Ridges and umbilics of a sampled smooth surface: a complete picture gearing toward topological coherence

Frédéric Cazals, Marc Pouget

INRIA

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Ridges and umbilics of a sampled smooth surface: a complete picture gearing toward topological coherence

Frédéric Cazals — Marc Pouget

N° 5294

Septembre 2004

Thème SYM

ISSN 0249-6399

INRIA
SOPHIA ANTIPOLIS

Ridges and umbilics of a sampled smooth surface: a complete picture gearing toward topological coherence

Frédéric Cazals , Marc Pouget

Thème SYM — Systèmes symboliques
Projet Geometrica

Rapport de recherche n° 5294 — Septembre 2004 — 36 pages

Abstract: Consider a smooth surface, and at each point which is not an umbilic, respectively paint in blue (red) anything related to the maximum (minimum) principal curvature. Given such a surface, a blue (red) ridge is a curve on the surface such that at each of its points, the principal blue (red) curvature has an extremum along its blue (red) curvature line. Ridges are curves of *extremal* curvature and therefore encode important informations used in segmentation, registration, matching and surface analysis.

Surprisingly, no method developed so far to report ridges from a mesh approximating a smooth surface comes with a careful analysis, which entails that one does not know whether the ridges are reported in a coherent fashion. This paper aims at bridging this gap with the following contributions. First, a careful analysis of the Acute rule —an orientation procedure used in most algorithms— is presented. Second, given a triangulation T approximating a smooth generic surface S, we present sufficient conditions on T together with a certified algorithm reporting ridges in a topologically coherent fashion. Third, we develop an algorithm and a filtering procedure aiming at reporting the most salient features of a coarse mesh T.

Key-words: Ridges, Umbilics, Sampled Smooth Surfaces, Crest Lines, Geometric Approximation.

Unité de recherche INRIA Sophia Antipolis
2004, route des Lucioles, BP 93, 06902 Sophia Antipolis Cedex (France)
Téléphone : +33 4 92 38 77 77 — Télécopie : +33 4 92 38 77 65

Extrêmes de courbure (ridges) et ombilics sur une surface lisse échantillonnée: une étude complète adaptée à la cohérence topologique

Résumé : Considérons une surface lisse, et en chaque point qui n'est pas un ombilic, peignons respectivement en bleu (rouge) tout ce qui dépend de la courbure principal maximum (minimum). Étant donnée une telle surface, un ridge bleu (rouge) est une courbe telle qu'en chacun de ces points, la courbure principale bleu (rouge) a un extremum le long de sa ligne de courbure bleu (rouge). Les ridges sont des lignes d'extrêmes de courbure et donc codent des informations importantes utilisées en segmentation, enregistrement, comparaison et analyse de surfaces.

Néanmoins, aucune méthode calculant les ridges à partir d'un maillage approchant une surface lisse ne propose une analyse détaillée, de telle sorte qu'il n'est pas possible se savoir si les ridges sont calculés de façon cohérente. Cet article corrige ce problème avec les contributions suivantes. Premièrement, une analyse soigneuse de l' "Acute rule" —une procédure d'orientation utilisée dans la plus part des algorithmes— est présentée. Deuxièmement, étant donnée une triangulation T approchant une surface lisse générique S, nous donnons des conditions suffisantes sur T, ainsi qu'un algorithme certifié calculant les ridges avec une topologie cohérente. Troisièmement, nous développons un algorithme et une procédure de filtrage pour le calcul des courbes les plus saillantes sur un maillage grossier.

Mots-clés : Extrêmes de courbure, Ombilics, Surfaces Lisses Échantillonnées, Lignes Saillantes, Approximation Géométrique.

onsider a smooth surface, and at each point which is not an umbilic, respectively paint in blue (red) anything related to the maximum (minimum) principal curvature. Given such a surface, a blue (red) ridge is a curve on the surface such that at each of its points, the principal blue (red) curvature has an extremum along its blue (red) curvature line. Ridges are curves of *extremal* curvature and therefore encode important informations used in segmentation, registration, matching and surface analysis.

Surprisingly, no method developed so far to report ridges from a mesh approximating a smooth surface comes with a careful analysis, which entails that one does not know whether the ridges are reported in a coherent fashion. This paper aims at bridging this gap with the following contributions. First, a careful analysis of the Acute rule —an orientation procedure used in most algorithms— is presented. Second, given a triangulation T approximating a smooth generic surface S, we present sufficient conditions on T together with a certified algorithm reporting ridges in a topologically coherent fashion. Third, we develop an algorithm and a filtering procedure aiming at reporting the most salient features of a coarse mesh T.

1 Introduction

1.1 Ridges and applications

Curves of extremal curvature on high quality meshes. Processing high quality meshes naturally call for differential geometry methods, and given a mesh approximating a smooth surface, estimates for the surface area, the tangent space, or the Weingarten map –which encodes information on the principal directions and curvatures— are by now well understood. However, beyond these first and second order operators, several applications such as parameterization and segmentation, registration [TG95, PAT00], shape matching [HGY$^+$99] or surface analysis [HGY$^+$99] benefit from the characterization of higher order properties and in particular the characterization of curves of *extremal* curvatures, which are precisely the so-called *ridges*.

Consider a smooth oriented surface, and at each point which is not an umbilic, respectively paint in blue (red) anything related to the maximum (minimum) principal curvature. Given such a surface, a blue (red) ridge is a curve on the surface such that at each of its points, the principal blue (red) curvature has an extremum along its blue (red) curvature line. Ridges entertain subtle and deep relationships to the medial axis and to salient features of a smooth surface, which is why ridges are so often called for in applied geometry. But despite their importance in applications, the calculation and the manipulation of ridges poses several major problems.

Stability issues. Ridges witness extrema of principal curvatures and their definition involves derivatives of curvatures, whence third order differential quantities. Moreover, the classification of ridges (as we shall see a given ridge can be classified as elliptic or hyperbolic) involves fourth order differential quantities, so that the precise definition of ridges requires C^4 differentiable surfaces. Therefore, the calculation of ridges from a mesh approximating a smooth surface is a difficult problem.

As a consequence, out of all the publications available about ridges, none attempts at precisely qualifying the quality of the output. In particular, when surfaces are fitted to points clouds or meshes, one never knows whether the ridges reported are those of the original data or artifacts of the surface fitted. A related issue is that the difficult predicates / operations of the algorithms developed are never analyzed.

Completeness issues. Applications using ridges often if not always do so in a partial fashion i.e. focus on a subset of ridges only. (As we shall see later, these are in general the elliptic ridges.) We illustrate this trend, which can probably be charged to the fact that a precise use of ridges has yet to be found, by the following two examples.

First consider a generic closed surface of genus zero —a topological sphere. Each such surface has at least four umbilics, each being traversed by either one or three ridges —we shall recall the precise relationship between ridges and umbilics in a forthcoming section. Reporting ridges therefore requires reporting and classifying umbilics, an issue always overlooked. A second illustration of the non-completeness of ridges is the following. It is well known that the boundary of the medial axis projects onto a subset of the so-called elliptic ridges of the surface. Therefore, using the medial axis to report ridges not only misses all the hyperbolic ridges, but might also miss elliptic ridges.

1.2 Contributions

For analytically known surfaces, ridges can be defined analytically, and at least theoretically, can be recovered using computer algebra and numerical techniques. However, the complexity of the calculations and the numerical drifts often hinder from doing so. Consequently, effective algorithms rely at some point on calculations performed on a mesh, and then a mandatory —if ridges of different colors are reported separately— operation is to find a coherent orientation between principal vectors (vectors associated with the principal directions) at the endpoints of an edge of the mesh. Interestingly, the corresponding operation, which we call the *Acute Rule* —it just consists of picking the orientation so that the two vectors make an acute angle— has never been analyzed so that one never knows whether the ridges reported are reliable or not.

In this context, our contributions are threefold:

- We present the first analysis of the Acute Rule, and highlight its subtle interplay with the geometry of ridges and principal foliations.
- Given a triangulation T approximating a smooth generic surface S, we present sufficient conditions on T so that the ridges reported from T are topologically coherent with those of S. Under these hypothesis, we design a certified algorithm.
- For the cases where the surface S is not generic or the triangulation T does not meet the sufficient conditions, we develop a filtering procedure which aims at reporting the more significant ridges only.

All in all, our focus in on recovering delicate smooth objects from discrete ones. A natural question is therefore whether this is worth the effort since a vivid trend over the past years has

been to investigate discrete concepts which are independent of any differential setting. Two such examples are the theory of normal cycle for the estimation of curvatures [CSM03], or the discrete Morse theory developed by R. Forman [For98]. We believe the answer is affirmative for two reasons.

First, when discrete concepts exist, the question of their convergence to the smooth counter-parts should always be addressed since the meaning of a calculation which does not yield the smooth value on a dense model is at best questionable. A well known example is the angular defect to estimate the Gauss curvature [BCM03], which in general does not allow to estimate the Gauss curvature of a smooth surface.

Second, for our concern, no discrete notion of principal foliation and ridges exist as of today. (Notice that a notion of principal foliation must be available if one wishes to talk about critical points along curvature lines.) Interestingly, for foliations, the simpler problem of computing the orbit structure of a gradient vector field has been addressed under two perspectives, namely simulation of differentiability [EHZ01] and discrete Morse theory [CCL03]. But the question of understanding the relationships between the two diagrams is open, a problem definitely simpler than that of understanding the orbit structure of principal direction fields.

1.3 Paper overview

In section 2, we precisely specify the problem addressed and the contributions made. Fundamentals on ridges are recalled in section 3, while previous contributions are discussed in section 4. The acute rule is discussed in section 5, the certified algorithm is developed in section 6, and the filtering procedures are discussed in section 7. Section 8 presents a method to report umbilics, while experimental results are presented in section 9.

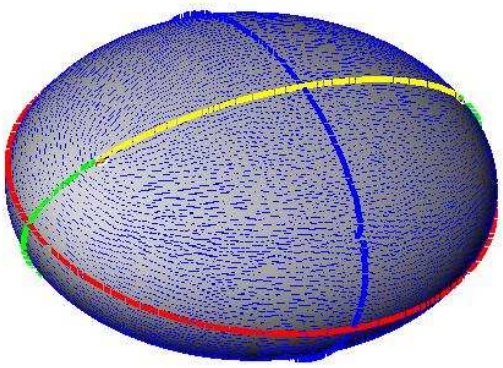

Figure 1: Umbilics, ridges, and principal blue foliation on the ellipsoid (10k points)

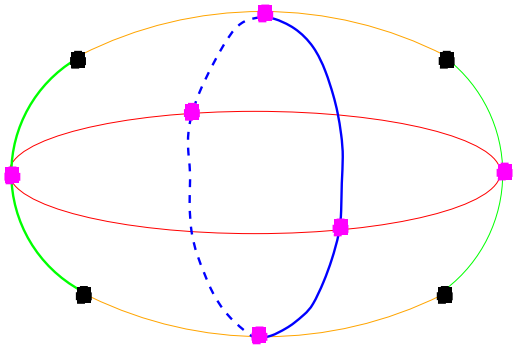

Figure 2: Schematic view of the umbilics and the ridges. Max of k_1: blue; Min of k_1: green; Min of k_2: red; Max of k_2: yellow

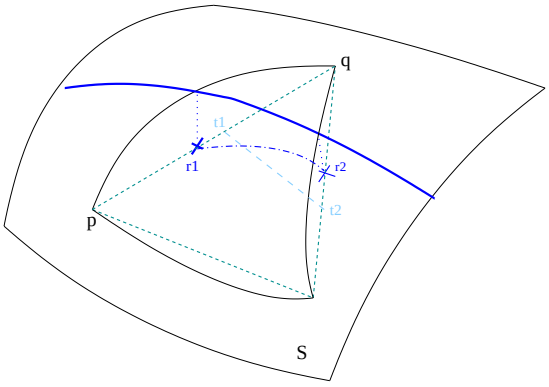

Figure 3: A ridge on a smooth surface, its image (pull-back) $[r_1, r_2]$ on an inscribed triangulated surface, and a straight segment $[t_1, t_2]$ isotopic to this pull-back in the triangle *pqs*

2 Problem overview and contributions

In this section, we define the problem addressed and list the contributions. Readers not familiar with the local theory of smooth surfaces might read section 3 first.

2.1 Ridges of a smooth surface

To state the problem addressed, let us define ridges informally. Given a smooth compact oriented surface S without boundary, a ridge is a curve drawn on S, which either a closed curve —homeomorphic to S^1— and free of umbilics, or an open curve —homeomorphic to the real line— connecting two umbilics. Let us define the set R_S of ridges as the set of these ridge loops, ridge segments and the umbilical points. For generic surfaces, the ridge loops and segments associated to the same principal curvature do not intersect.

A ridge actually consists of the set of points where a bivariate function $b(u,v)$ vanishes. In local coordinates, function $b(u,v)$ encodes the gradient of a principal curvature, and since at any point which is not an umbilic there are two principal curvatures, there are two functions $b_0(u,v)$ or $b_3(u,v)$ witnessing the ridges associated to the maximum (k_1) and minimum (k_2) principal curvatures. We shall assume that principal curvatures are ordered, i.e. $k_1 \geq k_2$, and shall paint in blue (red) objects related to k_1 (k_2). For example, we shall speak of red/blue curvature lines or ridges.

2.2 Problem addressed

Assume we are given a triangle mesh T providing a piecewise-linear approximation of a smooth surface S as an inscribed mesh —that is the vertices of T belong to S. Triangulation T might have been reconstructed from a point cloud [AB99, ACDL00, BC00] or might be the output of a meshing algorithm such as Chew's algorithm [Che93, BO03].

As indicated on Fig. 3, we aim at reporting the *pull-back* of the set R_S of ridges of S onto T. More precisely, we aim at reporting a set R_T^b (R_T^r) of blue (red) elements on T defined as

- a set of points on T corresponding to umbilics of S,
- polygonal simple curves on T corresponding to the blue (red) ridges of S.

From these, we shall say that:

Definition. 1 *Ridges are reported a in topologically coherent fashion provided that the set R_T^b (R_T^r) has the same topology than the set R_S^b (R_S^r); which means that the* push-forward *of R_T^b (R_T^r) should be isotopic to R_S^b (R_S^r) on S, or equivalently the* push-back *of R_S^b (R_S^r) should be isotopic to R_T^b (R_T^r) on T.*

For this problem to make sense, we consider on S, T and the ridge sets the topology induced by $\mathbb{R}^3$. Moreover we assume that the triangulation T and the surface S are isotopic. By now, several sufficient conditions for this property to hold are known. Conditions having a geometric flavor can be found in [ACDL00, APR03], while a more topological one can be found in [CCs04]. Intuitively, the mesh T has to be an embedded mesh lying in the tubular neighborhood of the surface S —which naturally is assumed to have a positive reach. (The reach is defined by inf lfs, with lfs the local feature size i.e. the distance to the medial axis of S.) We assume that the projection from T to S, which projects a point of T to its nearest point on S is an homeomorphism, inducing the isotopy between S and T as constructed in [APR03]. Therefore:

Definition. 2 *The* image *of an edge $[p,q]$ of T is the unique path defined on S by projecting each point of $[p,q]$. Similarly, the* image *of a triangle is its projection onto S.*

As suggested by the previous definition, we shall use the following abuses of terminology. When saying that "an edge is crossed by a ridge" or "a triangle contains an umbilic" we shall mean "the image of the edge is crossed by a ridge" or "the image of the triangle contains an umbilic". Equivalently, this also means that "an edge is crossed by the pull-back of a ridge" or "a triangle contains the pull-back of an umbilic".

Before proceeding, a comment is in order. We do not consider purple points —intersection between ridges of different colors— because the topology of the blue and red sets of ridges are processed separately. The incentive for ignoring purple points is that red and blue ridges are independent since functions b_0 and b_3 are so. Incidently, this assumption alleviates the constraint of reporting the correct topology of ridges around purple points, as depicted on Fig. 4.

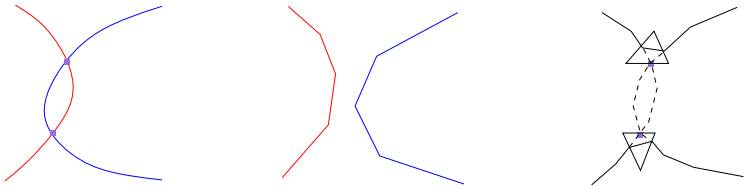

Figure 4: (a)Two ridges of different colors (b)Blue / red ridges reported independently: the topology of each ridge is respected, but that of the union is not (c)Ridges reported simultaneously by the Gaussian extremality: ridges are disconnected at purple points

2.3 Methodology

Consider a mesh T approximating of a smooth compact oriented generic surface S. The main issue tackled is to understand which properties T must have in order for one to report the ridges of S faithfully. We first recall the following:

Observation. 1 *At any point of a smooth oriented surface S which is not an umbilic, orienting the principal directions at a vertex is equivalent to choosing a unit maximal principal vector, and there are two such orientations. (The minimal principal vector is then uniquely defined to be the unit vector so that the basis (maximal vector, minimal vector, normal) is direct.) Orienting the principal directions of an edge of a triangulation T whose endpoints lie on S means orienting them at its two vertices, and there are four possibilities.*

As already noticed, a ridge is witnessed by the zero crossing of a bivariate function $b = 0$ — function b either refers to b_0 or b_3. Since we are given a triangulation T of S, using an idea reminiscent from Marching Lines and Marching cubes, it is natural to seek the zero crossings of function b along the edges of T. But in our case however, function b depends on the orientation of the principal directions [1], and by the above observation there are four possible orientations for the edge. Consider an edge e of the mesh T. In the case of a dense mesh T, and whenever the vertices of e are close on S, one expects the principal direction (maximum or minimum) at the vertices of e to be nearly aligned —at least far from umbilics, which motivates the following:

Definition. 3 (Acute Rule) *Orienting an edge with the* acute rule *consists of choosing maximal principal vectors at the two vertices so that they make an acute angle in the 3D space.*

Once an edge has been oriented, the signs of function b at the endpoints are known, and a zero crossing along the edge can be sought. Connecting such points provides pull-backs of ridge

[1] Functions b_0 and b_3 are coeffi cients of odd terms of the Monge form of the surface —i.e. coeffi cients of the Taylor expansion of the surface expressed locally as a bivariate height function, and their sign therefore depends upon the orientation of the principal frame. See section 3 for the details.

segments, and the only tasks to be taken care of are the processing of ridges near umbilics. From an algorithmic standpoint, using a triangulation to pull-back ridges provides an appealing solution. However, as we shall see, certifying that the results reported are correct subsumes pretty demanding properties on the triangulation.

2.4 Contributions

Having defined the problem addressed, we list the contributions. To do so, we define three different models for the surface S and the triangulation T:

Case I. Mesh T provides an approximation of a smooth compact oriented generic surface S. The Monge coefficients at each vertex can be computed analytically with arbitrary precision.

Case II. Mesh T provides an approximation of a smooth compact oriented generic surface S . However, the Monge patches of S at the vertices of T are approximated using an algorithm such as [CP03].

Case III. Mesh T is not fine enough for a surface S such as above, or approximates a surface which fails some of the requirements.

Having presented these models, we can now comment the three contributions.

Understanding the Acute rule. The A.R. is implicitly used in [Mor90, Mor96, TG95, OBS04a], yet, it has never been analyzed carefully and it its interplay with umbilics, principal foliations and ridges has never been addressed. As a consequence, one cannot specify whether the algorithms fail or not.

We provide a detailed analysis of this rule and highlight its subtle interplay with geometric properties of the surface.

A certified algorithm for case I. Using the analysis of the A.R. , we derive sufficient conditions on T so as to report ridges in a topologically coherent fashion. Under these conditions, we provide a provably correct algorithm which runs through three stages. First, umbilics are reported and classified. Second, ridges are computed and classified using the acute rule. Third, ridges are connected at umbilics with the correct topology.

A *wish for the best* algorithm for cases II and III. If local differential quantities are not known perfectly, if the mesh is not *dense* enough, or if surface S is not a generic surface or is not smooth, it is hopeless to expect reporting a configuration of umbilics and ridges matching that of a smooth generic surface. To face these situations we propose a filtering procedure aiming at reporting significant features of the surface.

3 A primer on ridges

We consider a smooth surface S, oriented, compact and without boundary, embedded in the Euclidean space E^3 equipped with the orientation of its world coordinate system —referred to as the *direct orientation* in the sequel.

First, recall that at each point of the surface which is not an umbilic, there are two orthogonal principal directions d_1, d_2 and two associated principal curvatures k_1 and k_2. These principal directions define two line or direction fields on S, one everywhere orthogonal to the other —so that it is sufficient to study only one of these. Each principal direction field defines lines of curvature which are integral curves of the corresponding principal field, and the set of all these lines defines the principal foliation. Following standard usage, we shall always sort principal curvatures, that is we will always assume $k_1 \geq k_2$. Moreover, objects related to the larger (smaller) principal curvature are painted in blue (red). For example, we shall speak of a blue curvature line or of the blue foliation.

At a point of S which is not an umbilic, the non oriented principal directions d_1, d_2 together with the normal vector n define two direct orthonormal frames. If v_1 is a unit vector of direction d_1 (we call it a maximal principal vector) then there exists a unique unit minimal principal vector v_2 so that (v_1, v_2, n) is direct, and the other possible frame is $(-v_1, -v_2, n)$. (*direct* must be understood with reference with the direct orientation of the world coordinate system mentioned above.) In such a coordinate system, S can be locally described as a Monge form:

$$z = \frac{1}{2}(k_1 x^2 + k_2 y^2) + \frac{1}{6}(b_0 x^3 + 3b_1 x^2 y + 3b_2 x y^2 + b_3 y^3) \tag{1}$$

$$+ \frac{1}{24}(c_0 x^4 + 4c_1 x^3 y + 6c_2 x^2 y^2 + 4c_3 x y^3 + c_4 y^4) + \ldots \tag{2}$$

Moreover, it should be noticed that switching from one of the two coordinate systems to the other reverts the sign of all the odd coefficients on the Monge form of the surface.

Having recalled the fundamental notions related to principal curvatures, let us get to ridges. Defining ridges precisely is a serious endeavor requiring technical notions from contact theory and singularity theory, and we refer the reader to standard textbooks [Por01, HGY+99], as well as to [CP04] for an overview. A blue (red) ridge of a smooth surface is a line on the surface such that at each of its points, the principal blue (red) curvature has an extremum along a blue (red) curvature line. Intuitively, the essence of ridges is best captured by looking at the Taylor expansion of a principal curvature along its corresponding line of curvature. Taking the example of the blue principal curvature, this Taylor expansion is given by [HGY+99]:

$$k_1(x) = k_1 + b_0 x + \frac{P_1}{2(k_1 - k_2)} x^2 + \ldots, \qquad P_1 = 3b_1^2 + (k_1 - k_2)(c_0 - 3k_1^3). \tag{3}$$

A blue ridge point is characterized by $b_0 = 0$, but as illustrated on Fig. 5, the sign of b_0 depends on the orientation of the curvature line. Moreover, the value of P_1 determines the type of the ridge: if $P_1 < 0$ ($P_1 > 0$) the ridge is called elliptic (hyperbolic). In between such regions, one finds isolated points called turning points characterized by $P_1 = 0$. From Eq. 3 —and its dual for k_2, it is also easily seen that an elliptic ridge corresponds to either a maximum of k_1 or a minimum of k_2. Similarly, an hyperbolic ridge corresponds to a minimum of k_1 or a maximum of k_2. The corresponding geometric interpretation when moving along a curvature line and crossing the ridge is recalled on Fig. 6.

To summarize, a ridge is distinguished by its color and its type. When displaying ridges, we shall adopt the following conventions:

- blue elliptic (hyperbolic) ridges are painted in blue (green),
- red elliptic (hyperbolic) ridges are painted in red (yellow).

At last, ridges displayed in black refer either to red or blue ridges.

From a topological standpoint, excluding umbilics, a ridge is either homeomorphic to a circle or to the real line, whence the following:

Definition. 4 *A ridge* segment *is defined as a ridge homeomorphic to the real line, and a ridge* loop *is defined as a ridge homeomorphic to a circle.*

Ridge loops correspond to closed ridges of a given color free of umbilic. Ridge segments connect two umbilics or twice the same one. At an umbilic, the color of ridges always changes —the ridge switching from a blue hyperbolic segment to red hyperbolic segment or vice-versa.

To finish up this review, let us recall the following generic properties:

- a ridge (segment or loop) contains an even number of turning points at which the ridge changes from elliptic to hyperbolic,
- umbilics are 1-ridge umbilics (also called hyperbolic umbilics) or 3-ridges umbilics (also called elliptic umbilics),
- ridges of the same color do not cross (they only meet at 3-ridges umbilics). Two ridges of different colors may cross at a so-called *purple point*.

Apart from the ridge type, we shall pay a special attention to connections between ridges at umbilics and turning points —but we shall not track the crossings of ridges at purple points. These notions are illustrated on the famous example of the ellipsoid on Figs. 1 and 2.

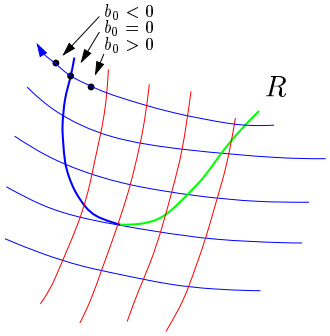

Figure 5: Variation of the b_0 coefficient and turning point of a ridge

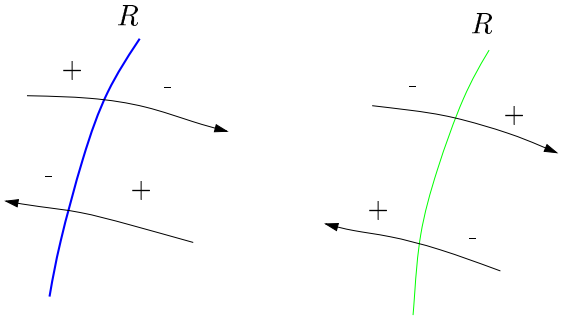

Figure 6: Classification of a blue ridge as elliptic (max of k_1, left), and hyperbolic (min of k_1, right) from the sign change of b_0

4 Previous work

4.1 Reporting umbilics

For parametric surfaces, Morris [Mor90] minimizes the function $k_1 - k_2$. Sander et al [SZ92] note that on noisy data given by range images, the determination of umbilics with a thresholding on the function $k_1 - k_2$ is not reliable. Then they propose to compute indices of the field on the neighborhood of each data point. If the index is not zero then there is a umbilic in the neighborhood and the index is known.

4.2 Reporting ridges

4.2.1 Extrema of curvatures from second order quantities

Ohtake et al [OB01] or Stylianou et al [SF00] enable rough detection of some ridges on triangulations with the computation of second order quantities. The input of these methods is a triangulated surface, principal curvatures and directions are computed at vertices with a surface fitting algorithm. Considering the one-ring of a vertex, they define two points on its boundary, which are assumed to be the intersection with the curvature line going through the center vertex. The center vertex is defined as ridge point if its principal curvature is extremal wrt these two points. Ridge points are connected by piecewise linear lines, decimated and smoothed.

Khaneja et al [KMG98] only use principal curvatures and discard directions. They use dynamic programming to generate a curve maximizing the principal curvature between two given points. These curves do not necessarily correspond to ridges and one has to define start and end points.

4.2.2 Extremum of curvatures from third order quantities

Parametric surfaces. The method developed by Morris [Mor90, Mor96] applies to a surface parameterized on a two dimensional domain. The parametric domain is triangulated and zero crossing are sought on edges using the A.R. . Ridges are defined in the parameter space as zero-sets of a function of the derivatives of the parameterization up to the third order. Orientation problems inherent to the A.R. are dealt with differently away from umbilics and at umbilics. In the former case, a locally non vanishing continuous vector field orienting the principal direction is used. In the later case, two principal vector fields vanishing on two different lines radiating from the umbilic are used. This enables to detect crossings assuming there is at most one such crossing on an edge. The precise position of the crossing point is found with a binary subdivision of the edge. Ridges are recovered by connecting ridge points on each triangle. Umbilics are in triangles with 1 or 3 crossings and can be localized minimizing the function $k_1 - k_2$. As this simple algorithm shall fail near umbilics, a special processing is used to correct the topology.

Implicit surfaces. The method developed by J.P.Thirion et al. [TG95] applies to implicit surfaces defined as the zero-set of a trivariate function. Ridges are defined as the intersection of the surface with another implicit surface. The intersection of two implicit surfaces can be extracted with their

marching line algorithm which gives a set of non-intersecting closed curves. The marching line algorithm relies upon the A.R. , and the problem of orientation is addressed in two ways. First, one can extract the curves associated to the Gaussian extremality $E_g = b_0 b_3$, but these curves are disconnected at umbilics and purple points. Second, one can extract blue and red ridges independently using a local orientation procedure. No attempt toward topological correctness is made, and the focus is rather on stable ridges —the so-called crest lines— and their application to matching. This algorithm is applied to 3D medical images, the implicit surface being defined by convolution of the discrete data with a Gaussian function. Belyaev et al [BPK98] improve the E_g marching line algorithm in an adaptive way with a refinement of the triangulation approximating the implicit surface near ridges.

However, the behavior of these algorithms near umbilics or the hypothesis to get a correct topology are not addressed comprehensively.

4.2.3 Focal surface based methods

The focal surfaces are the centers of principal curvature spheres. These surfaces have singularities along cuspidal edges corresponding to ridges —reciprocally all ridges have such a characterization. In [WB01] and [LA98] authors uses the "area degenerating" effect of the focal surface near its singularities and define a algorithm for polygonal surfaces. Only second order quantities are computed, extremal triangles are reported for visualization purpose.

4.2.4 Medial Axis based methods

It is well known that the boundary of the medial axis of a smooth surface projects onto selected elliptic ridges of the surface [CP04]. M.Hisada et al. [HBK02] use a Voronoi diagram for medial axis approximation and detect the corresponding elliptic ridge points on polygonal surfaces. They consider closed surfaces and the inner/outer medial axis (that is the part of the medial axis inside/outside the surface) and call ridges/ravines the corresponding points which are subsets of positive max of k_1/negative min of k_2.

This method apparently carries the advantage of not requiring the computation of high-order derivative on the surface. But the medial axis being a very unstable construction, the method cannot provide satisfactory results unless special care is devoted to filtering. (Medial axis filtering procedures have recently been developed [eRS04, CL04], but these procedures have not been applied to robust ridge detection.)

A second drawback of this methods concerns the type of ridges reported. Reporting the medial axis using a Voronoi diagram relies upon the convergence of empty balls (encoded in the Voronoi diagram) to maximal balls (whose centers are on the medial axis). Therefore, while Voronoi diagrams could certainly be made local enough so as to report all elliptic ridges, a major difficulty arises for hyperbolic ones. Since for such ridges the intersection of the osculating sphere and the surface does not reduce to a point, no corresponding empty ball exists in Voronoi.

5 Zero crossings of b_0 and b_3 and the acute rule

Recall that we wish to report pull-backs of ridges by tracking the zero crossings of the b_0 and b_3 coefficients of the Monge form of S using the acute rule. We aim at providing sufficient conditions on T for the acute rule to yield reliable results.

5.1 The acute rule in action: examples

Before providing a formal description of the A.R. and to see which conditions on the triangulation T we shall need, we illustrate the interplay of the A.R. with the geometry of ridges and foliations.

A.R. and ridge crossing. Consider Fig. 7. Traveling from p to q on the image of the edge and following the lines of curvature with the indicated orientation results in going toward the ridge, crossing it, and getting away from it. the ridge crossing is witnessed by a sign change of b_0, which is detected iff the orientations of the principal directions are as indicated —or both in opposite direction. This orientation is provided by the A.R. , but also by *continuity* as follows. Starting from one endpoint and propagating the orientation of the first principal vector by continuity along the image of the edge results in the orientation of the second principal vector provided by the A.R. .

A second illustrative example is depicted on Fig. 8. As opposed to the previous one, the lines of curvature going through (rq) does not cross the ridge. But the sign of b_0 remains the same after the crossing. The crossing is thus witnessed by the sign change of b_0 which is again detected provided the principal directions at p and q are oriented as indicated.

A.R. and transversality of foliations. Consider the image of an edge. Call the edge transverse if its image intersects each line of curvature transversally. Figures 9 and 10 feature two transverse edges intersecting a ridge R at a point —spotted by the green dot. As illustrated by these two examples, the acute rule gives the correct orientation in the first case only. Consider now Figures 11 and 12 , which feature two non transverse intersections. Here again, the acute rule provides the correct orientation in the first case but not the second one. (Notice on these examples that orienting the edge by continuity requires going through a point belonging to a curvature line tangent to the image of the edge.)

The failure cases of the A.R. come from a large deviation in the $3D$ space of the maximal principal directions between the endpoints of the edge. The deviation of the maximal principal direction along a curvature line on the surface has two components which are extrinsic (the normal curvature) and intrinsic (the geodesic curvature). The denser the mesh, the shorter the edges of the triangulation and the smaller the extrinsic deviation. But the situation is different for the geodesic curvature which is high near umbilics, but might also be large away from umbilics as depicted on Fig. 12 –the left two curvature lines. In such cases, the acute rule may yield erroneous orientations and no guarantee can be provided.

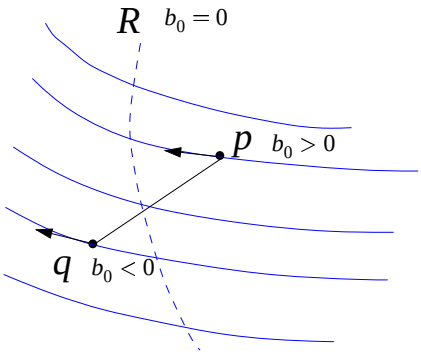

Figure 7: The image of an edge, a blue ridge and the blue principal foliation

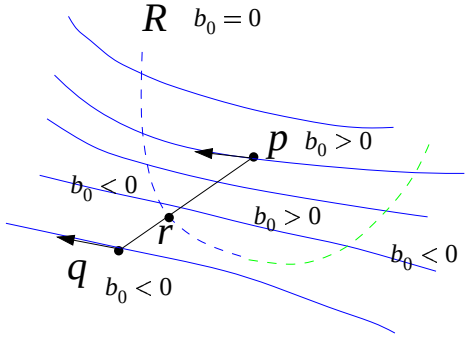

Figure 8: The image of an edge, a blue/green ridge and the blue principal foliation

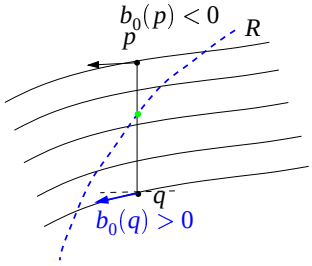

Figure 9: Edge and foliation are transverse, orientation is correct

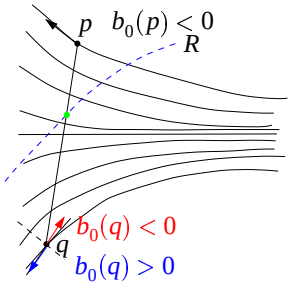

Figure 10: Edge and foliation are transverse, orientation is erroneous

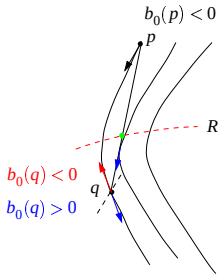

Figure 11: Edge and foliation are not transverse, orientation is correct

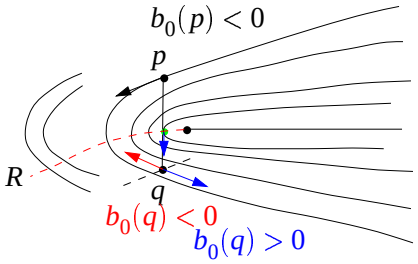

Figure 12: Edge and foliation are not transverse, orientation is erroneous

5.2 The acute rule

5.2.1 Simplicity hypothesis for smooth surfaces

Following a well established trend and to keep the description of the algorithm tractable without having to consider degenerate situations, we shall require genericity conditions.

Hypothesis. 1 *The vertices and edges of T are assumed to meet the following* generic *conditions:*

S1 *the maximal principal directions at the two vertices of an edge are not at right angle;*

S2 *the image of the edges does not contain any umbilic;*

S3 *no ridge goes through a vertex of T;*

S4 *ridges intersect the images of edges transversally.*

Second, we shall assume the mesh is dense enough so as to guarantee that:

D1 *the image of a triangle onto the surface contains at most one umbilic;*

D2 *no ridge is included in the image of a triangle.*

5.2.2 Detecting a ridge crossing with the acute rule

Equipped with the previous hypothesis, we shall derive the applicability condition of the A.R. .

Consider an edge of T. By hypothesis S1 —principal vectors are not orthogonal, orienting an edge with the A.R. is always possible. There are actually two such possible orientations and only one once an orientation is chosen at an end point.

Consider now a simple curve C homeomorphic to the line-segment $[0,1]$ drawn on S. If C does not contain any umbilic —by hypothesis S2, one can define along C two continuous unit vector fields which orient the maximum direction field. Once a unit maximum vector has been chosen at one endpoint of C, we call the orientation induced by the corresponding vector field at the other endpoint the *orientation by continuity*. Replacing C by the image of an edge yields the following:

Definition. 5 *The orientation of the principal directions of an edge given by the A.R. is called* correct *(erroneous) if it coincides (differs) with (from) one of the two orientations by continuity.*

Having discussed these orientation issues, we finally raise the observation used to track the zero crossings of b_0 and b_3 along an edge:

Observation. 2 *Let C be the image of an edge* $[p,q]$. *Assume the simplicity hypothesis S3 and S4 hold and that C crosses at most one blue ridge. If the orientation of the principal directions of the edge is correct then a blue ridge crosses the image of the edge iff* b_0 *changes sign:* $b_0(p)b_0(q) < 0$.

5.2.3 Erroneous orientations

As illustrated above, the A.R. does not always yield the correct orientation of an edge. To see which consequences this may have, let us study the decisions induced by an erroneous orientation assuming each edge is intersected by at most one ridge of each color. We discuss the result induced by the A.R. and compare it to decisions made using the Gaussian extremality $E_g = b_0 b_3$:

- no ridge crosses the edge: b_0 and b_3 report one blue and one red ridge crossings, whereas E_g reports the correct answer;
- only one blue (resp red) ridge is crossing the edge: b_0 and b_3 report one red (resp blue) ridge crossing —that is colors are swapped, E_g also reports one crossing —of un-specified color;
- one blue and one red ridge cross the edge: b_0 and b_3 report no crossing, E_g also reports no crossing.

In conclusion, neither the tandem b_0, b_3 nor E_g are able to report the correct answer an all cases. The only advantage of E_g is that it never reports crossings when there is none. But it misses the two crossings if there are two, does not specify the color of the ridge, and more consequently is unable to preserve the topology of ridges near purple points —see Fig. 4(c).

5.3 Finding the type of a ridge

Once a ridge has been reported, one needs to classify it as elliptic or hyperbolic. As recalled in section 3, the type of a ridge can be computed from 4th order differential quantities at a ridge point using Eq. (3), which also has the geometric interpretation recalled on Fig. 6. A natural question is therefore: can the type be determined geometrically with only third order quantities —a procedure likely to be more stable than the previous one?

Consider an edge along which the sign of b_0 changes. As illustrated on Fig. 13, the knowledge of the principal directions at the edge endpoints together with the location of the ridge point r falls short from providing enough information to state the ridge type.

In [OBS04a] the following heuristic is used: the ridge is tagged as elliptic if the principal vectors u_1 and u_2 at the endpoints v_1 and v_2) are such that $b_0(u_1)(u_1.(v_2 - v_1)) > 0$ and $b_0(u_2)(u_2.(v_1 -$

$v_2)) > 0$. This rule takes into account the principal directions at the endpoints of the edge but not local information on the ridge itself. This method implicitly assumes that orienting the principal direction with a vector making a acute angle with the edge leads the curvature line towards the ridge. As an example, the rule fails at vertex v_2 of Fig. 14 and misclassifies the second situation of Fig. 13. Such a situation is likely to occur when an edge is almost parallel to the ridge.

To accommodate the configurations of Fig. 13, we propose the following geometric criterion. Consider a triangle crossed by a ridge segment $[r_1 r_2]$. The idea is to use the direction information given by $[r_1 r_2]$ to distinguish between the two types —see Fig. 5.3.

The sign of b_0 for a maximal principal vector pointing towards the ridge segment $[r_1, r_2]$ for the triangle is defined as the sign appearing at least at two vertices. If this sign is positive then there is a maximum of k_1 and the ridge is elliptic else it is hyperbolic. Let (v_1, v_2, v_3) be a triangle, r_1 and r_2 the ridge points on the edges $[v_1, v_2]$ and $[v_1, v_3]$. As S is oriented, assume (v_1, v_2, v_3) is direct. Then the sign of $b_0(v_1)$ for an orientation pointing toward the ridge segment $[r_1, r_2]$ is that of:

$$sign(b_0(u_1)) \det(u_1, r_2 - r_1, n)$$

with u_1 any of the two orientations of $d_1(v_1)$ and $b_0(u_1)$ the associated b_0 value. The sign of $b_0(v_2)$ (resp. $b_0(v_3)$ for an orientation pointing toward the ridge segment $[r_1, r_2]$ is that of:

$$-sign(b_0(v_2)) \det(u_2, r_2 - r_1, n) \qquad \text{(resp. } -sign(b_0(u_3)) \det(u_3, r_2 - r_1, n))$$

with u_2 (resp. u_3) any of the two orientations of $d_1(v_2)$ (resp. $d(v_3)$).

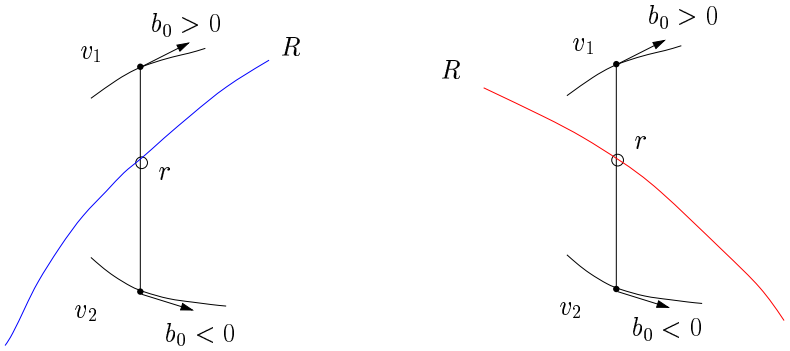

Figure 13: An elliptic (left), and an hyperbolic (right) ridge crossing the image of an edge of the triangulation T

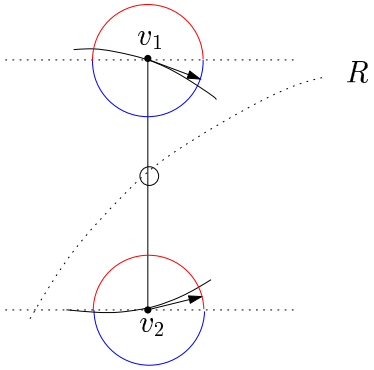

Figure 14: The tagging rule of [OBS04a]

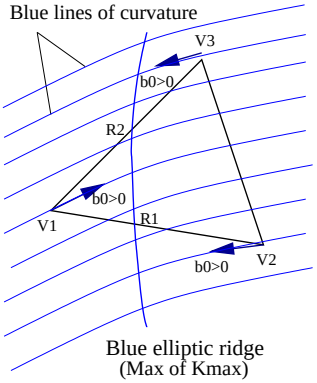

Figure 15: Determining the type of a ridge segment using third order properties

6 Algorithm for case I

An algorithm reporting a set of ridges R_T^b, R_T^r is called certified if the topology of R_T^b, R_T^r match that of R_S^b, R_S^r as specified by definition 1. The goal of this section is to design sufficient conditions on T so as to have certified algorithms.

6.1 Certified algorithms and sufficient hypothesis

For the surface, the setting assumed is that of case I, that is the surface is generic and differential quantities are known exactly. For the triangulation T, we assume that all generic and density conditions of section 5.2.1 are met, and also assume that each edge of the triangulation is oriented with the A.R. . Under these assumptions, we provide three sets of additional sufficient conditions on T, that is conditions coming along with a certified algorithm.

Although the previous assumptions are unrealistic in practice, the conditions developed highlight the subtle interplay of the A.R. with the geometry of the ridges and foliations of the surface. As explained in section 2.2 blue and red ridges are processed separately so we focus on blue ridges.

The structure of each section is as follows. First, we discuss the hypothesis made on the triangulation T. Second, we provide an algorithm using these hypothesis. Third, we prove that under the hypothesis, the ridges computed have the correct topology.

6.2 A sufficient yet unrealistic condition

Comment. We start with the most ideal assumptions on T. These assumptions are so strong that they also define implicitly the location of umbilics which do not have to be computed a priori.

Hypothesis 1.

1. each edge is intersected by at most one (blue) ridge,
2. each edge (oriented with the acute rule) has a correct orientation,
3. only ridge segments connected to an umbilic are crossing the triangle containing this umbilic, and they cross it at most twice.

Algorithm. First we detect crossings of ridges on edges according to Observation 2: a ridge crossing r is drawn on the edge $[p,q]$ if $b_0(p)b_0(q) < 0$. If a triangle has 1 or 3 ridge crossings, it contains a 1 or 3 ridges umbilic. We define the corresponding point of R^b_T as the center of inertia of the triangle. We connect this point to the ridge crossings on edges with a segment. If a triangle has exactly 2 ridge crossings, we connect these points with a segment. The polygonal lines of R^b_T are composed of these segments.

Proof of correctness. Hypothesis 1.1 and 1.2 imply that the detection of ridge crossings on edges in correct and that each triangle has 0,1,2 or 3 edges intersected by a ridge:

- If there is one ridge crossing in a triangle, then a ridge segment ends in this triangle which must contain a 1-ridge umbilic (Fig. 16).
- If there are two ridge crossings in a triangle, then the same ridge is crossing the triangle (Fig. 17). Otherwise two ridge segments were ending in the triangle which must contain two umbilics, this contradicts hypothesis D1.
- If there are three ridge crossings then either three different ridge segments end in the triangle which must contain a 3-ridge umbilic; or two ridge segments end in the triangle, one of them connected at both ends to a 3-ridge umbilic (Fig. 18). The other possibilities cannot occur due to hypothesis 1.3: either a single ridge segment crosses the triangle thrice, or there is one ridge segment connected to a 1-ridge umbilic and another ridge crossing the triangle but not connected to the umbilic (Fig. 19).

Hence a one-to-one correspondence between umbilics and centroids of 1 and 3 ridge crossings triangles is given. The topology of each ridge pulled-back on T is preserved in each triangle. Due to hypothesis D2, no ridge is missing. In conclusion the set R^b_T has the same topology than the pull-back of R^b_S.

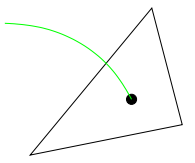

Figure 16: One ridge crossing

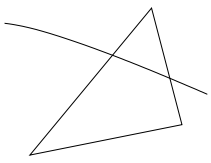

Figure 17: Two ridge crossings

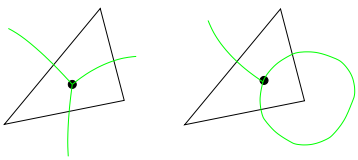

Figure 18: Three ridge crossings

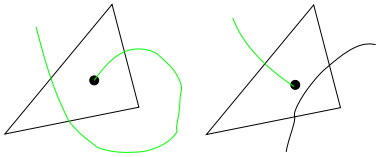

Figure 19: Three ridge crossings not allowed

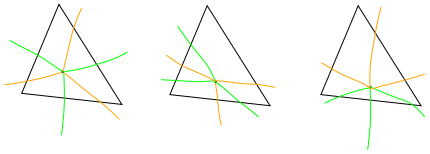

Figure 20: A three-ridges umbilic within a triangle. The first two umbilics are symmetric, the third is unsymmetric

6.3 Improvement at umbilics

Comment. The assumptions used above are not realistic for several reasons. First, at umbilics, the maximal principal direction makes half a turn, so that at least one edge of the triangle containing the umbilic is likely to have erroneous orientation —see Fig. 12 for an index $1/2$ umbilic. Moreover, at a 3-ridge umbilic, requiring that the three ridges intersect the three edges of its triangle is a strong requirement —see Fig. 20 for a symmetric 3-ridge umbilic, and is actually impossible for a unsymmetric 3-ridge umbilic.

Hence we need to extend the neighborhood of an umbilic with a patch (i.e. a topological disk as defined in section 8.1) with more boundary edges so as to subdivide the deviation of the direction field and separate the ridge crossings.

Hypothesis 2.

1. triangles containing umbilics are detected,
2. topological disk patches made of triangles are given for each triangle containing an umbilic so that:
   (a) these patches do not overlap,
   (b) there is no ridge segment or loop included in any patch,
   (c) only ridge segments connected to an umbilic are crossing the border of this umbilic patch and they cross it at most twice.
   (d) all edges outside and on the border of the patches:
      i. have a correct orientation,
      ii. are intersected by at most one ridge.

Algorithm. First the points corresponding to umbilics in R_T^b are obviously the centroid of the triangles containing an umbilic. Second, we detect crossings of ridges on edges outside and on the border of the patches. Outside patches, if a triangles have 2 ridge crossings, we connect them with a segment. On the boundary of a patch, there are 1 or 3 crossings for a 1 or 3 ridges umbilic, and these crossings are connected to the umbilic by polygonal lines inside the patch (for a 3 ridge umbilic, these lines should not cross). The polygonal lines of R_T^b are composed of all these segments.

Proof of correctness. Hypothesis 2.2.d implies that the detection of ridge crossings outside and on the border of patches is correct. On this part of T there is no umbilic, hence a triangle has 0 or 2 ridge crossings, if any then the same ridge is crossing the triangle. The analysis of the ridge crossings on the patches is similar to that of the 1 or 3 ridge crossings triangles of the last section. Due to hypothesis 2.2.c, only ridges connected to the umbilic of the patch can cross its border.

If the patch contains a 1-ridge umbilic, the ridge segment connected to this umbilic must cross the border of the patch (hypothesis 2.2.b). If this ridge segment crosses the border more than once, then either it is connected twice to the 1-ridge umbilic or it crosses at least thrice the border —since it has to leave the region as indicated on Fig. (19); this could not occur since the umbilic would no longer be a 1-ridge umbilic or hypothesis 2.2.c would not be satisfied.

If the patch contains a 3-ridge umbilic, it has 2 or 3 different ridge segments connected to the umbilic and crossing the border thrice.

In conclusion, a patch has 1 or 3 ridge crossings on its border if it contains a 1 or 3-ridge umbilic. Due to hypothesis D2 and 2.2.b, no ridge is missing. The topology of ridges inside patches and in triangles outside patches is preserved, so that the set R_T^b has the same topology than the pull-back of R_S^b.

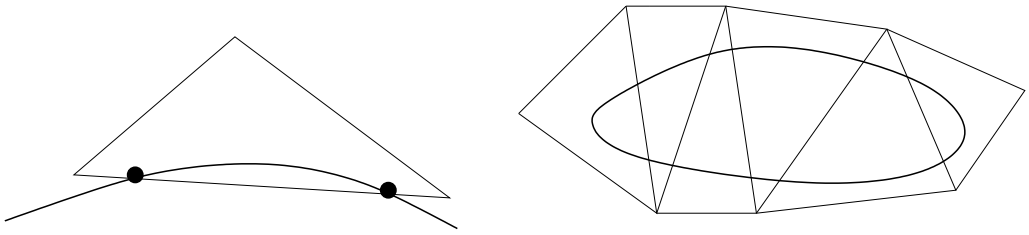

Figure 21: A double crossing and a closed ridge with only double crossings that cannot be detected

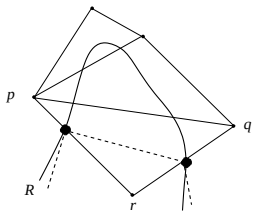

Figure 22: Double crossing with an edge correctly processed: the ridge (solid curve) is shortened (dashed polyline)

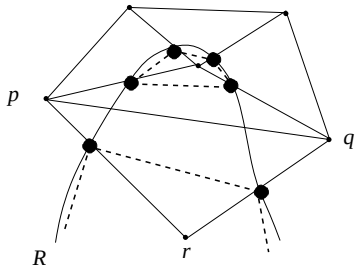

Figure 23: Double crossing with an edge leading to an erroneous topology: the ridge (solid curve) is split (dashed polyline)

6.4 Double crossings

Comment. A ridge can intersect twice the same edge for example if it is almost tangent to the edge (Fig. 21). Such a double crossing of the same ridge is not detected with the sign of b_0 since it changes twice along the edge. Hence, the polygonal ridges constructed as previously explained ignore double crossings. How do these double crossings affect the ridges topology? To answer this question, one needs to compare the pull-back of a ridge on T and the computed ridge.

A first necessary condition for a ridge loop is that the loop must not retract to a single point —because it will not be detected at all!— as illustrated on Fig.21. A second condition is that the topology is not modified if one can "retract" the actual ridge on the computed one without enclosing vertices of T. As an illustration, this is possible on Fig. 22 but not on Fig. 23.

Hypothesis 3. Hypothesis 2 is assumed except 2.2.d.ii which is weakened: edges outside patches can be intersected at most twice by the same ridge. Moreover, the double crossings do not modify the topology of this ridge.

Algorithm. The method is the same that for hypothesis 2.

Proof of correctness. Double crossings are not detected and a ridge with such double crossings is *retracted* as indicated on Fig. 22. (Note that such retractions do not preserve intersection, but ridges of a single color do not intersect.) Hence the topology of each ridge is preserved and so is the topology of the set R^b_T.

6.5 Algorithm for case I

In this section we recall all steps of the algorithm proved to be able to report the ridges with a correct topology with the assumption of hypothesis 3. Moreover we specify the ridge types of the polygonal lines.

Detect umbilics and patches Run the algorithm of section 8.

Processing an edge outside patches Processing an edge requires three steps:

1. The edge $[p,q]$ is oriented with the acute rule.
2. A blue ridge crossing r is detected if $b_0(p)b_0(q) < 0$. The position of r on the edge is computed by linear interpolation:
$$r = \frac{|b_0(q)|p + |b_0(p)|q}{|b_0(q)| + |b_0(p)|} \tag{4}$$
3. We associate to the point r the differential quantities interpolated as above from the vertices p and q. For example the type elliptic or hyperbolic of r is given by the sign of $P_1(r)$.

Processing a triangle outside patches If there are two ridge crossings on two edges of a triangle, the two points are connected by a segment. The type of the segment can be computed in two ways:

- with fourth order properties, if the type of both ridge crossings is the same it defines the type of the segment. If types are distinct, then the segment is defined as a turning point segment;
- with third order properties as explained in section 5.3.

Notice though, that while the topology of the ridges is certified, the type of the ridges is not.

Processing patches around umbilics There is one or three ridge crossings on edges on the border of a patch. We join them to the centroid of the triangle containing the umbilic with polygonal lines drawn on T. The type of these polygonal lines is hyperbolic.

7 Algorithm for cases II and III

7.1 Algorithm for case II

Recall that in this case, S is supposed to be generic but the Monge coefficients at the vertices of T are approximated. Moreover T may not be dense enough to satisfy hypothesis D1 and D2. We can still run the algorithm described previously. In the sequel, we analyze the possible outcomes of the algorithm.

Detection of umbilics: a necessary condition for correctness for a compact surface S without boundary is that the sum of indices of umbilics matches the Euler characteristic of S.

Processing edges: some erroneous orientations may occur, this implies that some ridge crossings are added or missed (cf. 5.2.3).

Processing triangles: due to erroneous ridge crossings, some triangles can have only 1 or 3 ridge crossings. This may prevent some ridges to reach their umbilics. The topology of ridges can also be affected by non-retractable double crossings.

7.2 Case III: Filtering sharp ridges and crest lines

For real world applications dealing with coarse meshes, or meshes featuring degenerate regions or sharp features, one cannot expect a configuration of umbilics and ridges matching that of a smooth generic surface. For example, if the principal curvatures are constant —which is the case on a plane or a cylinder, then all points are ridge points.

As an alternative to ridges, some applications focus on the so-called crest-lines [PAT00]. A crest line is an elliptic ridge which is a maximum of $\max(|k_1|,|k_2|)$, and one may see these lines as the visually most salient curves on a surface. (Notice that these lines do not cross each other and avoid umbilics.)

In a broader setting, an appealing notion is that of *sharp* ridge or *prominent* ridge. Since ridges are witnessed by zero crossings of b_0 and b_3, one can expect erroneous detections as long as these coefficients remain small. In order to select the most prominent ridge points, we can focus on points where the variation of the curvature is fast along the curvature line. From Eq. (3), these points are characterized by a high absolute value of $P/(k_1-k_2)$ since when b_0 is zero the leading term is that of x^2.

Observation. 3 *At a ridge point, the second derivative of k_1 along its curvature line satisfies:*

$$k_1''(0) = P_1/(k_1-k_2).$$

Filtering ridge points using this observation retains more information than focusing on crest lines since hyperbolic ridges might be sharp also. This observation might find applications in surface registration.

Using the previous observation, one can define the *sharpness of a ridge segment* as the integral of the absolute value of $P_1/(k_1-k_2)$ along the ridge. This gives a scale-independent threshold and an associated sharpness-filter. It should be noticed this filter is different from the *strength of a ridge segment* as defined in [OBS04b], which is the integral of the curvature along the line.

8 Umbilic detection

We want to detect triangles of T containing an umbilic. The method combines a minimization and an index computation on the neighborhood of each triangle of T. The size of the neighborhood is

the only parameter of the algorithm. There is no guaranty that this method gives a correct answer, nevertheless it gives satisfaction in practice.

8.1 Finding patches around triangles

Given a triangle t, we aim at defining a collection of triangles around it so that this collection defines a topological disk on the triangulation T. To do so, the most natural way consists of using the successive rings of triangles around t. Let us define the 0-ring neighbors of a triangle as the triangle itself. The k-th ring neighbors is defined recursively from the $(k-1)$-th ring neighbors by adding to the $(k-1)$-th ring neighbors the triangle incidents to this set. However, the k-th ring neighbors may not form a topological disk as indicated on Fig. 24.

To get around this difficulty and starting from the patch consisting of the triangle t, we iteratively construct a patch P by the following algorithm. Each triangle incident to one or two edges of the boundary of P is placed into a priority queue —the grade being the distance between the centroid of this triangle to that of t. Then, patch P is enlarged with the triangle t' having the least grade provided $P \cup t'$ remains a topological disk. If triangle t' is stitched to the patch, its neighbors are inserted in the queue —if they are not already in it. The process stops as soon as the least distance is more than some threshold.

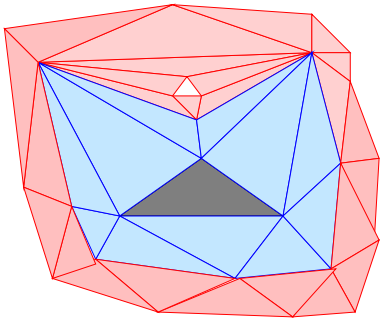

Figure 24: The two-rings triangles may not form a topological disk

8.2 Reporting umbilics

We aim at identifying some triangles of the mesh T as generic umbilics of index $\pm 1/2$. Notice that on a smooth surface, the function $k_1 - k_2$ is always positive, and vanishes only at umbilics. To use this criterion, we define the value of $k_1 - k_2$ for a triangle as the arithmetic mean of the values at its vertices. The detection proceeds in three steps combining the minimization of the function $k_1 - k_2$ and an index computation:

1. Compute a patch around each triangle;
2. Select a triangle if it minimizes the function $k_1 - k_2$ amongst all the triangles of its patch;

3. Compute the index of the principal directions on the contour of the patches of triangles selected in step 2.

The first step is parameterized by a positive number defining the *size* of the patch, which is defined as a multiple of the greatest grade of the one-ring triangles —grade bearing the meaning of the previous subsection. The second step is straightforward. The third step requires following the contour of each patch, orienting the maximal principal directions so that two consecutive of them make an angle smaller than $\pi/2$, and computing the index by adding the angle deviations of the corresponding maximal principal vectors. (See the acute rule in section 5.2 for more details on the orientation of two consecutive principal directions.) We then keep only triangles with an index $\pm 1/2$. Theoretically, the computation of the index of a direction field at a point on a manifold needs the use of a chart [BG88, chap.7,p.260]. Here, we assume that the projection on the tangent plane of the studied vertex is a chart, the index is computed for principal vectors projected on this plane. From a theoretical perspective, if the triangulated surface T is inscribed in a smooth generic closed surface, then the sum of indices of umbilics equals the Euler characteristic of the surface.

The Monge coefficients of the umbilical triangle (defined as the arithmetic mean of that of its vertices) should be close to those of the umbilic it is identified to. Hence we can use them to decide further knowledge of the umbilic type. For example, the sign of $S = (b_0 - b_2)b_2 - b_1(b_1 - b_3)$ which is a third order quantity should also give the index of the umbilic. It is likely to be less accurate than our computation of step 3 which uses only second order quantities. Other invariants of third order decide the type 1-ridge or 3-ridge and the symmetry of ridges at the umbilic (cf. [CP04]).

9 Experimental results

Experimental setup. We conducted experiments on smooth surfaces meshed with the algorithm described in [BO03], and on coarse models reconstructed from scanned models. In both cases, the Monge forms at the vertices were estimated with the fitting algorithm [CP03].In the case of analytically known surfaces, estimating the Monge forms instead of computing them formally introduces numerical errors penalizing the algorithm. In any case and as discussed above, two additional sources of errors are due to (i)erroneous orientations induced by the Acute Rule (ii)double crossings modifying the topology.

Smooth surfaces. The first test surface is the ellipsoid of Fig.1, where the algorithm reports perfectly the well-known patterns of umbilics and ridges. The second test surface is a blend between two ellipsoids displayed on Fig.25 and defined by the following equation:

$$1 - \exp(-0.7(\frac{x^2}{0.15^2} + \frac{y^2}{0.25^2} + \frac{z^2}{0.35^2})) + \exp(-0.7(\frac{(x-0.25)^2}{0.1^2} + \frac{(y-0.1)^2}{0.2^2} + \frac{(z-0.1)^2}{0.3^2})) = 0.$$

The umbilic detection algorithm gives good results for a size of the patch around a vertex 2 or 3 times the size of its 1-ring. Indeed, the greater the patch, the fewer the number of points collected

with the minimization step. More importantly, the boundary of large patches are not too close from umbilics, which as explained in section 6.3 favors a correct computation of the index. As an example on this surface, 14 points are detected by the minimization algorithm —using a patch size 3 times the 1-ring size, while six umbilics of index $+1/2$ and two of index $-1/2$ are reported by the index computation. Notice that this complies with the Euler characteristic. On this model, apart from isolated ridge segments of erroneous type, the ridges reported look convincing! The topology near umbilics is easily corrected, and one can also observe purple points and turning points.

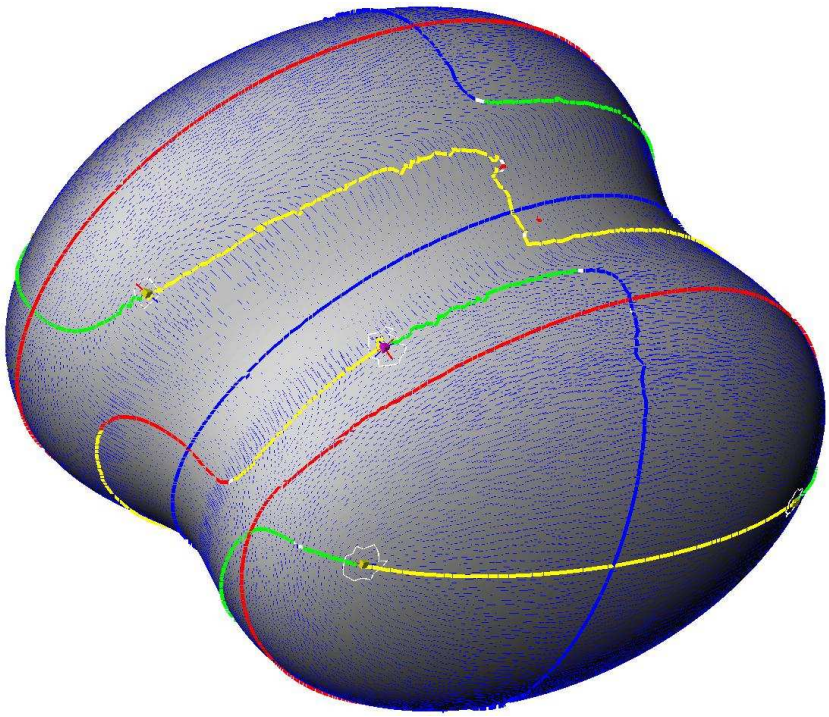

Figure 25: Implicit blending of two ellipsoids (40k points)

Coarse meshes. Figure 26 features a coarse mesh of a mechanical part with first all crests lines, second crest lines filtered with their strength and third filtered with their sharpness. The whole configuration of crests is very noisy, but while the strength-filtered crest lines remain affected by noise on cylindrical parts, the sharpness-filtered crests actually appear on salient features of the model as expected.

This example is especially interesting since it features flat and cylindrical regions. Each point on such region has constant thus critical principal curvatures and is therefore a ridge point. A strength of our approach is to provide a simple filtering procedure to get rid of these ridge points. This example also calls for a comment on methods aiming at reporting ridges after to have performed an

interpolation / approximation of the model. Again, if the model features flat or cylindrical regions, such algorithms report many insignificant ridges —that would also have to be filtered.

The whole process of this mesh takes about 10 seconds on a 2GHz PC.

Figure 27 features the David model (380k pts) processed in 2 minutes. The sharpness filter is used, but the weight filter gives similar results on this model more generic than the mechanical one.

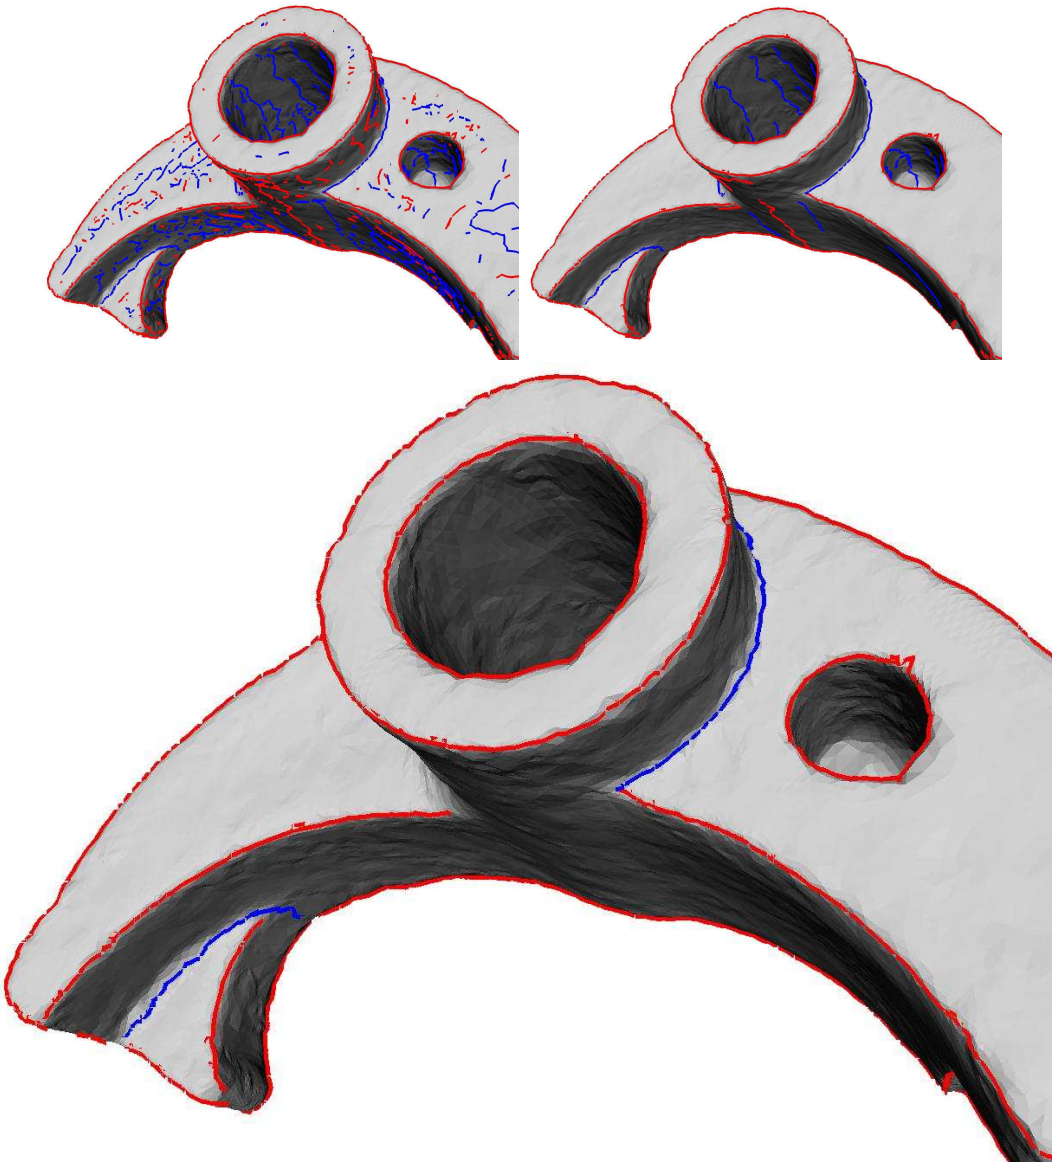

Figure 26: Mechanical part (37k pts): All crest lines, crests filtered with the strength and crests filtered with the sharpness. Notice that any point on a flat or cylindrical part lies on two ridges, so that the noise observed on the left two Figs. is unavoidable. It can however easily be filtered out using Observation 3.

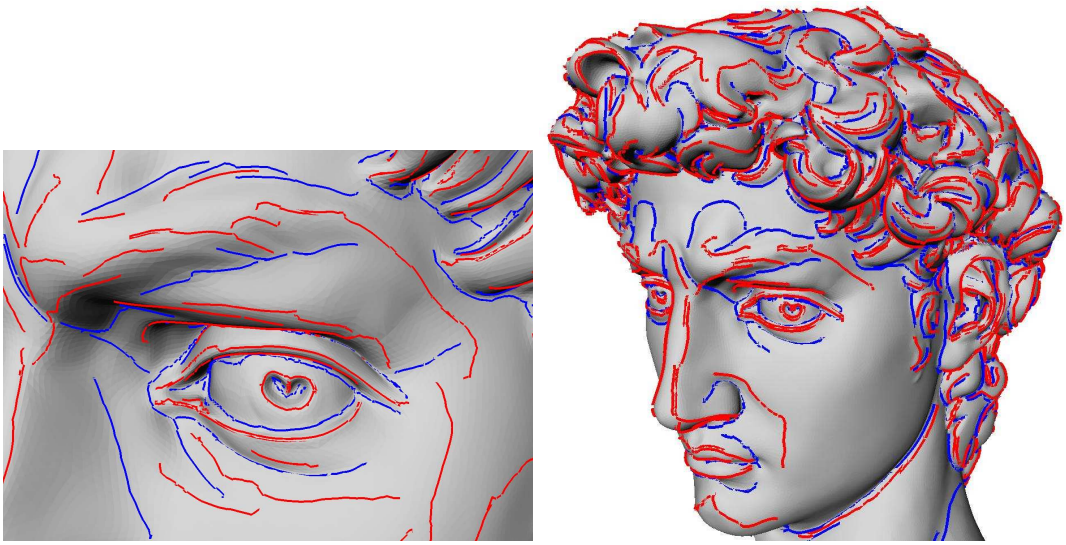

Figure 27: Filtered crest lines on a 380k pts model

10 Conclusion

In this paper, we present the first detailed analysis of the Acute Rule, an orientation procedure for principal vectors present in all algorithms reporting ridges of different colors separately. This analysis is used to derive sufficient conditions on a triangulation T approximating a smooth compact generic surface S so as to report the ridges of S in a topologically coherent fashion. At last, we develop an algorithm and a filtering procedure aiming at reporting the most salient features of a coarse mesh T, and provide experimental results booth for sampled smooth surfaces and coarse meshes. In spite of these contributions, several challenging issues remain.

First, although we provide sufficient conditions for ridges to be reported correctly, these conditions are not constructive and computing compatible triangulations is a open question. As observed when discussing the acute rule, solving this question is likely to require developing an algorithm for computing principal foliations, another challenging task. Apart from these *smooth realm* related issues, developing alternative methods based on more discrete concepts is another open research avenue. There again, one could design discrete foliations and discrete ridges. For example, extremality coefficients derived from the Normal Cycle framework could yield interesting concepts.

Acknowledgments. Steve Oudot is acknowledged for providing his meshing software.

References

[AB99] Nina Amenta and Marshall Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.*, 22(4):481–504, 1999.

[ACDL00] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. In *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, pages 213–222, 2000.

[APR03] N. Amenta, T. Peters, and A. Russell. Computational topology: ambient isotopic approximation of 2-manifolds. *Theoretical Computer Science*, 305:3–15, 2003.

[BC00] Jean-Daniel Boissonnat and Frédéric Cazals. Smooth surface reconstruction via natural neighbour interpolation of distance functions. In *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, pages 223–232, 2000.

[BCM03] V. Borrelli, F. Cazals, and J-M. Morvan. On the angular defect of triangulations and the pointwise approximation of curvatures. *CAGD*, 20(6), 2003.

[BG88] M. Berger and B. Gostiaux. *Differential geometry: Manifolds, Curves and Surfaces*. Springer, 1988.

[BO03] J.-D. Boissonnat and S. Oudot. Provably good surface sampling and approximation. In *Symp. on Geometry Processing*, 2003.

[BPK98] A. Belyaev, A. Pasko, and T Kunii. Ridges and ravines on implicit surfaces. In *IEEE Computer Graphics International*, pages 530–535, 1998.

[CCL03] F. Cazals, F. Chazal, and T. Lewiner. Molecular shape analysis based upon the morse-smale complex and the connolly function. In *ACM Symposium on Computational Geometry*, 2003.

[CCs04] F. Chazal and D. Cohen-steiner. A condition for isotopic approximation. In *ACM Symposium on Solid Modeling*, 2004.

[Che93] L. P. Chew. Guaranteed-quality mesh generation for curved surfaces. In *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, pages 274–280, 1993.

[CL04] F. Chazal and A. Lieutier. Stability and homotopy of a subset of the medial axis. In *ACM Symp. Solid Modeling and Applications*, 2004.

[CP03] F. Cazals and M. Pouget. Estimating differential quantities using polynomial fitting of osculating jets. In *Symp. on Geometry Processing*, 2003.

[CP04] F. Cazals and M. Pouget. Smooth surfaces, umbilics, lines of curvatures, foliations, ridges and the medial axi: a concise overview. Technical Report 5138, INRIA, 2004.

[CSM03] D. Cohen-Steiner and J.-M. Morvan. Restricted delaunay triangulations and normal cycle. In *ACM Symposium on Computational Geometry*, 2003.

[EHZ01] Herbert Edelsbrunner, John Harer, and Afra Zomorodian. Hierarchical morse complexes for piecewise linear 2-manifolds. In *Symposium on Computational Geometry*, pages 70–79, 2001.

[eRS04] F. Chazal et R. Soufflet. Stability and finiteness properties of medial axis and skeleton. *J. of Control and Dynamical Systems*, 10(2), 2004.

[For98] R. Forman. Morse theory for cell complexes. *Advances in Mathematics*, 134:90–145, 1998.

[HBK02] M. Hisada, A. Belyaev, and T. L. Kunii. A skeleton-based approach for detection of perceptually salient features on polygonal surfaces. In *Computer Graphics Forum*, volume 21, pages 689–700, 2002.

[HGY+99] P. W. Hallinan, G. Gordon, A.L. Yuille, P. Giblin, and D. Mumford. *Two-and Three-Dimensional Patterns of the Face*. A.K.Peters, 1999.

[KMG98] N. Khaneja, M. I. Miller, and U. Grenander. Dynamic programming generation of curves on brain surfaces. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 20, pages 1260–1265, 1998.

[LA98] G. Lukács and L. Andor. Computing natural division lines on free-form surfaces based on measured data. In M. Daehlen, T. Lyche, and L. Schumaker, editors, *Mathematical Methods for Curves and Surfaces II*, pages 319–326. Vanderbilt University Press, 1998.

[Mor90] R. Morris. *Symmetry of Curves and the Geometry of Surfaces: two Explorations with the aid of Computer Graphics*. Phd Thesis, 1990.

[Mor96] R. Morris. The sub-parabolic lines of a surface. In Glen Mullineux, editor, *Mathematics of Surfaces VI, IMA new series 58*, pages 79–102. Clarendon Press, Oxford, 1996.

[OB01] Y. Ohtake and A. Belyaev. Automatic detection of geodesic ridges and ravines on polygonal surfaces. *The Journal of Three Dimensional Images*, 15(1):127–132, 2001.

[OBS04a] Y. Ohtake, A. Belyaev, and H-P. Seidel. Ridge-valley lines on meshes via implicit surface fitting. In *ACM Siggraph*, 2004.

[OBS04b] Y. Ohtake, A. Belyaev, and H. P. Siedel. Ridge-valley lines on meshes via implicit surface fitting. In *Siggraph'04*. To appear, 2004.

[PAT00] X. Pennec, N. Ayache, and J.-P. Thirion. Landmark-based registration using features identified through differential geometry. In I. Bankman, editor, *Handbook of Medical Imaging*. Academic Press, 2000.

[Por01] I. Porteous. *Geometric Differentiation (2nd Edition)*. Cambridge University Press, 2001.

[SF00] G. Stylianou and G. Farin. Crest line extraction from 3d triangulated meshes. *NSF/DoE Lake Tahoe Workshop on Hierarchical Approximation and Geometrical Methods for Scientific Visualization*, 2000.

[SZ92] P. Sander and S. Zucker. Singularities of principal direction fields from 3D images. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 14(3):309–317, 1992.

[TG95] J.-P. Thirion and A. Gourdon. Computing the differential characteristics of isointensity images. *Computer Vision and Image Understanding*, 61(2):190–202, 1995.

[WB01] K. Watanabe and A.G. Belyaev. Detection of salient curvature features on polygonal surfaces. In *Eurographics*, 2001.

Contents

Unité de recherche INRIA Sophia Antipolis
2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Unité de recherche INRIA Futurs : Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)
Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifi que
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)
Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)
Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)
Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
http://www.inria.fr
ISSN 0249-6399