Machine Vision and Applications (1997) 10: 17–26

Machine Vision and Applications

A comparison of local surface geometry estimation methods

Alan M. McIvor, Robert J. Valkenburg

Imaging and Sensing Team, Industrial Research Limited, P.O. Box 2225, Auckland, New Zealand

Received: 10 March 1996 / Accepted: 10 September 1996

Abstract. The performance of several methods for estimating local surface geometry (the principal frame plus the principal quadric) are examined by applying them to a suite of synthetic and real test data which have been corrupted by various amounts of additive Gaussian noise. Methods considered include finite differences, a facet based approach, and quadric surface fitting. The nonlinear quadric fitting method considered was found to perform the best but has the greatest computational cost. The facet based approach works as well as the other quadric fitting methods and has a much smaller computational cost. Hence it is the recommended method to use in practice.

Key words: Principal quadric – Geometry – Estimation – Evaluation

1 Introduction

This paper presents a comparative analysis of several techniques for determining the local surface geometry of smooth surfaces from 3D scanner data. That is, techniques for determining the principal frame and principal quadric at a point on a smooth surface. Many people consider it important to evaluate the performance of computer vision algorithms. For example, (Jain and Dorai 1995) advocate determining the *admissibility* of algorithms, i.e., under which circumstances an algorithm is the best one for a given task. A comparison of techniques based on surface fitting for local geometry estimation methods is given in (Flynn and Jain 1989). Other related surveys include (Hoover et al. 1995; Jiang et al. 1995) and (Trucco and Fisher 1995), which present studies of range image segmentation techniques.

This paper assumes that surface position samples (in a fixed world coordinate frame) are provided by a sensor in the form of 3D data points. A variety of sensors exist which provide such data, for example (Llario and Martinez 1992; Shirai 1992; Stahs and Wahl 1992; Jarvis 1993; Reid and Brady 1994; Araki et al. 1995; DePiero and Trivedi 1995; McIvor and Valkenburg 1995; Sato 1995). It is useful to estimate the local geometry of surfaces because this provides features for tasks such as surface segmentation (Yu et al. 1994; Masuda and Yokoya 1995; Song and Lee 1995) and recognition of objects (Soucy and Laurendeau 1995).

At a point on a smooth surface, the *normal curvature* in a given direction (in the tangent plane) is the curvature of the planar curve formed by the intersection with the surface of the plane defined by that direction and the surface normal. The *principal directions* are the directions in which the normal curvature has minimum and maximum values. The associated values of the normal curvature are called the *principal curvatures*. The *principal coordinate frame* $\mathbf{X} = \begin{bmatrix} X & Y & Z \end{bmatrix}^T$ has its X and Y axes aligned with the principal directions, and the Z axis aligned with the surface normal.

In the principal coordinate frame, the *principal quadric* or osculating quadric is a second order description of the surface

$$Z = aX^2 + cY^2 \tag{1}$$

with

$$a = \kappa_1/2 \tag{2}$$

$$c = \kappa_2/2 \tag{3}$$

where κ_1 and κ_2 are the principal curvatures (Guggenheimer 1977; Koenderink 1990).

The principal coordinate frame is well-defined except at umbilic points, which are points at which the principal curvatures are equal. At umbilic points, all tangent directions are principal, and any two orthogonal directions can be used. In general, if a surface contains umbilic points, it is not possible to choose principal directions in a globally consistent way.

Let $\mathbf{X}_w$ be the coordinates of a point in the global coordinate frame in which the data is obtained, and let

Correspondence to: Dr Alan M. McIvor (a.mcivor@irl.cri.nz)

$\mathbf{X}$ be the coordinates of that point in the principal coordinate frame associated with a given point p. Then

$$\mathbf{X} = \mathbf{R}(\mathbf{X}_w - \mathbf{X}_w^p) \tag{4}$$

where $\mathbf{X}_w^p$ is the coordinates of p in the global coordinate frame, and $\mathbf{R}$ is a rotation matrix (called the *attitude matrix*). Note that the surface normal (in the global coordinate frame) is given by the bottom row of the attitude matrix.

Shape descriptors can be derived from the principal curvatures. The Gaussian (K) and mean (H) curvatures are particularly useful, and are defined by

$$K = \kappa_1 \kappa_2 \tag{5}$$

$$H = (\kappa_1 + \kappa_2)/2 \tag{6}$$

1.1 Differential geometry approach

This section shows how the principal frame and principal quadric can be determined using the differential properties of the surface. Let $\mathbf{u} = \begin{bmatrix} u \; v \end{bmatrix}^T$ be the domain of a 2D parameterisation of the surface in a local vicinity of the point of interest. The surface *normal* is given by

$$\mathbf{n} = \pm \frac{\partial \mathbf{X}}{\partial u} \times \frac{\partial \mathbf{X}}{\partial v} / \| \frac{\partial \mathbf{X}}{\partial u} \times \frac{\partial \mathbf{X}}{\partial v} \| \tag{7}$$

The sign ambiguity is resolved by a choice of sidedness for the surface. For situations where the surface forms the boundary of an object, the outward pointing normal is often used.

The *first fundamental form* of the surface is defined by

$$\mathbf{G} = \begin{bmatrix} g_{11} & g_{12} \\ g_{21} & g_{22} \end{bmatrix} = \begin{bmatrix} \frac{\partial \mathbf{X}}{\partial u} \cdot \frac{\partial \mathbf{X}}{\partial u} & \frac{\partial \mathbf{X}}{\partial u} \cdot \frac{\partial \mathbf{X}}{\partial v} \\ \frac{\partial \mathbf{X}}{\partial v} \cdot \frac{\partial \mathbf{X}}{\partial u} & \frac{\partial \mathbf{X}}{\partial v} \cdot \frac{\partial \mathbf{X}}{\partial v} \end{bmatrix} \tag{8}$$

and the *second fundamental form* of the surface is defined by

$$\mathbf{D} = \begin{bmatrix} d_{11} & d_{12} \\ d_{21} & d_{22} \end{bmatrix} = \begin{bmatrix} \mathbf{n} \cdot \frac{\partial^2 \mathbf{X}}{\partial u^2} & \mathbf{n} \cdot \frac{\partial^2 \mathbf{X}}{\partial u \partial v} \\ \mathbf{n} \cdot \frac{\partial^2 \mathbf{X}}{\partial v \partial u} & \mathbf{n} \cdot \frac{\partial^2 \mathbf{X}}{\partial v^2} \end{bmatrix} \tag{9}$$

A tangent vector $\dot{\mathbf{X}}$ at the point and its associated pull-back into the parameterisation domain $\dot{\mathbf{u}}$ are related by

$$\dot{\mathbf{X}} = \mathbf{A}\dot{\mathbf{u}} \tag{10}$$

where

$$\mathbf{A} = \begin{bmatrix} \frac{\partial \mathbf{X}}{\partial u} & \frac{\partial \mathbf{X}}{\partial v} \end{bmatrix} \tag{11}$$

The normal curvature, κ_n, of the surface in the direction $\dot{\mathbf{X}}$ is given by

$$\kappa_n = \frac{\dot{\mathbf{u}}^T \mathbf{D} \dot{\mathbf{u}}}{\dot{\mathbf{u}}^T \mathbf{G} \dot{\mathbf{u}}} \tag{12}$$

The principal curvatures are the solutions to

$$|\mathbf{G}|\kappa_n^2 - \alpha \kappa_n + |\mathbf{D}| = 0 \tag{13}$$

where

$$\alpha = g_{11}d_{22} + g_{22}d_{11} - 2g_{12}d_{12} \tag{14}$$

and the principal directions of curvature satisfy

$$(\mathbf{D} - \kappa_i \mathbf{G})\dot{\mathbf{u}}_i = 0 \tag{15}$$

The Gaussian curvature satisfies

$$K = \frac{|\mathbf{D}|}{|\mathbf{G}|} \tag{16}$$

and the mean curvature satisfies

$$H = \frac{\alpha}{2|\mathbf{G}|} \tag{17}$$

1.2 Other approaches

The approaches considered in this paper are based on estimating the derivatives in Eqns. (7)–(9). Other approaches to the problem of determining the local surface geometry have been considered.

One such approach is given in (Yoshimi and Tomita 1994). Here it is assumed that surface normal estimates have been calculated or obtained directly. Then a local region about a point is defined as being those neighbouring points at which the angle between the two surface normals is less than a given constant. Except when the Gaussian curvature is zero, this region is a single closed region with an elliptic boundary. The major and minor axis of the ellipse forming the boundary are the principal directions, and the radii are related to the principal curvatures. The surface geometry calculation algorithm based on this approach can also handle surfaces with zero Gaussian curvature.

Instead of using a surface description based on results from differential geometry, as described above, an alternative is to consider the local variation in position (Liang and Todhunter 1990; Berkmann and Caelli 1994).

Consider the *local (first order) surface covariance matrix*

$$\mathbf{C}_I = \frac{1}{n} \sum_{i=1}^{n} (\mathbf{X}_i - \bar{\mathbf{X}})(\mathbf{X}_i - \bar{\mathbf{X}})^T \tag{18}$$

where

$$\bar{\mathbf{X}} = \frac{1}{n} \sum_{i=1}^{n} \mathbf{X}_i \tag{19}$$

$\mathbf{X}_i$ is the ith point in the neighbourhood being considered, and n is the number of points in the neighbourhood.

Two of the eigenvectors of $\mathbf{C}_I$ define the plane which minimises, in the least squares sense, the orthogonal distance from all points to that plane. This plane can be considered as being analogous to the tangent plane. Consequently, the other eigenvector forms the equivalent of the surface normal. The “topology preserving principal” guides the choice of eigenvectors to use.

The normal vectors for all of the points can then be projected onto the tangent plane (given a coordinate system by the eigenvectors). The 2-dimensional covariance matrix associated with these points, $\mathbf{C}_p$, can then be determined. The 2 eigenvalues and eigenvectors of this matrix can be related to surface shape. Planar regions project to a single point. A locally developable surface projects to a straight line.

1.3 Outline

The rest of the paper is organised as follows. Sect. 2 contains a description of each estimation method considered. Sect. 3 contains a description of the experimental evaluation system, and the performance results for each estimation method. Sect. 4 compares the performance of the various methods, and discusses their attributes. Sect. 5 concludes the paper with a discussion of future work on applying the analysis to other techniques and extending the number of surfaces considered.

2 Details of methods

This section contains a description of the estimation methods being tested.

All of the quadric fitting methods make use of a *rotated principal quadric*, which is the principal quadric Eq. (1) expressed in a coordinate frame related to the principal frame by a rotation about the surface normal. That is, a *rotated principal frame*, $\hat{\mathbf{X}} = \left[\hat{X}\ \hat{Y}\ \hat{Z}\right]^T$, related to the principal frame $\mathbf{X}$ by

$$\mathbf{X} = \begin{bmatrix} \cos\theta & \sin\theta & 0 \\ -\sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \hat{\mathbf{X}} \tag{20}$$

The rotated principal quadric has the form

$$\hat{Z} = a'\hat{X}^2 + b'\hat{X}\hat{Y} + c'\hat{Y}^2 \tag{21}$$

The principal curvatures are the eigenvalues of the second fundamental form and the principal directions are the eigenvectors (do Carmo 1976, pg 144). For a rotated principal quadric, the second fundamental form is

$$\mathbf{D} = \begin{bmatrix} 2a' & b' \\ b' & 2c' \end{bmatrix} \tag{22}$$

An eigenvalue decomposition can thus be used to solve for a, b, and θ. Explicitly, this gives

$$a = \frac{(a'+c') + ((a'-c')^2 + b'^2)^{1/2}}{2} \tag{23}$$

$$c = \frac{(a'+c') - ((a'-c')^2 + b'^2)^{1/2}}{2} \tag{24}$$

$$\theta = 1/2 \tan^{-1}(b'/(a'-c')) \tag{25}$$

Note that θ is determined only up to an ambiguity of π due to the symmetry of the principal quadric.

The transformation from the world coordinate frame to a rotated principal frame is

$$\hat{\mathbf{X}} = \hat{\mathbf{R}}(\mathbf{X}_w - \mathbf{X}_w^p) \tag{26}$$

One useful rotated principal frame is defined by choosing $\hat{\mathbf{R}}$ to be

$$\hat{\mathbf{R}} = \left[\mathbf{r}_1\ \mathbf{r}_2\ \mathbf{r}_3\right]^T \tag{27}$$

$$\mathbf{r}_3 = \mathbf{n} \tag{28}$$

$$\mathbf{r}_1 = \frac{(\mathbf{I} - \mathbf{n}\mathbf{n}^T)\mathbf{e}_1}{\|(\mathbf{I} - \mathbf{n}\mathbf{n}^T)\mathbf{e}_1\|} \tag{29}$$

$$\mathbf{r}_2 = \mathbf{n} \times \mathbf{r}_1 \tag{30}$$

where $\mathbf{n}$ is the surface normal (in world coordinates). This physically represents the rotation which aligns $\hat{X}$ with the projection of X_w onto the tangent plane defined by $\mathbf{n}$.

The reason for introducing rotated principal quadrics is that they separate the problem of estimating the principal quadric into two subproblems: estimating the surface normal, and estimating the principal directions (within the tangent plane). The surface normal, and thus $\hat{\mathbf{R}}$, is always well-defined (since it is assumed that the surface is smooth). Hence the first subproblem always has a unique solution. The parameters of the rotated principal quadric are also well-defined. The angle θ is well-defined except at umbilic points, where an arbitrary choice such as 0 can be made. When considered as a single optimisation problem using the principal quadric directly, there is no local minima for the attitude matrix at umbilic points.

In the implementation of all of the methods based on quadric fitting, the quadric was fitted over 7×7 neighbourhoods. Similarly, the facet based approaches used 7×7 neighbourhoods. Operators defined over 3×3 methods were used in the finite difference based methods.

2.1 Quadric fitting

The rotated principal quadric Eq. (21) forms the basis of the most commonly used approach to estimation (Ferrie et al. 1990; Sander and Zucker 1990; Monga et al. 1991; Ferrie et al. 1993a). The procedure implemented and tested here, which is typical of this approach, is as follows:

1. Fit a plane to the patch of $\mathbf{X}_w$ data points using total least squares. Use the plane's normal as an estimate of the surface normal at the centre of the patch.
2. Construct $\hat{\mathbf{R}}$ from the surface normal estimate using Eqns. (27)–(30).
3. Map the data into the $\hat{\mathbf{X}}$ coordinate frame using Eq. (26).

4. Fit the rotated principal quadric given in Eq. (21) to the mapped data.

5. Determine the coefficients – a, b, θ – of the principal quadric and the transformation from $\hat{\mathbf{X}}$ to $\mathbf{X}$ using Eqns. (23)–(25).

6. Estimate the attitude matrix by concatenating the transformations.

7. Calculate the curvature parameters required from the quadric parameters.

In terms of the parameters of the rotated principal quadric, the Gaussian and mean curvatures are

$$K = 4a'c' - b'^2 \tag{31}$$
$$H = a' + c' \tag{32}$$

2.2 Extended quadric fitting

The method described in the previous section depends on the plane fitting by total least squares accurately estimating the surface normal. No further improvement is made in the quadric fitting stage. Inaccuracies in the surface normal estimate can be allowed for by fitting the quadric

$$\hat{Z} = a'\hat{X}^2 + b'\hat{X}\hat{Y} + c'\hat{Y}^2 + d'\hat{X} + e'\hat{Y} \tag{33}$$

instead of the rotated principal quadric, Eq. (21).

For such a quadric, the surface normal at points in the mapped coordinate frame is

$$\mathbf{n}(\hat{\mathbf{X}}) \propto \begin{bmatrix} -2a'\hat{X} - b'\hat{Y} - d' \\ -b'\hat{X} - 2c'\hat{Y} - e' \\ 1 \end{bmatrix} \tag{34}$$

and, in particular, the surface normal estimate at the origin (i.e., the point at which the principal quadric is being estimated) is

$$\mathbf{n}(0) = \frac{1}{(1 + d'^2 + e'^2)^{1/2}} \begin{bmatrix} -d' \\ -e' \\ 1 \end{bmatrix} \tag{35}$$

The rotation needed to align $\hat{Z}$ with the new surface normal estimate can thus be calculated.

The Gaussian and mean curvature estimates are

$$K = \frac{4a'c' - b'^2}{(1 + d'^2 + e'^2)^2} \tag{36}$$

$$H = \frac{a' + c' + c'd'^2 + a'e'^2 - b'd'e'}{(1 + d'^2 + e'^2)^{3/2}} \tag{37}$$

It is useful to interpret fitting Eq. (33) as fitting Eq. (21) where the origin of the latter is allowed to move to some other point on the surface patch. That is, assuming the surface normal generated by the total least squares surface fitting is not necessarily an estimate of the surface normal at the origin but of the surface normal at some other place on the surface patch. The point on the fitted surface at which the surface normal corresponds to that generated by the total least squares plane fitting is given by

$$\begin{bmatrix} 2a' & b' \\ b' & 2c' \end{bmatrix} \begin{bmatrix} \hat{X} \\ \hat{Y} \end{bmatrix} = - \begin{bmatrix} d' \\ e' \end{bmatrix} \tag{38}$$

A solution is only guaranteed when $K \neq 0$. However, when $K = 0$, one or both of the principal curvatures are zero, and the total least squares estimate of the surface normal can be expected to be orthogonal to the corresponding principal direction. Thus, a solution will still exist.

2.3 Full quadric fitting

The quadrics being fitted in the previous two sections can be extended to include all terms of order two or less. That is, fitting the quadric

$$\hat{Z} = a'\hat{X}^2 + b'\hat{X}\hat{Y} + c'\hat{Y}^2 + d'\hat{X} + e'\hat{Y} + f' \tag{39}$$

The addition of the zero-order term f' allows for noise in the data point at the centre of the patch, i.e., the centre data point is not constrained to lie on the fitted surface.

2.4 Iterative quadric fitting

The method described in Sect. 2.2 gives a new estimate of the surface normal, which is used to generate a new estimate of the principal frame. The data can then be mapped into this new coordinate frame and a new quadric fitted. This should give a better estimate of the principal quadric.

The iteration stopping criteria used was to stop when the incremental change in the direction of the normal fell below a small tolerance level. Convergence was an issue for surface data that contains a significant amount of noise. The processing was also stopped if there was no convergence after 10 iterations. The stability of the system could probably be improved by adding the constraint that the origin of the quadric being fitted remains close to the origin of the data patch.

A similar iterative scheme is discussed in (Fitzgibbon and Fisher 1993), where the model used is

$$\hat{\mathbf{X}} = \hat{\mathbf{R}}(\mathbf{X}_w - \mathbf{X}_w^p) + \mathbf{t} \tag{40}$$
$$\hat{Z} = a'\hat{X}^2 + c'\hat{Y}^2 \tag{41}$$

The surfaces representable with this model are a strict superset of the surfaces representable by Eq. (4) and Eq. (33). The extra parameters, though, generate problems as mentioned in (Fitzgibbon and Fisher 1993).

2.5 Quadric fitting by nonlinear optimisation

The principal quadric is defined by the transformation from the world coordinate frame to the principal coordinate frame, Eq. (4), and the parameters of the principal quadric in this frame, Eq. (1). In Sect. 2.4, several adhoc schemes are discussed for iteratively estimating the principal quadric. A more rigorous approach is to pose the problem as

$$\min_{\hat{\mathbf{R}},a,c,f} \|\hat{Z} - (a\hat{X}^2 + c\hat{Y}^2 + f)\|^2 \tag{42}$$

where

$$\hat{\mathbf{X}} = \hat{\mathbf{R}}(\mathbf{X}_w - \mathbf{X}_w^p) \tag{43}$$

Specifying the attitude matrix requires three parameters, e.g., Euler angles. The parameter f is added to accommodate noise in the surface data point at the centre of the patch. Thus, this is a 6-dimensional nonlinear optimisation problem.

An improved formulation is to use the optimisation criteria

$$\min_{\hat{\mathbf{R}},a,b,c,f} \|\hat{Z} - (a\hat{X}^2 + b\hat{X}\hat{Y} + c\hat{Y}^2 + f)\|^2 \tag{44}$$

where

$$\hat{\mathbf{X}} = \hat{\mathbf{R}}(\mathbf{X}_w - \mathbf{X}_w^p) \tag{45}$$

As discussed in the introduction to this section, Eq. (44) is more stable at umbilic points. Moreover, while both Eq. (42) and Eq. (44) are partially linear problems, Eq. (44) has 4 linear and 2 nonlinear parameters, whereas Eq. (42) has 3 linear and 3 nonlinear parameters[1]. Special techniques for separable least squares problems can be used to solve these problems (Seber and Wild 1989), with Eq. (44) requiring less computational effort.

2.6 Finite differences

The surface normal, curvature, etc., are defined in terms of derivatives of positions on the surface as a function of a two dimensional parameterisation. Such a parameterisation is provided by the data acquisition procedure. Thus, estimates of the surface normal, etc., can be generated by using finite differences to estimate these derivatives. A simple scheme based on 3×3 operators was implemented for testing.

The optimal selection of derivative operators for estimating surface properties is considered in (Weiss 1994; Stoddart et al. 1994; Stoddart et al. 1995). A method for selecting the parameters of the optimal operators on the basis of available a priori information is also presented therein.

[1] The nonlinear parameters represent the rotations.

2.7 Smoothed finite differences

Section 2.6 analysed a very simple finite difference based scheme. One way in which the performance can be improved is to use more complex finite difference operators that, for example, use larger neighbourhoods for their basis. Another approach is to smooth the data before the finite difference operators are used. This section analyses such a method.

The data in the neighbourhood of a point is smoothed using a Gaussian kernel. The weighting factor used depends on the 3D distance of neighbours from the point, as opposed to the separation in terms of grid indexes. That is, the weighting kernel is

$$\exp(-\|\mathbf{X}_w(u+i, v+j) - \mathbf{X}_w(u,v)\|^2/(2\omega^2)) \tag{46}$$

This approach has an advantage in that it rejects outliers because, by definition, they are a long way from the point and hence are given a low weighting. Similarly, it also naturally inhibits smoothing over occlusion boundaries.

Another approach to smoothing data whilst still preserving boundaries is presented in (Cai 1989). A related operator is described in (Boulanger 1994) which weights neighbouring data points by their geodesic distance on the lattice formed by the data points. Note that one of the effects of smoothing will be that the estimated curvatures will be of smaller magnitude than the actual curvatures.

In the implementation, the Gaussian smoothing used had a ω of 10mm and was truncated at 25 pixels from the centre. The same finite difference operators as in Sect 2.6 were used.

2.8 Facets

Each of the components of the 3D data can individually be viewed as a single valued function of the data grid indices. Thus, they can be treated like images and a facet model approach (Haralick and Shapiro 1992) can be used to fit a parameterised function to them. The derivatives of position with respect to the data grid can then be estimated from these fitted functions. This was implemented by fitting quadratic patches

$$\hat{\mathbf{X}}_w = \mathbf{p}_1 u^2 + \mathbf{p}_2 uv + \mathbf{p}_3 v^2 + \mathbf{p}_4 u + \mathbf{p}_5 v + \mathbf{p}_6 \tag{47}$$

using least squares.

For several types of scanner, e.g., (Trucco et al. 1994; DePiero and Trivedi 1995), the data grid corresponds directly to the X and Y components of 3D position. That is, the data is actually generated in the form of a depth map. In these cases, it is only necessary to fit one such facet surface.

With other scanners, 3D data is generated at each data point which satisfies some constraints. For example, in a temporally encoded structured light system, the data grid is the pixel grid of a camera. Hence, when the data is presented in the camera coordinate system, it must satisfy the collinearity equations which describe the perspective projection. This can also be used to reduce the number of facet models that need to be fitted to one. The facet model $Z = f(u, v)$, where u, v are the image plane coordinates, can be fitted using the least squares error measure

$$E = \sum_{(u,v)\in\mathcal{N}} (X - uf(u,v)/f_u)^2 + (Y - vf(u,v)/f_v)^2 + (Z - f(u,v))^2 \quad (48)$$

In addition, with a temporally encoded structured light system, the data is collected by a camera, and the "raw" data for each pixel is an identifier for the stripe sequence observed there. Given calibration information, the identifiers can be converted into 3D data points. Hence, another possibility is to fit a single facet model to the stripe identifiers. However, it can be shown that even for a planar surface, the stripe identifier is a rational function of image plane coordinates (McIvor and Valkenburg 1995). Thus, facet models more complex than low order polynomials will be needed. Note that stripe identifier images are banded due to the discrete number of stripes projected, and the facet model fitting is closely related to the substripe estimation needed to get more precise stripe identifier information.

In many cases, orthonormal polynomials are used in the implementation of the facet fitting algorithm (Flynn and Jain 1989). However, at many points, the data neighbourhood will be incomplete due to occlusions, etc. Over these reduced neighbourhoods, the orthonormal polynomials are no longer orthogonal, and hence it is necessary to revert to the ordinary least squares approach.

3 Experimental method and results

This section of the paper describes the performance of the estimation methods presented in Sect. 2 when applied to a selection of synthetic 3D data, as a function of added Gaussian noise. This analysis approach was taken because it means that the ground truth is known exactly, and the performance is isolated from miscalibration errors, etc. The maximum measurement noise level at which reliable principal quadric estimates can be generated can also be determined.

The synthetic data was generated from planar, cylindrical, and spherical surfaces. The Gaussian noise is independent and identically distributed, and added to all three components of position. That is, the noisy image data, $\mathbf{X}_w$, is generated from the synthetic image data, $\bar{\mathbf{X}}_w$, by

$$\mathbf{X}_w = \bar{\mathbf{X}}_w + \mathbf{n} \quad (49)$$

where

$$\mathbf{n} \sim \mathrm{N}(0, \sigma^2\mathbf{I}_3) \quad (50)$$

Table 1. True values of Gaussian and mean curvature for the 3 test surfaces

Surface	Gaussian curvature	mean curvature
plane	0	0
cylinder	0	−3.9062E-3
sphere	6.1035E-5	−7.8125E-3

The noisy data points are thus uniformly distributed in 3-space about the true position. This is a good first approximation to most sensors.

The plane used in the experiments is

$$0.25\bar{X}_w + 0.25\bar{Y}_w - \bar{Z}_w + 1000 = 0 \quad (51)$$

The cylinder used in the experiments is

$$\bar{X}_w^2 + (\bar{Z}_w - 1000)^2 - 128^2 = 0 \quad (52)$$

The sphere used in the experiments is

$$\bar{X}_w^2 + +\bar{Y}_w^2 + (\bar{Z}_w - 1000)^2 - 128^2 = 0 \quad (53)$$

True values of Gaussian and mean curvature for these surfaces are given in Table 1.

Sampling of these surfaces was simulated by using the preimage of a 2D planar grid under perspective projection. This simulates the data capture for sensors such as temporally encoded structured light systems which use a camera to collect the raw data. The camera coordinate system used coincides with the world coordinate system in which the equations of the surfaces are expressed above. The 3D data was projected onto the camera's image plane assuming that the horizontal and vertical focal lengths were 1000 pixels. Therefore, data samples are approximately 1mm apart on the surfaces.

Output from an estimation method applied to a noise corrupted image was evaluated by calculating the spatial sample average and standard deviation of the angle between the true and estimated surface normals, and of the estimated Gaussian and mean curvatures, over a disk shaped region centred in the images, of radius 100. The nonlinear quadric fitting method is an exception. For this method, 32×32 central subimages were used with statistics collected at all points.

4 Discussion

This section discusses the results from the experimental evaluation of the estimation methods. Some of the results are presented in Figs. 1–3, where they are plotted as a function of the standard deviation, σ, of the noise added to the synthetic images[2]. Complete results are available in (McIvor and Valkenburg 1996).

[2]Note that, in the figures, the absolute value of the curvature averages is plotted since a logarithmic scale is used.

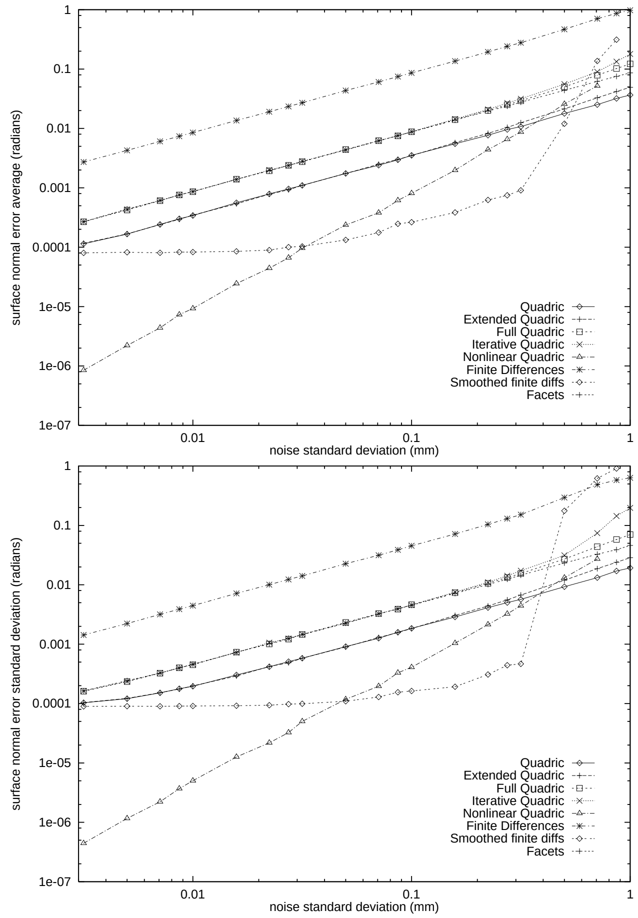

Fig. 1. Surface normal estimation performance for plane

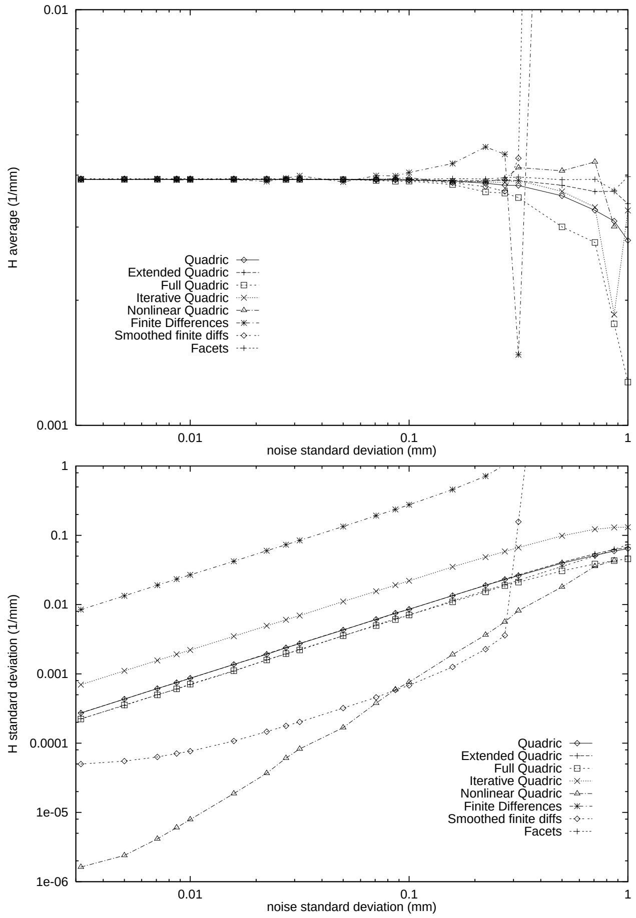

Fig. 2. Mean curvature estimation performance for cylinder

The angle between the true and estimated surface normal is in the range $[0, \pi)$. Hence, a uniform distribution of such angles would have an average value of $\pi/2$ radians. This can be used as a performance benchmark. Furthermore, since the angles are always positive, the average value can only be zero if the estimator is unbiased and has zero variance.

Taking 0.1 radians (i.e., approximately 5°) as the breakpoint at which the bias is considered unacceptable, then all methods have acceptable surface normal error bias below $\sigma = 0.1$. Similarly, below $\sigma = 0.1$, the bias in the curvature estimates is also small. Above $\sigma = 0.1$, the performance of each estimator is very similar for all 3 surfaces. This is presumably because the variation in the underlying surface becomes insignificant compared to the variation caused by the noise. Note that, except for a few cases, the bias is less than 1.96 times the estimator standard deviation. That is, the true value is always within a 95% confidence interval centred around the estimator average value.

One explanation for the bias in the curvature estimates generated by the quadric fitting methods is that cylindrical and spherical patches cannot be represented exactly by a quadric. The least squares fitted quadric tends to have smaller magnitude curvature, and a position offset.

The surface normal performance of each estimator is very similar for all surfaces. The standard deviations of both curvature estimates is also almost independent of the surface type. However, results from a larger set of surfaces are needed before the hypothesis that the performance is independent of the surface can be accepted.

All five quadric fitting methods use a surface normal estimate generated by fitting a plane to the data by total least squares. For the quadric method, this is the final surface normal estimate. For the other methods it is only an initial estimate. The surface normal estimate selected by the total least squares fit is the unit eigenvector corresponding to the smallest eigenvalue of the design matrix. If the surface curvature is high, or the data spread perpendicular to the true surface normal is large compared to the patch size (because of the sampling noise), the selected eigenvector will not be the correct one and the principal quadric estimate will be no good. The breakdown point can be estimated as the noise level at which 3σ equals the patch half size. This is $\sigma = 1$ for the experiments since 7×7 patches were used.

mated values of Gaussian curvature. The comparison of mean curvatures is inconclusive.

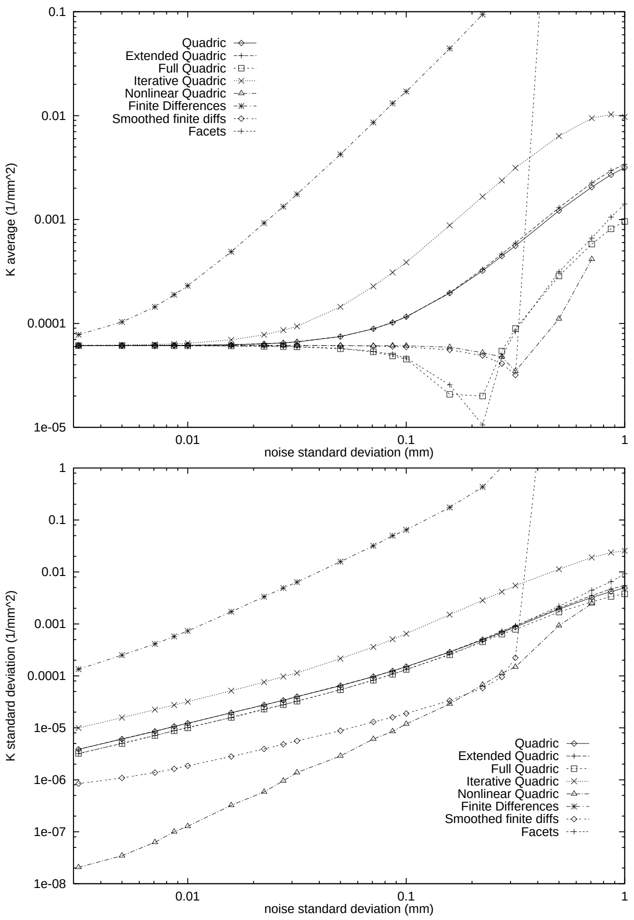

Fig. 3. Gaussian curvature estimation performance for sphere

The quadric, extended quadric, and full quadric methods are closely related, only differing in the number of terms used in the quadric model fitted. The performance of the curvature estimates generated by these methods is quite similar. However, the standard deviation of the surface normal estimates generated by the extended quadric and full quadric methods is smaller than those generated by the quadric fitting methods, for the cylinder and sphere (i.e., when at least one principal curvature is nonzero). This is because the extra quadric terms in the extended and full quadric fitting methods can correct errors in the surface normal estimated by the total least squares planar fit, as shown in Sect. 2.2. The central point of a data patch (i.e., the point at which the principal quadric is being estimated) is displaced from the true underlying surface by noise. The zero-order term in the full quadric method compensates for this, and allows selection of a quadric that passes through the (estimated) corresponding true surface point. The other two quadric fitting methods do not have a zero-order term and hence are forced to select a quadric that passes through the noisy data point. Consequently, the fitted quadrics have higher curvature magnitude to compensate. This conjecture is clearly supported by a comparison of the estimated values of Gaussian curvature. The comparison of mean curvatures is inconclusive.

The iterative quadric fitting method performs worse than all of the three non-iterative versions. This is probably due to the instability encountered in the implementation. The standard deviation of the curvature estimates generated by the nonlinear quadric fitting method is several orders of magnitude less than the standard deviation of the estimates from the other methods. This suggests that if a stable iterative method can be found, it would result in an improved estimator which does not have the computational complexity of the nonlinear method.

The formulation of the quadric fitting approaches is only approximate because the error criteria used assumes that all errors are in the direction of the normal at the centre of the data patch. It would require a much more complex formulation to handle, say, the model where the errors are assumed to be samples from a 3D measurement error distribution.

Not surprisingly, the finite difference approach has the worst performance of all methods. However, the finite difference method with pre-smoothing has performance equal to or better than all other methods, except the non-linear quadric fitting, at low noise levels. The smoothing can be viewed as a low-pass filter that rejects the noise until it is so large (at $\sigma = 0.3$) that the underlying surface is extremely distorted. Above this noise level, the performance is no better than finite differences alone. The smoothing is also known to introduce a bias in the curvature estimates (an underestimation of the magnitude), although this is not evident for the three surfaces considered. More work is needed to investigate the effects of smoothing.

The facet approach gives very similar performance to the full quadric fitting approach.

The three surfaces used in the experimental evaluation are all symmetric. Hence the performance of methods using total least squares surface normal estimates is probably over-stated.

The computational complexity of the nonlinear quadric fitting method is largest. The method requiring the least computation is the facet method since the design matrix depends on the sampling grid and thus its decomposition only needs to be computed once.

The performance statistics for real data acquired with a temporally encoded structured light system are presented in Tables 2 and 3. The sensor has a σ value of 0.5mm. The true position and orientation of the surfaces was estimated by least squares fitting a plane (resp. cylinder) to all data points. The correct mean curvature of the cylinder was determined to be -0.0132. These results agree with the simulation results, although they are slightly higher. This is due, in part, to residual errors in the surface position estimates as a result of unmodelled lens distortions, etc.

Since points are often classified on the basis of the sign of the Gaussian and mean curvatures, it is use-

Table 2. Statistics for real data taken from a planar surface with a temporally encoded structured light system

	surface normal error		Gaussian curvature		mean curvature	
	ave	SD	ave	SD	ave	SD
Quadric	0.1181	0.0708	0.0014	0.0019	−0.0004	0.0502
Extended Quadric	0.1169	0.0715	0.0014	0.0020	−0.0004	0.0506
Full Quadric	0.1166	0.0715	−0.0001	0.0011	−0.0004	0.0339
Iterative Quadric	0.1169	0.0715	0.0014	0.0020	−0.0004	0.0506
Nonlinear Quadric	0.1124	0.0052	−0.0001	0.0006	0.0015	0.0270
Finite Differences	0.1359	0.0836	0.0012	0.0373	−0.0006	0.2425
Smoothed finite diffs	0.0072	0.0037	0.0000	1.2E-6	−0.0001	0.0010
Facets	0.1180	0.0708	−0.0012	0.0023	−0.0004	0.0340

Table 3. Statistics for real data taken from a cylindrical surface with a temporally encoded structured light system

	surface normal error		Gaussian curvature		mean curvature	
	ave	SD	ave	SD	ave	SD
Quadric	0.0934	0.0866	0.0098	0.0288	−0.0126	0.1514
Extended Quadric	0.0858	0.0742	0.0099	0.0306	−0.0129	0.1538
Full Quadric	0.0878	0.0803	−0.0016	0.0164	−0.0116	0.1106
Iterative Quadric	0.0863	0.0754	0.0099	0.0309	−0.0130	0.1542
Nonlinear Quadric	0.0524	0.0285	−0.0003	0.0051	−0.0121	0.0463
Finite Differences	0.3053	0.1965	0.0586	0.5080	−0.0111	0.7308
Smoothed finite diffs	0.0086	0.0059	0.0000	0.0005	−0.0140	0.0090
Facets	0.0873	0.0806	−0.0086	0.0273	−0.0136	0.1115

ful to look at the probability of getting the sign wrong (when the true value is non-zero). Figure 4 shows the probability of getting the mean curvature sign wrong for the cylinder. Similiar results are obtained for the sphere's Gaussian and mean curvatures. Only the nonlinear quadric fitting and smoothed least squares methods keep the probability below 5% for all values of noise standard deviation below 0.1. The basic quadric fitting method and the facet method are both less likely to get the sign of the sphere's Gaussian curvature wrong than the extended and full quadric methods, when the noise standard deviation is above 0.03. Below this, they all behave similarly. The probability is always less than 0.5 because the estimator average always has the correct sign. Also, as noted above, the true value is almost always within the 95% confidence interval of the estimator average.

(Flynn and Jain 1989) also looks at the performance of several methods for estimating the surface curvature, including the quadric fitting method considered here. Synthetic plane, cylindrical, and spherical surfaces are considered, but only the averages of the estimator outputs are reported. They also consider the effects of quantisation within the data generation process, and concluded that this is a major source of errors. Accuracies of 10% for non-zero curvatures were recorded. (Hilton et al. 1995) attempts to model the noise distribution in curvature estimates given a model of the noise in the range data.

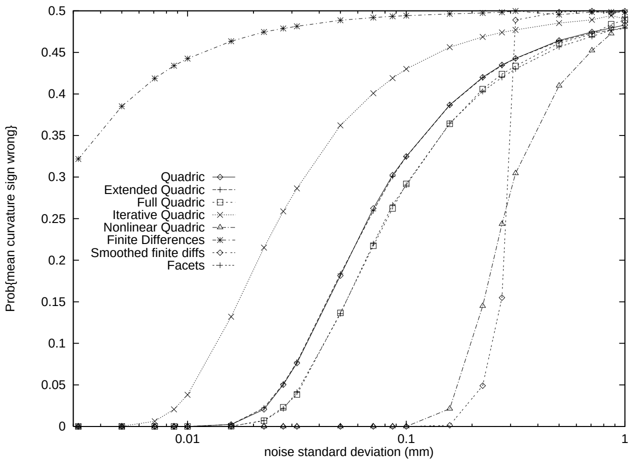

Fig. 4. The probability of getting the mean curvature sign wrong for the cylinder

5 Conclusion

This paper is about measuring the performance of methods for estimating the local surface geometry from 3D data. Several methods, including those in common use, are introduced and then the results are presented from applying them to a set of synthetic and real data. The performance was measured by the average and standard deviation of the angle between the true and estimated surface normals, and the average and standard deviation of the Gaussian and mean curvatures.

As expected, the finite difference scheme had very poor performance. However, the performance of the finite difference scheme with pre-smoothing indicates that such an approach may generate a useful method.

Many of the estimators are 2-step methods based on estimating the surface normal and then fitting a quadric patch. The number of terms used in the quadric seemed to make little difference to the curvature estimate performance, although having first and zero order terms improves the surface normal estimates. The nonlinear quadric fitting method showed that much better performance is possible, although at much greater computational cost. A simple iterative scheme was tried but this proved unstable and hence not an improvement over the non-iterative schemes. Other approaches, e.g., (Fitzgibbon and Fisher 1993), have proved to be more stable.

Similar performance to the quadric fitting methods was obtained from a facet based approach applied to the 3D data. This method has the computational advantage that the design matrix is dependent on the data grid not the data, and hence only needs to be computed once when the algorithm is applied over a whole data set.

Future work will include looking at other methods, such as using higher order surface models, e.g., (Flynn and Jain 1988), and using *curvature consistency* (Ferrie et al. 1993b). Another approach would be to use low-order surface models, but couple this with spatial data weighting to reflect the fact that the principal quadric is not necessarily a good second order approximation over a finite surface patch. Another direction for future work will be to consider other surfaces, especially non-symmetric ones. Also the effects of different values of surface curvature need to be explored to test the tentative hypothesis that performance is independent of the underlying surface. Finally, the optimal selection of parameters, e.g., patch size, needs investigation.

Acknowledgement. The authors acknowledge the support of the New Zealand Foundation of Science, Research and Technology, under contract CO8410, Objective 2.

References

Araki K, Shimizu M, Noda T, Chiba Y, Tsuda Y, Ikegaya K, Sannonmiya K, Gomi M (1995) A high-speed and continuous 3D measurement system. Machine Vision and Applications 8:79–84

Berkmann J, Caelli T (1994) Computation of surface geometry and segmentation using covariance techniques. IEEE Transactions on Pattern Analysis and Machine Intelligence 16:1114–1116

Boulanger P (1994) Multiscale edge detection based on a new geometrically intrinsic filter. In: El-Hakim SF (ed) Videometrics III. SPIE, pp 264–278

Cai LD (1989) A small leakage model for diffusion smoothing of range data. In: 11th IJCAI. Detroit pp 1585–1590

DePiero FW, Trivedi MM (1995) 3-D computer vision using structured light: Design, calibration and implementation issues. Advances in Computers 43:243–278

do Carmo MP (1976) Differential Geometry of Curves and Surfaces. Prentice-Hall, Englewood Cliffs

Ferrie FP, Lagarde J, Whaite P (1990) Recovery of volumetric object descriptions from laser rangefinder images. In: ECCV90. pp 387–396

Ferrie FP, Lagarde J, Whaite P (1993a) Darboux frames, snakes, and super-quadrics: Geometry from the bottom up. IEEE Transactions on Pattern Analysis and Machine Intelligence 15:771–784

Ferrie FP, Mathur S, Soucy G (1993b) Feature extraction for 3-D model building and object recognition. In: Jain AK, Flynn PJ (eds) Three-Dimensional Object Recognition Systems. Elsevier Science Publishers, Amsterdam pp 57–113

Fitzgibbon AW, Fisher RB (1993) Invariant fitting of arbitrary single-extremum surfaces. In: Illingworth J (ed) British Machine Vision Conference. University of Surrey, Guildford, U. K., BMVA Press, Sheffield pp 569–578

Flynn PJ, Jain AK (1988) Surface classification: Hypothesis testing and parameter estimation. In: Computer Vision and Pattern Recognition. IEEE, pp 261–267

Flynn PJ, Jain AK (1989) On reliable curvature estimation. In: CVPR. pp 110–116

Guggenheimer HW (1977) Differential Geometry. Dover, Mineola

Haralick RM, Shapiro LG (1992) Computer and robot vision. Addison Wesley, Reading, Mass.

Hilton ADM, Illingworth J, Windeatt T (1995) Statistics of surface curvature estimates. Pattern Recognition 28:1201–1221

Hoover A, Jean-Baptiste G, Jiang X, Flynn PJ, Bunke H, Goldof D, Bowyer K (1995) Range image segmentation: The user's dilemma. In: IEEE International Symposium on Computer Vision. Coral Gables, Florida pp 323–328

Jain AK, Dorai C (1995) Keynote address — Practicing vision: Integration, evaluation and applications. In: ACCV'95 Second Asian Conference on Computer Vision. Singapore

Jarvis R (1993) Range sensing for computer vision. In: Jain AK, Flynn PJ (eds) Three-Dimensional Object Recognition Systems. Elsevier Science Publishers, Amsterdam pp 17–56

Jiang XY, Hoover A, Jean-Baptiste G, Goldgof D, Bowyer K, Bunke H (1995) A methodology for evaluating edge detection techniques for range images. In: ACCV'95 Second Asian Conference on Computer Vision. Nanyang Technological University, Singapore pp II–415 – II–419

Koenderink JJ (1990) Solid Shape. MIT Press, Cambridge, Mass.

Liang P, Todhunter JS (1990) Representation and recognition of surface shapes in range images: A differential geometric approach. Computer Vision, Graphics, and Image Processing 52:78–109

Llario V, Martinez AB (1992) Active methods for obtaining depth maps. In: Torras C (ed) Computer Vision: Theory and Industrial Applications. Springer-Verlag, Berlin Heidelberg New York Ch. 3, pp 97–134

Masuda T, Yokoya N (1995) A robust method for registration and segmentation of multiple range images. Computer Vision and Image Understanding 61:295–307

Monga O, Ayache N, Sander P (1991) Modeling uncertainty for estimating local surface geometry. In: 7th Scandinavian Conference on Image Analysis. Aalborg pp 403–410

McIvor AM, Valkenburg RJ (1995) Calibrating a structured light system. In: Image & Vision Computing New Zealand. Industrial Research Limited, Lincoln, Canterbury pp 167–172

McIvor AM, Valkenburg RJ (1996) A comparison of local surface geometry estimation methods. Report 520 Industrial Research Limited Available as ftp://ftp.vision.irl.cri.nz/pub/doc/1996/irl-report-520.ps

Reid ID, Brady JM (1994) Range image feature extraction and representation. In: Cameron S, Probert P (eds) Advanced Guided Vehicles — Aspects of the Oxford AGV project. World Scientific, Singapore Ch. 9, pp 141–162

Sander PT, Zucker SW (1990) Charting surface structure. In: ECCV90. pp 418–426

Sato Y (1995) Active rangefinding and recognition with Cubiscope. In: ACCV'95 Second Asian Conference on Computer Vision. Singapore pp I–369 – I–372

Seber GAF, Wild CJ (1989) Nonlinear Regression. Wiley, New York

Shirai Y (1992) 3D computer vision and applications. In: 11th International Conference on Pattern Recognition. International Association for Pattern Recognition, IEEE Computer Society Press, pp 236–245

Song B, Lee S (1995) On the range image segmentation techique based on local correlation. In: ACCV'95 Second Asian Conference on Computer Vision. Singapore pp I–528 – I–532

Soucy M, Laurendeau D (1995) A dynamic integration algorithm to model surfaces from multiple range view. Machine Vision and Applications 8:53–62

Stahs T, Wahl F (1992) Fast and versatile range data acquisition in a robot work cell. In: IEEE/RSJ International Conference on Intelligent Robots and Systems. pp 1169–1174

Stoddart AJ, Illingworth J, Windeatt T (1994) Optimal parameter selection for derivative estimation from range images. In: Hancock E (ed) BMVC94: Proceedings of the 5th British Machine Vision Conference. BMVA Press, Sheffield pp 165–174

Stoddart AJ, Illingworth J, Windeatt T (1995) Optimal parameter selection for derivative estimation from range images. Image and Vision Computing 13:629–635

Trucco E, Fisher RB (1995) Experiments in curvature-based segmentation of range data. IEEE Transactions on Pattern Analysis and Machine Intelligence 17:177–182

Trucco E, Fisher RB, Fitzgibbon AW (1994) Direct calibration and data consistency in 3-D laser scanning. In: Hancock E (ed) BMVC94: Proceedings of the 5th British Machine Vision Conference. BMVA Press, Sheffield pp 489–498

Weiss I (1994) High-order differentiation filters that work. IEEE Transactions on Pattern Analysis and Machine Intelligence 16:734–739

Yoshimi T, Tomita F (1994) Robust curvature vectors calculation from range data using ISL method. In: MVA'94: IAPR Workshop on Machine Vision Applications. Kawasaki, Japan pp 506–509

Yu X, Bui TD, Krzyzak A (1994) Robust estimation for range image segmentation and reconstruction. IEEE Transactions on Pattern Analysis and Machine Intelligence 16:530–538

Alan M. M^{c}Ivor is a research scientist at Industrial Research Limited, where his interests include computer vision, communication networks, and SGML. He is currently working on techniques for interpreting 3D data. Previous positions include 3 years at the British Petroleum Research Centre developing control systems for remotely operated submersible vehicles, and a post-doc at Oxford University working in Computer Vision. He has a PhD and SM from Harvard University in Engineering and a BE (first class honours) in Engineering Science from Auckland University.

Bob Valkenburg has been a research scientist for six years. He is project leader for the Videometrics project and specialises in mathematical analysis, high speed computing and hardware design. He is currently working towards a PhD (Engineering) at Auckland University. BE (Electronics), ME (Electronics), DipSci (Maths), PGDipSci (Maths) Auckland University.

This article was processed by the author using AMM's LaTeX document class which is an approximation to the document style *cljour2* from Springer-Verlag.