Shape Invariants and Principal Directions from 3D Points and Normals

George Kamberov

Gerda Kamberova

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 is presented. 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 approach to the estimation of the differential invariants in computer vision is to reconstruct an approximation of the surface and its parameterization, and then derive the differential invariants using classical differential geometry. First, range sensors, stereo, MRI, or CT techniques are used to collect a sufficiently dense set of sampled surface points; second, an approximate parameterization 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 approach is to forgo the 3D model, and 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.

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 $S = \mathbf{f}(M)$,

$$d\mathbf{f}_p : T_p(M) \to T_{f(p)}(S) := df_p\,(T_pM) \subset \boldsymbol{R}^3.$$

where $T_p(M)$ denotes the tangent plane to the abstract surface M. The tangent plane $T_p(M)$ is the linear space, that best approximates the surface M at the point 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.

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 $\mathrm{II}(\cdot,\cdot)$ is represented by a diagonal matrix

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

where $\mathrm{II}(\mathbf{e_1}, \mathbf{e_2}) = \mathrm{II}(\mathbf{e_2}, \mathbf{e_1}) = 0$, and $\mathrm{II}(\mathbf{e_j}, \mathbf{e_j}) = \kappa_j$, $j = 1, 2$. 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*

$$d\mathbf{N} = -H d\mathbf{f} + \omega, \qquad (3)$$

where H *is the mean curvature, and* ω *is a* $\boldsymbol{R}^3$*-valued 1-form,*

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

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

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

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

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*

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

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*

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

Furthermore, $\omega(\mathbf{u})$ *is colinear 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

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

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*

$$K = H^2 - \lambda^2. \qquad (10)$$

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

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

4 Implementation and Results

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. compute 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 real stereo data.

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

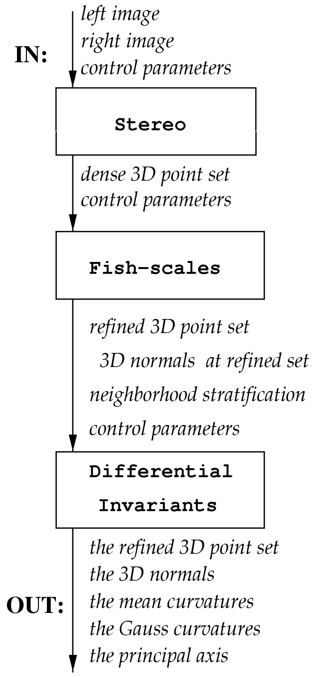

Figure 1: Reconstruction of 3D points and their differential invariants based on stereo.

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 .

The input to the Differential invariants module (Figure 1) 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 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 adaptively 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. We calculate 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 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.

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 stereo images of a surface, goes through stereo reconstruction, Gauss map computation, and finally differential invariants recovery, and produces 3D sample points from the surface, the mean and Gauss curvatures, and the principal axes at these points. Figure 2 shows one input image, the recovered 3D points, and the mean and the Gauss curvature recovered. The curvature data are displayed as shaded 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,