

Principal curvatures from the integral invariant viewpoint.

Helmut Pottmann, Johannes Wallner, Yong-Liang Yang, Yu-Kun Lai, and Shi-Min
Hu

Project Area(s):

Industrielle Geometrie - Algorithmische Differentialgeometrie

Institut für Geometrie

Principal curvatures from the integral invariant viewpoint

Helmut Pottmann^{1*} and Johannes Wallner²
Yong-Liang Yang³, Yu-Kun Lai³, and Shi-Min Hu³

¹ *TU Wien* ² *TU Graz* ³ *Tsinghua University, Beijing*

Abstract

The extraction of curvature information for surfaces is a basic problem of Geometry Processing. Recently an integral invariant solution of this problem was presented, which is based on principal component analysis of local neighbourhoods defined by kernel balls of various sizes. It is not only robust to noise, but also adjusts to the level of detail required. In the present paper we show an asymptotic analysis of the moments of inertia and the principal directions which are used in this approach. We also address implementation and, briefly, robustness issues and applications.

Key words: integral invariants, principal curvatures, robustness

1 Introduction

The role of differential geometry in the investigation of curves and surfaces is a very important one, and geometry processing tasks frequently require information about properties which for smooth surfaces are obtained by differentiation – normal vectors, curvatures, principal directions. Also the global understanding of shapes benefits from differential entities, as exemplified by the network of principal curvature lines, and by the crest lines of a surface. References on these topics include Alliez et al. (2003), Cazals and Pouget (2005), Hildebrandt et al. (2005), Ohtake et al. (2004), and Yoshizawa et al. (2005).

* corresponding author. Address: Wiedner Hauptstr. 8–10/104, A-1040 Wien, Austria. Tel. (+43)1-58801-11310. Fax. (+43)1-58801-11399. email pottmann@geometry.tuwien.ac.at.



Fig. 1. Feature extraction on multiple scales using principal component analysis on ball neighbourhoods of different radii. Darker regions are classified as features on all scales, lighter shaded regions correspond to features extracted at only one or two scales. The images from left: Ravines – Original data — Ridges.

Real-world data, which are e.g. obtained by laser scanning, frequently exhibit too much noise for straightforward numerical differentiation to make sense. Thus, the use of differential invariants requires data smoothing and de-noising prior to computation. This may be done in a global way via appropriate geometric flows (cf. the work by Bajaj and Xu (2003), Clarenz et al. (2004b), and Osher and Fedkiw (2002)). Local methods, using smooth approximations of the data in an appropriate neighbourhood, are presented by Cazals and Pouget (2003), Goldfeather and Interrante (2004), Ohtake et al. (2004), Taubin (1995), and Tong and Tang (2005). In either case, the preservation of features which may not be considered as noise is not an easy task and requires especially adapted algorithms.

Classical differential geometry cannot be used directly for frequently occurring data types like triangle meshes. This led to the development of *discrete differential geometry*. This is a highly interesting and practically useful area which gained increasing attention over the past years. It offers both precise statements on the given geometry and elegant extensions of the classical theory. For an introduction with a focus on Computer Graphics, see Desbrun et al. (2005). As to the topic of the present paper, it is possible to employ discrete differential geometry concepts to extract curvatures at multiple scales: It can be done by a multiresolution mesh representation. This method requires

certain changes of the underlying geometry, but there are others which do not need this complication. Handling noisy data is possible (cf. Hildebrandt and Polthier (2004)), but is neither the main intent nor strength of discrete differential geometry.

Recently Yang et al. (2006) have presented a solution of the problem of robust and multiscale computation of principal curvatures via *integral invariants* obtained by integration over local neighbourhoods. This is an approach initiated by Manay et al. (2004) and Clarenz et al. (2004b,a). The present paper serves as a theoretical foundation for the work by Yang et al. (2006), which presents numerical and experimental results.

Our method is based on principal component analysis (PCA) of local neighbourhoods defined via kernels of variable size. It leads in a natural way to families of geometry descriptors dependent on a kernel radius r (the ‘scale’), and which for smooth surfaces converge to the curvatures of the surface in question, as $r \rightarrow 0$. The kernel radius serves as an approximate threshold size which distinguishes noise from features. The use of integration rather than differentiation has a smoothing effect.

The aim of the present paper is not to present yet another method for curvature estimation, but rather to investigate a new tool which takes a more *global view*. It computes integrated curvature-like quantities over the chosen neighbourhoods. Only for very small kernel radii the invariants computed in this way are closely related to curvatures. Their behavior on larger scales is different, but consistent and useful for applications.

1.1 Prior work on integral invariants, principal curves and feature extraction

While PCA has been used to obtain shape characteristics for a very long time (see e.g. Taubin and Cooper (1992) for applications in Computer Vision), local integral invariants are a rather new topic in geometric computing. Manay et al. (2004) investigate integral invariants for curves in the plane and show their superior performance on noisy data, especially for the reliable retrieval of shapes from geometric databases. A special case of an integral invariant, defined for 2D curves or 3D surfaces, has been used by Connolly (1986) for molecular shape analysis. In the case of a smooth surface Φ which has the property that it occurs as the boundary of a suitable domain D , Connolly considers the surface area $A_s(r, \mathbf{p})$ of the spherical patch neighbourhood

$$N_s(r, \mathbf{p}) := D \cap S(r, \mathbf{p})$$

(see Fig. 2); here \mathbf{p} is a point of Φ and $S(r, \mathbf{p}) = \{\mathbf{x} \mid \|\mathbf{x} - \mathbf{p}\| = r\}$ denotes the sphere of radius r centered at \mathbf{p} . Equation (2), which describes the relation of




Fig. 2. From left to right: Kernel ball $B(r, \mathbf{p})$ and surface Φ ; ball neighbourhood $N_b(r, \mathbf{p})$; sphere neighbourhood $N_s(r, \mathbf{p})$; surface patch neighbourhood $N_p(r, \mathbf{p})$; spherical intersection curve $c_s(r, \mathbf{p})$.

the function $A_s(r, \mathbf{p})$ to mean curvature in a precise way, has been derived by Cazals et al. (2003). Gelfand et al. (2005) use the volume $V_b(r, \mathbf{p})$ of the ball neighbourhood

$$N_b(r, \mathbf{p}) := D \cap B(r, \mathbf{p})$$

to obtain a geometry descriptor useful for finding correspondences in matching problems. Here $B(r, \mathbf{p}) = \{\mathbf{x} \mid \|\mathbf{x} - \mathbf{p}\| \leq r\}$ is the ball with radius r and center \mathbf{p} . Both functions A_s and V_b turn out to be related to mean curvature (cf. Hulin and Troyanov (2003); Cazals et al. (2003); Gelfand et al. (2005)):

$$V_b(r, \mathbf{p}) = \frac{2\pi}{3}r^3 - \frac{\pi}{4}H(\mathbf{p})r^4 + O(r^5), \quad (1)$$

$$A_s(r, \mathbf{p}) = 2\pi r^2 - \pi H(\mathbf{p})r^3 + O(r^4). \quad (2)$$

Here $H(\mathbf{p})$ is the mean curvature in the point \mathbf{p} . Most closely related to the present work are the papers of Clarenz et al. (2004a,b), where integral invariants are employed for feature detection. The authors consider the surface patch neighbourhood

$$N_p(r, \mathbf{p}) = \Phi \cap B(r, \mathbf{p})$$

and perform PCA on this patch. They show that the distance between \mathbf{p} and the barycenter of $N_p(r, \mathbf{p})$ is related to mean curvature, and discuss some scaling properties of the eigenvalues coming up in PCA. The eigenvalues resulting from PCA on the set $N_p(r, \mathbf{p})$ have also been used by Pauly et al. (2003) for multi-scale feature extraction on point-sampled surfaces.

Globally defined features computable from principal directions are the principal curvature lines. These curves proved to be useful for the detection of special shapes, and also for such diverse applications as the work on non-isotropic remeshing by Alliez et al. (2003), and on line-art rendering of smooth surfaces by Hertzmann and Zorin (2000). Global features whose definition relies on even higher order differential invariants (like derivatives of principal curvatures) also carry important shape information: For instance, the curves known as feature lines and crest lines received a lot of attention in the geometry processing community (see Cazals and Pouget (2005), Clarenz et al. (2004b), Hildebrandt et al. (2005), Ohtake et al. (2004), Pauly et al. (2003),

and Yoshizawa et al. (2005)).

1.2 Contributions and overview

The methods analyzed in this paper were presented for the first time in a short paper by Yang et al. (2006), which also contains numerical results and comparison with other methods. The present paper is an extension of that work, and provides proofs and theoretical background.

The main idea of our analysis is to compute quantities like volumes and covariance matrices of small sets associated with points on a smooth surface Φ by way of Taylor expansion; and to express the coefficients of these expansions in terms of Taylor coefficients of Φ . As the latter Taylor coefficients carry curvature information, it is possible to give asymptotic relations between the geometry descriptors computed here and the curvatures of Φ .

The geometry descriptors computed by our methods make sense for noisy surfaces as well, even if the analysis is based on smooth surfaces. This is because they are computed by integration (i.e., averaging), and therefore the values we obtain for a noisy surface Φ' agree with the values of some hypothetical smooth surface Φ which approximates Φ' . When we compute a ‘curvature at scale r ’ for Φ' , we actually compute that curvature for the surface Φ ; and it is the surface Φ which the analysis in the present paper applies to. A thorough robustness analysis of geometry descriptors obtained in this way is given in the forthcoming paper by Pottmann et al. (2006), while the experimental validation of our claims regarding robustness is presented by Yang et al. (2006). Our paper is organized as follows:

- Section 2 presents a thorough study of PCA of four types of neighbourhoods (see Fig. 2) which can be defined by means of a kernel ball or kernel sphere. The precise relation of the quantities obtained by PCA to the curvatures of the underlying surface is obtained by an asymptotic analysis as the kernel radius tends to zero.
- Integral invariants lead to the definition of principal curvatures at a given resolution level r which are consistent with the classical theory ($r \rightarrow 0$). This is discussed in Sec. 3. The multiscale behaviour of these modified principal curvatures and comparison of this method with others is the topic of the short paper by Yang et al. (2006), as is the computation of principal curves at a given scale and multi-scale feature extraction. We only briefly discuss these applications here.
- Section 4 discusses implementation issues. We use FFT and hierarchical refinement of sphere triangulations.

2 Principal component analysis of local neighbourhoods

Principal component analysis of a point set A requires the computation of its volume $V_A = \int_A d\mathbf{x}$, its barycenter $\mathbf{s}_A = \frac{1}{V_A} \int_A \mathbf{x} d\mathbf{x}$, and its covariance matrix

$$J_A = \int_A (\mathbf{x} - \mathbf{s}_A) \cdot (\mathbf{x} - \mathbf{s}_A)^T d\mathbf{x} = \int_A \mathbf{x}\mathbf{x}^T d\mathbf{x} - V_A \mathbf{s}_A \mathbf{s}_A^T. \quad (3)$$

We use column vector notation, so $\mathbf{x}\mathbf{x}^T$ is a 3 by 3 matrix of rank 1. If A is considered as a subset of Euclidean space, then $d\mathbf{x}$ is the usual volume element. If A is contained in a smooth (rectifiable) surface, then $d\mathbf{x}$ means the surface area element. If A is contained in a smooth (rectifiable) curve, then $d\mathbf{x}$ is the arc length element. When we use an orthonormal coordinate system of eigenvectors, J_A takes the form $J_A = \text{diag}(J_{xx}, J_{yy}, J_{zz})$, where J_{xx} , J_{yy} and J_{zz} are the eigenvalues of J_A . The principal moments of inertia of the set A then have the form $J_1 = J_{yy} + J_{zz}$, $J_2 = J_{zz} + J_{xx}$, $J_3 = J_{xx} + J_{yy}$. In the following text we don't refer to principal moments, but only to eigenvalues of the covariance matrix.

In the cases studied by the present paper, the set A is a small set near a certain point \mathbf{p} . The point \mathbf{p} is contained in the surface under consideration and signifies the location where we want to extract geometry descriptors. The size of A is defined by a certain kernel radius. We will usually write $A(r, \mathbf{p})$ to indicate the dependence on the point and the radius, and we will give the first terms in the Taylor expansions of the functions $r \mapsto V_A(r, \mathbf{p})$, $r \mapsto \mathbf{s}_A(r, \mathbf{p})$, and $r \mapsto J_A(r, \mathbf{p})$.

2.1 The principal frame and the simplification of volume integrals

For theoretical investigations it is convenient to work in a coordinate system associated with a point on the surface. We repeat well known results here – they can be found e.g. in do Carmo (1976). For any given point of a C^2 surface there is the so-called principal frame, with respect to which the surface is realized as the graph of the function

$$z(x, y) = \frac{1}{2}(\kappa_1 x^2 + \kappa_2 y^2) + O(\rho^3) \quad (\rho^2 = x^2 + y^2). \quad (4)$$

Here κ_1 and κ_2 denote the principal curvatures of the surface in the point of interest, which is located at the origin of the principal frame. The vector $(0, 0, 1)$ is a normal vector of the surface, and we use the convention that whenever the surface occurs as boundary of a domain D , it points towards the inside of D .

We use the notation $\rho = \sqrt{x^2 + y^2}$ also later in the text. Further, we let $\mathbf{x} = (x, y, z)$. Consider now the general volume integral

$$I(r, f) = \int_{\|\mathbf{x}\| \leq r, z \geq z(x, y)} f(\mathbf{x}) d\mathbf{x}, \text{ where } f : \mathbb{R}^3 \rightarrow \mathbb{R}^k. \quad (5)$$

The value of $I(r, f)$ then is contained in \mathbb{R}^k , depending on the function f . The next theorem shows how to approximately compute such expressions. We use the notation $B(r)^+$ for the half ball $x^2 + y^2 + z^2 = \rho^2 + z^2 \leq r^2, z \geq 0$.

Theorem 1 *If the function f is of magnitude $O(\rho^k z^l)$, then its integral $I(r, f)$ over the domain $x^2 + y^2 + z^2 \leq r^2, z \geq z(x, y)$ can be expressed as*

$$I(r, f) = \tilde{I}(r, f) + O(r^{k+2l+5}), \text{ where}$$

$$\tilde{I}(r, f) = \int_{B(r)^+} f(\mathbf{x}) d\mathbf{x} - \int_{x^2+y^2 \leq r^2} \left(\int_{z=0}^{\frac{1}{2}(\kappa_1 x^2 + \kappa_2 y^2)} f(\mathbf{x}) dz \right) dx dy. \quad (6)$$

Proof. First, the integral $I(r, f)$ is approximated by

$$I(r, f) \approx \int_{B(r)^+} f(\mathbf{x}) d\mathbf{x} - \int_{x^2+y^2 \leq r^2} \left(\int_{z=0}^{z=z(x, y)} f(\mathbf{x}) dz \right) dx dy. \quad (7)$$

The error we make is the volume integral of f over that part \tilde{A} of the original domain which lies outside the sphere $x^2 + y^2 + z^2 = r^2$, but still inside the cylinder $x^2 + y^2 = r^2$. The integral $\int_{\tilde{A}} f$ is bounded by $\text{Vol}(\tilde{A}) \cdot \max_{\tilde{A}} |f|$. The set \tilde{A} has circumference $\approx 2\pi r$ and in view of (4) its extent in the z direction is $O(r^2)$. It has been shown by Hulin and Troyanov (2003) that the radial thickness of \tilde{A} is of magnitude $O(r^3)$, so $\text{Vol}(\tilde{A}) \sim r \cdot r^2 \cdot r^3 = r^6$. Equation (4) implies that within the domain \tilde{A} , z is of magnitude r^2 . With $f \sim \rho^k \cdot z^l$, we now have an upper bound of the form $|\int_{\tilde{A}} f| \leq O(r^6 \cdot r^k (r^2)^l)$.

We further approximate the integral by neglecting the term $O(r^3)$ in the expression $z(x, y)$, i.e., by replacing the surface by its osculating paraboloid. The error we make here is given by the integral of f over the layer which is bounded by the graph of the function $z(x, y)$ given by (4) and the graph of the same function without the $O(r^3)$ term. I.e., the layer has a thickness of magnitude $O(r^3)$. With $z(x, y) = O(\rho^2)$ we make an error of magnitude $\int_{\phi=0}^{2\pi} \int_{\rho=0}^r O(\rho^3) \cdot O(\rho^k (\rho^2)^l) \cdot \rho d\rho d\phi = O(r^{k+2l+5})$.

To sum up, the two modifications of the integral $I(r, f)$ lead to errors of magnitudes r^{2k+l+6} and r^{2k+l+5} , respectively. This implies that $I(r, f) - \tilde{I}(r, f) = O(r^{2k+l+5})$. \square

2.2 Principal component analysis of the ball neighbourhood

We consider the surface Φ as the boundary of the domain D and perform PCA on the ball neighbourhood $N_b(r, \mathbf{p}) := B(r, \mathbf{p}) \cap D$. The volume $V(r, \mathbf{p})$ of $N_b(r, \mathbf{p})$ is the integral of the constant function 1, which is of magnitude $O(\rho^0 z^0)$. We use (6) to approximate the volume and get precisely the known result of Equation (1), with $H = (\kappa_1 + \kappa_2)/2$. For the barycenter \mathbf{s}_b , we have to compute $\int x d\mathbf{x}$, $\int y d\mathbf{x}$, and $\int z d\mathbf{x}$. The functions x, y are of magnitude $O(\rho^1 z^0)$, whereas the z coordinate function is of magnitude $O(\rho^0 z^1)$. Consequently,

$$\begin{aligned} \mathbf{s}_b(r, \mathbf{p}) &= \frac{1}{V_b(r, \mathbf{p})} (\tilde{I}(r, \mathbf{x}) + O(r^6)) \\ &= \left(\frac{2\pi}{3} r^3 - \frac{\pi}{4} H(\mathbf{p}) r^4 + O(r^5) \right)^{-1} \left(\begin{bmatrix} 0 \\ 0 \\ \pi r^4/4 \end{bmatrix} + O(r^6) \right). \end{aligned} \quad (8)$$

(in this computation, many terms evaluate to zero because of symmetry). Division of power series yields

$$\mathbf{s}_b(r, \mathbf{p}) = \left[0, \quad 0, \quad \frac{3}{8} r + \frac{9}{64} H(\mathbf{p}) r^2 \right]^T + O(r^3). \quad (9)$$

Now we consider the covariance matrix. The following table shows the order $O(\rho^k z^l)$ of magnitude of the functions we are going to integrate:

f	x^2	y^2	z^2	xy	xz	yz
(k, l)	$(2, 0)$	$(2, 0)$	$(0, 2)$	$(2, 0)$	$(1, 1)$	$(1, 1)$

We approximate the integrals over those functions according to (6): In any case the error is at most $O(r^7)$, while $\tilde{I}(r, xy) = \tilde{I}(r, xz) = \tilde{I}(r, yz) = 0$ because of symmetry. Thus, the covariance matrix reads

$$\begin{aligned} J(r, \mathbf{p}) &= \int \mathbf{x}\mathbf{x}^T d\mathbf{x} - V_b \mathbf{s}_b \mathbf{s}_b^T = \text{diag} \left(\tilde{I}(r, x^2), \tilde{I}(r, y^2), \tilde{I}(r, z^2) \right) + O(r^7) \\ &\quad - \text{diag} \left(0, 0, \frac{\pi}{4} r^4 \left(\frac{3}{8} r + \frac{9}{64} H(\mathbf{p}) r^2 \right) \right) + O(r^9). \end{aligned} \quad (10)$$

For this formula we have used that the product $V_b \mathbf{s}_b$ is already known from Equation (8). This leads to the following result:

Theorem 2 *The eigenvalues $M_{b,i}(r, \mathbf{p})$ of the covariance matrix of the set*

$N_b(r, \mathbf{p})$ have the following Taylor expansion

$$M_{b,i}(r, \mathbf{p}) = \widetilde{M}_{b,i}(r, \mathbf{p}) + O(r^7) \quad (i = 1, 2, 3), \quad \text{where}$$

$$\widetilde{M}_{b,1}(r, \mathbf{p}) = \frac{2\pi}{15}r^5 - \frac{\pi}{48}(3\kappa_1(\mathbf{p}) + \kappa_2(\mathbf{p}))r^6, \quad (11)$$

$$\widetilde{M}_{b,2}(r, \mathbf{p}) = \frac{2\pi}{15}r^5 - \frac{\pi}{48}(\kappa_1(\mathbf{p}) + 3\kappa_2(\mathbf{p}))r^6, \quad (12)$$

$$\widetilde{M}_{b,3}(r, \mathbf{p}) = \frac{19\pi}{480}r^5 - \frac{9\pi}{512}(\kappa_1(\mathbf{p}) + \kappa_2(\mathbf{p}))r^6. \quad (13)$$

Here $\kappa_1(\mathbf{p})$ and $\kappa_2(\mathbf{p})$ are the principal curvatures of the surface which is used in the definition of the ball neighbourhood.

Proof. For symmetric matrices J and \widetilde{J} in general, the difference of eigenvalues is bounded by $\|J - \widetilde{J}\|$, where $\|\cdot\|$ is any matrix norm. When computed with respect to the principal frame associated with the point \mathbf{p} , the covariance matrix $J(r, \mathbf{p})$ differs from a diagonal matrix $\widetilde{J}(r, \mathbf{p})$ only by an error of magnitude $O(r^7)$, as shown by (10). It follows that the eigenvalues $M_{b,i}$ coincide with the diagonal elements of $\widetilde{J}(r, \mathbf{p})$, up to $O(r^7)$. The precise values of these diagonal elements are found by computing the integrals $\widetilde{I}(r, x^2), \dots$ according to (6), which is elementary but tedious. From (10) we see that $M_{b,1} = \widetilde{I}(r, x^2) + O(r^7)$ and $M_{b,2} = \widetilde{I}(r, y^2) + O(r^7)$. Interestingly

$$\widetilde{I}(r, z^2) = \frac{2}{15}\pi r^5 \quad (14)$$

does not contain any curvature information. In combination with the other terms in (10), we get $M_{b,3}$. \square

Theorem 3 *We consider the situation of Theorem 2. As $r \rightarrow 0$, the eigenvectors $\mathbf{e}_{b,1}(r, \mathbf{p})$ and $\mathbf{e}_{b,2}(r, \mathbf{p})$ of the covariance matrix which correspond to the eigenvalues $M_{b,1}(r, \mathbf{p})$ and $M_{b,2}(r, \mathbf{p})$ converge to the principal directions $\mathbf{e}_1, \mathbf{e}_2$ associated with κ_1 and κ_2 , provided $\kappa_1 \neq \kappa_2$. The eigenvector $\mathbf{e}_{b,3}(r, \mathbf{p})$ converges towards the surface normal vector \mathbf{n} . For $\mathbf{e}_{b,1}(r, \mathbf{p})$ and $\mathbf{e}_{b,2}(r, \mathbf{p})$, this convergence is linear, whereas for $\mathbf{e}_{b,3}$ it is quadratic:*

$$\angle(\mathbf{e}_{b,1}(r, \mathbf{p}), \mathbf{e}_1(\mathbf{p})), \angle(\mathbf{e}_{b,2}(r, \mathbf{p}), \mathbf{e}_2(\mathbf{p})) \sim \frac{r}{\kappa_1(\mathbf{p}) - \kappa_2(\mathbf{p})}, \quad (15)$$

$$\angle(\mathbf{e}_{b,3}(r, \mathbf{p}), \mathbf{n}(\mathbf{p})) \sim r^2. \quad (16)$$

Proof. The statement about convergence of eigenvectors follows from the ‘‘sine theta’’ theorems for perturbation of eigenvectors of Hermitean matrices due to Davis and Kahan (1970). If λ_i is an eigenvalue of the matrix J , and ε is the minimum distance from the remaining eigenvalues, then a perturbation of size h causes a change in the eigenvector \mathbf{e}_i of magnitude h/ε , provided $h \ll \varepsilon$. We apply this result to $J = \text{diag}(\widetilde{M}_{b,1}, \widetilde{M}_{b,2}, \widetilde{M}_{b,3})$, with $\widetilde{M}_{b,i}$ from Equations

(11)–(13). Then $h = O(r^7)$. In the case $i = 1, 2$, we have $\varepsilon \sim r^6(\kappa_1 - \kappa_2)$, and in the case $i = 3$, we have $\varepsilon \sim r^5$. \square

The eigenvalues $M_{b,1}(r, \mathbf{p})$ $M_{b,2}(r, \mathbf{p})$ do not feature the principal curvatures in their dominant terms, but their difference does:

$$M_{b,2}(r, \mathbf{p}) - M_{b,1}(r, \mathbf{p}) = \frac{\pi}{24}(\kappa_1(\mathbf{p}) - \kappa_2(\mathbf{p}))r^6 + O(r^7). \quad (17)$$

2.3 Principal component analysis of the sphere neighbourhood

In this section we consider PCA of the sphere neighbourhood $N_s(r, \mathbf{p}) = S(r, \mathbf{p}) \cap D$ in a way analogous to the previous section which dealt with the ball neighbourhood. As to notation, see Section 1.1 and Figure 2.

Theorem 4 *The barycenter $\mathbf{s}_s(r, \mathbf{p})$ of the sphere neighbourhood $N_s(r, \mathbf{p})$ is expressed in the principal frame as*

$$\mathbf{s}_s(r, \mathbf{p}) = \left[0, 0, \frac{1}{2}r + \frac{\kappa_1(\mathbf{p}) + \kappa_2(\mathbf{p})}{8}r^2\right]^T + O(r^3). \quad (18)$$

For the sphere neighbourhood, the eigenvalues $M_{s,i}$ of the covariance matrix have the form $M_{s,i} = \widetilde{M}_{s,i}(r, \mathbf{p}) + O(r^6)$, where

$$\widetilde{M}_{s,1}(r, \mathbf{p}) = \frac{2\pi}{3}r^4 - \frac{\pi}{8}(3\kappa_1(\mathbf{p}) + \kappa_2(\mathbf{p}))r^5, \quad (19)$$

$$\widetilde{M}_{s,2}(r, \mathbf{p}) = \frac{2\pi}{3}r^4 - \frac{\pi}{8}(\kappa_1(\mathbf{p}) + 3\kappa_2(\mathbf{p}))r^5, \quad (20)$$

$$\widetilde{M}_{s,3}(r, \mathbf{p}) = \frac{\pi}{6}r^4 - \frac{\pi}{8}(\kappa_1(\mathbf{p}) + \kappa_2(\mathbf{p}))r^5. \quad (21)$$

Proof. Any integral $I(r, f)$ over the ball neighbourhood can be written as an iterated integral:

$$I(r, f) = \int_{D \cap B(r, \mathbf{p})} f(\mathbf{x}) d\mathbf{x}, \quad I'(\rho, f) = \int_{D \cap S(\rho, \mathbf{p})} f(\mathbf{x}) d\mathbf{x} \implies I(r, f) = \int_{\rho=0}^r I'(\rho, f) d\rho.$$

This implies the differential relation

$$I'(r, f) = \frac{d}{dr}I(r, f). \quad (22)$$

Thus we can compute the surface area $A_s(r, \mathbf{p})$, the integral $I'(\mathbf{x})$ used for the barycenter, and the matrix $I'(\mathbf{xx}^T)$ used for the covariance matrix, simply by differentiating $V_b(r, \mathbf{p})$, $I(r, \mathbf{x})$, and $I(r, \mathbf{xx}^T)$, respectively. We get the known

expression for the surface area A_s given by (2), and further

$$A_s \mathbf{s}_s = \begin{bmatrix} 0 \\ 0 \\ \pi r^3 \end{bmatrix} + O(r^5), \quad \int_{D \cap S(r, \mathbf{p})} \mathbf{x} \mathbf{x}^T d\mathbf{x} = \begin{bmatrix} d\widetilde{M}_{b,1}/dr & 0 & 0 \\ 0 & d\widetilde{M}_{b,2}/dr & 0 \\ 0 & 0 & \frac{2}{3}\pi r^4 \end{bmatrix} + O(r^6).$$

For the lower right corner of the matrix $I'(\mathbf{x} \mathbf{x}^T)$, we have differentiated (14). By dividing $I'(\mathbf{x})$ by the Taylor series of the surface area A_s , we get (18). Now that the barycenter is known, we are able to compute the covariance matrix: $J'(r, \mathbf{p}) = I'(\mathbf{x} \mathbf{x}^T) - A_s \mathbf{s}_s \mathbf{s}_s^T = \text{diag}(\widetilde{M}_{s,1}, \widetilde{M}_{s,2}, \widetilde{M}_{s,3}) + O(r^6)$. This completes the proof. \square

2.4 Principal component analysis of the spherical intersection curve

We now consider the intersection curve $c_s(r, \mathbf{p}) = S(r, \mathbf{p}) \cap \Phi$, which is defined as the intersection of a kernel sphere of radius r , centered in the point \mathbf{p} , with the smooth surface Φ . Principal component analysis of the spherical intersection curve is included here for the sake of completeness. It is not as robust against perturbations as PCA of the ball or sphere neighbourhoods $N_b(r, \mathbf{p})$ and $N_s(r, \mathbf{p})$. The analysis consists mainly of computations not all of which are repeated here. We first give a cylinder coordinate representation $(\phi, \rho(\phi), z(\phi))$ of the intersection curve, where $x = \rho \cos \phi$ and $y = \rho \sin \phi$.

The intersection of the plane $x : y = \cos \phi : \sin \phi$ with the surface of (4) results in the curve $z(\rho, \phi) = \frac{1}{2}\kappa_n(\phi)^2 \rho^2 + O(\rho^3)$, where the expression

$$\kappa_n(\phi) = \kappa_1 \cos^2 \phi + \kappa_2 \sin^2 \phi \quad (23)$$

has an interpretation as the normal curvature of the surface Φ associated with the direction $x : y = \cos \phi : \sin \phi$ (cf. do Carmo (1976)). Intersection of that curve $z = z(\rho, \phi)$, $\phi = \text{const}$, with the sphere $\rho^2 + z^2 = r^2$ requires some elementary computations and leads to the following parameterization of the intersection curve:

$$c_s : \quad \rho(\phi) = r - \frac{1}{8}\kappa_n(\phi)^2 r^3 + O(r^4), \quad z(\phi) = \kappa_n(\phi)^2 r^2 + O(r^3). \quad (24)$$

The arc length differential in cylinder coordinates reads $(ds/d\phi)^2 = \rho^2 + \rho_\phi^2 + z_\phi^2$. Here the subscript ϕ indicates differentiation. We get $s_\phi^2 = r^2 + \frac{1}{4}(\kappa_{n,\phi}^2 - \kappa_n^2) + O(r^5)$. Taking the square root via the binomial series yields $s_\phi = r + \frac{1}{8}(\kappa_{n,\phi}^2 - \kappa_n^2)r^3 + O(r^4)$. When we expand this expression and integrate, we can compute

the arc length $L_c(r, \mathbf{p})$ of the spherical intersection curve $c_s(r, \mathbf{p})$:

$$\begin{aligned} L_c(r, \mathbf{p}) &= \oint ds = \int_{\phi=0}^{2\pi} \frac{ds}{d\phi} d\phi \\ &= 2\pi r + \frac{\pi}{32}(\kappa_1(\mathbf{p})^2 - 10\kappa_1(\mathbf{p})\kappa_2(\mathbf{p}) + \kappa_2(\mathbf{p})^2)r^3 + O(r^4). \end{aligned} \quad (25)$$

Barycenter \mathbf{s}_c and covariance matrix J_c of the curve c_s are subsequently found by integration:

$$\mathbf{s}_c = \frac{1}{L_c} \int \mathbf{x}(\phi) \frac{ds}{d\phi} d\phi, \quad J_c = \frac{1}{L_c} \int \mathbf{x}(\phi) \mathbf{x}(\phi)^T \frac{ds}{d\phi} d\phi.$$

After the substitution $x(\phi) = \rho(\phi) \cos \phi$ and $y(\phi) = \rho(\phi) \sin \phi$, these integrations are elementary and will not be carried out in detail. The results are similar to the ball and sphere neighbourhood case:

Theorem 5 *The barycenter $\mathbf{s}_c(r, \mathbf{p})$ of the spherical intersection curve $c_s(r, \mathbf{p})$ is expressed in the principal frame as $\mathbf{s}_c(r, \mathbf{p}) = [0, 0, \frac{1}{4}(\kappa_1(\mathbf{p}) + \kappa_2(\mathbf{p}))r^2]^T + O(r^3)$. The covariance matrix of the curve $c_s(r, \mathbf{p})$ has the eigenvalues $M_{c,i}(r, \mathbf{p})$ ($i = 1, 2, 3$) with*

$$\begin{aligned} M_{c,1} &= \pi r^3 + \frac{\pi}{64}((\kappa_1(\mathbf{p}) - \kappa_2(\mathbf{p}))^2 - 12\kappa_2(\mathbf{p})(\kappa_1(\mathbf{p}) + \kappa_2(\mathbf{p})))r^5 + O(r^6), \\ M_{c,3} &= \frac{\pi}{16}(\kappa_1(\mathbf{p}) - \kappa_2(\mathbf{p}))^2 r^5 + O(r^6). \end{aligned}$$

The formula for $M_{c,2}$ is the one for $M_{c,1}$ with indices 1 and 2 exchanged.

2.5 PCA of the surface patch neighbourhood

We now discuss principal component analysis of the surface patch neighbourhood $N_p(r, \mathbf{p})$, thereby providing more precise estimates than those given by Clarenz et al. (2004b). While the surface patch neighbourhood is of equal geometric interest as the ball and sphere neighbourhoods, it has been shown by the numerical experiments of Yang et al. (2006) that it is not as robust against noise. This is only to be expected, since a variation of the surface causes a variation in the surface area which can be bounded only if we know bounds on the derivatives of the surface.

Like in the case of the spherical intersection curve c_s , the mathematics consists almost entirely of computations, which are not repeated here. We again refer to the cylinder coordinate representation $x = \rho \cos \phi$, $y = \rho \sin \phi$, $z = z(\rho \cos \phi, \rho \sin \phi)$, of the surface (4). The surface area element is given by $dA = (1 + (dz/dx)^2 + (dz/dy)^2)^{1/2} dx dy = (\rho + \frac{1}{2}\rho^3(\kappa_1^2 \cos^2 \phi + \kappa_2^2 \sin^2 \phi) + O(\rho^5)) d\rho d\phi$. The surface area is computed as the integral $\int_{\phi=0}^{2\pi} \int_{\rho=0}^{\rho(\phi)} dA$, where $\rho(\phi)$ is taken

from (24). The result is

$$A_p(r, \mathbf{p}) = \pi r^2 + \frac{\pi}{32}(\kappa_1(\mathbf{p}) - \kappa_2(\mathbf{p}))^2 r^4 + O(r^5). \quad (26)$$

The remaining computations of barycenter and principal components are similar to the curve case:

Theorem 6 *The barycenter $\mathbf{s}_p(r, \mathbf{p})$ of the spherical path neighbourhood $N_p(r, \mathbf{p}) = B(r, \mathbf{p}) \cap \Phi$ is expressed in the principal frame as $\mathbf{s}_p(r, \mathbf{p}) = [0, 0, \frac{1}{8}(\kappa_1(\mathbf{p}) + \kappa_2(\mathbf{p}))r^2]^T + O(r^3)$. The covariance matrix of the spherical patch neighbourhood has the eigenvalues $M_{p,i}$ ($i = 1, 2, 3$), with*

$$\begin{aligned} M_{p,1}(r, \mathbf{p}) &= \frac{\pi}{4}r^4 + \frac{\pi}{192}\left(\kappa_2(\mathbf{p})^2 - 3\kappa_1(\mathbf{p})^2 - 6\kappa_1(\mathbf{p})\kappa_2(\mathbf{p})\right)r^6 + O(r^7), \\ M_{p,3}(r, \mathbf{p}) &= \frac{\pi}{192}\left(3\kappa_1(\mathbf{p})^2 + 3\kappa_2(\mathbf{p})^2 - 2\kappa_1(\mathbf{p})\kappa_2(\mathbf{p})\right)r^6 + O(r^7). \end{aligned}$$

The formula for $M_{p,2}$ is the one for $M_{p,1}$ with indices 1 and 2 exchanged.

3 Geometry descriptors from integral invariants

For many geometry processing tasks it is important to compute numerical geometry descriptors which act as a low pass filter, i.e., ignore details which are smaller than a certain threshold size. One way to define such descriptors is described below.

3.1 Principal curvatures at a given scale

The relationship between covariance matrices and principal curvatures leads to the definition of *principal curvatures at scale r* , where r is the kernel ball radius. This is done by ignoring the $O(r^7)$ -terms in (11), (12) and solving for κ_1, κ_2 . They are estimators for the actual values of the principal curvatures, and are denoted by $\kappa_{b,1}, \kappa_{b,2}$, where the index b stands for ‘ball neighbourhood’:

$$\begin{aligned} \kappa_{b,1}(r, \mathbf{p}) &:= \frac{6}{\pi r^6}(M_{b,2}(r, \mathbf{p}) - 3M_{b,1}(r, \mathbf{p})) + \frac{8}{5r}, \\ \kappa_{b,2}(r, \mathbf{p}) &:= \frac{6}{\pi r^6}(M_{b,1}(r, \mathbf{p}) - 3M_{b,2}(r, \mathbf{p})) + \frac{8}{5r}. \end{aligned} \quad (27)$$

Theorem 4 likewise makes it possible to define principal curvature estimators $\kappa_{s,1}$ and $\kappa_{s,2}$, using the eigenvalues $M_{s,1}$ and $M_{s,2}$. And of course also the eigenvalues computed from the curve c_s or the surface patch N_p can be used

as well. A comparison of this method to define curvatures with the method of normal cycles is furnished by Figure 5.

The arithmetic mean $(\kappa_{b,1}(r, \mathbf{p}) + \kappa_{b,2}(r, \mathbf{p}))/2$ is a possible definition of the mean curvature estimator at scale r , but also other formulae (e.g., the ones involving barycenters, volumes, or areas) can be used to define a mean curvature at scale r : We may, for instance, let

$$H_{ball}(r, \mathbf{p}) = \frac{4}{\pi r^4} \left(\frac{2\pi}{3} r^3 - V_b(r, \mathbf{p}) \right), \quad H_{sphere}(r, \mathbf{p}) = \frac{1}{\pi r^3} (2\pi r^2 - A_s(r, \mathbf{p})). \quad (28)$$

These definitions are based on Equations (1) and (2), and illustrated by Fig. 6. Any of these different definitions of principal curvatures or mean curvature at scale r yields a family of geometry descriptors which approximates curvatures as the kernel radius tends to zero.

An important property of principal curvatures at a given scale is that features of the surface under investigation which are smaller than the kernel radius are considered as noise and are in general smoothed away. When we use principal curvatures at various different scales r_1, r_2, \dots in order to recognize features on a surface, it may happen that some features are detected only for one r_i , whereas others (the *persistent* ones) are detected for several radii r_j . This helps to distinguish between unimportant and important features of given data. Examples are given by Yang et al. (2006). Figure 5 which computes one principal curvature at two different scales also illustrates this fact. An example of a computation where various kernel sizes are employed in order to detect ridges and ravines is shown by Fig. 1.

3.2 *Principal curves on multiple scales*

Similar to curvature estimators, we can also get estimators for principal directions. PCA of a local neighbourhoods defined by a kernel of size r yields principal directions *at scale* r . These directions are in general not tangential to the given surface, since they are computed in a global way from the chosen neighbourhood. In fact, they should not follow details which are small compared to the kernel size. In order to obtain vector fields which are tangent to the given surface, such directions have to be projected onto the surface, but then they are no longer orthogonal. This loss of a prominent geometric property is actually not important and rather an advantage when we now integrate these two vector fields to obtain *principal curves at the chosen scale* r (see Figures 3 and 4 for an illustration).

Applications of principal curves at a given resolution r include remeshing for the purpose of constructing a quad-dominant mesh with planar faces approximating a given surface, as presented by Liu et al. (2006).

Due to the robustness of these principal curves, it seems natural to use such curves for detecting special shapes such as pipe surfaces, rotational surfaces, canal surfaces and developable surfaces (see Fig. 3). Non-planar developable surfaces can be characterized by the fact that exactly one principal curvature vanishes; they have one family of straight principal curvature lines. Recognizing and especially reconstructing developable surfaces from measurement data is not an easy task, as discussed by Peternell (2004). Classical principal curvature lines can hardly be employed for computations where robustness is essential. However it turns out that principal curvatures at a larger scale r are useful for such purposes. This is illustrated by Fig. 3: Apart from smoothing effects near sharp edges, we can nicely observe one family of nearly straight principal curves. They are good candidates for rulings in an approximating surface. The ‘D-form’ model of Fig. 3, right, is a triangle mesh generated from a laser scan of a paper model. For more facts on D-forms, the interested reader is referred to pp. 317, 401, and 418 of the monograph by Pottmann and Wallner (2001), and also to Bobenko and Izmistiev (2005).

3.3 Comparison with other methods, robustness, and multiscale behaviour

This paper provides theoretical background and proofs for the material presented by Yang et al. (2006). For this reason, Section 3.3 is rather short, as numerical results have already been presented in that paper. They compare the principal curvatures obtained via the ball and sphere neighbourhoods with other methods:

- Principal components of the patch neighbourhood according to Clarenz et al. (2004b) and Pauly et al. (2003);
- The method of normal cycles of Cohen-Steiner and Morvan (2003);
- A fitting method (osculating jets by Cazals and Pouget (2003)).

Yang et al. (2006) report that PCA of the patch neighbourhood is quite sensitive to noise, normal cycles less so. PCA of the ball and sphere neighbourhoods compare favourably with these three methods. The multiscale behaviour of these different methods is good for low noise levels, while local fitting methods apparently have defects at coarse scales.

The present paper shows a comparison with the method of normal cycles in Figure 5, by marking in dark regions of high principal curvatures – this is a simple way of defining features. For this particular data set it is obvious that the PCA methods of the present paper detect more features than the method of normal cycles.

The multiscale behaviour of integral invariants defined via the ball and sphere neighbourhoods is illustrated in Figures 6 and 5. The latter illustrates the

maximum principal curvature at scale r via grey values, the former shows level curves of a mean curvature at scale r . Figure 4 shows how principal curves defined at a coarser scale exhibit much more stable behaviour than curves defined at a finer scale – features smaller than the kernel balls tend to be considered as noise.

4 Implementation

4.1 Ball neighbourhood computations via FFT

Integral invariants defined via the ball neighbourhood can easily be computed by convolution, which means that FFT can be employed. In more detail, this is done as follows: Using the terminology from above, D is the domain whose boundary is the surface under investigation. In order to do PCA on the ball neighbourhood $N_b(r, \mathbf{p}) = B(r, \mathbf{p}) \cap D$, we compute the integrals $I(r, \mathbf{p}, 1)$, $I(r, \mathbf{p}, \mathbf{x})$, $I(r, \mathbf{p}, \mathbf{x}\mathbf{x}^T)$ defined by

$$I(r, \mathbf{p}, f) = \int_{N_b(r, \mathbf{p})} f(\mathbf{x}) d\mathbf{x}. \quad (29)$$

So the first argument of the function I is the kernel radius, the second argument is the point \mathbf{p} under investigation, and the third argument is the function we integrate. The definition of the same quantities by Equation (5) uses a coordinate system where the point \mathbf{p} under investigation is the origin, and the z axis is orthogonal to the surface. However, it does not matter which coordinate system we use for computing the eigenvalues of the covariance matrix of $N_b(r, \mathbf{p})$, as long as we consistently use the same coordinate system while we work with the point \mathbf{p} .

We use the notation $1_D(\mathbf{x})$ for the indicator function of the domain D , which is defined by $1_D(\mathbf{x}) = 1$ if $\mathbf{x} \in D$ and $1_D(\mathbf{x}) = 0$ otherwise. Likewise, we define $1_B(\mathbf{x})$ to be the indicator function of the ball with center \mathbf{o} and radius r . Apparently, the definitions of $I(r, \mathbf{p}, 1), \dots$ can be rewritten as

$$\begin{aligned} I(r, \mathbf{p}, 1) &= \int_{\mathbb{R}^3} 1_D(\mathbf{x}) 1_B(\mathbf{p} - \mathbf{x}) d\mathbf{x}, \\ I(r, \mathbf{p}, \mathbf{x}) &= \int_{\mathbb{R}^3} 1_D(\mathbf{x}) (-\mathbf{p} - \mathbf{x}) 1_B(\mathbf{p} - \mathbf{x}) d\mathbf{x}, \\ I(r, \mathbf{p}, \mathbf{x}\mathbf{x}^T) &= \int_{\mathbb{R}^3} 1_D(\mathbf{x}) (\mathbf{p} - \mathbf{x})(\mathbf{p} - \mathbf{x})^T 1_B(\mathbf{p} - \mathbf{x}) d\mathbf{x}. \end{aligned}$$

Note that we write $\mathbf{p} - \mathbf{x}$ instead of $\mathbf{x} - \mathbf{p}$, which is justified by the symmetries present. We use convolution notation, i.e., $(f * g)(\mathbf{p}) = \int f(\mathbf{x}) g(\mathbf{p} - \mathbf{x}) d\mathbf{x}$. Then

the formulas above are converted into

$$\begin{aligned}
I(r, \mathbf{p}, 1) &= (1_D * 1_B)(\mathbf{p}) \\
I(r, \mathbf{p}, \mathbf{x}) &= (1_D * (1_B \cdot f))(\mathbf{p}), \quad \text{where } f(\mathbf{x}) = -\mathbf{x}, \\
I(r, \mathbf{p}, \mathbf{xx}^T) &= (1_D * (1_B \cdot g))(\mathbf{p}), \quad \text{where } g(\mathbf{x}) = \mathbf{xx}^T.
\end{aligned}$$

This means that we may compute all necessary information by convolving the indicator function 1_D with the 10 different functions $1_B(x, y, z)$, $-x \cdot 1_B(x, y, z)$, $-y \cdot 1_B(x, y, z)$, $-z \cdot 1_B(x, y, z)$, $x^2 \cdot 1_B(x, y, z)$, $y^2 \cdot 1_B(x, y, z)$, $z^2 \cdot 1_B(x, y, z)$, $xy \cdot 1_B(x, y, z)$, $xz \cdot 1_B(x, y, z)$, $yz \cdot 1_B(x, y, z)$.

For the actual numerical computation, the indicator function 1_D is represented as an occupancy voxel grid, as discussed by Gelfand et al. (2005). Computing convolutions via FFT means that we compute the result of convolution for all points in a cube-shaped domain at once. In order to optimize computation time, we need to cover the boundary of D (i.e., the surface Φ to be analyzed) by a suitable collection of cubes, so as not to compute too many values irrelevant for curvature information at the boundary of the domain D .

Model	Number of triangles	time for PCA ball	time for PCA sphere	Figure number
pillow	24576	23.3 s	8.1 s	*2 and *3
dragon	209227	83.9 s	39.1 s	5 and *4, *5
bunny	69451	50.1 s	15.1 s	6

Table 1

Computation times for PCA. A star before a figure number refers to figures in the paper Yang et al. (2006).

Model	No. of triangles	type of PCA	time for PCA	time for curves	Figure number
horse	114560	ball	46.7s	56.6s	4 and *6, left
pipe	30560	ball	39.7s	5.5s	3, left
D-form	16384	sphere	4.8s	2.8s	3, right

Table 2

Computation times for principal curves. A star before a figure number refers to figures in the paper Yang et al. (2006).

4.2 Sphere neighbourhood computations

Here we use a more geometric method, which is based on an almost uniform multilevel discretization of the sphere $S(r, \mathbf{p})$. By discrete integration over the mesh faces, we can compute any integral invariant which is defined via the sphere neighbourhood.

To accelerate the computation of the sphere neighbourhood $D \cap S(r, \mathbf{p})$, we use a 4-level hierarchical representation of the sphere. It has a highly regular triangulation with 218 faces at the coarsest level. The next finer level is obtained by splitting each triangle into 4 sub-triangles, where new vertices get projected onto the sphere. During refinement, we maintain a list of points inside D . A specified face F at a given level is put into this list, if all 3 vertices are in D . Otherwise, we split the face into 4 sub-faces and judge them one by one in the next level. If we reach the finest discretization level, we use the barycenter of the face to judge inside-outside information. After a hierarchical traversal, we get a list of inside faces which may belong to different discretized sphere levels. PCA is based on the union of these inside faces.

Table 1 gives some computation times of PCA on ball and sphere neighbourhood, respectively. We have done the experiments on a 2.8 GHz PC with 2GB RAM. Table 2 shows the computation times for principal curves.

5 Conclusion

Principal component analysis of local neighbourhoods defined by a ball kernel or spherical kernel allows us to define principal directions and principal curvatures on a given scale. We have shown an asymptotic analysis of the principal moments and directions of inertia, and their relation to principal curvatures. Geometry descriptors derived in this way are robust against noise and exhibit multiscale behaviour suitable for applications like feature extraction and computation of principal curves.

Acknowledgments

This research was supported by grants No. 16002-N05 and S9207 of the Austrian Science Fund (FWF).

References

- Alliez, P., Cohen-Steiner, D., Devillers, O., Lévy, B., Desbrun, M., 2003. Anisotropic polygonal remeshing. *ACM Trans. Graphics* 22 (3), 485–493, Proc. SIGGRAPH.
- Bajaj, C., Xu, G., 2003. Anisotropic diffusion on surfaces and functions on surfaces. *ACM Trans. Graphics* 22 (1), 4–32.
- Bobenko, A., Izvestiev, I., 2005. Alexandrov’s theorem, weighted Delaunay triangulations, and mixed volumes. preprint.
URL <http://arxiv.org/math/0609447>
- Cazals, F., Chazal, F., Lewiner, T., 2003. Molecular shape analysis based upon the Morse-Smale complex and the Connolly function. In: Proc. Symp. Comp. Geometry. pp. 351–360.
- Cazals, F., Pouget, M., 2003. Estimating differential quantities using polynomial fitting of osculating jets. In: Kobbelt, L., Schröder, P., Hoppe, H. (Eds.), Symp. Geometry processing. Eurographics, pp. 177–178.
- Cazals, F., Pouget, M., 2005. Topology driven algorithms for ridge extraction on meshes. Tech. Rep. 5526, INRIA.
URL <http://www.inria.fr/rrrt/rr-5526.html>
- Clarenz, U., Rumpf, M., Schweitzer, M. A., Telea, A., 2004a. Feature sensitive multiscale editing on surfaces. *Vis. Computer* 20 (5), 329–343.
- Clarenz, U., Rumpf, M., Telea, A., 2004b. Robust feature detection and local classification for surfaces based on moment analysis. *IEEE Trans. Vis. Comp. Graphics* 10 (5), 516–524.
- Cohen-Steiner, D., Morvan, J.-M., 2003. Restricted Delaunay triangulations and normal cycle. In: SCG ’03: Proc. 19th symposium on Computational geometry. ACM Press, pp. 312–321.
- Connolly, M., 1986. Measurement of protein surface shape by solid angles. *J. Mol. Graphics* 4 (1), 3–6.
- Davis, C., Kahan, W. M., 1970. The rotation of eigenvectors by a perturbation. III. *SIAM J. Numer. Anal.* 7, 1–46.
- Desbrun, M., Grinspun, E., Schröder, P., 2005. Discrete differential geometry: An applied introduction. SIGGRAPH Course Notes.
- do Carmo, M., 1976. *Differential Geometry of Curves and Surfaces*. Prentice-Hall.
- Gelfand, N., Mitra, N. J., Guibas, L. J., Pottmann, H., 2005. Robust global registration. In: Desbrun, M., Pottmann, H. (Eds.), Symp. Geometry Processing. Eurographics, pp. 197–206.
- Goldfeather, J., Interrante, V., 2004. A novel cubic-order algorithm for approximating principal direction vectors. *ACM Trans. Graphics* 23 (1), 45–63.
- Hertzmann, A., Zorin, D., 2000. Illustrating smooth surfaces. In: SIGGRAPH ’00. pp. 517–526.
- Hildebrandt, K., Polthier, K., 2004. Anisotropic filtering of non-linear surface features. *Computer Graphics Forum* 23 (3), 391–400, Proc. Eurographics.
- Hildebrandt, K., Polthier, K., Wardetzky, M., 2005. Smooth feature lines on

- surface meshes. In: Desbrun, M., Pottmann, H. (Eds.), *Symp. Geometry Processing*. Eurographics, pp. 85–90.
- Hulin, D., Troyanov, M., 2003. Mean curvature and asymptotic volume of small balls. *Amer. Math. Monthly* 110, 947–950.
- Liu, Y., Pottmann, H., Wallner, J., Yang, Y., Wang, W., 2006. Geometric modeling with conical meshes and developable surfaces. *ACM Trans. Graphics* 25 (3), 681–689, Proc. SIGGRAPH.
- Manay, S., Hong, B.-W., Yezzi, A. J., Soatto, S., 2004. Integral invariant signatures. In: Pajdla, T., Matas, J. (Eds.), *Computer Vision — ECCV 2004*, Part IV. Springer, pp. 87–99.
- Ohtake, Y., Belyaev, A., Seidel, H.-P., 2004. Ridge-valley lines on meshes via implicit surface fitting. *ACM Trans. Graphics* 23 (3), 609–612, Proc. SIGGRAPH.
- Osher, S., Fedkiw, R., 2002. *The Level Set Method and Dynamic Implicit Surfaces*. Springer.
- Pauly, M., Keiser, R., Gross, M., 2003. Multi-scale feature extraction on point-sampled geometry. *Computer Graphics Forum* 22 (3), 281–289, Proc. Eurographics.
- Peternell, M., 2004. Developable surface fitting to point clouds. *Comput. Aided Geom. Des* 21, 785–803.
- Pottmann, H., Wallner, J., 2001. *Computational Line Geometry*. Springer.
- Pottmann, H., Wallner, J., Huang, Q., Yang, Y.-L., 2006. Integral invariants for robust geometry processing. preprint, TU Wien.
URL <http://dmg.tuwien.ac.at/wallner/iirgp.pdf>
- Taubin, G., 1995. Estimating the tensor of curvature of a surface from a polyhedral approximation. In: *ICCV '95: Proceedings of the Fifth International Conference on Computer Vision*. IEEE Computer Society, pp. 902–907.
- Taubin, G., Cooper, D. B., 1992. Object recognition based on moment (or algebraic) invariants. In: Mundy, J. L., Zisserman, A. (Eds.), *Geometric invariance in Computer Vision*. MIT Press, pp. 375–397, Proceedings of the DARPA-ESPRIT workshop held in Reykjavik, March 25–28, 1991.
- Tong, W.-S., Tang, C.-K., 2005. Robust estimation of adaptive tensors of curvature by tensor voting. *IEEE Trans. PAMI* 27 (3), 434–449.
- Yang, Y.-L., Lai, Y.-K., Hu, S.-M., Pottmann, H., 2006. Robust principal curvatures on multiple scales. In: Polthier, K., Sheffer, A. (Eds.), *Symp. Geometry processing*. Eurographics, pp. 223–226.
URL <http://dmg.tuwien.ac.at/wallner/sgp06-electronic.pdf>
- Yoshizawa, S., Belyaev, A., Seidel, H.-P., 2005. Fast and robust detection of crest lines on meshes. In: *SPM '05: Proceedings of the 2005 ACM symposium on Solid and Physical Modeling*. ACM Press, pp. 227–232.




Fig. 3. Surface interrogation with principal curves at a certain scale r , computed for models which contain developable surfaces. Left: pipe surfaces which contain cylindrical parts. Right: D-form. The figures show also one of the kernel balls employed in the computation of principal curves.




Fig. 4. Principal curves at different scales, computed by PCA of the ball neighbourhood. The different kernel sizes are shown by one kernel ball each.




Fig. 5. Grey coded values of the maximal principal curvature computed (a) by the method of normal cycles by Cohen-Steiner and Morvan (2003); (b) as the value $\kappa_{b,1}$ according to Equation (27), by PCA of the ball neighbourhood; (c) as the value $\kappa_{s,1}$, by PCA of the sphere neighbourhood. In each the kernel size is shown.




Fig. 6. Level sets of various ‘mean curvatures at scale r ’, computed for the Stanford bunny dataset. Figures (a) and (b) show the quantity H_{sphere} , which is defined in Equation 28. Figures (c) and (d) show H_{ball} . The kernel radius is visualized by one kernel ball per image.