Principal Frame and Principal Quadric Estimation

Dr Alan M. McIvor and Robert J. Valkenburg

Machine Vision Team, Industrial Research Limited
P.O. Box 2225, Auckland
New Zealand
Ph: +64 9 3034116, Fax: +64 9 3028106
E-mail: a.mcivor@irl.cri.nz, b.valkenburg@irl.cri.nz
URL: http://kauri.auck.irl.cri.nz/

Abstract

The principal frame and principal quadric form a fundamental description of a smooth surface at a point. This paper is concerned with estimating these from the 3D data acquired by a sensor such as a structured light system. Eight different estimation methods are described and a comparison of their performance is presented.

1 Introduction

Sensors such as temporally encoded structured light systems [1] and laser scanners [2] can generate dense 3D position samples from the surfaces of objects. The basic properties of the underlying surface at a data point are described by the osculating paraboloid. This is the local quadratic approximation to the surface which has second order contact. Moreover, it is possible to define a unique coordinate representation of the osculating paraboloid, except for certain symmetrical surfaces. This representation is called the *principal quadric* when it is defined in the *principal coordinate frame*.

The principal coordinate frame at a point on a smooth surface is defined as follows. One coordinate is chosen in the direction of the surface normal[1]. The intersection of any plane containing the surface normal with the surface is a (planar) curve. The curvature of this curve is the *normal curvature* of the surface in the direction of the plane. The maximum and minimum values of normal curvature are the *principal curvatures* and their associated directions, which are orthogonal, are the *principal directions*. In the principal frame, the principal quadric is given by

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

with

$$\begin{aligned} a &= \kappa_1/2 & (2)\\ c &= \kappa_2/2 & (3) \end{aligned}$$

where κ_1 and κ_2 are the principal curvatures [3, 4]. The representation in terms of the principal frame and principal curvatures (which is an eigenvalue representation) is "usually

[1] In all that follows this will be assumed to be the third or Z coordinate.

the most obvious one. The principal directions tend to coincide with visually obvious symmetries" [4, Pg 228].

Consider any coordinate frame which has its third coordinate axis in the direction of the surface normal (but not necessarily having the others axes aligned with the principal directions). Such a coordinate frame will be called a *rotated principal frame*. In a rotated principal frame, the osculating paraboloid has the representation

$$\hat{Z} = a'\hat{X}^2 + b'\hat{X}\hat{Y} + c'\hat{Y}^2 \tag{4}$$

The coefficients of the principal quadric can be obtained from these by

$$a = \frac{(a'+c') + ((a'-c')^2 + b'^2)^{1/2}}{2} \tag{5}$$

$$c = \frac{(a'+c') - ((a'-c')^2 + b'^2)^{1/2}}{2} \tag{6}$$

and the coordinate transformation from the rotated principal frame coordinates $\hat{\mathbf{X}}$ to the principal frame coordinates $\mathbf{X}$ is

$$\mathbf{X} = \begin{bmatrix} \cos\theta & \sin\theta & 0 \\ -\sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \hat{\mathbf{X}} \tag{7}$$

where

$$\theta = \frac{1}{2}\operatorname{atan2}(b', a' - c') \tag{8}$$

There are many different ways in which the principal quadric can be estimated from dense 3D surface point data. This paper briefly describes 8 such methods and compares their performance. The methods are described in Section 2 and the comparison results are presented in Section 3.

2 Estimation methods

The 3D surface point data is provided such that it is indexed on a two dimensional grid, e.g., the camera pixel array in a temporally encoded structured light system. The continuous analogue of this is a chart mapping $\mathbb{R}^2$ to $\mathbb{R}^3$. The coefficients of the coordinate transformation from the data frame to the principal frame, and the principal curvatures, can all be defined in terms of derivatives (up to second order) of this mapping. The first estimation method, `Finite Differences`, consists of approximating these derivatives using finite difference operators. The second, `Smoothed finite diffs`, augments these operators with Gaussian pre-smoothing.

The next three methods all are based on partitioning the problem into two subproblems:

1. Estimate a rotated principal frame, and map the data into the rotated frame;
2. Fit a quadric to the mapped data, and calculate the parameters of the principal quadric from the fitted quadric.

All 3 methods use the same technique for the first subproblem. Total least squares is used to find a planar approximation to the data set, and the normal to this plane is used as an estimate of the surface normal, $\mathbf{n}$. Given this, the transformation from data coordinates $\mathbf{X}_w$ to the rotated principal frame coordinates is

$$\hat{\mathbf{X}} = \hat{\mathbf{R}}(\mathbf{X}_w - \mathbf{X}_w^p) \tag{9}$$

where $\mathbf{X}_w^p$ is the point at which the principal quadric is being estimated (the centre of the data patch), and

$$\hat{\mathbf{R}} = \begin{bmatrix} \mathbf{r}_1 & \mathbf{r}_2 & \mathbf{r}_3 \end{bmatrix}^T \tag{10}$$

$$\mathbf{r}_3 = \mathbf{n} \tag{11}$$

$$\mathbf{r}_1 = \frac{(\mathbf{I} - \mathbf{n}\mathbf{n}^T)\mathbf{e}_1}{\|(\mathbf{I} - \mathbf{n}\mathbf{n}^T)\mathbf{e}_1\|} \tag{12}$$

$$\mathbf{r}_2 = \mathbf{n} \times \mathbf{r}_1 \tag{13}$$

All 3 methods involve fitting a quadric by least squares for the second subproblem. They differ by which low order terms they include in the quadric, as follows:

- `Quadric`

$$\hat{Z} = a'\hat{X}^2 + b'\hat{X}\hat{Y} + c'\hat{Y}^2 \tag{14}$$

- `Extended Quadric`

$$\hat{Z} = a'\hat{X}^2 + b'\hat{X}\hat{Y} + c'\hat{Y}^2 + d'\hat{X} + e'\hat{Y} \tag{15}$$

- `Full Quadric`

$$\hat{Z} = a'\hat{X}^2 + b'\hat{X}\hat{Y} + c'\hat{Y}^2 + d'\hat{X} + e'\hat{Y} + f' \tag{16}$$

At the origin, i.e., where the principal quadric is being estimated, the normal to the surface patch fitted in the `Quadric` method always lies in the $\hat{Z}$ axis direction. The other two methods have enough flexibility to allow the normal direction to vary. This new surface normal estimate can be used to form a new rotated principal frame estimate. The data can be mapped into this new frame, and a new quadric fitted. The process can be iterated in this way, thereby allowing for initial estimates in the surface normal direction. The `Iterative Quadric` method is formed by applying this updating to (15).

The 4 quadric fitting approaches just described can be viewed as linearised solutions to the non-linear optimisation problem of minimising the error between a data set and a quadric. Another approach is to use a non-linear least squares technique. The `Nonlinear Quadric` method uses such a technique to solve the problem

$$\min_{\hat{\mathbf{R}},a,b,c,f} \|\hat{Z} - (a\hat{X}^2 + b\hat{X}\hat{Y} + c\hat{Y}^2 + f)\|^2 \tag{17}$$

where

$$\hat{\mathbf{X}} = \hat{\mathbf{R}}(\mathbf{X}_w - \mathbf{X}_w^p) \tag{18}$$

The last method, `Facets`, is based on fitting quadratic facets individually to each component of the surface position, as a function of the data set index. The quadratic functions can be differentiated to obtain the parameters needed to estimate the principal quadric.

3 Evaluation

The performance of the estimation methods described in Section 2 was compared by applying them to a test suite of synthetic 3D data sets. Each data set was generated by taking data from a smooth "base" surface and adding Gaussian noise. The Gaussian noise was independently and identically distributed over all 3 components of position. A range of variances was considered. Three different base surface were used: a plane, a cylinder, and a sphere.

The output from an estimation method applied to a noise corrupted data set 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 the central region of the data set. The Gaussian curvature (the product of the principal curvatures) and the mean curvature (the average of the principal curvatures) were recorded instead of the principal curvatures because they are more commonly used in subsequent processing. Representative results are shown in Figures 1–4. Complete results are given in [5].

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.

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 very 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). The performance of the iterative quadric fitting method is 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. The smoothing is 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.

4 Conclusion

The performance of a number of methods for estimating the principal quadric was evaluated by applying them to a collection of test data sets. Statistics on the surface normal estimate errors, and the Gaussian and mean curvature estimate errors, for each method were thereby generated, as a function of the amount of noise in the data set.

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 two step methods based on estimating the surface normal and then fitting a quadric patch. The number of low order 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., [6], 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.

References

1. Alan M. McIvor. *Determining surface geometry using structured light.* In Image & Vision Computing New Zealand, pages 83–88, Lincoln, Canterbury, August 1995.

2. Ray Jarvis. *Range sensing for computer vision.* In Anil K. Jain and Patrick J. Flynn, editors, Three-Dimensional Object Recognition Systems, volume 1 of *Advances in Image Communications*, pages 17–56. Elsevier Science Publishers, Amsterdam, 1993.

3. Heinrick W. Guggenheimer. *Differential Geometry.* Dover, Mineola, 1977.

4. Jan J. Koenderink. *Solid Shape.* MIT Press, Cambridge, Mass., 1990.

5. Alan M. McIvor and Robert J. Valkenburg. *A comparison of local surface geometry estimation methods.* Report 520, Industrial Research Limited, February 1996. Available as ftp://ftp.vision.irl.cri.nz/pub/doc/1996/irl-report-520.ps.

6. Andrew W. Fitzgibbon and Robert B. Fisher. *Invariant fitting of arbitrary single-extremum surfaces.* In John Illingworth, editor, British Machine Vision Conference, pages 569–578, Sheffield, September 1993. University of Surrey, Guildford, U. K., BMVA Press.

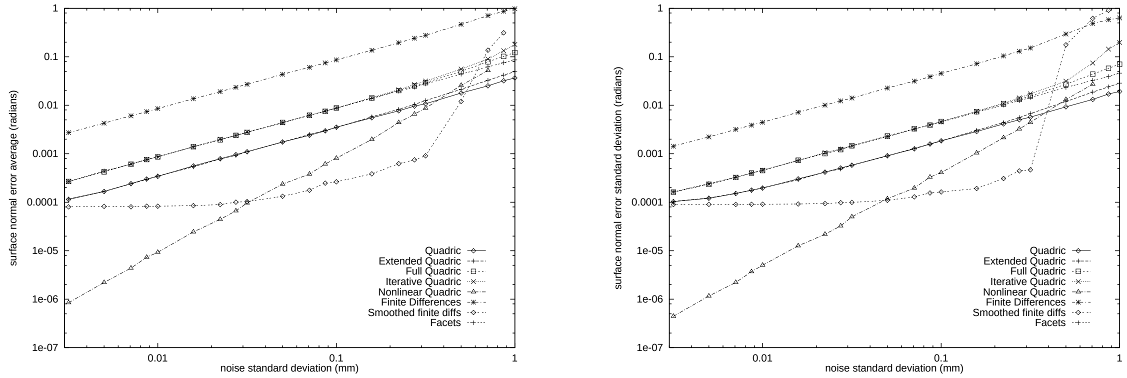

Figure 1: Surface normal estimation performance for plane.

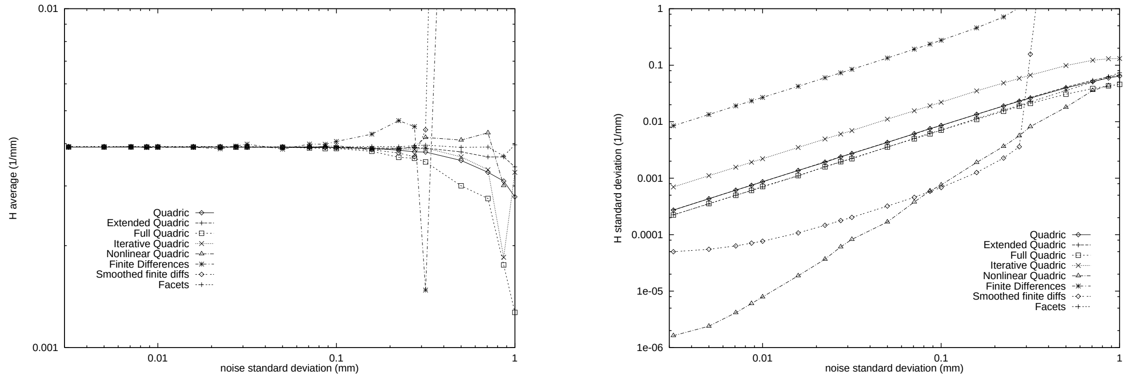

Figure 2: Mean curvature estimation performance for cylinder.

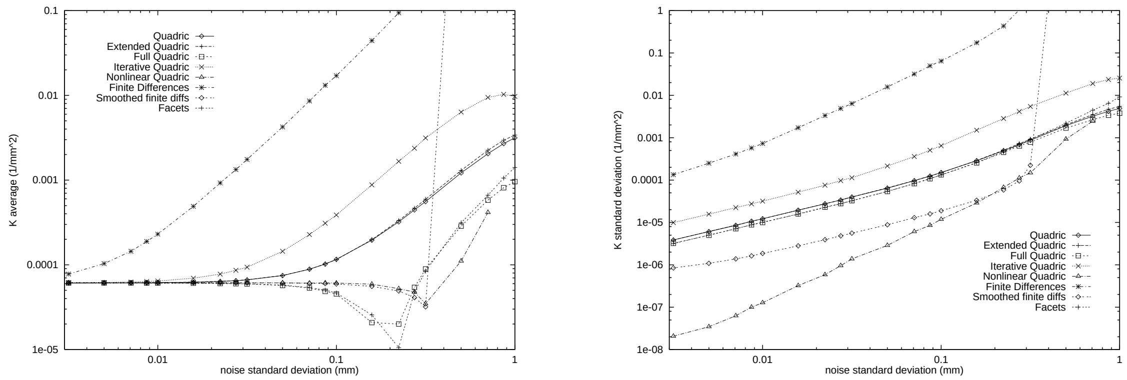

Figure 3: Gaussian curvature estimation performance for sphere.

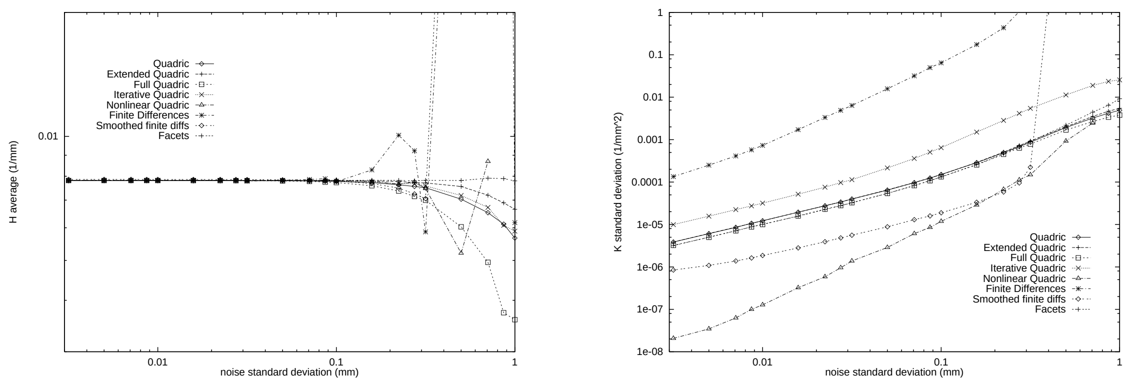

Figure 4: Mean curvature estimation performance for sphere.