

Robust Principal Curvatures on Multiple Scales

Yong-Liang Yang¹ Yu-Kun Lai¹ Shi-Min Hu¹ Helmut Pottmann²

¹Tsinghua University, Beijing ²Vienna University of Technology.

Abstract

Geometry processing algorithms often require the robust extraction of curvature information. We propose to achieve this with principal component analysis (PCA) of local neighborhoods, defined via spherical kernels centered on the given surface Φ . Intersection of a kernel ball B_r or its boundary sphere S_r with the volume bounded by Φ leads to the so-called ball and sphere neighborhoods. Information obtained by PCA of these neighborhoods turns out to be more robust than PCA of the patch neighborhood $B_r \cap \Phi$ previously used. The relation of the quantities computed by PCA with the principal curvatures of Φ is revealed by an asymptotic analysis as the kernel radius r tends to zero. This also allows us to define principal curvatures “at scale r ” in a way which is consistent with the classical setting. The advantages of the new approach are discussed in a comparison with results obtained by normal cycles and local fitting; whereas the former method somewhat lacks in robustness, the latter does not achieve a consistent behavior at features on coarse scales. As to applications, we address computing principal curves and feature extraction on multiple scales.

1. Introduction

Differential geometry plays a central role in the analysis of curves and surfaces. Local investigations frequently use differential invariants such as curvatures, but also the global understanding of shapes can benefit from differential geometric entities. Since differentiation is very sensitive to noise, the use of differential invariants requires data smoothing and de-noising prior to computation. This can be done in a global way via appropriate geometric flows [CRT04] or locally, using smooth approximations of the data in an appropriate neighborhood [CP03, GI04, Tau95, TT05]. In both cases, the preservation of features which may not be considered as noise is not an easy task and requires especially adapted algorithms. A related difficulty is the suppression of those details in the geometry which lie below the scale one is interested in. Whereas classical differential geometry cannot be used directly for such important objects as meshes, discrete differential geometry is extending the theory to the discrete setting [CSM03, DGS05, HP04]. Some discrete operators work on larger local neighborhoods as well and thus are defined on multiple scales. Though handling noisy data is possible [HP04], this is not the main intent and strength of discrete differential geometry.

The present paper pursues an approach via *integral invariants* obtained by integration over local neighborhoods.



Figure 1: Feature extraction on multiple scales using PCA on ball neighborhoods 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.

These neighborhoods are constructed by means of balls, whose radius r defines the scale on which one is working. The origin of this method is work on molecular shape analysis [Con86], on 2D shape matching [MHYS04], and feature extraction [CRT04]. Consider a domain D , its boundary surface Φ , a point $\mathbf{p} \in \Phi$, the ball $B_r(\mathbf{p})$ with radius r and

center \mathbf{p} , and the boundary sphere $S_r(\mathbf{p}) = \partial B_r(\mathbf{p})$. We define *ball and sphere neighborhoods* as $N_b^r(\mathbf{p}) := D \cap B_r(\mathbf{p})$ and $N_s^r := D \cap S_r(\mathbf{p})$, resp. Clarenz et al. [CRT04] and Pauly et al. [PKG03] consider the *surface patch neighborhood* $N_p^r(\mathbf{p}) = \Phi \cap B^r(\mathbf{p})$ and perform PCA on this patch. It turns out that PCA of the patch neighbourhood is less robust against noise than PCA of the ball and sphere neighborhoods.

Contributions and overview. The main contributions of our short paper are: (i) We present the results of a thorough study of PCA on all three neighborhoods described above. (ii) Our analysis also yields definitions of principal curvatures at scale r which are consistent with the classical theory. However, note that PCA yields an integrated quantity, which is not the same as a curvature estimate at a single point. (iii) In Sec. 3, we compare our approach with local fitting [CP03] and normal cycles [CSM03]. (iv) As applications, we study principal curves and feature extraction on multiple scales (Sec. 4). Details and further applications will be provided in a forthcoming paper.

2. Principal component analysis of local neighborhoods

Principal component analysis of a point set A means to compute its barycenter $\mathbf{s} := (\int_A \mathbf{x} d\mathbf{x}) / (\int_A d\mathbf{x})$ and covariance matrix

$$J(A) := \int_A (\mathbf{x} - \mathbf{s}) \cdot (\mathbf{x} - \mathbf{s})^T d\mathbf{x}. \quad (1)$$

If A is not full-dimensional but contained in a smooth surface, (1) is understood as a surface integral. The principal directions and principal components of A are given by the eigenvectors \mathbf{e}_i and the corresponding eigenvalues λ_i ($i = 1, 2, 3$) of $J(A)$.

In order to derive results on the principal components of the neighborhoods defined above, we use the principal frame of the surface Φ at \mathbf{p} as coordinate frame and approximate Φ up to second order by a paraboloid Π ,

$$\Pi: z = \frac{1}{2}(\kappa_1 x^2 + \kappa_2 y^2). \quad (2)$$

Here, κ_1, κ_2 denote the principal curvatures at \mathbf{p} . In order to compute the first two terms in the Taylor expansion of moments with respect to the radius r , it is sufficient to work with Π instead of Φ . By the symmetry of the paraboloid Π , the barycenters of all neighborhoods lie on the z -axis, and the eigenvectors \mathbf{e}_i of the corresponding covariance matrices are parallel to the coordinate axes. In the following, we present some results of [PHYK05].

PCA of the ball and sphere neighborhoods. Consider the barycenters \mathbf{s}_b^r and \mathbf{s}_s^r of the ball and sphere neighborhoods $N_b^r = B_r(\mathbf{p}) \cap D$ and $N_s^r = S_r(\mathbf{p}) \cap D$, respectively. Their signed distances d_b^r and d_s^r from the surface are related to the surface's *mean curvature* H via

$$d_b^r = \frac{3}{8}r + \frac{9H}{64}r^2 + O(r^3), \quad d_s^r = \frac{1}{2}r + \frac{H}{4}r^2 + O(r^3). \quad (3)$$

Two eigenvectors of the covariance matrix are approximately tangent to the surface, the corresponding eigenvalues M_{b1}^r, M_{b2}^r (ball) and M_{s1}^r, M_{s2}^r (sphere) read

$$M_{bi}^r = \frac{2\pi}{15}r^5 - \frac{\pi}{48}(2\kappa_i + \kappa_1 + \kappa_2)r^6 + O(r^7), \quad (4)$$

$$M_{si}^r = \frac{2\pi}{3}r^4 - \frac{\pi}{8}(2\kappa_i + \kappa_1 + \kappa_2)r^5 + O(r^6). \quad (5)$$

An important consequence of equations (4) and (5) is that we can *define* principal curvatures on scale r which, as r tends to zero, converge to the actual principal curvatures in the classical sense. We may use the ball neighborhood and compute curvatures $\kappa_{b1}^r, \kappa_{b2}^r$ from M_{b1}^r and M_{b2}^r via (4), or use the sphere neighborhood and compute $\kappa_{s1}^r, \kappa_{s2}^r$ from M_{s1}^r and M_{s2}^r via (5):

$$\kappa_{b1}^r := \frac{6}{\pi r^6}(M_{b2}^r - 3M_{b1}^r) + \frac{8}{5r}, \quad (6)$$

$$\kappa_{s1}^r := \frac{1}{\pi r^5}(M_{s2}^r - 3M_{s1}^r) + \frac{4}{3r}, \quad (7)$$

and analogous expressions for $\kappa_{b2}^r, \kappa_{s2}^r$. Our tests showed a very similar and useful behavior of both types of principal curvatures.

PCA of the patch neighborhood. Similar results hold for the patch neighborhood N_p^r , but unfortunately here we have higher sensitivity to noise (see [PHYK05] and the discussion below). Properties of N_p^r not contained in Clarenz et al. [CRT04] are the following: The surface area PA^r of N_p^r is related to the principal curvatures κ_1, κ_2 of Φ at \mathbf{p} by

$$PA^r = \pi r^2 + \frac{\pi}{32}(\kappa_1 - \kappa_2)^2 r^4 + O(r^5). \quad (8)$$

The two eigenvalues of $J(N_p^r)$ which correspond to the approximate principal directions satisfy

$$M_{pi}^r = \frac{\pi}{4}r^4 + \frac{\pi}{192}(\kappa_1^2 + \kappa_2^2 - 4\kappa_i^2 - 6\kappa_1\kappa_2)r^6 + O(r^7).$$

3. Comparison of principal component analysis with local fitting and normal cycles

The computation of integral invariants based on the ball neighborhood can efficiently be done by means of the fast Fourier transform, whereas for PCA of the sphere neighborhood, we use a geometric method based on an almost uniform multilevel discretization of the sphere S_r [PHYK05]. We have run extensive tests in order to compare PCA of different neighborhoods with the main methods used in geometry processing: normal cycles [CSM03] and local fitting. As representative of the many available methods for local fitting we employ osculating jets [CP03].

Robustness. In order to avoid effects caused by a complicated geometry, we illustrate sensitivity to noise by means of a rather simple surface (Fig. 2). A summary of results concerning the principal curvature directions is depicted in Fig. 3. PCA of the patch neighborhood [CRT04, PKG03] is quite sensitive to noise, normal cycles less so. The effect of




Figure 2: Principal curvature lines (left) and directions of minimum principal curvature (middle) for a smooth surface ('pillow'); noisy surface (right) used for the robustness test.

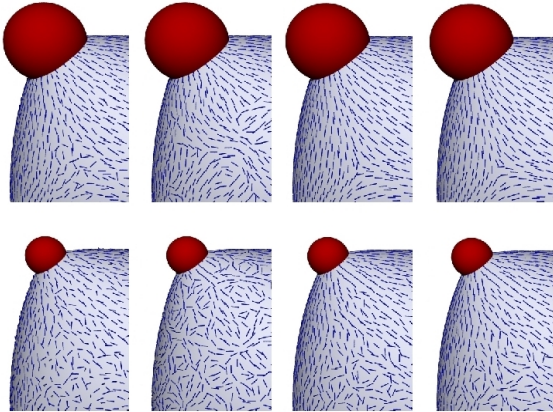


Figure 3: Directions of minimum principal curvature computed from the noisy surface of Fig. 2, right, but shown over the smooth surface for clarity. From left: normal cycles, PCA (patch), osculating jet, PCA (ball).

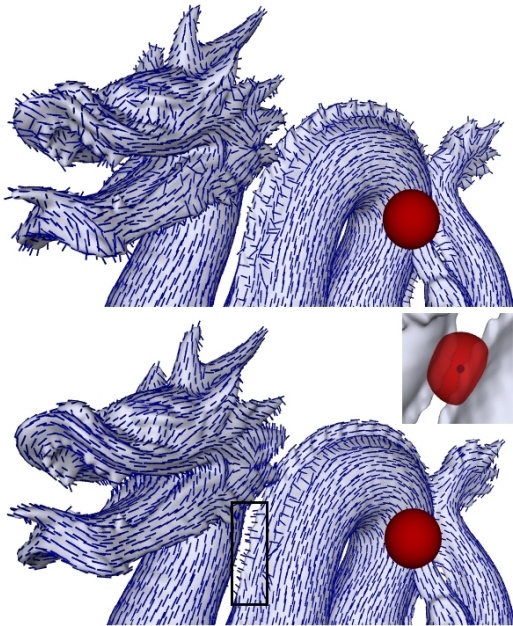


Figure 4: Directions of minimum principal curvature obtained with osculating jets (top) and PCA(ball) (bottom). Larger kernels may have extraneous intersections (see detail) which – if not detected – cause unwanted effects.

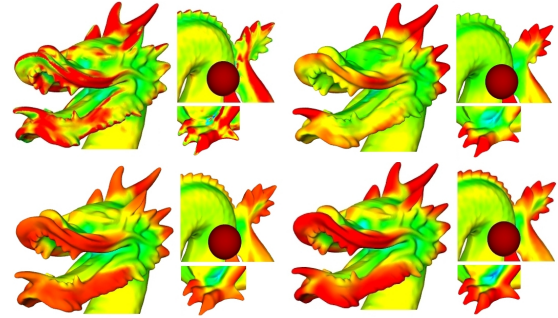


Figure 5: Maximum principal curvature obtained with osculating jets (top left), normal cycles (top right), PCA/sphere (bottom left), and PCA(ball) (bottom right). The jet method exhibits inconsistencies, normal cycles miss some features.

noise removal by averaging over a greater domain is apparently much more pronounced for PCA of ball and sphere neighborhoods than for the patch neighborhood. The normal cycle method performs in between. For simple surfaces, fitting methods behave like PCA of ball/sphere neighborhoods.

Multi-scale behavior. At a low noise level all PCA based methods respond very well to an increase of the kernel radius and exhibit the expected smoothing and simplification effect. This is also true for normal cycles, although their feature detection capability appears to be somewhat weaker. Local fitting methods have defects when it comes to judging curvature at coarse scales (see Figs. 4 and 5). While using a larger neighborhood has some smoothing effect on the local fitting surface Φ_l , one evaluates curvature of Φ_l at the center of the neighborhood. Here, the sensitivity of classical curvatures to small changes becomes crucial and explains why the overall picture of the extracted quantities for a large neighborhood is by far less smooth than for PCA based methods. Apparently the ball and sphere PCA are most robust to noise and exhibit the desired scaling behavior.

4. Principal curves and feature extraction on multiple scales

The eigenvectors \mathbf{e}_1^r , \mathbf{e}_2^r of the covariance matrices of local neighborhoods serve as principal directions of the given surface at the chosen scale r . They are not exactly tangential to the given surface. In fact they should not follow details which are small compared to r . In order to define principal curves on Φ we project \mathbf{e}_1^r , \mathbf{e}_2^r onto Φ and integrate the resulting tangential vector fields. The direction of this projection shall be given by the eigenvector \mathbf{e}_3^r , which estimates the surface normal. Note that the projected directions are usually not orthogonal to each other, which is actually a numerical advantage for integration. By guiding the whole projection procedure with the directions arising from PCA, principal curves become less sensitive to local surface deviations. Figure 6 illustrates their behavior for different scales. The more we increase r , the better these curves follow the

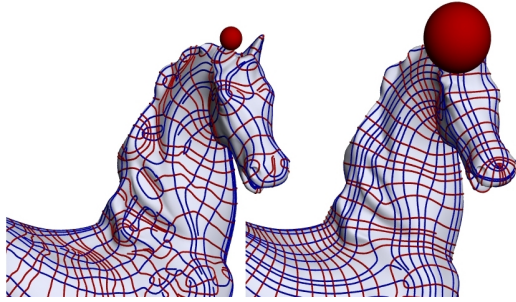


Figure 6: Principal curves for larger kernel radius (right) better follow the global shape.

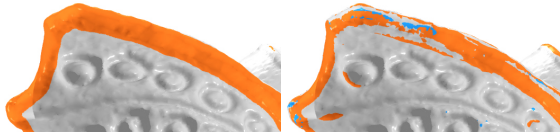


Figure 7: Feature extraction on a coarse scale with PCA/ball (left) and osculating jet (right).




Figure 8: Feature extraction using principal curvatures obtained with PCA on the sphere neighborhood.

global geometry of the surface. We believe that this tool will be very valuable for global shape understanding (see e.g. [LPW*06]).

Feature regions on 3D models are characterized by at least one large principal curvature. We have shown how to define and compute principal curvatures κ_1^r, κ_2^r on multiple scales, e.g. based on formulae (6) or (7). Hence, we also have a simple tool for robust multi-scale feature extraction: On a given scale r , feature regions are first filtered out by $\max(|\kappa_1^r|, |\kappa_2^r|) > T_c$, where T_c is a threshold on curvatures; T_c can be derived from a few user-specified feature and non-feature samples based on a simple statistical model. Using the sign of the dominant principal curvature, the extracted feature regions are further classified into ridges and valleys.

Feature regions extracted at different scales have been combined into the single image of Fig. 1: The dark shaded regions are the persistent features, and lighter shaded regions correspond to features extracted at only one or two scales. See also Figures 7 and 8.

Acknowledgements. This research was supported by the Austrian Science Fund (FWF) under Grant No. P16002-N05 and by the National Science Foundation of China under Grant No. 60225016. We gratefully acknowledge fruitful discussions with Johannes Wallner and his help with the preparation of the final version.

References

- [Con86] CONNOLLY M.: Measurement of protein surface shape by solid angles. *J. Mol. Graphics* 4 (1986). 1
- [CP03] CAZALS F., POUGET M.: Estimating differential quantities using polynomial fitting of osculating jets. In *Symp. Geometry Processing* (2003), pp. 177–178. 1, 2
- [CRT04] CLARENZ U., RUMPF M., TELEA A.: Robust feature detection and local classification for surfaces based on moment analysis. *IEEE TVCG* (2004). 1, 2
- [CSM03] COHEN-STEINER D., MORVAN J. M.: Restricted Delaunay triangulations and normal cycle. In *ACM Symp. Comp. Graphics* (2003), pp. 312–321. 1, 2
- [DGS05] DESBRUN M., GRINSUN E., SCHRÖDER P.: *Discrete Differential Geometry: An Applied Introduction*. SIGGRAPH Course Notes, 2005. 1
- [GI04] GOLDFEATHER J., INTERRANTE V.: A novel cubic-order algorithm for approximating principal direction vectors. *ACM TOG* 23, 1 (2004), 45–63. 1
- [HP04] HILDEBRANDT K., POLTHIER K.: Anisotropic filtering of non-linear surface features. *Computer Graphics Forum* 23, 3 (2004), 391–400. 1
- [LPW*06] LIU Y., POTTMANN H., WALLNER J., WANG W., YANG Y.: Geometric modeling with conical meshes and developable surfaces. *ACM TOG* 25, 3 (2006). 4
- [MHYS04] MANAY S., HONG B.-W., YEZZI A. J., SOATTO S.: Integral invariant signatures. In *Proc. European Conf. Computer Vision* (2004), pp. 87–99. 1
- [PHYK05] POTTMANN H., HUANG Q.-X., YANG Y.-L., KÖLPL S.: Integral invariants for robust geometry processing. *Geometry Preprint* 146, TU Wien, 2005. 2
- [PKG03] PAULY M., KEISER R., GROSS M.: Multi-scale feature extraction on point-sampled geometry. *Computer Graphics Forum* 22, 3 (2003), 281–289. 2
- [Tau95] TAUBIN G.: Estimating the tensor of curvature of a surface from a polyhedral approximation. In *Proc. Int. Conf. Computer Vision* (1995). 1
- [TT05] TONG W.-S., TANG C.-K.: Robust estimation of adaptive tensors of curvature by tensor voting. *IEEE Pattern Anal. Machine Intell.* 27, 3 (2005), 434–449. 1