|\mathbf{c}_i - \bar{\mathbf{c}}_n\|^2 \leq \alpha^2(\check{\mathbf{c}}_n + \acute{\mathbf{c}}) \tag{4}$$

where $\bar{r}_n$, etc, are estimate means, $\check{r}_n$, etc, are estimate variances, $\acute{r}$, etc, are measurement error variances, and α is a chosen significance level. The statistics are simple, uniformly weighted, estimators which are updated recursively after each point is added. More details are given in [8].

A Monte Carlo simulation was used to measure the performance as a function of system noise. The number of points in the largest extracted region, and its estimated radius are shown in Figure 3. At small noise standard deviation levels, approximately a quarter of the underlying spherical region is extracted in the largest spherical patch (there were 45225 points total). Accurate radius estimates are also obtained; the systematic bias is due to the smoothing inherent in the principal quadric estimation. In the central part of the noise standard deviation range, around the value 0.1, the size of the extracted (connected) regions gets smaller but good estimates of the radius (and centre) are still generated. Above noise levels of 1.0, the performance is poor by all measures.

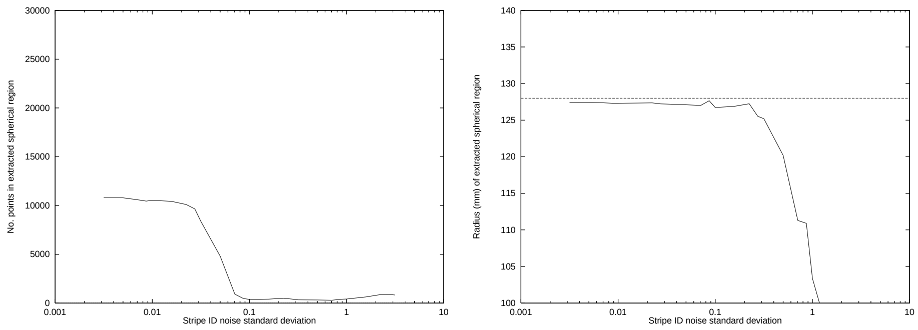

Fig. 3. Performance of the region growing based spherical region extraction method, as measured by the number of data points in the extracted spherical region (left), and the radius of the extracted spherical region (right). The true radius was 128mm.

3.2 Unsupervised Learning

The Basic Isodata Procedure [11, pg 201] was used to cluster the data points into regions of constant principal curvature.

The clustering algorithm requires the number of clusters, c, to be specified as input. This, however, is not known. It is determined by repeating the clustering for $c = 1, 2, 3, \ldots$, and looking at how the sum-of-squared-error criterion $J_e(c)$, used by the clustering algorithm to determine the optimal partitioning, varies with c. Termination tests as defined in [11, pg 242] are not appropriate here because the data often consists of one or more clusters against a background scattering of "noise" points. This leads to small decreases in $J_e(c)$ as the "noise" points are partitioned into small clusters of their own. Hence a suitable criterion is the smallest value of c which satisfies

$$\frac{J_e(c+1)}{J_e(c)} > 1 - \epsilon \tag{5}$$

where a good value for ϵ was determined to be 0.01.

A Monte Carlo simulation was used to measure the performance as a function of system noise. The number of points in the largest extracted region, and the estimated radius are shown in Figure 4. These results show that the Bayesian classification method is significantly more robust at high noise levels than the region growing based method.

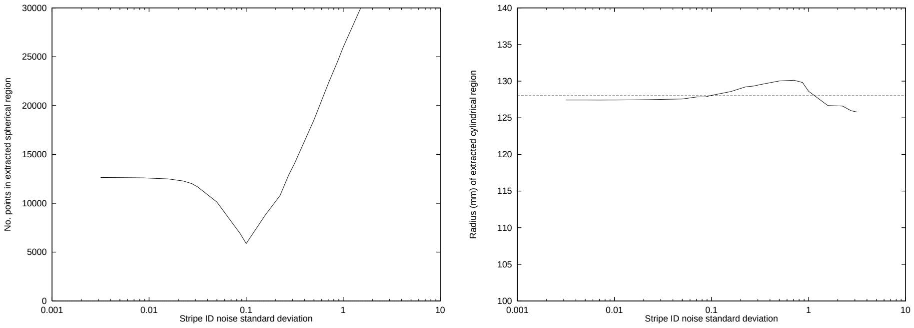

Fig. 4. Performance of the unsupervised learning based spherical region extraction method, as measured by the number of data points in the extracted spherical region (left), and the radius of the extracted spherical region (right). The true radius was 128mm.

4 Cylindrical Surface Patches

Given a point and the principal quadric there, the parameters of the cylinder it must lie on can be uniquely determined as follows (for the ith data point):

1. Radius estimate r_i

$$r_i = \frac{1}{2|H_i|} \tag{6}$$

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

$$\mathbf{d}_i = \pm\mathbf{r}_i \tag{7}$$

   where $\mathbf{r}_i$ is the principal direction in which the principal curvature is 0.
3. Unique point on axis $\mathbf{p}_i$

$$\mathbf{x}_a = \mathbf{x}_i + \frac{1}{2H_i}\mathbf{n}_i \tag{8}$$

$$\mathbf{p}_i = \mathbf{x}_0 + (\mathbf{I} - \bar{\mathbf{d}}_n\bar{\mathbf{d}}_n^T)\mathbf{x}_a \tag{9}$$

   where $\mathbf{n}_i$ is the surface normal, $\bar{\mathbf{d}}_n$ is the current estimate of the axis direction, and $\mathbf{x}_0$ is an arbitrary fixed point.

All points on a cylindrical surface are parabolic points (i.e., the Gaussian curvature is zero) with non-zero mean curvature. That is, one but not both principal curvatures is zero. This is used to segment out potentially cylindrical regions for further segmentation into cylindrical regions. As with spherical surface segmentation, region growing and unsupervised learning approaches were implemented and their performance measured using Monte Carlo simulation [8]. Results for the region growing method are shown in Figure 5. The underlying cylindrical region contained 123392 points, so the extracted regions contain approximately

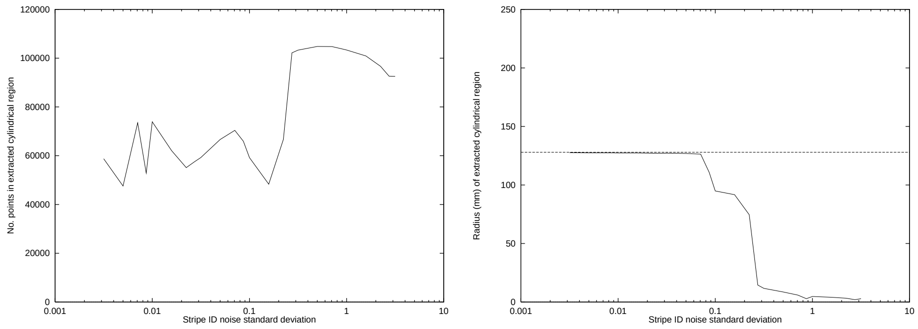

Fig. 5. Performance of the region growing based cylindrical region extraction method, as measured by the number of data points in the extracted cylindrical region (left), and the radius of the extracted cylindrical region (right). The true radius was 128mm.

half. Results for the unsupervised learning approach are shown in Figure 6. For low stripe id noise standard deviation, almost all points in the cylinder are extracted and a good estimate of the cylinder's radius is produced. However, there is a sharp drop-off in performance between standard deviation levels of 0.01 and 0.1, after which only cylinder fragments are extracted.

The performance on real data is illustrated in Figure 7, which shows the largest cylindrical region extracted. The radius of the extracted cylinder was 39.50mm compared to the actual value of 37.76mm estimated by other means.

5 Conclusion

This paper has considered the recovery of shapes from a collection of standard surface classes. The starting point is a set of data points that correspond to a smooth surface patch, i.e., a preliminary segmentation along occlusion boundaries and surface folds has been performed. The methods developed for shape recovery use the information in the principal quadric at surface points, as well as the point location information. This has resulted in robust methods that have two phases. Firstly, points that could not possibly be on such a surface because of their curvature properties are eliminated. Then, the remaining points are examined in more detail to identify the parameters of the surface that they are on. Region growing and clustering (i.e., unsupervised learning) techniques are used to segment the points into connected sets corresponding to surfaces of the various types. The methods were tested on real and synthetic data. The latter was used to measure the performance as a function of system noise.

Three classes of surface type are considered: planar, spherical, and cylindrical surfaces. Planar surfaces are characterised by containing planar umbilic points. Such points are detected by thresholding the principal curvatures. Then the resulting point set is segmented into connected components using a labelling algorithm.

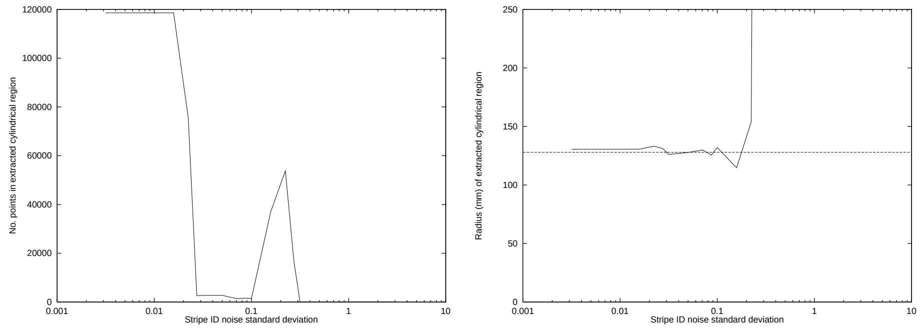

Fig. 6. Performance of the unsupervised learning based cylindrical region extraction method, as measured by the number of data points in the extracted cylindrical region (left), and the radius of the extracted cylindrical region (right). The true radius was 128mm.

Spherical surfaces contain only umbilic points. Such points are detected by thresholding the difference between the principal curvatures. The points then need to be segmented into connected regions containing points all with the same principal curvature. Region growing and clustering algorithms are considered for this task. The clustering method was found to be the more robust of the two.

Cylindrical surfaces contain only parabolic points. They are detected by thresholding the smaller magnitude principal curvature. The other principal curvature gives the cylinder radius, and its principal direction is parallel to the cylinder axis. They need to be grouped on the basis of their radius and cylinder axis (direction and position) to segment out connected cylindrical regions. Both region growing and clustering methods were considered for doing the segmentation. The region growing method was found to be more reliable than the clustering method.

All of the methods have a number of parameters that control their operation and whose value must be selected. The initial point rejection phase requires principal curvature thresholds. The region growing method requires initial estimates of the surface parameters, variance values for them, and a significance value for generating the validation gate. The clustering method requires a significance level for deciding how many clusters are present. Appropriate values for some of these are data dependent. Of those that are, the initial surface parameter estimates are easily generated from the starting point or a small patch around it. The others are the thresholds for the principal curvatures, which clearly are dependent on the noise in the stripe id measurements.

Acknowledgement

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

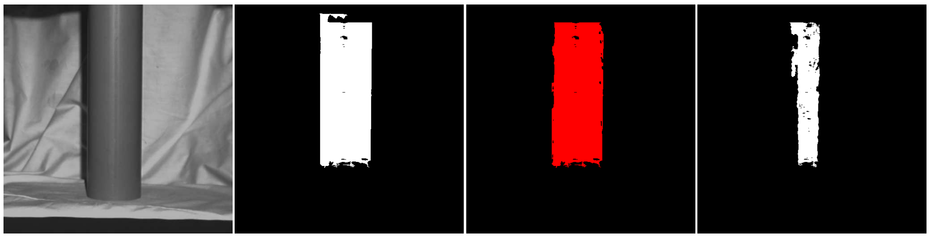

Fig. 7. Example of cylindrical region extraction. From left to right: An image of a scene containing a piece of pipe; A mask indicating the largest connected region of parabolic points; The largest cylindrical region extracted by region growing; The largest cylindrical region extracted by Bayesian classification.

References

1. Alan M. McIvor and Robert J. Valkenburg. A comparison of local surface geometry estimation methods. *Machine Vision and Applications*, 10(1):17–26, 1997.
2. Frank P. Ferrie, Jean Lagarde, and Peter Whaite. Darboux frames, snakes, and super-quadrics: Geometry from the bottom up. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(8):771–784, August 1993.
3. Demetri Terzopoulos, J. Platt, A. Barr, and K. Fleischer. Elastically deformable models. *Computer Graphics*, 21(4):205–214, 1987.
4. Xiaoyi Jiang and Horst Bunke. Fast segmentation of range images into planar regions by scan line grouping. *Machine Vision and Applications*, 7:115–122, 1994.
5. Frank Pipitone. Rapid recognition of elementary surface shapes in cluttered range images using tripod operators. In *MVA'94 IAPR Workshop on Machine Vision Applications*, pages 21–25, Kawasaki, Dec 13–15 1994.
6. Narayan Sriranga Raja and Anil K. Jain. Obtaining generic parts from range images using a multiview representation. *CVGIP: Image Understanding*, 60(1):44–64, July 1994.
7. B. Parvin and G. Medioni. B-rep object description from multiple range views. *International Journal of Computer Vision*, 20(1/2):81–112, October 1996.
8. Alan M. McIvor and Peter T. Waltenberg. Recognition of simple curved surfaces from 3D surface data. IRL Report 690, Industrial Research Limited, March 1997. URL ftp://ftp.vision.irl.cri.nz/pub/doc/1997/irl-report-690.ps.
9. Kenichi Kanatani. *Geometric Computation for Machine Vision*. Number 37 in The Oxford Engineering Science Series. Claredon Press, Oxford, 1993.
10. T. Pavlidis. *Algorithms for graphics and image processing*. Computer Science Press, Springer-Verlag, Berlin – Heidelberg, 1982.
11. R. O. Duda and P. E. Hart. *Pattern Classification and Scene Analysis*. Wiley, New York, 1973.