Shape Invariants and Principal Directions from 3D Points and Normals

George Kamberov

Gerda Kamberov

Department of Computer Science
Stevens Institute of Technology
Hoboken, NJ 07030, USA

Department of Computer Science
Hofstra University
Hempstead, NY 11549, USA

ABSTRACT

A new technique for computing the differential invariants of a surface from 3D sample points and normals. It is based on a new conformal geometric approach to computing shape invariants directly from the Gauss map. In the current implementation we compute the mean curvature, the Gauss curvature, and the principal curvature axes at 3D points reconstructed by area-based stereo. The differential invariants are computed directly from the points and the normals without prior recovery of a 3D surface model and an approximate surface parameterization. The technique is stable computationally.

Keywords: shape, mean curvature, gauss curvature, principal directions

1 Introduction

The differential invariants of a surface provide a complete description of its shape. They are used in surface matching, in localizing important regions in surfaces, providing descriptors and measurements in various applications.

One set of approaches in computer vision to the estimation of the differential invariants is to reconstruct an approximation of the surface and its parameterization, and then derive the differential invariants using classical differential geometry. In these approaches, first, range sensors, stereo, MRI, or CT techniques are used to collect a sufficiently dense set of sampled surface points; second, an approximate parametrization is derived from an estimated 3D surface model; and third, the differential invariants are computed from the approximated surface. The 3D model is obtained by applying a marching cubes method, a Delaunay triangulation, or some model fitting or smoothing technique (some perennial favorites include Gaussian convolution, or wavelets expansions). See [4, 5, 8, 10, 13]. An important caveat concerning the derivation of invariants from 3D models is that the models do not come with robust and general error estimates, in particular, for the curvature estimates obtained from them. Yet, another set of approaches forgo the 3D model, and attempt to extract the curvature invariants directly from range, stereo, or photometric data [4]. Both approaches use classical differential geometry in the recovery of the differential invariants (including taking second order derivatives, and solving general characteristic polynomials). An additional source of errors is the computational instability of the methods. For example, to compute the principal curvature vectors and the principal curvatures, the methods rely on diagonalizing general symmetric matrices (in fact the operators are often only close to symmetric due to noise and round off errors). The standard diagonalization routines introduce additional errors.

The classical shape invariants, principal curvatures, mean curvature, and Gauss curvature, belong to the realm of metric geometry.

We present new theoretical results that utilize the conformal structure in the derivation of the differential invariants. In the implementation we, first, use stereo to reconstruct 3D points; second, recover the Gauss map; and third, apply the new theory to recover the differential invariants. Our approach is based on

the realization that: (i) Unless the surface is minimal, the Gauss map completely determines it up to scale and translation. (ii) Furthermore, in all cases (including minimal surfaces) the differential invariants of the surface can be extracted by, first, estimating the mean curvature, and second, computing the Gauss curvature and principal axes by diagonalizing a special trace-free, symmetric matrix. The theory we present for the extraction of the mean curvature does not require a diagonalization, also, for the recovery of the rest of the invariants, the special type of the matrix allows us to diagonalize it to find the principal curvature vectors and Gauss curvature without resorting to solving general characteristic polynomials. This results in improved stability of our approach because we can reduce the number of nonstable nonlinear operations, including taking of square roots, for example.

The new methods require tools for extracting the Gauss map and the conformal structure (the notion of angles) on a surface in $\boldsymbol{R}^3$. Fortunately such tools exist. Our current implementation uses the fish-scales method developed in [8].

The rest of the paper is organized as follows: basic definitions and preliminaries are reviewed in Section 2; the new theoretical results are discussed in Section 3; implementation, results, and evaluation are given in Section 4.

2 Theoretical Background

In this section we give some basic definitions from differential geometry and introduce the notation. The reader is referred to [3] for in depth introduction to the subject.

Surface parameterizations are a convenient tool for analyzing surface properties.

A parameterized surface: We think of a surface, S, in space as a vector-valued map, $\mathbf{f}$, from some two-dimensional domain M into Euclidean three space:

$$\mathbf{f} : M \to \boldsymbol{R}^3, \quad S = \mathbf{f}(M).$$

The domain M is often chosen to be a planar region endowed with some coordinates (u, v) but one can use any smooth 2D manifold.

The differential and the tangent plane: The differential, $d\mathbf{f}_p$, of $\mathbf{f}$ at a point $p \in M$ is a linear map that maps tangent vectors to tangent vectors, i.e., if $\mathbf{u}$ is the velocity (tangent) vector to a curve in M, $d\mathbf{f}_p(\mathbf{u})$ is the velocity (tangent) vector to the image of that curve in $S = \mathbf{f}(M)$,

$$d\mathbf{f}_p : T_p(M) \to T_{f(p)}(S)$$

where $T_p(M)$ and $T_{f(p)}(S)$ denote the tangent planes to the abstract surface M and to the image surface S, respectively. The plane $T_p(M)$ to a surface M at a point p is the linear space, that best approximates the surface at p. A choice of local coordinates in M defines a basis in the tangent plane $T_p(M)$. It is customary to omit the subscript p when discussing the differential or the tangent plane, and so we do, but this should not cause any confusion. It should be understood that all statements are local, i.e. apply to a neighborhood of a point.

The map $\mathbf{f}$ is an immersion if its differential $d\mathbf{f}$ is an isomorphism. In particular, if $\mathbf{f}$ is an immersion, S has a well-defined tangent plane at each point, and a normal at each point if M is oriented.

Oriented surfaces: In this paper, we consider only oriented surfaces, i.e. there is a consistent way of identifying positively oriented frames in the tangent plane.

The Gauss map $\mathbf{N}$ (the surface normal), the Gauss curvature, the mean curvature, and all other differential invariants are expressed in terms of the map $\mathbf{f}$ and its derivatives.

The Gauss map: If M is an abstract oriented two dimensional manifold then the value of the Gauss map at a point $p \in M$ is defied by

$$\mathbf{N} = \frac{1}{\|d\mathbf{f}(\mathbf{v_1}) \times d\mathbf{f}(\mathbf{v_2})\|} d\mathbf{f}(\mathbf{v_1}) \times d\mathbf{f}(\mathbf{v_2})$$

where $(\mathbf{v}_1, \mathbf{v}_2)$ is a positively oriented frame of the tangent plane $T_p(M)$. Here $\times$ is the usual cross product in $\boldsymbol{R}^3$.

In particular, if M is a planar domain with a fixed coordinate system (u, v), then

$$d\mathbf{f} = \frac{\partial \mathbf{f}}{\partial u} du + \frac{\partial \mathbf{f}}{\partial v} dv,$$

and the Gauss map is the vector-valued function

$$\mathbf{N} = \frac{1}{\|\frac{\partial \mathbf{f}}{\partial u} \times \frac{\partial \mathbf{f}}{\partial v}\|} \frac{\partial \mathbf{f}}{\partial u} \times \frac{\partial \mathbf{f}}{\partial v}. \tag{1}$$

In general, it is convenient to think of the Gauss map as a map from M to the unit sphere, $\mathbf{S^3}$, $\mathbf{N} : M \to \mathbf{S^3} \subset \boldsymbol{R}^3$.

The fish-scales method designed by Šára and Bajcsy in [8] gives estimates of the Gauss map and the conformal structure.

A conformal structure and a complex structure induced by a parameterization:

A conformal structure on a surface is a choice of angles between tangent vectors. On an oriented surface, a conformal structure is equivalent to defining the operation, J, of rotating tangent vectors by ninety degrees counterclockwise in the tangent plane. This operation is also called a complex structure. A surface parameterization, $\mathbf{f} : M \to \boldsymbol{R}^3$, defines a complex structure J_f on the domain M. Indeed, let $\mathbf{v}$ be a vector tangent to M at some point $p \in M$, then $J_f(\mathbf{v})$ is the unique vector tangent to the domain satisfying

$$d\mathbf{f}(J_f(\mathbf{v})) = \mathbf{N} \times d\mathbf{f}(\mathbf{v}).$$

Thus the defining relation for the complex structure J_f is

$$d\mathbf{f} \circ J_f = \mathbf{N} \times d\mathbf{f}. \tag{2}$$

The differential invariants mean curvature, Gauss curvature, principal axes and principal curvatures: Recall that the second fundamental form of $\mathbf{f}$ is a symmetric quadratic form defined by $\mathrm{I\!I}(\mathbf{u},\mathbf{v}) = - < d\mathbf{N}(\mathbf{u})|d\mathbf{f}(\mathbf{v}) >$ where $< \cdot|\cdot >$ is the Euclidean scalar product in $\boldsymbol{R}^3$. At every point $p \in M$ there exists a positively oriented orthonormal frame $\{\mathbf{e_1}, \mathbf{e_2} = J_f(\mathbf{e_1})\}, \|d\mathbf{f}(\mathbf{e}_i)\| = 1$, of $T_p(M)$ in which the symmetric quadratic form $\mathrm{I\!I}(\cdot,\cdot)$ is represented by a diagonal matrix

$$\begin{pmatrix} \mathrm{I\!I}(\mathbf{e}_1,\mathbf{e}_1) & \mathrm{I\!I}(\mathbf{e}_1,\mathbf{e}_2) \\ \mathrm{I\!I}(\mathbf{e}_2,\mathbf{e}_1) & \mathrm{I\!I}(\mathbf{e}_2,\mathbf{e}_2) \end{pmatrix}$$

where $\mathrm{I\!I}(\mathbf{e_1},\mathbf{e_2}) = \mathrm{I\!I}(\mathbf{e_2},\mathbf{e_1}) = 0$, and $\mathrm{I\!I}(\mathbf{e_j},\mathbf{e_j}) = \kappa_j$, $j = 1, 2$. The vectors $\mathbf{e_1}$ and $\mathbf{e_2}$ are called principal curvature vectors, they define the *principal axes*, and the numbers κ_1, κ_2 are the *principal curvatures*. The *mean curvature*, H is the average of the principal curvature, and the *Gauss curvature* is the product of the principal curvatures. The curvatures are classical geometrical invariants [7].

3. Conformal Method for Computing the Differential Invariants: Theory

We now present the theoretical results for computing the mean curvature, Gauss curvature, and the principal axes from the Gauss map and the conformal structure. The key theoretical result is the following theorem. All proofs are in the Appendix.

Theorem 1 *Let $\mathbf{N}$ be the Gauss map of a parameterized surface $\mathbf{f} : M \to \boldsymbol{R}^3$ and let J_f be the induced complex structure. If $\mathbf{f}$ is twice continuously differentiable then the differential $d\mathbf{N}$ of the Gauss map satisfies*

$$d\mathbf{N} = -H d\mathbf{f} + \omega, \tag{3}$$

where H is the mean curvature, and ω is a $\boldsymbol{R}^3$-valued one form from the tangent plane to the Euclidean three space,

$$\omega : T(M) \to \boldsymbol{R}^3, \tag{4}$$

such that, for every vector $\mathbf{v}$ tangent to the domain M the image $\omega(\mathbf{v})$ satisfies

$$\omega(\mathbf{v}) \perp \mathbf{N} \tag{5}$$

$$\omega\left(J_f(\mathbf{v})\right) = -\mathbf{N} \times \omega(\mathbf{v}). \tag{6}$$

Therefore we have the following corollaries expressing the differential invariants in terms of the Gauss map and the complex structure.

Corollary 3.1 *(Mean curvature) Let $\mathbf{N}$ be the Gauss map of a parameterized surface $\mathbf{f} : M \to \boldsymbol{R}^3$ and let J_f be the induced complex structure. If $\mathbf{f}$ is twice continuously differentiable, then*

$$-H d\mathbf{f} = \frac{1}{2}\left(d\mathbf{N} - \mathbf{N} \times d\mathbf{N} \circ J_f\right) \tag{7}$$

Thus if we have estimates for $d\mathbf{f}$, the Gauss map and the complex structure we can estimate the mean curvature directly from (7).

Corollary 3.2 *(Principal axes) Let $\mathbf{N}$ be the Gauss map of a parameterized surface $\mathbf{f} : M \to \boldsymbol{R}^3$, and let J_f be the induced complex structure. If $\mathbf{f}$ is twice continuously differentiable, then*

$$\omega = \frac{1}{2}\left(d\mathbf{N} + \mathbf{N} \times d\mathbf{N} \circ J_f\right) \tag{8}$$

Furthermore, $\omega(\mathbf{u})$ is collinear to $d\mathbf{f}(\mathbf{u})$ if and only if the vector $\mathbf{u}$ is collinear to a principal curvature vector. Thus the quadratic form $< \omega(\cdot)|d\mathbf{f}(\cdot) >$ is symmetric and trace-free (i.e., has zero trace), and its eigenvalues are precisely $\pm\frac{1}{2}(\kappa_1 - \kappa_2)$, where κ_1 and κ_2 are the principal curvatures.

We can estimate the principal curvature vectors by solving

$$\frac{1}{2}\left(d\mathbf{N}(\mathbf{u}) + \mathbf{N} \times d\mathbf{N}(J_f(\mathbf{u}))\right) = \lambda\, d\mathbf{f}(\mathbf{u}) \tag{9}$$

for the scalar λ and the vector $\mathbf{u}$. This amounts to diagonalizing a symmetric trace free matrix representing the quadratic form $< \omega(\cdot)|d\mathbf{f}(\cdot) >$. The diagonalization of such matrices is more stable than the diagonalization of general matrices.

Corollary 3.3 *(Gauss curvature:) Let $\mathbf{N}$ be the Gauss map of a parameterized surface $\mathbf{f} : M \to \boldsymbol{R}^3$ and let J_f be the induced complex structure. Let H be the mean curvature. Let $\mathbf{f}$ be twice continuously differentiable, ω be the one form defined in (4), and λ^2 be the sum of the squares of the eigen values of the quadratic form, $< \omega(\cdot)|d\mathbf{f}(\cdot) >$. Then, the Gauss curvature, K, satisfies*

$$K = H^2 - \lambda^2. \tag{10}$$

Equation (10) gives a stable method for computing the Gauss curvature K. We do not need to diagonalize the quadratic form matrix. To compute λ^2, we can chose any orthonormal basis of the tangent plane to the surface in $\mathbf{R^3}$, then we represent the quadratic form $< \omega(\cdot)|d\mathbf{f}(\cdot) >$ as matrix A, and set λ^2 as follows

$$A = \begin{pmatrix} a & b \\ b & -a \end{pmatrix},\ \lambda^2 = a^2 + b^2. \tag{11}$$

4 Implementation and Results

Our goal in this paper is to illustrate that the new methods are well suited to handle dense 3D point data generated from stereo. Thus, in this first implementation, we do not care about the real-time issue. In the future we plan to exploit the spatial continuity and reduce the number of redundant calculations.

The inputs to our system are pairs of stereo images, and the outputs are the 3D points and the differential invariants of the surface at these points. For the recovery of the differential invariants we use the conformal method based on the theory presented in the paper. The images are processed by area-based stereo to produce a cloud of 3D samples from a surface. These points are processed by a simplified version of the fish-scales technique introduced in [8] to compute the surface normals and the neighborhood stratification which are then processed by a third module implementing a discretized version of the conformal method. All computations in the third module are local, but to account for the noise in the data, we have taken multiple measurements, i.e. measure differential invariants not in a single direction, but in all available directions (from a point to all its neighbors).

We have tested the system on various surfaces with known differential invariants (including catenoids, spheres, cylinders), and on stereo data.

Figure 1 represents a diagram of the complete system. The particular implementation of the stereo

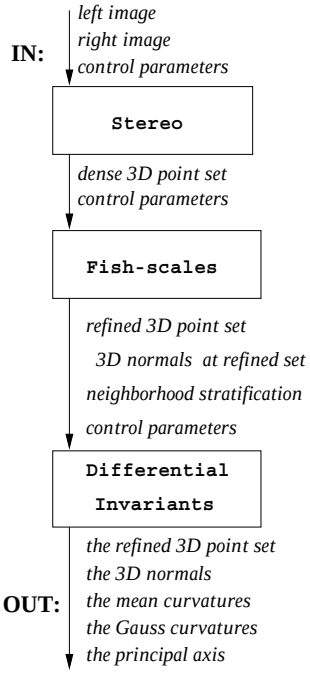

Figure 1: Reconstruction of 3D points and their differential invariants based on stereo.

module is discussed in [1].

We use a fish-scales method for recovery of the Gauss map and the neighborhood stratification from the 3D noisy unorganized set of points. The fish-scales method is introduced in [8], and is used there as a step in the recovery of 3D surface model .

In an implementation we have to deal with discrete versions of the objects defined in the theory section. For each neighborhood, a scale unit is selected adaptivly based on the geometry of the discrete neighborhood. Also, in an ideal continuous world, directional derivatives could be calculated exactly in any direction. Since we are dealing with noisy, discrete data from stereo, calculating differential invariants based on one direction may be catastrophic. Thus we choose to calculate the differential invariants at a point based on as many directions that the neighborhood allows (i.e. from the center point to its neighbors, depending on the size of the neighborhood). In theory the differential invariants should be independent of the direction chosen for the computation. In the real world, because of the discretization and the noise, we obtain different sample values of mean and Gauss curvatures for different directions. We use the multiple sample to our advantage: statistically estimate the differential invariants based on the multiple measurements. The outer `for` loop of `DifferentialInvariants` includes a step that calculates the sample variance of the mean curvature at a point

$$S(\mathbf{H}) = \frac{1}{m-1}\sum_{u}(H_u - H)^2, \qquad (12)$$

where H is the estimated mean curvature, u ranges over all directions we selected based on the neighborhood $\mathcal{U}$, and m is the number of directions selected. The sample variance, $S(\mathbf{H})$ is an unbiased estimate of the variance of a sample from a normal distribution, a model that we chose here for the matter of convenience. We use the sample variance to derive confidence intervals for H. In a final step, we reject the differential invariants of those points for which the the confidence intervals for H are not tight enough at the selected confidence level.

Adaptively, for each neighborhood, we set the local scale unit to be the minimum of two values: the radius, `r`, of the sphere inscribed inside the convex hull of the neighborhood, and an input control scale parameter.

The input to the procedure is the Gauss map `N`, the neighborhood stratification `U`, and an input scale controlling the maximum step size in discrete approximations. We use the conformal method to recover the mean and Gauss curvatures, and the principal axes.

In the current implementation, since the neighborhood is not very densely populated, we use all available directions. In applications with densely populated neighborhoods, a random sample of an appropriate size would suffice. We use the mean square error criterion, and a location data model with Normal sampling distribution. While this model appears to give satisfactory results, its choice is a matter of mathematical convenience at this point.

First, we show results which illustrate the application of the conformal method to data from stereo. Next, for the purpose of evaluation, we show how the method works on randomly selected points from synthetic surfaces.

Recall Figure 1, the system starts with 2D images of a surface, and produces 3D sample points from the

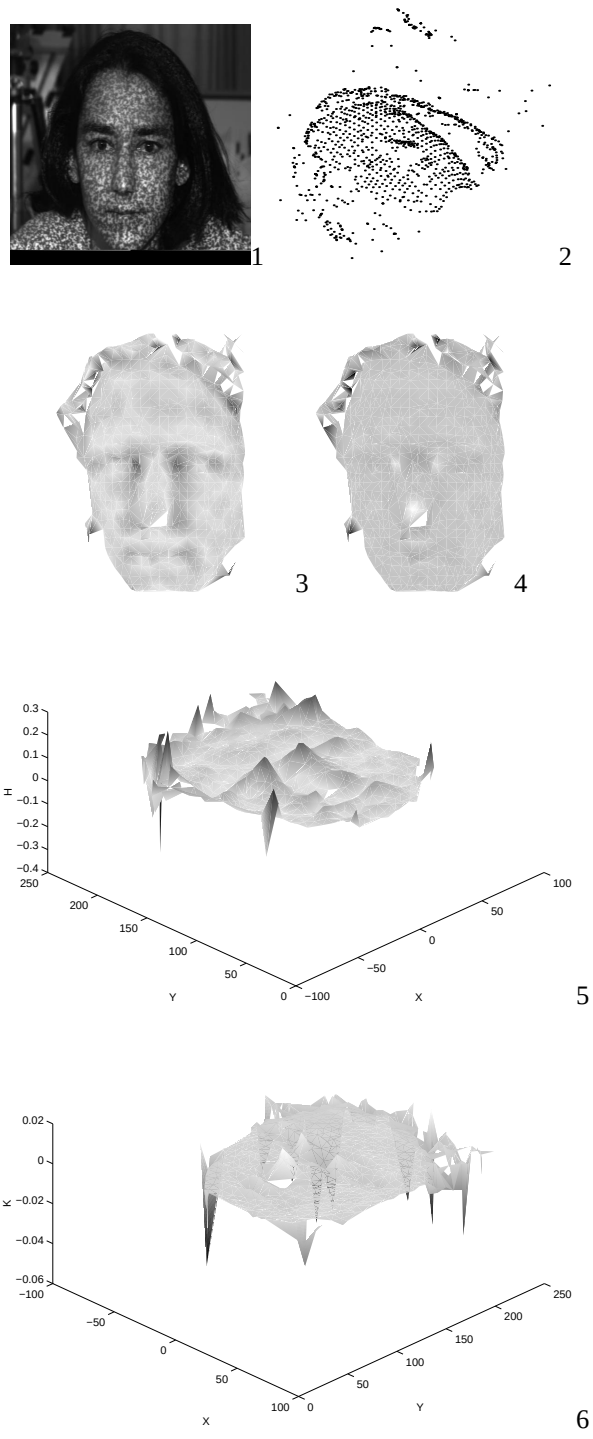

Figure 2: Stereo data, "face". (1) The face(one image from the stereo pair). (2) The reconstructed 3D points, from stereo. Two views of the recovered mean curvature surface (x, y, H), images (3) and (5), and two views of the recovered Gauss curvature surface (x, y, K), images (4) and (6).

surface and the differential invariants at these points. Figure 2 shows the mean and the Gauss curvature surfaces as triangulated meshes. For example, the point (x, y, H), or (x, y, K), on the curvature surface corresponds to a recovered 3D point (x, y, z) from the real face, and the mean curvature at that point is H, or K. The surface height, i.e. H, or K, values, are used to color the point (x, y, H), or (x, y, K).

Figure 3 shows the principal axes at each of the reconstructed points. A trained eye, can spot easily the lines of curvature that are formed by properly aligned principal axes. See the area of the chin, the boundary line going from the chin towards the right ear, and the areas of the eyebrows, for example.

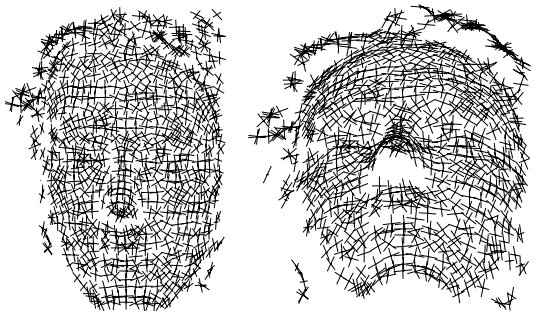

Figure 3: Stereo data "face": the principal axes attached at the 3D points. The 3D data set is shown from different view points. We can identify well-formed curvature lines in nonplanar regions (on the chin, the eyebrow areas, the cheeks)

We present results illustrating the performance of the conformal method on inputs from simulated surfaces for which the precise values of the invariants are known.

In addition, for a comparison, we calculated the mean curvatures on the same data sets using the classical definitions [3], i.e. using the eigenvalues of the second fundamental form (Section 2). Methods for calculating differential invariants in computer vision applications rely mostly on these definitions. They do not use the closed expressions provided by our conformal method.

The mean curvature is selected for the comparison since in our method the Gauss curvature is derived using the mean curvature (i.e. the mean curvature is calculated first).

We show results for three different surfaces here: a sphere (constant mean and Gauss curvatures, all vectors are principal axes), a catenoid (0 mean curvature, negative Gauss curvature, and well defined principal axes), and a cylinder (constant mean curvature, 0 Gauss curvature, and well defined principal axes).

In each of the three cases, we sample uniformly, at random the 3D surface (using the analytic surface representation). We use the exact normals (for the sphere and the cylinder), and normals recovered by the fish-scales method (for the catenoid). The neighborhood

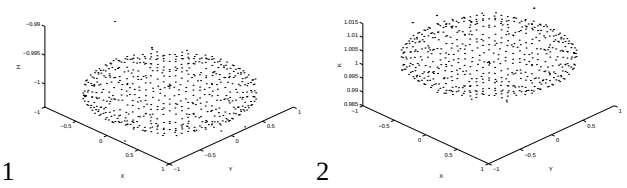

1 2

Figure 4: Sphere: (1) and (2) are the calculated points of the mean curvature surface, (x, y, H), and the Gauss curvature surface, (x, y, K), respectively.

stratification is obtained by a Delaunay triangulation.

The simulated input data to the conformal method could be separated in two groups: (i) input data that represent exact points and normals from a synthetic surface (a sphere and a cylinder); (ii) input data that represent approximated points and normals from the synthetic surface (catenoid).

Sphere. We sample randomly points from the unit sphere, (x, y, z), $y > 0$

$$x^2 + y^2 + z^2 = 1.$$

The surface is oriented by the outward normal. The counterclockwise direction in the tangent plane is selected to be the positive orientation (this is a standard choice in computer vision and computer graphics). Under this choice of orientation, the unit sphere has a constant mean curvature -1 and a constant Gauss curvature 1. *We remind the reader that in standard differential geometry texts, [3], the orientation is selected in such a way that the mean curvature of the sphere is positive, i.e.* 1*; for that orientation the inward normals to the surface are selected.* The principal curvatures of the sphere are both equal to one, so the sphere does not have well-defined principal curvature vector fields.

Figure 4 shows the recovered mean and Gauss curvatures.

In this experiment 97.75% of the mean curvature values are within 0.0022 from the true mean curvature value -1; 97.7% of the Gauss curvature values are within 0.0044 from the true value of the Gauss curvature 1. Note also that since the points on the surface are elliptic, by using linear interpolation in intermediate steps we always approximate surface points by points inside the sphere, which explains the underestimation for the mean curvature (i.e. the values are within 0.0022 of the true values and less than it), and overestimating the Gauss curvature.

Table 5, part A, contains statistics of the absolute mean curvature errors for the conformal and the standard methods.

Cylinder. We sample the half cylinder, (x, y, z),

$$x^2 + z^2 = 1,\ 0 \leq y \leq 1.$$

The surface oriented the by the outward normals. Under this choice of orientation, the cylinder has a constant mean curvature -0.5, and a constant Gauss curvature 0. The principal curvatures of the cylinder are equal to -1 and 0, respectively; the principal axes are such that one of the axes is always collinear with the generating straight line, in this case collinear with the y axes.

Test	MIN	MAX	MEAN	STD
A.1	0	0.094175	0.001761	0.009923
A.2	0.000134	0.006289	0.001470	0.000448
A.3	0.883836	0.995832	0.937364	0.019062
A.4	0.000268	0.012616	0.002944	0.000899
B.1	0	0.00019	0.00000	0.00001
B.2	0.00000	0.00202	0.00032	0.00026
B.3	0.43624	0.49910	0.47750	0.013258
B.4	0.0000	0.00002	0.000000	0.000002
C.1	0	0.000763	0.000142	0.000138
C.2	0.000025	0.126203	0.010550	0.018231
C.3	0.000001	0.002044	0.000297	0.000335

Figure 5: Statistics related to the simulated examples. A: for the sphere. B: for the cylinder. C: for the catenoid. The statistics in each case are for: (1), the sample variance, $S(H)$ used for rejection of outliers; (2), the absolute mean curvature error for the conformal method; (3), the absolute mean curvature error for the standard method; (4), the absolute Gauss curvature error for the conformal method. For the sphere and the cylinder (A and B), we use exact points and normals. For the catenoid (C) , the sampled points are piped through the fish-scales first, and thus the differential invariants are recovered from approximated points and normals. See discussion regarding the catenoid in the text.

Figure 6 shows mean and Gauss curvatures, and the principal axes recovered.

In this experiment 93.5% of the mean curvature values are within 0.0005 from the true mean curvature value $-0,5$; more than 99% of the Gauss curvature values are within 0.00001 from the true value of the Gauss curvature value statistics of the absolute mean curvarure errors, and the sample variances, $S(\mathbf{H})$, used for rejection of outlayers the table, Figure 5. In this case the mean curvature computed by the standard method has large residual error.

In this experiment the orientation of the principal axes was recovered correctly.

Catenoid. In this example we want to illustrate the sensitivity of the conformal method to noise in estimates of the Gauss map. We sample the catenoid

$$\begin{aligned} x &= -\cosh(t)\cos(\theta) \\ y &= t \\ z &= -\cosh(t)\sin(\theta) \end{aligned}$$

in the interval $-1 \leq t \leq 1$, $\pi/3 \leq \theta \leq 2\pi/3$. The surface orientation is by the outward normals. Under this choice of orientation, the catenoid has a constant mean curvature 0, and negative Gauss curvature.

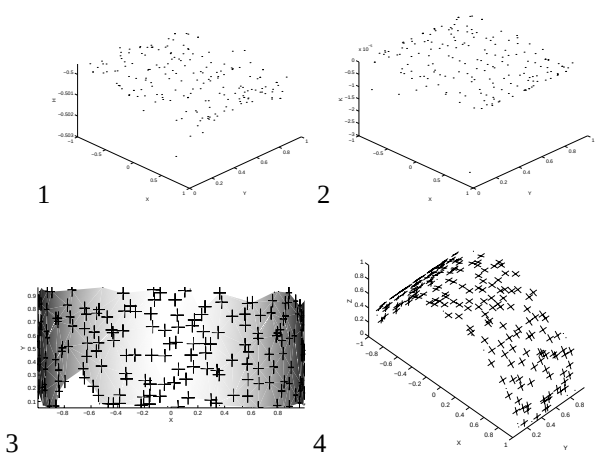

Figure 6: Cylinder. (1) and (2) the calculated points of the mean curvature surface, (x, y, H), and the Gauss curvature surface, (x, y, K). (3) the principal axes superimposed on the cylindrical surface which is shaded according to depth values (i.e. z). (4) the principal axes at the sampled points, shown from a different viewpoint.

In this experiment we use the refined 3D point set and the Gauss map output by the fish-scales method, not exact points and normals from the synthetic surface.

Figure 7 shows mean and Gauss curvatures, and principal axes recovered.

In this experiment 99% of the mean curvature values are within 0.032 from 0. See statistics given in the table, Figure 5. For this one particular example the standard method produces results closer to zero. Although we include the error statistics for completeness, note that we do not know the mean curvature of the model surface presented by the approximated points and normals. In this sense one can use these statistics to gauge the performance of the complete systems, but can not use the statistics to compare the performance of the conformal v.s. the standard method for estimating mean curvature. Note that the estimated Gauss normal still exhibits the rotational symmetry of the catenoid. The conformal method produced correct curvature lines indicated be the proper alignment of the recovered principal axes Figure 7, subfigure 5.

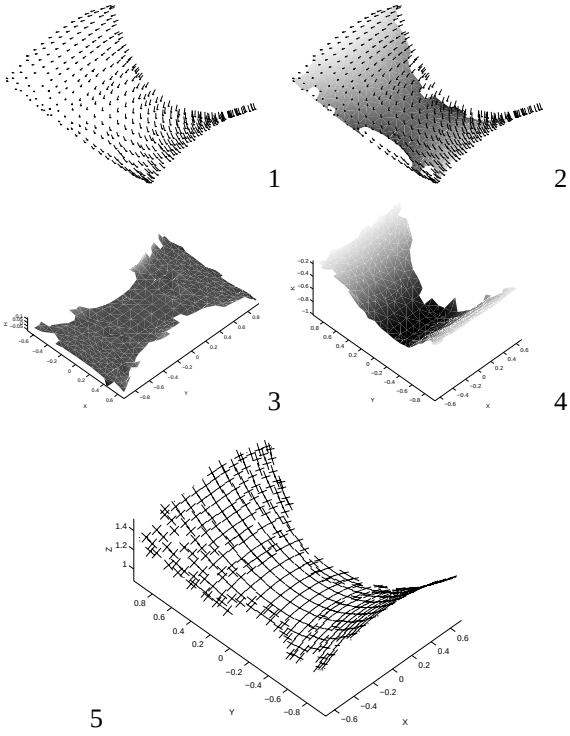

Figure 7: Catenoid: (1)Gauss map; (2) The Gauss map and the part of recovered surface after the rejection of outliers based on the sample variance; (3) and (4) the recovered mean curvature surface, (x, y, H), and the Gauss curvature surface, (x, y, K); (5) the principal axes superimposed on the catenoid surface.

Conclusions

We present a new conformal method for calculating the differential invaraints of a surface from 3D points and normals. We derive the theory, and implement the algorithms. We compute a measure of the precision/reliability of the recovered invariants in terms of the variance of the mean curvature (propagated from the variability of the local neighborhood). We show the results produced by the conformal method for surfaces with known differential invariants (a sphere and a cylinder), for surfaces with approximately known differential invariants (catenoid), and for surfaces with unknown invariants.

In the first group, sphere and cylinder, we evaluate the results of the conformal method in terms of the absolute mean curvature and Gauss curvature errors, and in terms of the angle between the principal axes and the known constant axes direction (i.e. the 'y' axes). We also compared the results obtained by the conformal method with results obtained using the classical definitions directly. The conformal method is more stable.

The second group of results (that includes the catenoid) illustrates the effect of the Gauss map errors on the differential invariant. More careful study in this direction is needed (including noise model propagation from normals to differential invariants), and is subject of our future research.

The third group of results (that includes the face) illustrates the performance of the conformal method on 3D points recovered from stereo, and normals recovered by the fish-scales method. Results here are evaluated only qualitatively and support the usefulness of the conformal approach in deriving differential invariants from real stereo data.

Appendix

The fastest method to prove Theorem 1 and its corollaries is to use quaternionic calculus but here we provide proofs along the lines of classical differential geometry.

Recall the definitions of the differential invariants from Section 2. Note that e_1 and e_2 are principal curvature directions iff

$$d\mathbf{f}(\mathbf{e_1}) \perp d\mathbf{f}(\mathbf{e_2}) \quad (13)$$

$$d\mathbf{N}(\mathbf{e_1}) = -\kappa_1 d\mathbf{f}(\mathbf{e_1}) \quad (14)$$

$$d\mathbf{N}(\mathbf{e_2}) = -\kappa_2 d\mathbf{f}(\mathbf{e_2}) \quad (15)$$

Proof of Theorem 1 The form of equations (14) and (15) suggests that the one form $d\mathbf{N}$ can be represented as

$$d\mathbf{N} = A\, d\mathbf{f} + \omega \quad (16)$$

for some coefficient A and some $\boldsymbol{R}^3$-valued one-from ω. We decide to look for a form ω satisfying the condition

$$\omega(J_f(\mathbf{u})) = -\mathbf{N} \times \omega(\mathbf{u}). \quad (17)$$

This choice for ω can be motivated by the decomposition of symmetric tensors into diagonal and trace-free components. A direct way to motivate our choice is to notice that the form $d\mathbf{f}$ satisfies $d\mathbf{f}(J_f(\mathbf{u})) = \mathbf{N} \times d\mathbf{f}(\mathbf{u})$. That is, $d\mathbf{f}$ relates a counter-clockwise rotation by ninety degrees in $T_p(M)$ to a counterclockwise rotation by ninety degrees around the axes $\mathbf{N}$ in $\boldsymbol{R}^3$. On the other hand, the condition (17) guarantees that the form ω relates a counter-clockwise rotation by ninety degrees in $T_p(M)$ to a clockwise rotation by ninety degrees around the axes $\mathbf{N}$ in $\boldsymbol{R}^3$. The representation (16) accounts for the possibility that the one-form $d\mathbf{N}$ may be a combination of forms which rotate in different directions around the N axes. From (16), (14), and (15) we obtain $\omega(\mathbf{e_i}) = (\kappa_i + A)\, d\mathbf{f}(\mathbf{e_i}), \quad i = 1, 2$. These identities show that (5) holds. Furthermore combining the identities with (17) and $\mathbf{e_2} = J_f(\mathbf{e_1})$ we obtain

$$\begin{aligned} \omega(\mathbf{e_2}) &= (\kappa_2 + A)\, d\mathbf{f}(\mathbf{e_2}) \\ -\mathbf{N} \times \omega(\mathbf{e_1}) &= \mathbf{N} \times (\kappa_2 + A)\, d\mathbf{f}(\mathbf{e_1}) \\ -\mathbf{N} \times (\kappa_1 + A)\, d\mathbf{f}(\mathbf{e_1}) &= \mathbf{N} \times (\kappa_2 + A)\, d\mathbf{f}(\mathbf{e_1}) \end{aligned}$$

The last identity implies that the tangential vector $(\kappa_1 + \kappa_2 + 2A)\, d\mathbf{f}(\mathbf{e_1})$ is colinear to the normal $\mathbf{N}$. This can only happen if it is the zero vector in $\boldsymbol{R}^3$, that is,

$$A = -\frac{1}{2}(\kappa_1 + \kappa_2) = -H.$$

Proof of Corollaries 3.1,3.2,3.3. From (3), (2), and (17) we get

$$\begin{aligned} \mathbf{N} \times d\mathbf{N} \circ J_f &= -H\mathbf{N} \times (\mathbf{N} \times d\mathbf{f}) - \mathbf{N} \times (\mathbf{N} \times \omega) \\ &= H\, d\mathbf{f} + \omega. \end{aligned}$$

The identities (7) and (8) follow directly from (3) and the identity $\mathbf{N} \times d\mathbf{N} \circ J_f = H\, d\mathbf{f} + \omega$. Rewriting (3) in the form $\omega = d\mathbf{N} + H\, d\mathbf{f}$ we conclude that $\omega(\mathbf{u})$ is colinear to $d\mathbf{f}(\mathbf{u})$ if and only if the later is colinear to $d\mathbf{N}(\mathbf{u})$, that is, if and only if $\mathbf{u}$ is parallel to a principal curvature vector. Furthermore, from equations (14) and (15) we obtain

$$\omega(\mathbf{e_1}) = \frac{k_2 - k_1}{2} d\mathbf{f}(\mathbf{e_1}), \omega(\mathbf{e_2}) = \frac{k_1 - k_2}{2} d\mathbf{f}(\mathbf{e_2})$$

References

[1] R. Bajcsy, R. Enciso, G. Kamberova, L. Nocera, R. Šara, "3D Reconstruction of Environments for Virtual Collaboration", *proc. 4th IEEE Workshop on Allications of Computer Vision*, Princeton, NJ, Oct 1998.

[2] P. Besl and R. Jain. Invariant surface characteristics and 3d object recognition in range images. *CVGIP*, 33:33–80, 1986.

[3] M. DoCarmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, 1976.

[4] F. Devernay, "Computing Differential Properties of 3-D Shapes from Stereoscopic Images without 3-D Models", INRIA, RR-2304, Sophia Antipolis, 1994.

[5] J. Feldmar and N. Ayache. *Registration of Smooth Surfaces Using Differential Properties*. Number 801. Springer-Verlag, 1994.

[6] K. Ikeuchi and M. Herbert. Spherical Representations: from EGI to SAI CMU Technical report: CMU-CS-95-197 (1995)

[7] J. Koenderink. *Solid Shape* MIT press, 1990

[8] R. Šára and R. Bajcsy, "Fish-Scales: Representing Fuzzy Manifolds," *Proc. Int. Conference on Computer Vision*,Bombay, India,Narosa Publishing House, 1998.

[9] H. Shum, M. Herbert, and K. Ikeuci. On 3d shape synthesis. In *Proc. Image Under. Workshop*, volume 2, pages 1103–1112, 1996.

[10] G. Taubin. Estimating the tensor of curvature of a surface from a polyhedral approximation. In *Proc. ICCV*. IEEE Comp. Soc. Press, 1995.

[11] R. Woodham. Photometric method for determining surface orientation from multiple images. *J. Opt. Eng.*, 19(1):139–144, 1980.

[12] R. Woodham. Gradient and curvature from the photometric-stereo method, including local confidence estimation. *Jour. Opt. Soc. of Am. A*, 11(11):3050–3068, 1984.

[13] P. Yuen et al. Curvature and Torsion Feature Extraction from 3-D Meshes and Multiple Scales . *IEE Proc. Vis. Image Signal Process*, volume 147, No 5, 454–462 (2000).