



Topics in Surface Discretization

David Cohen-Steiner

► To cite this version:

David Cohen-Steiner. Topics in Surface Discretization. Computational Geometry [cs.CG]. Ecole Polytechnique X, 2004. English. NNT: . tel-00832502

HAL Id: tel-00832502

<https://theses.hal.science/tel-00832502>

Submitted on 10 Jun 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THESE

présentée par

David Cohen-Steiner

pour obtenir le titre de
DOCTEUR DE L'ECOLE POLYTECHNIQUE

Spécialité : Algorithmique

QUELQUES PROBLÈMES LIÉS À LA DISCRÉTISATION DES SURFACES

Date de soutenance : 21 Janvier 2004

Composition du Jury :

Jean-Marc Steyaert	PRESIDENT
Herbert Edelsbrunner	RAPPORTEUR
Pierre Pansu	RAPPORTEUR
Jean-Daniel Boissonnat	EXAMINATEUR
Jean-Marie Morvan	EXAMINATEUR
Gert Vegter	EXAMINATEUR
Claude Puech	EXAMINATEUR INVITE

Thèse préparée au sein du projet **Géométrica - INRIA Sophia-Antipolis**

THESE

présentée par

David Cohen-Steiner

pour obtenir le titre de
DOCTEUR DE L'ECOLE POLYTECHNIQUE

Spécialité : Algorithmique

QUELQUES PROBLÈMES LIÉS À LA DISCRÉTISATION DES SURFACES

Date de soutenance : 21 Janvier 2004

Composition du Jury :

Jean-Marc Steyaert	PRESIDENT
Herbert Edelsbrunner	RAPPORTEUR
Pierre Pansu	RAPPORTEUR
Jean-Daniel Boissonnat	EXAMINATEUR
Jean-Marie Morvan	EXAMINATEUR
Gert Vegter	EXAMINATEUR
Claude Puech	EXAMINATEUR INVITE

Thèse préparée au sein du projet **Géométrica - INRIA Sophia-Antipolis**

Remerciements

En premier lieu, je voudrais remercier Jean-Daniel Boissonnat pour la qualité de son encadrement. Pendant ces trois ans, il m’a à la fois aiguillé sur des pistes fécondes, et laissé la liberté nécessaire à l’apprentissage de la recherche. Faire une thèse sous sa direction a été pour moi un grand plaisir, et une grande chance.

Si j’ai eu cette chance, c’est en grande partie grâce à Frédéric Cazals, qui a été mon premier contact dans le projet PRISME et m’a donné l’envie de venir y faire une thèse. Pour cela et pour le reste, merci.

Un grand merci aussi à Pierre Alliez, sans qui une partie importante de cette thèse, le quatrième chapitre de la deuxième partie, n’aurait pas vu le jour. Je tiens également à remercier les autres personnes qui ont contribué à ce travail : Mathieu Desbrun, Oliviers Devillers, et Bruno Lévy.

Deux mathématiciens m’ont beaucoup apporté durant ces trois années : Jean-Marie Morvan aura été, entre autre, mon mentor en géométrie différentielle et en théorie de la mesure géométrique. Je voudrais le remercier de m’avoir communiqué son goût pour ces sujets passionnants, ainsi qu’une partie de son savoir. Frédéric Chazal m’a aussi considérablement aidé, notamment pour la démonstration de l’isotopie donnée dans la dernière partie. Cette thèse leur doit beaucoup.

Elle doit également beaucoup à Mariette Yvinec et à Sylvain Pion qui, outre le soutien amical qu’ils m’ont apporté, m’ont bien souvent tiré des traquenards informatiques dans lesquels j’étais tombé. Par ailleurs, la première partie de cette thèse est le fruit d’une collaboration avec Mariette et Eric Colin de Verdière, que je voudrais remercier particulièrement. Merci aussi à Olivier Devillers pour m’avoir initié à ce jeu passionnant qu’est le poker, et à Monique Teillaud

pour sa bonne humeur et son sens de l'humour.

Les doctorants de l'ex-projet PRISME m'ont beaucoup aidé dans la réalisation de ce travail, par leur soutien scientifique, technique, et moral. Un grand merci donc à Christophe, Frank, Julia, Laurent, Luca, Marc, Marie, Philippe, Pierre-Marie, Raphaëlle, Steve, et Thomas.

Et, bien sûr, je ne remercierai jamais assez Agnès pour son soutien sans faille et son constant dévouement. Cela aurait été beaucoup plus difficile sans elle.

Enfin, outre Jean-Daniel et Jean-Marie, je remercie les autres membres du jury pour le temps qu'ils ont consacré à ce texte, ainsi que pour leurs remarques constructives : Claude Puech, Gert Vegter, Herbert Edelsbrunner, Jean-Marc Steyaert, et Pierre Pansu.

Sur un plan plus personnel, je voudrais remercier certaines personnes qui, si elles lisent ces lignes —même si c'est assez improbable, se reconnaîtront...

Contents

Introduction	7
I Conforming Delaunay triangulations in 3D	13
1.1 The algorithm	17
1.1.1 Definitions and notations	17
1.1.2 Protecting balls	17
1.1.3 Center-points, h -points, p -points, and SOS -points	18
1.1.4 The “split-on-a-sphere” strategy	19
1.1.5 The protection procedure	20
1.1.6 The whole algorithm	20
1.2 Proof of the algorithm	21
1.2.1 Properties maintained in the algorithm	21
1.2.2 Termination proof	24
1.3 Construction of the protecting balls	27
1.4 Improvements	28
1.4.1 Speeding up the protection procedure	28
1.4.2 Restricting the area where balls are required	29
1.5 Experimental results	31
1.6 Conclusion	32
II Estimation of the curvature tensor via the normal cycle	35
2 Elementary presentation of normal cycles theory	41
2.1 Statement of the main theorem	43

2.1.1	Curvature measures	43
2.1.2	Anisotropic curvature measures	45
2.1.3	Theorem	46
2.2	Normal cycles and curvature measures	47
2.2.1	A first approach	47
2.2.2	Background	48
2.2.3	Smooth case	51
2.2.4	Convex case	53
2.2.5	Polyhedral case	54
2.3	The second fundamental form via the normal cycle	58
2.3.1	Two more 2-differential forms	59
2.3.2	Smooth case	59
2.3.3	Polyhedral case	60
2.4	Sketch of proof of the theorem	61
2.5	Results	62
3	Approximation of normal cycles	69
3.1	Background on geometric measure theory	70
3.2	Geometric measure theory and curvature measures	74
3.2.1	Normal cycle of a geometric subset of $\mathbb{E}^n$	74
3.2.2	Curvature measures from normal cycles	76
3.2.3	Second fundamental form of geometric sets	76
3.3	An approximation result	78
3.3.1	Strongly close hypersurfaces	79
3.3.2	A homotopy between normal cycles	80
3.3.3	Approximation of curvature measures	83
3.3.4	The case of triangulations in 3-space	86
4	Application to anisotropic polygonal remeshing	95
4.0.5	Previous Work	96
4.0.6	Contributions	97
4.0.7	Overview	98
4.1	Principal Direction Fields	100
4.1.1	Robust 3D Curvature Tensor Estimation	100

4.1.2	Flattening the Curvature Tensor Field	102
4.1.3	Tensor Field Smoothing	103
4.1.4	Tensor Field Umbilic Points	103
4.1.5	Taking Care of Features	105
4.2	Resampling	105
4.2.1	Curve-based Sampling for Anisotropic Areas	106
4.2.2	Point-based Sampling in Spherical Areas	110
4.3	Meshing	111
4.3.1	Vertex Creation	111
4.3.2	Edge Creation	112
4.3.3	Polygon Creation	112
4.4	Results and Discussion	114
4.5	Conclusions and Future Work	117

III Meshing implicit surfaces with certified topology 119

1.6	A condition for isotopic meshing	123
1.6.1	A glimpse at stratified Morse theory	123
1.6.2	Collapses	128
1.6.3	Main result	129
1.6.4	Proof of the homeomorphy	132
1.6.5	Proof of the isotopy	134
1.7	Algorithm	138
1.7.1	PL watersheds	139
1.7.2	Main algorithm	142
1.8	Conclusion	146

Conclusion 149

References 150

Introduction

In the computer graphics community, it is sometimes argued that after sound, image, and video, the next digital media revolution will be about 3-dimensional geometry. While this prediction may seem somewhat optimistic, one has to acknowledge that shapes are becoming more and more widespread in computer science. The success of video games, as well as the one of special effects in motion pictures provide the most prominent examples. In addition to these applications in entertainment, geometry is also crucial in many industrial sectors. A considerable fraction of real world objects are produced from geometric models designed using CAGD systems. Also, it is often useful to produce a computer model of a real world shape. Examples include reverse engineering of mechanical parts, digitalization of art works, finite element simulations, GIS, surgery simulation, oil reservoirs modelling, or protein engineering.

Dealing with shapes on a computer requires the ability to build and process numeric, that is discrete models of shapes. By now, discretization of signals such as sound and image, and to a smaller extent video, is a fairly well-mastered area. Indeed, powerful theoretical tools such as Fourier transform, Shannon's sampling theory or wavelets lead to efficient algorithms to discretize signals and to manipulate them. For geometric data, the situation is quite different. Because shapes are not signals, that is real functions, but rather sets, the tools above-mentioned cannot be applied, at least directly. While some attempts were made to adapt signal processing tools to treat geometric data, a fully satisfying framework for shape discretization still remains elusive.

There is no consensus on the way a shape should be encoded on a computer. Instead, several discrete models of shapes are commonly used, each having its pros and cons :

- Point sets arise naturally in applications such as reverse engineering, where the shape of interest is only known through a finite number of points measured on its surface. Point sets,

even huge, can be rendered efficiently using splatting techniques. Also, efficient algorithms are available to compute geometric structures associated with point sets, such as the Voronoi diagram. The Voronoi diagram of a sampling of a surface gives much information about the surface itself. For that reason, in order to study a surface, it is often useful to build a sampling of that surface. On the other hand, point sets are not really surface models, as they do not encode the topology of the surface they lie on.

- Implicit surfaces encode surfaces as zero-sets of functions having an analytic expression. They are particularly well suited for certain tasks. In particular, they are a convenient tool for morphing or for shape animation. They also are the basis of level-sets methods. In surface modeling, they are used to build complex shapes by blending elementary shapes. Modeling with implicit surfaces has the advantage to allow for topological changes without any special care. As a drawback, rendering them directly requires ray-tracing and is thus computationally demanding. Moreover, their manipulation can lead to difficult computer algebra issues and robustness problems.
- Meshes are probably the most popular way to encode shapes. Surfaces meshes are by far the preferred representation for rendering, at least up to a reasonable number of polygons, since they can be processed very efficiently by graphic cards. Discrete conformal parameterizations lead to satisfying texture mapping algorithms for surfaces represented by meshes. Surface meshes also are a convenient data structure for many geometric algorithms, allowing for instance to “walk” on the surface easily. Finally, they form the basis of splines and subdivision surfaces, which are the main surface models used in CAGD. Volume meshes are also ubiquitous, since they are essential components in finite elements simulations.
- Volumetric images arise in applications where shapes are acquired by tomography, such as medical imaging. Being -3D- images, they can be processed with techniques from functional analysis, but they are not very well handled by more geometric methods, due to the intensity of the noise they contain. Also, high resolution volumetric images can be expensive to process, since their size is cubical in their spatial resolution. Let us mention an interesting intermediate between meshes and images, the so-called geometry images [65]. The idea is to model shapes by coarse -and possibly curved- meshes, whose polygons are endowed with textures encoding the detailed geometry of the shape. Typically, the value of these textures

give x,y,z coordinates of surface points, or coordinates of the normal vectors. This representation enables the use of image processing technique at a small scale, while keeping the combinatorial structure of a mesh at a larger scale.

Shape discretization raises numerous problems. At the heart of these problems lie geometric approximation questions : given a discretization of a shape, to which extent can one recover the geometry of the original shape? How can one discretize shapes while controlling the approximation error? A first class of problems posed by surface discretization is concerned with the conversion between discrete models of shapes. Conversions are useful because they permit to take advantage of the particular strengths of either one of the models.

- The problem of surface reconstruction [39] consists in finding a mesh approximating a given surface, knowing only a sampling of that surface. Its main application is reverse engineering. Ideally, the approximating mesh should have the same topology as the original surface, and be geometrically close to it. Most surface reconstruction algorithms from computational geometry are based on the Delaunay triangulation of the sample points. The latest among these are provably correct, that is guarantee the topology and the geometry of the output mesh provided the initial surface is smooth and densely enough sampled. The key fact ensuring these guarantees is the convergence of a certain subset of the Voronoi vertices of the sample points towards the medial axis of the complement of the original surface. From a certain point of view, these results can be considered as a geometric analog of Shannon's theory, since they give a sampling condition under which a given shape can be recovered. However, they are not fully satisfactory. Indeed, the sampling condition requires that the local density of points should be at least a constant times the curvature of the surface. As a consequence, this sampling condition cannot be met -by finite point clouds- when the surface contains a sharp edge, even almost flat.
- Passing from a point cloud to an implicit surface is also useful. Several techniques are available for that purpose. Most of them first compute a decomposition -or rather a covering- of space induced by the points. The desired implicit function is then defined by gluing functions defined on each element of the covering. The coverings used can be derived from the Voronoi diagram of the points [15], or from an octree [95]. A notable exception are the point set surfaces, which do not use any covering of space.

One may also want to convert an implicitly defined surface into a point cloud. In a popular method [35], an initial sampling is first computed, and the sample points then mimic the motion of particles repelling each other while staying on the surface. When equilibrium is reached, the resulting point cloud is evenly distributed on the surface. Usually, the repelling force is driven by the curvature of the surface, so that the density of the final sampling adapts to the curvature of the surface. Another technique aims at building a sampling matching the conditions ensuring correct surface reconstruction [16].

Finally, it is sometimes needed to build a mesh approximating an implicit surface. If guarantees on the geometric distortion are often desired, the major specific problem here is to recover the topology of the implicit surface correctly. We do not elaborate further as part III is entirely devoted to this problem.

- Building efficient meshes of known shapes is a major area of research. Efficiency can have different meanings according to the problem studied. For surface meshes, the goal is to optimize the ratio between geometric distortion and number of mesh cells. Efficient surface meshes can be stored, transmitted or processed at lower cost. Several measures of geometric distortion are commonly used : Hausdorff distance, closeness between the normal vectors, L^p norms, or volume enclosed by the surface and its approximating mesh. It is known from approximation theory that the shape of cells in an efficient mesh is given by the curvature of the surface to be meshed [34]. In Part II chapter 4 of this thesis, we address the problem of computing an efficient mesh of a surface described by a high resolution mesh.

For volume meshes, efficiency usually refers to their main application, namely the finite elements method. In this context, a mesh is used to define a finite-dimensional vector space of a given space of real functions defined on the volume to be meshed. These functions are then approximated by elements of the finite-dimensional vector space and partial differential equations are replaced by finite-dimensional analogs. The goal here is to optimize the ratio between the dimension of the approximation space and the quality of the approximation. In this respect, Delaunay meshes have good theoretical properties and are often preferred [110]. Also, in many cases, the phenomenon studied is not uniform in space, due to the presence of media having different characteristics. Given the interfaces between these media, usually described by a mesh, it can thus be useful to build a 3-dimensional Delaunay mesh containing the interface mesh as a sub-complex. As an example, when modeling the air flow around an airplane wing, it is needed that the boundary of the wing is represented as a sub-complex

of the mesh used for simulation. The problem of building such a mesh is addressed in part I.

Another direction of research related to the problem of shape discretization is the extension of notions from the continuous world to the discrete one. Indeed, smooth objects, such as surfaces, have been extensively studied by many generations of mathematicians. There is thus an extremely well developed theoretical framework to deal with such objects. By contrast, the study of discrete objects such as meshes has received much less attention, leaving much room for further developments. In this respect, a particularly exciting area of research consists in finding discrete counterparts of objects classically defined only in the smooth setting. As an example, smooth Morse theory involves derivative computations and thus do not generalize to the discrete setting directly. Still, the critical point theory for meshes developed by T. Banchoff [8] provides a satisfying discrete counterpart of it. In this case, both the discrete and the smooth theory actually are special cases of a more general one, stratified Morse theory. This ensures structural consistency, which is often sought after : the discrete concept should satisfy the same theorems as the smooth one. Discrete conformal parameterizations are another example of unified framework for both the discrete and the smooth case. In addition to being consistent, definitions of discrete analogs of smooth concept should also lead to efficient algorithm to compute them. Successful examples in this respect include the PL Morse complex by Edelsbrunner [46] et al. or Forman's combinatorial Morse theory [51]. Finally, an extension of a quantitative notion such as curvature to say meshes should also satisfy approximation results. For instance, the curvature of mesh approximating a smooth surface -in a sense to be made precise- should be close to the curvature of the surface. Indeed, this is useful for estimation purposes : if a surface is only known through an approximating mesh, this is the condition under which the curvature of the surface can be recovered by computing the one of the mesh. In Part II, we propose a definition of curvature for meshes that respects the three criteria we just mentioned : consistency, low computational cost, and good behavior under approximation.

This thesis is organized in three almost independent parts. The first one describes an algorithm devoted to the problem of conforming Delaunay triangulations, which we already mentioned. The goal is to compute a Delaunay mesh containing a given set of polygonal constraints in 3D as sub-complex. This is the first practical algorithm which is provably correct without any restriction on the input constraints. Implementation has shown that the algorithm works well in practice. This part has been published in the proceedings of SOCG 2002, and will

appear in the special issue of CGTA devoted to this conference.

The second part is concerned with curvature estimation from meshes. Our contribution is twofold. First, using the theory of normal cycles, we give an extension of the curvature tensor to a broad class of surfaces, including smooth and piecewise linear ones. More precisely, we associate with each region of a surface a tensor which, in the smooth case, equals the integral of the curvature tensor over that region. Our main result is that under certain assumptions, this tensor behaves nicely under approximation. In particular, a mesh approximating a smooth surface well will have curvature tensors close to the ones of the smooth surface. Thus, when computed a mesh approximating a smooth surface, our curvature tensor yields a reliable estimator of the curvature tensor of the smooth surface. This estimator also proves to work well in practice and is straightforward to compute. We actually prove the general approximation result not only for surfaces, but for hypersurfaces of any dimension. This work on the curvature tensor has been published in the proceedings of SOCG 2003. Second, we describe an algorithm for anisotropic polygonal remeshing of triangulated surfaces. Using the algorithm mentioned above, we first estimate the principal curvatures and directions at each vertex of the input surface. We then trace a net of lines of curvatures on the surface by numeric integration. Finally, we derive an anisotropic mesh of the input surface from this net. With an appropriate choice of the spacing between lines of curvature, the resulting mesh is actually close to being optimal for approximation purposes. This work on remeshing has been published in the proceedings of SIGGRAPH 2003.

In the last part, we give a new algorithm for implicit surface polygonalization. To the best of our knowledge, this is the first algorithm that recovers the topology of the implicit surface in a provably correct way. We actually show that the polygonalization we output is always isotopic to the input implicit surface. Instead of sticking to the precise geometry of the implicit surface, we try to capture its topology with the least computational cost. For that reason, we believe our algorithm should be efficient in practice, though we did not implement it yet by lack of time.

Part I

Conforming Delaunay triangulations in 3D

Introduction

In the following, the term *faces* denotes objects in 3D space which are either 0-dimensional faces called vertices, 1-dimensional faces called edges or 2-dimensional faces called 2-faces. The vertices are just points, the edges are straight line segments, and the 2-faces are polygonal regions possibly with holes and isolated edges or vertices included in their interior. A piecewise linear complex, called for short PLC, is a finite set $\mathcal{C}$ of faces such that:

- the boundary of any face of $\mathcal{C}$ is a union of faces of $\mathcal{C}$;
- the intersection of any two faces of $\mathcal{C}$ is either empty or a union of faces of $\mathcal{C}$.

A triangulation $\mathcal{T}$ is said to *conform* to a PLC $\mathcal{C}$ if any face of $\mathcal{C}$ is a union of faces of $\mathcal{T}$. In this paper, we propose an algorithm which, given a PLC $\mathcal{C}$, finds a set of points $\mathcal{P}$ whose Delaunay triangulation conforms to $\mathcal{C}$. The set $\mathcal{P}$ includes the vertices of $\mathcal{C}$ and a certain number of additional points which are usually called Steiner points.

This question is motivated by problems in mesh generation and geometric modeling: in these fields, it is crucial to decompose the space into a set of simplices which conforms to a given PLC, with the additional restriction that the shape of the cells must satisfy certain properties. Delaunay triangulations present several features (see, *e.g.*, [17]) which can be exploited to solve this problem, and many mesh generation algorithms make use of this concept.

The problem of computing a conforming 2D Delaunay triangulation was solved by Saalfeld [103] and Edelsbrunner and Tan [44]. The algorithm by Edelsbrunner and Tan [44] guarantees an $O(n^3)$ bound on the number of generated Steiner vertices, if n is the size of the input. Most of the further works on the subject are based on the Delaunay refinement approach pioneered by Ruppert [101] and Chew [31]. Shewchuk [109] gave an algorithm in 3D which builds a conforming Delaunay triangulation under restrictive conditions on the angles of the PLC. Murphy, Mount, and Gable [94] found a solution which works under no restriction, but produces far too many points in practice. The main interest of their paper is to show the existence of a conforming Delaunay triangulation with a finite set of vertices for any 3D PLC.

Our algorithm uses the Delaunay refinement approach. Initially, the set $\mathcal{P}$ is the set of vertices of the complex $\mathcal{C}$. Points are then added to $\mathcal{P}$ until each edge and each face of the complex $\mathcal{C}$ is a union of simplices which are in the Delaunay triangulation of $\mathcal{P}$.

The main difficulty with such a strategy is to ensure termination. Indeed, it is known that sharp edges and corners may induce cascading additions of Steiner points. To avoid this effect, we first define a protected area around edges and vertices of the PLC with a special refinement process. Outside the protected area, the PLC can be refined using Ruppert's process and the interaction between refinements in both areas can be controlled. Murphy, Mount, and Gable use a similar approach. The main difference with our work lies in the definition of the protected area. In our case, this area adapts to the local geometry of the input PLC.

The algorithm is presented in Section 1.1 and proved to be correct in Section 1.2. In Section 1.3, we present the details of the construction of the initial protected area, skipped in Section 1.1. Section 1.4 presents some refinements to improve the running time of the algorithm and to lower the number of vertices in the output conforming triangulation. At last, we end with experimental results in Section 1.5.

1.1 The algorithm

After a few definitions, we describe the protected area (Subsections 1.1.2 and 1.1.3). We then define the refinement process used for this area (Subsections 1.1.4 and 1.1.5). Finally, we describe the main procedure and summarize the whole algorithm.

1.1.1 Definitions and notations

The *circumball* of a segment ab is the ball admitting the segment ab as diameter. The *circumball* of a triangle abc is the ball admitting the circumscribing circle of abc as great circle.

An edge (resp. a triangle) is said to have the *Gabriel property* if its circumball contains no point of $\mathcal{P}$ in its interior. A point in the interior of the circumball of an edge (resp. a triangle) is said to *encroach upon* this edge (resp. this triangle).

In the following, we note $\text{bd}(B)$ the boundary of a ball B , $\text{int}(B)$ the interior of B and $\text{circum}(ab)$ (resp $\text{circum}(abc)$) the circumball of the segment ab (resp. of the triangle abc).

1.1.2 Protecting balls

The *1-skeleton* Sk of the complex $\mathcal{C}$ is the union of the 0- and 1-dimensional faces of $\mathcal{C}$. The protected area is defined by means of a set $\mathcal{B}$ of closed balls, called protecting balls, satisfying the following requirements:

- i. the union of the balls in $\mathcal{B}$ covers the 1-skeleton Sk of the complex $\mathcal{C}$;
- ii. the balls are centered on points which are in Sk ;
- iii. if two balls intersect, their centers belong to the same edge of the complex $\mathcal{C}$;
- iv. if a face of $\mathcal{C}$ intersects a ball, then it contains the center of this ball;
- v. the intersection of any three balls in $\mathcal{B}$ is empty;
- vi. any two balls are not tangent;
- vii. the center of any ball is inside no other ball.

(i) and (iv) imply that any vertex in $\mathcal{C}$ is the center of a ball in $\mathcal{B}$. We show in Section 1.3 how to build a set of balls satisfying these requirements. Furthermore, in Section 1.4, we show that there is in fact no need to cover all the edges.




Figure 1.1: The situation in the neighborhood of a ball B , incident to three other balls B_1, B_2 and B_3 . There are two faces in the complex, limited by three edges, in the plane of the figure. Point h_i is added on the radical plane of B and B_i . p -points a, b, c , and d belong to the boundary of two balls and to a face, they are therefore also inserted in $\mathcal{P}$. Incident to o are four right-angled triangles (e.g., oh_2a) and two isosceles triangles (e.g., oab). The shield edges are ab and cd .

1.1.3 Center-points, h -points, p -points, and SOS-points

We describe here a few subsets of points, included in the balls of $\mathcal{B}$, that we need to add first in the set $\mathcal{P}$. See Figure 1.1.

Let B be a ball in $\mathcal{B}$ with center o . Let $\mathcal{B}_B$ be the set of balls in $\mathcal{B}$ that intersect B . By condition (v), the intersections of B with the elements of $\mathcal{B}_B$ are disjoint.

We first add the center o of B . Such a point will be called a *center-point*. Then, for each element B_i of $\mathcal{B}_B$, consider the radical plane of B and B_i . It intersects the line joining the centers of B and B_i at a point h_i , which is on an edge of $\mathcal{C}$ by condition (iii). The point h_i is added to the set $\mathcal{P}$. Such points will be called *h -points*.

By condition (iv), any face of $\mathcal{C}$ which intersects $B \cap B_i$ contains the centers of B and B_i , and thus can be either the edge including the segment oo_i (o_i is the center of B_i) or a 2-face incident this edge. For each 2-face F of $\mathcal{C}$ intersecting $B \cap B_i$, we add to $\mathcal{P}$ the intersection

points of F with the circle $\text{bd}(B) \cap \text{bd}(B_i)$. We called those points *p-points*.

Consider the plane Q of a 2-face of $\mathcal{C}$ intersecting B (and thus containing o). The edges of $\mathcal{C}$ split the disk $Q \cap B$ into one or several sectors. We focus on sectors which are included in $\mathcal{C}$. The *p-points* further split these sectors in subsectors. We call *right-angled subsectors* the subsectors limited by an edge of $\mathcal{C}$ and a *p-point* and *isosceles subsectors* the subsectors limited by two *p-points*.

If some isosceles subsectors form an angle $\geq \pi/2$, we add some points on their bounding circular arcs to subdivide them in new subsectors forming an angle $< \pi/2$. For reasons that will be clear in Subsection 1.1.4, these points are called *SOS-points*. The new subsectors with angle $< \pi/2$ are still called *isosceles subsectors*.

Center-points and *h-points* are the only categories of points added in the interior of protecting balls. *p-points* and *SOS-points* lie on the boundaries of protecting balls. *SOS-points* belong to a single protecting ball while *p-points* belong to the intersection of two balls.

Isosceles subsectors are defined by the center o of a ball B and by two points a and b (either *p-points* or *SOS-points*) on $\text{bd}(B)$. Line segments such as ab , joining two points that define an isosceles subsector, are called *shield edges*. In the following, triangles defined by center-points and shield edges such as oab are referred to as *isosceles triangles*. Triangles spanned by a center-point, a *h-point* and a *p-point* on the boundary of some right-angled subsector are referred to as *right-angled triangles*.

Definition 1 *The protected area is the union of the isosceles and right-angled triangles. See the dark gray area in Figure 1.1. In particular, the protected area is included in the union of the protecting balls.*

Definition 2 *The unprotected area is the complex $\mathcal{C}$, minus the protected area.*

1.1.4 The “split-on-a-sphere” strategy

During the process, it will be necessary to split shield edges. Since we do not want to add more points inside the balls in $\mathcal{B}$, we use a special treatment to split such a shield edge, called the “split-on-a-sphere” strategy (SOS for short). See Figure 1.2.

Let ab be a shield edge to be split, in a ball B . We distinguish two cases: a and b are both *SOS-points* and belong to a single ball B , or at least one of these two points (for example a) is a *p-point* and belongs also to another ball B' .




Figure 1.2: *The SOS strategy: We split the shield edge ab by inserting the point c on the boundary of the ball.*

If a and b belong only to B , let c be the midpoint of the shortest geodesic arc ab on $\text{bd}(B)$. To refine edge ab , we add c to $\mathcal{P}$ and replace the shield edge ab by two shield edges ac and cb .

If a is a p -point belonging to $\text{bd}(B) \cap \text{bd}(B')$, the idea is quite similar; however, if we do not take care, the SOS strategy could lead to cascading insertions of points, because refining an edge on B would lead to refinement of an edge on B' , and so on. We thus use a strategy “à la Ruppert” [101], using circular shells. We consider the length of the segment ab , divided by two, and round it to the nearest distance d which is of the form 2^k , $k \in \mathbb{Z}$ (the unit distance has been chosen arbitrarily at the beginning of the algorithm). Let c be the point of the shortest geodesic arc ab on $\text{bd}(B)$ at distance d from a . We split the shield edge ab using the point c .

In both cases, the added point c belongs to the category of SOS-points. Note that, due to the SOS refinement strategy, the protected and unprotected areas, still defined as in Subsection 1.1.3, will slightly evolve during the algorithm. Each SOS refinement increases the protected area and decreases the unprotected area.

1.1.5 The protection procedure

This procedure adds some points to set $\mathcal{P}$ to ensure that shield edges and isosceles triangles have the Gabriel property. It uses recursively the SOS strategy and works as follows: While there is an encroached shield edge ab or an encroached isosceles triangle oab , refine the edge ab using the SOS strategy.

1.1.6 The whole algorithm

Let us recall that the algorithm works by adding points to set $\mathcal{P}$. We note $Dt_3(\mathcal{P})$ the 3D Delaunay triangulation of points in $\mathcal{P}$. For each plane Q of a 2-face in $\mathcal{C}$, we note $Dt_2(\mathcal{P} \cap Q)$ the 2D Delaunay triangulation of points in $\mathcal{P} \cap Q$. These triangulations are updated upon each

insertion of a point in $\mathcal{P}$.

The algorithm performs the initialization step and the main procedure described below.

The Initialization Step:

- Construct and initialize the protected area (as described in 1.1.2 and 1.1.3);
- execute the protection procedure.

We will see later that the Delaunay triangulation of $\mathcal{P}$ conforms to the part of $\mathcal{C}$ which is inside the protected area. Because the algorithm maintains the Gabriel property of shield edges, in each plane Q of a 2-face F of $\mathcal{C}$, the 2D triangulation $Dt_2(\mathcal{P} \cap Q)$ conforms to the shield edges in this plane and thus to the unprotected part F_u of F . The main procedure ensures that the triangles of $Dt_2(\mathcal{P} \cap Q)$ included in F_u appear in the 3D triangulation $Dt_3(\mathcal{P})$.

The Main Procedure:

The Main Procedure consists in executing the following loop: While there is a triangle T in the 2D Delaunay triangulation $Dt_2(\mathcal{P} \cap Q)$ of the plane Q of a 2-face F of $\mathcal{C}$ such that:

- a. T is included in the unprotected part F_u of F ,
- b. T does not appear in $Dt_3(\mathcal{P})$,

refine T trying to insert its circumcenter c , that is:

- if c encroaches upon no shield edge, insert it;
- otherwise, split all the shield edges encroached upon by c using the SOS strategy, and then execute the protection procedure.

1.2 Proof of the algorithm

Two steps are involved for the proof of this algorithm. First, we prove invariants of the algorithm concerning the positions of the points added and the Gabriel property of some triangles and edges. After that, we are able to prove termination.

1.2.1 Properties maintained in the algorithm

Lemma 1 *At the beginning (and the end) of each execution of the main loop, the shield edges have the Gabriel property.*




Figure 1.3: The circumcenter p of a triangle T lies in the unprotected area.

Proof. Indeed, this is true before the first execution of the main loop, because the protection procedure, which has just been executed, ensures this property; for the same reason, this also holds after an execution of the loop leading to the split of shield edges. At last, a circumcenter is inserted in $\mathcal{P}$ only if it does not violate this property. $\square$

In the following, we define an *added* circumcenter to be a circumcenter inserted in the set $\mathcal{P}$, and a *rejected* circumcenter to be a circumcenter considered in the algorithm but not inserted because it encroaches upon some shield edge.

Lemma 2 *Any circumcenter (added or rejected) considered by the algorithm lies in the unprotected area, outside the protecting spheres. In particular, no point is added inside the protecting spheres after the initialization step, and $\mathcal{P}$ is included in $\mathcal{C}$.*

Proof. Let T be a triangle whose circumcenter is considered at some step of the algorithm. T lies in the unprotected area, and belongs to the 2D Delaunay triangulation $Dt_2(\mathcal{P} \cap Q)$ of the plane Q of some 2-face in $\mathcal{C}$. Let p be the circumcenter of T . Assume for contradiction that p lies outside the unprotected area. Let m be a point in T . Since shield edges enclose the connected component of the unprotected area which contains T , the segment pm must intersect a shield edge ab . The vertices a and b cannot be inside $\text{circum}(T)$ because T belongs to $Dt_2(\mathcal{P} \cap Q)$. Hence (Figure 1.3), triangle T belongs to the circumball of ab , which is impossible by Lemma 1.

Moreover, since the circumballs of shield edges cover the intersection of the unprotected area with the protecting balls (see Figure 1.4), any added circumcenter is actually outside the protecting spheres. $\square$




Figure 1.4: The intersection of the unprotected area with the union of protecting balls is included in the circumballs of shield edges.

Proposition 3 *At the beginning (and the end) of each execution of the main loop, the isosceles triangles have the Gabriel property.*

Proof. The proposition is obvious after the initialization step because the protection procedure is called and enforces the Gabriel property of isosceles triangles. For the same reason, it is also the case when a circumcenter has just been rejected because it encroaches upon some shield edge.

It remains to see that this proposition is still true when a circumcenter has just been inserted: such a circumcenter lies outside the protecting spheres (by Lemma 2) and outside the circumball of any shield edge (otherwise it is not inserted in $\mathcal{P}$). Let ab be such a shield edge, belonging to ball B . We note that the boundaries of B , $\text{circum}(ab)$, and $\text{circum}(oab)$ belong to a pencil of spheres. Because the angle $\widehat{aob}$ is smaller than $\pi/2$, we have $\text{circum}(oab) \subset \text{circum}(ab) \cup B$ (Figure 1.5). The result follows. $\square$




Figure 1.5: The balls B , $\text{circum}(oab)$, and $\text{circum}(ab)$.

Lemma 4 *Let B be a ball with center o , and p be a point on the boundary of B . If, at some stage of the algorithm, the segment op is encroached upon, the encroaching point is a h -point h_i on the radical plane of B and B_i , and p belongs to $\text{bd}(B) \cap \text{int}(B_i)$.*

Proof. The circumball of op is inside B . Therefore, op can only be encroached upon by a vertex in this ball, and not by the center of B , hence only by a h -vertex in B . Suppose that op is encroached upon by a vertex h_i , belonging to B and B_i . The encroachment condition can be rewritten $\widehat{oh_i p} > \pi/2$. Because points q in $\text{bd}(B)$ that satisfy $\widehat{oh_i q} > \pi/2$ lie in $\text{bd}(B) \cap \text{int}(B_i)$, p belongs to $\text{int}(B_i)$. $\square$

Proposition 5 *At each stage of the algorithm, the right-angled triangles have the Gabriel property.*

Proof. Suppose that a right-angled triangle $oh_j p$ does not have the Gabriel property at some stage of the algorithm: h_j is on the radical plane between B and B_j , and p is on the boundary of B and B_j . Because the circumball of $oh_j p$ is the circumball of op , by Lemma 4, the encroaching point is a h -point, and p has to belong to the interior of a third ball B_i , which is impossible by condition (v). $\square$

Center points and h -points cut the edges of $\mathcal{C}$ in subedges. Note that Proposition 5 implies that these subedges are edges of $Dt_3(\mathcal{P})$.

1.2.2 Termination proof

Proposition 6 *The protection procedure always terminates.*

The proof is a straightforward consequence of the following lemma.

Lemma 7 *For each call to the protection procedure, there exists $\theta > 0$ such that no isosceles triangle with angle at the center of the ball less than θ will be split.*

Proof. Let oab be an isosceles triangle with shield edge ab in a protecting ball B . We consider in turn three kinds of possible encroaching points: points on the boundary of B (case 1), points in the interior of B (case 2), and points outside B (case 3). In each case k , we prove the existence of a value θ_k , such that neither oab nor ab can be encroached upon by a point of type k if $\widehat{aob} < \theta_k$.

Recall that the three balls B , $\text{circum}(ab)$ and $\text{circum}(oab)$ belong to a pencil of spheres. Because the angle $\widehat{aob}$ is smaller than $\pi/2$, we have $\text{circum}(oab) \subset B \cup \text{circum}(ab)$ and $\text{circum}(ab) \cap B \subset \text{circum}(oab)$ (see Figure 1.5). Therefore, it is enough to check that points on the boundary of B or outside B (cases 1 and 3) do not encroach upon ab and that points in B (case 2) do not encroach upon oab .

1. For a plane Q of a 2-face of $\mathcal{C}$ intersecting B , we consider the circle bounding $B \cap Q$ and we denote by $S(Q, B)$ the union of arcs on this circle spanned by the isosceles triangles in Q . Notice that all the SOS-points inserted on B are located on such a set $S(Q, B)$.

If Q is the plane containing oab , no point of $S(Q, B)$ encroaches upon ab . If Q' is another plane, the distance between $S(Q, B)$ and $S(Q', B)$ is strictly positive, so there is a value $\theta_1(B, Q, Q')$ such that ab is not encroached upon by a point on $S(Q', B)$ if $\widehat{aob} < \theta_1$. Setting $\theta_1 = \min\{\theta_1(B, Q, Q')\}$ achieves the proof of case 1.

2. The only points in a ball B which can encroach upon an isosceles triangle oab in B are the h -points in B . Suppose that a point h_i (on the radical plane of B and B_i) encroaches upon oab .

If h_i is in the plane Q of oab , we prove that encroachment is not possible. Indeed, if h_i encroaches upon oab , h_i encroaches either upon oa or upon ob . Thus a or b would belong to $\text{bd}(B) \cap \text{int}(B_i)$, by Lemma 4, which is impossible because a and b are either p -points or SOS-points.

Let us now deal with the case where h_i does not belong to the plane Q . Let $c \in S(Q, B)$; c does not belong to B_i , for otherwise h_i would belong to Q . Let us prove that h_i is not in the closed ball $\text{circum}(oc)$. If h_i is in the interior of $\text{circum}(oc)$, this means that oc is encroached upon by h_i , hence, by Lemma 4, c belongs to $\text{int}(B_i)$, which is not the case. Similarly, if h_i is on the boundary of $\text{circum}(oc)$, c belongs to B_i .

Hence, the distance between h_i and the ball $\text{circum}(oc)$ is strictly positive. Let $\delta(B, Q, h_i)$ be the minimum (strictly positive) of this distance for $c \in S(Q, B)$. Let $\delta'(B, \theta)$ be the Hausdorff distance between $\text{circum}(oc)$ and $\text{circum}(oa'b')$ where $oa'b'$ is an isosceles triangle with a' and b' on $\text{bd}(B)$, axis oc and $\widehat{a'ob'} = \theta$. As $\delta'(B, \theta)$ goes to 0 when θ goes to 0, there exists $\theta_2(B, Q, h_i)$ such that $\delta'(B, \theta) < \delta(B, Q, h_i)$ for any $\theta < \theta_2(B, Q, h_i)$. It follows that oab cannot be encroached upon by h_i if $\widehat{aob} < \theta_2(B, Q, h_i)$. Setting $\theta_2 = \min\{\theta_2(B, Q, h_i)\}$ achieves the proof of case 2.

3. Consider now the case where edge ab is encroached upon by a point p outside the ball B . At each call of the protection procedure, the set of points outside the protecting spheres is fixed. Also, the distance between two sets $S(Q_1, B_1)$ and $S(Q_2, B_2)$ which do not share a p -point is bounded from below. Thus, there is a value θ'_3 such that, if $\widehat{aob} < \theta'_3$, edge ab cannot be encroached upon by p except if p belongs to $S(Q, B')$ where Q is the plane of

oab and B' intersects B . Therefore, the only case remaining to be considered is the case where a is a p -point in $Q \cap \text{bd}(B) \cap \text{bd}(B')$ and ab is encroached upon by a point p of $S(Q, B')$. However, in this case, we split edges incident to a using circular shells. Hence, after a few splits, the edges incident to a will have the same lengths and will be unable to encroach upon each other. Therefore, we get a value $\theta_3 \leq \theta'_3$ satisfying the desired requirement.

□

Theorem 8 *The algorithm terminates, and, once it is the case, the Delaunay triangulation of $\mathcal{P}$ conforms to the complex $\mathcal{C}$.*

Proof. It is sufficient to prove that the main procedure terminates: indeed, once it is the case, Propositions 3 and 5 show that the Delaunay triangulation of $\mathcal{P}$ conforms to the protected area of $\mathcal{C}$, and the fact that the algorithm ends precisely means that the Delaunay triangulation of $\mathcal{P}$ also conforms to the unprotected area of $\mathcal{C}$. We prove the termination of the main procedure by proving first that the number of added circumcenters is finite and second that the number of shield edges encroached upon by rejected circumcenters is finite. Because the protection procedure is already known to terminate, these two facts imply the termination of the main procedure.

By construction of the protecting spheres, the unprotected area is a disjoint union of plane regions. Let F_u be such a region. As previously noticed, owing to the SOS strategy, these unprotected regions slightly evolve during the algorithm; however, they are always shrinking. Consequently, the distance between F_u and the other regions as well as the distance between F_u and the set of center-points and h -points added in the interior of the protecting balls can be bounded from below by a constant δ_F . Let T be a triangle in F_u whose circumcenter has to be inserted in $\mathcal{P}$ and let C_T be the circumcircle of T . As T does not belong to $Dt_3(\mathcal{P})$, its circumball $\text{circum}(T)$ contains a point in $\mathcal{P}$ which is not in the plane of F_u . Such a point can be inside a protecting ball (a center-point or a h -point), on the boundary of a protecting ball (and thus on the boundary of another region), or an added circumcenter (in another region by Lemma 2). Therefore $\text{circum}(T)$ either contains a point added in the interior of a protecting sphere or intersects another unprotected region, and the radius of C_T is thus larger than δ_F . Because T belongs to the 2D Delaunay triangulation in the plane of F_u , C_T encloses no point of $\mathcal{P}$. The area of F_u being finite, this shows that the number of added circumcenters is bounded.




Figure 1.6: *The shortest shield edge ab which may be encroached upon by a rejected circumcenter p .*

Let us now show that the total number of edges encroached upon by rejected circumcenters is finite. For this purpose, consider a shield edge encroached upon by the center p of a circumcircle C in a region F_u . C being empty and of radius larger than δ_F , it is easy to show that the shield edge has length at least $\delta_F \sqrt{2}$ (see Figure 1.6). Thus the number of those edges is finite. $\square$

1.3 Construction of the protecting balls

We have to build the set $\mathcal{B}$ of protecting balls satisfying the conditions described in Subsection 1.1.2. The efficiency of the algorithm really depends on this construction: the less balls there are, the less points will be produced in $\mathcal{P}$.

Definition 3 *Let $\mathcal{C}$ be a PLC. The local feature size of a point p with respect to $\mathcal{C}$ is the distance between p and the union of faces of $\mathcal{C}$ that do not contain p .*

Let $lfs(p)$ denote the local feature size of point p with respect to the PLC which is given as input of the algorithm. We address the following construction of the enclosing balls. Let α be a real, $0 < \alpha < \frac{1}{2}$ (typically $\alpha = 0.4$).

First, for each vertex v of the PLC, construct a ball of radius $\alpha \cdot lfs(v)$.

Then, on each edge e , do the following. While e is not completely covered by balls, consider a maximal open line segment a_1a_2 in e and outside the union of the balls in the current set $\mathcal{B}$. Point a_i ($i = 1, 2$) is an intersection of ball B_i (with center o_i and radius r_i) with edge e . We will insert a ball between B_1 and B_2 . Let o be the midpoint of a_1a_2 . Insert a new ball B in $\mathcal{B}$, of

center o and radius r , with:

$$r = \min \left\{ \alpha \cdot lfs(o), oa_1 + \frac{r_1}{2}, oa_2 + \frac{r_2}{2} \right\}.$$

To ensure condition (vi), if $r = oa_1$, we replace r by $(1 - \varepsilon)r$ where ε is a small positive constant.

Lemma 9 *This construction terminates.*

Proof. Consider an edge e , whose vertices have just been protected by two spheres. Let A be the union of the (open) line segments which are in e minus the union of the current set of balls. Call A_0 the set A just after the protection of the endpoints of e . The distance $d = \min\{lfs(p) \mid p \in \overline{A_0}\}$ is strictly positive (the lfs function is continuous on $\overline{A_0}$, and lfs does not vanish on $\overline{A_0}$). The insertion of a new ball:

- either increases by one the number of connected components of A and decreases the measure of A by at least $2(1 - \varepsilon) \cdot \alpha \cdot d$ (hence this case can happen only a finite number of times),
- or decreases by one the number of connected components of A (without increasing the measure of A).

The result follows. □

Conditions (i), (ii), (iv), (vi) and (vii) are obviously satisfied. (iii) follows from the fact that if two points o and o' do not belong to the same edge, oo' is larger than or equal to $lfs(o)$ and $lfs(o')$. If two balls B and B' , centered at o and o' with radii r and r' , are in $\mathcal{B}$, then $r < \frac{1}{2}lfs(o)$ and similarly for r' . Thus $r + r' < oo'$, hence the balls cannot intersect.

(v) is also true. Indeed, if three balls intersect, their centers must be vertices of a triangle in $\mathcal{C}$. But it follows from our construction that two balls centered on vertices of the PLC cannot intersect because $\alpha < \frac{1}{2}$.

Hence we have:

Proposition 10 *This construction of $\mathcal{B}$ is correct.*

1.4 Improvements

1.4.1 Speeding up the protection procedure

The following proposition shows that when the protection procedure is called from the main procedure, there is no need to check whether isosceles triangles have the Gabriel property.

Proposition 11 *After the initialization process, enforcing Gabriel property for shield edges in the protection procedure is enough to ensure Gabriel property for isosceles triangles.*

Proof. Upon termination of the initialization step, all isosceles triangles have the Gabriel property. Suppose that, at some stage of the algorithm, a point encroaches upon some isosceles triangle oab without encroaching ab . Let B be the ball containing oab . Since $\text{circum}(oab)$ is included in the union of B and $\text{circum}(ab)$ (Figure 1.5), the encroaching point must be inside B .

Hence it is sufficient to show that no isosceles triangle is encroached upon by a vertex inside its protecting ball during the algorithm. By contradiction, let $T = oab$ be the first isosceles triangle encroached upon by a vertex in B . Since no point is inserted inside the balls during the main procedure, T must be a triangle which results from the splitting of some triangle $T' = oac$. The encroaching point can thus only be a h -point h_i lying inside B . Arguing that $\text{circum}(oab)$, $\text{circum}(oac)$, and $\text{circum}(oa)$ belong to a sphere pencil and comparing their radii, we deduce (Figure 1.7) that $\text{circum}(oab) \subset \text{circum}(oac) \cup \text{circum}(oa)$. However, h_i does not belong to $\text{circum}(oac)$ because $T' = oac$ was not encroached upon by h_i , nor to $\text{circum}(oa)$ (by Lemma 4). Therefore h_i does not belong to $\text{circum}(oab)$, which yields the contradiction. $\square$




Figure 1.7: $\text{circum}(oab) \subset \text{circum}(oac) \cup \text{circum}(oa)$.

1.4.2 Restricting the area where balls are required

In 1.1.2, the set $\mathcal{B}$ is constructed so that the balls cover the whole 1-skeleton Sk of $\mathcal{C}$. We explain here that this is not always necessary. Indeed, the balls are introduced to avoid troubles with small angles; they are thus not required at places where faces intersect with an angle large

enough. This remark enables to put less balls in $\mathcal{B}$, hence to reduce the size of the output $\mathcal{P}$. We first describe the modification in the construction of the balls, and then prove that, despite this slight modification, the algorithm is still correct.

Let $e = o_1o_2$ be an edge of the PLC so that all angles between faces incident to e are $\geq \pi/2$. We modify the algorithm in the following way. Still construct balls B_1 and B_2 centered at the vertices o_1 and o_2 . In $\mathcal{P}$, insert o_1 , o_2 , and the two intersections p_1 and p_2 of e with the boundaries of B_1 and B_2 .

Consider p_1p_2 as a shield edge in the main procedure. In other words, whenever this edge would be encroached upon by the insertion of a point in $\mathcal{P}$, split this edge in the middle, to keep it protected at each stage of the algorithm. The original edge of $\mathcal{C}$ is thus not in the protected area, but the process is exactly like in the standard algorithm.

There are only minor modifications for the proof of the algorithm. The unprotected area is still bounded with shield edges. The proof of termination of the protection procedure is analogous: Lemma 7 can be adapted without difficulty to show that there also exists a length $\delta > 0$ such that the protection procedure never splits a shield edge which is a part of an edge and with length less than δ . The only difficulty is to show the following proposition.

Proposition 12 *The modified version of the main procedure always terminates.*

Proof. Let F_u be a region, in a plane Q , incident to edge e . The distance between F_u and the regions non-incident to e as well as the distance between F_u and the set of center-points and h -points outside Q can be bounded from below by a constant $\delta_F > 0$. Let p be the circumcenter of a triangle T in F_u , added to $\mathcal{P}$. We will show that the circumball of T cannot contain a vertex of another face incident to e , which implies that the radius of this circumball is larger than δ_F , like in the proof of Theorem 8.

Suppose for contradiction that T is encroached upon by a point p' of $\mathcal{P}$ on a face incident to e . Necessarily, because the angles of the faces of $\mathcal{C}$ are obtuse at e , the circumball of T must intersect e . Let a and b be the intersection points of the boundary of $\text{circum}(T)$ with e . Let $a'b'$ be the unique shield edge included in e which is intersected by $\text{circum}(T)$. (The uniqueness follows from the fact that points in $\mathcal{P}$, like a' and b' , cannot lie in $\text{circum}(T)$.) Let H be the plane orthogonal to F_u and containing e , and H^+ be the half-space bounded by H and not containing T . Clearly, $\text{circum}(T) \cap H^+ \subseteq \text{circum}(ab) \cap H^+ \subseteq \text{circum}(a'b') \cap H^+$ (see Figure 1.8). The point p' is in $\text{circum}(T) \cap H^+$, hence in $\text{circum}(a'b')$, which means that p' encroaches upon the shield edge $a'b'$ and yields the contradiction.




Figure 1.8: In the half-space H^+ (above the edge e in the figure), the part of $\text{circum}(T)$ is included in the part of $\text{circum}(ab)$ which is, in turn, contained in the part of $\text{circum}(a'b')$.

The remaining part of the proof of termination of the main procedure is exactly the same as in the proof of Theorem 8. $\square$

1.5 Experimental results

The algorithm has been implemented and tested using the Computational Geometry Algorithms Library CGAL¹. Results for several models are displayed in Table 1.1 and Figures 1.9, 1.10, 1.11, and 1.12.

Table 1.1 gives for each model, the number of vertices of the input PLC (*nb input vertices*), the number of 2-faces to which the Delaunay triangulation of input vertices does not conform (*non Delaunay faces*), and the number of vertices of the conforming output triangulation (*nb output vertices*). In those examples and in most cases, the number of vertices in the output conforming triangulation and the number of input vertices are in a ratio comprised between 3 to 1 and 10 to 1.

The running times, measured on a PC with 500Mhz processor, do not include the computations of local feature size values, because the current implementation uses a very slow brute force algorithm for it. We are currently designing a data structure devoted to speed up these computations.

¹<http://www.cgal.org/>

Table 1.1: *Experimental data.*

	geological data	triceratops	umbrella
nb input vertices	7566	2832	16
nb non Delaunay faces	1045	2194	5
nb output vertices	25793	27947	122
running time (s)	83	570	0.7

1.6 Conclusion

We have presented an algorithm for computing a conforming Delaunay triangulation of any three-dimensional piecewise linear complex. The most important innovation, compared to the paper by Murphy et al. [94], is to enclose critical places by balls whose radii fit the local complexity of the complex, with the use of the local feature size. Our experimental results show that it is valuable in practice. The algorithm could be easily modified to guarantee in the resulting mesh the Gabriel property for any triangle included in a constraint. The next step currently under work is to investigate how conforming meshes with guarantees on the shape and size of the elements can be obtained. Several questions remain open: we did not try to find the time complexity of our algorithm. It would also be interesting, as in [44] in the plane, to find a bound on the output depending on the size of the initial complex and/or (like in [101]) the lfs function.




Figure 1.9: *Detail of a geological formation (Courtesy of T-surf and Mr. Reinsdorff). Solid line segments stand for shield edges.*




Figure 1.10: *Umbrella. Solid line segments stand for shield edges.*




Figure 1.11: *Triceratops*.




Figure 1.12: *Detail of the triceratops. Solid line segments stand for shield edges.*

Part II

Estimation of the curvature tensor via the normal cycle

Introduction

In many applications such as surface segmentation, anisotropic remeshing [2] or non-photorealistic rendering, a key step is to estimate the curvature tensor of a smooth surface knowing only a polyhedral approximation of it. A lot of efforts have been devoted to this problem, leading to several estimators, see [96] or [86] for a detailed survey. Most popular methods rely on either one of the following three approaches.

The first one is the quadric fitting approach, where the estimated curvature tensor is the one of the quadric that best fits the sample points locally [104]. Usually, such algorithms first estimate the tangent plane at the vertex at which the curvature is to be estimated. Then a set of nearby vertices is selected, and their projections p_i on the estimated tangent plane are computed, as well corresponding heights h_i . Finally, the quadratic function f defined on the estimated tangent plane that minimizes the least square error between the $f(p_i)$ and the h_i is determined, and the output is the curvature tensor of the graph of f at the considered point. Some multi-pass variants of this method use the estimated quadric to deduce a hopefully better estimation of the tangent plane, and repeat the whole procedure based on this estimation, until convergence is reached [86]. The use of higher degree polynomials instead of quadratic functions is also of interest [26]. The main advantages of quadric fitting are its relative robustness against noise, and the fact that it is *meshless* : it can proceed point clouds of which no triangulation is available. Still, one should be able to select a set of neighboring points for each sample, which is not that much easier than triangulating them, at least locally. The main drawback of this approach is that for some specific configurations of neighboring points, namely when the h_i lie on a conic, there can be a whole space of quadratic functions achieving a perfect fit. This results in a very high sensitivity of the estimation in the vicinity of these configurations, and thus in a poor accuracy [26].

Another general technique, closer to classical differential geometry, first consists in estimating the curvatures of well-chosen cross-sections of the surface. In most cases, these sections are defined by planes spanned by the considered point, and two neighbors of that point in the triangulation. Some methods discard triples of points that are not sufficiently aligned [29]. The curvature of corresponding cross-sections are usually estimated by the inverse of the radius of the circle that passes through the considered triple of points [85], though other choices are possible [78]. From these curvatures and an estimate of the normal vector, the normal curvatures

of the surface can be estimated using Meusnier’s theorem. Finally, the estimated curvature tensor is the one that best accounts for the previously estimated normal curvatures [119, 37]. The main shortcomings of this general approach are the lack of analysis of the quality of the obtained estimators, and also the lack of solid theoretical foundations.

Let us mention for sake of completeness a third class of methods, relying on covariance matrices. An example of algorithm falling in this class is [10]. It proceeds by estimating normal vectors at each vertex neighboring a given vertex p , and project these vectors on the estimated tangent plane at p . The estimated curvature tensor at p is defined as being the inertia matrix of the obtained set of 2-vectors. Other methods output the inertia matrix of sets of vectors obtained by different means [127]. Algorithms based on the covariance matrix approach are meshless and seem robust against noise but, again, no analysis of the estimation errors is available, and it seems that these algorithms are not well suited for irregular samplings.

In this work, we propose a sound approach to curvature estimation. Building upon the theory of normal cycles from differential geometry, we define curvature tensors for a general class of surfaces, including smooth and polyhedral ones. More precisely, we associate with each region a tensor which in the smooth case is the average of the curvature tensor over this region. The curvature tensor of a polyhedral approximation of a smooth surface then provides an estimator of the one of the smooth surface. Our definition can be viewed as an anisotropic generalization of what is usually called discrete curvatures.

Besides the aforementioned definition, our main contribution is a bound on the difference between the estimated curvature and the actual one. This bound holds whenever the polyhedral approximation is *closely inscribed* in the smooth surface, which is a rather mild assumption. In particular, restricted Delaunay triangulations can be shown to fulfill this property for sufficiently dense samplings. This latter case is of great importance in practice. Indeed, in most cases, the only available data on an object is a set of points sampled on its surface. A possible approach is then to apply a reconstruction algorithm to get a triangulated surface, and estimate the curvature of the object from that triangulation. As most popular reconstruction algorithms, at least among the Delaunay-based ones, return a triangulation close to the Delaunay triangulation of the samples restricted to the object, the special case of restricted Delaunay triangulations is particularly worth studying. For these triangulations, the bound we obtained implies that our estimator converges linearly with respect to the sampling density, under a mild local uniformity

condition on the sampling. This latter result actually holds for any closely inscribed triangulation in which all triangles have bounded aspect ratio. To the best of our knowledge, only weaker results have been obtained in the past [87]. Our result can be viewed as a quantitative version of a theorem obtained by J.Fu [52] for gaussian and mean curvatures.

This part of the thesis is organized in three chapters. In the first one, we give an elementary presentation of the theory of normal cycles, and how they can be used to define curvature for both smooth and polyhedral surfaces in a unified way. We also state the main theorem precisely. The second chapter, which is more mathematical, is devoted to the proof of the theorem. We will actually give a more general result, applying to a broad class of objects, from which follow corollaries corresponding to specific cases, such as the one of restricted Delaunay triangulations. Finally, the third chapter shows a practical application of curvature estimation, which we already briefly mentioned : the anisotropic remeshing of triangulated surfaces.

Chapter 2

Elementary presentation of normal cycles theory

Introduction

The curvature of smooth surfaces has been extensively studied during the past centuries. In a certain sense, it is defined by differentiating the considered surface twice. This definition, obviously, does not generalize to less than twice differentiable surfaces, and *a fortiori* not to piecewise linear surfaces. As shown in figure 2.1 in the case of curves, the very concept of pointwise curvature does not even make sense for the class of piecewise linear objects. Indeed, at any point lying in the interior of an edge, the curvature is 0, whereas at a vertex, it seems infinite. This problematic situation is easily overcome by shifting from the pointwise point of view to the *measure theoretic* one : instead of considering curvatures at a given point, one should consider integrals of curvature over a given region. For instance, if the curvature at a given point of a polygonal line is not geometrically relevant, it is intuitively clear that the total amount of curvature in the region **B** (in bold in figure 2.1) is β , the angle between the normals at the two endpoints of **B** .

The function that associates to each region B the number β is the simplest example of what are called *curvature measures*. Curvature measures are an appropriate way to deal with the curvature of non necessarily smooth objects. For surfaces, which have a richer geometry, one can define two curvature measures : the gaussian curvature measure, and the mean curvature measure.




Figure 2.1: *What is the curvature of a polygonal line?*

These two curvature measures give a rather coarse description of the geometry of the surface. In a certain sense, they tell how much curvature there is in the considered region, but not in which direction the curvature is. Typically, they contain enough information to recover the radius of a cylinder, but not the direction of its axis. We will show how this problem can be solved by defining a new type of curvature measures, which are tensor-valued : the *anisotropic curvature measures*.

Also, a crucial issue for our purpose of curvature estimation is the one of approximation : when the considered piecewise linear surface is a good approximation of a smooth surface, are respective curvature measures close to each other? Our main theorem states that under certain reasonable assumptions, the answer is yes.

Curvature measures have various interpretations. For convex objects, they are involved in the formula giving the volume of an offset of the object as a function of the offset parameter, as we will see. They can also be defined in the framework of integral geometry [128]. These interpretations, however, do not seem to lead easily to the definition of anisotropic curvature measures, nor to the mentioned approximation result, except for the convex case [67]. To reach this finer level of understanding, one needs to rely on a powerful tool lying at the crossroads of geometric measure theory and differential geometry : the theory of normal cycles.

This chapter is organized in four sections. We first introduce some notations and state the theorem (section 2.1). Then we present the theory of normal cycles (section 2.2) and how they can be used to define anisotropic curvature measures (section 2.3), which is a contribution of this thesis. A rough idea of the proof of the theorem, intended for the not too mathematically

inclined reader, is given in section 2.4.

2.1 Statement of the main theorem

In the sequel, we denote by M a surface in the three dimensional oriented euclidean space $\mathbb{R}^3$. We assume for simplicity that M is the boundary of some compact set $V \subset \mathbb{R}^3$.

2.1.1 Curvature measures

Let us first recall some basic definitions and notations in the case where M is smooth. A good reference for these is [9]. The unit normal vector at a point $p \in M$ pointing outward V will be referred to as $n(p)$. Note that M is thereby oriented. Given a vector v in the tangent plane $T_p M$ to M at p , the derivative of $n(p)$ in the direction v is orthogonal to $n(p)$ as $n(q)$ has unit length for any $q \in M$. The derivative $D_p n$ of n at p thus defines an endomorphism of $T_p M$, known as the Weingarten endomorphism, or shape operator ¹. The Weingarten endomorphism can be shown to be symmetric ; the associated quadratic form is called the second fundamental form. Eigenvectors and eigenvalues of the Weingarten endomorphism are respectively called principal directions and principal curvatures. Both principal curvatures can be recovered from the trace and determinant of $D_p n$, also called mean ² and gaussian curvature at p . Figure 2.1.1 shows the geometric meaning of the second fundamental form at a point p : applied to a unit vector $\vec{v}$ in the tangent plane at p , it yields the signed curvature of the section of the surface by the plane spanned by $n(p)$, $\vec{v}$, and passing through p . Principal directions, displayed in bold, correspond to the values of $\vec{v}$ where the second fundamental form is maximal or minimal. According to the sign of the gaussian curvature, one gets three different cases, respectively depicted in figure 2.1.1 : elliptic (positive), parabolic (zero), and hyperbolic (negative).

As mentioned in the introduction, our result does not involve curvatures at a single point, but rather *curvature measures*, which we define here :

Definition 4 *The gaussian curvature measure of M , ϕ_V^G , is the function that associates with every (Borel³) set $B \subset \mathbb{R}^3$ the quantity :*

$$\phi_V^G(B) = \int_{B \cap M} G(p) dp$$

¹for some reason, most authors add a minus sign in the definition of the Weingarten endomorphism.

²the mean curvature is usually defined as the half trace of the Weingarten endomorphism.

³this restriction is very weak : virtually all sets one can encounter are Borel sets.




Figure 2.2: The second fundamental form

where $G(p)$ is the gaussian curvature of M at point p . Similarly, we define the mean curvature measure ϕ_V^H by :

$$\phi_V^H(B) = \int_{B \cap M} H(p) dp$$

$H(p)$ being the mean curvature of M at point p .

Corresponding objects can be defined for triangulated surfaces. Assume now that V is a polyhedron with vertex set P and edge set E .

Definition 5 The discrete gaussian curvature measure of M , ϕ_V^G , is the function that associates with every (Borel) set $B \subset \mathbb{R}^3$ the quantity :

$$\phi_V^G(B) = \sum_{p \in B \cap P} g(p) \quad (2.1)$$

where $g(p)$ is the angle defect of M at point p , that is 2π minus the sum of angles between consecutive edges incident on p . Similarly, we define the discrete mean curvature measure ϕ_V^H by :

$$\phi_V^H(B) = \sum_{e \in E} \text{length}(e \cap B) \beta(e) \quad (2.2)$$

$|\beta(e)|$ being the angle between the normals to the triangles of M incident on e . The sign of $\beta(e)$ is chosen to be positive if e is convex and negative if it is concave.

In section 2.2 we will see where these formulas come from and why we use the same notation for continuous and discrete curvature measures.

2.1.2 Anisotropic curvature measures

In the case where M is smooth, the second fundamental form of M associates with each point $p \in M$ a 2×2 symmetric bilinear form on $T_p M$, denoted by $H_V(p)$. The 2×2 symmetric bilinear form on $T_p M$ having the same eigenvectors as $H_V(p)$ but with swapped eigenvalues will be denoted $\tilde{H}_V(p)$. As we will see, we will need to extend these bilinear forms to 3×3 symmetric bilinear forms, which we call $\bar{H}_V(p)$ and $\bar{\tilde{H}}_V(p)$. We do so by setting $\bar{H}_V(X, Y) = \bar{\tilde{H}}_V(X, Y) = 0$ whenever X or Y is orthogonal to $T_p M$. In other words, applying one of these 3×3 form to a couple of 3-vectors X and Y amounts to applying the corresponding 2×2 form to the projections of X and Y on $T_p M$. A form related to $\bar{H}_V$ has already been considered by Taubin in [119]. We now introduce two matrix valued measures which are in some sense anisotropic versions of curvature measures :

Definition 6 *The anisotropic curvature measures $\bar{H}_V$ and $\bar{\tilde{H}}_V$ associate with every (Borel) set the 3×3 symmetric bilinear form :*

$$\begin{aligned}\bar{H}_V(B) &= \int_{B \cap M} \bar{H}_V(p) dp \\ \bar{\tilde{H}}_V(B) &= \int_{B \cap M} \bar{\tilde{H}}_V(p) dp\end{aligned}$$

The above definition is the reason why we extended $\tilde{H}_V(p)$ and $H_V(p)$ to $\mathbb{R}^3$: indeed, it would have been impossible to integrate any of these forms directly, since the plane where they are defined depends on the considered point p . Again, corresponding objects can be defined in the polyhedral case. If V is a polyhedron, we define the discrete anisotropic curvature measures by :

Definition 7

$$\begin{aligned}\bar{\tilde{H}}_V(B) &= \sum_{e \in E} \beta(e) \text{length}(e \cap B) \vec{e} \otimes \vec{e} \\ \bar{H}_V(B) &= \sum_{e \in E} \frac{\text{length}(e \cap B)}{2} [(\beta(e) - \sin \beta(e)) \vec{e}^+ \otimes \vec{e}^+ \\ &\quad + (\beta(e) + \sin \beta(e)) \vec{e}^- \otimes \vec{e}^-]\end{aligned}$$

where $\vec{e}$ denotes a unit 3-vector with the same direction as edge e , and e^+ and e^- respectively denote the normalized sum and difference of unit normal vectors to triangles incident on e . If u and v are two vectors, $u \otimes v$ is the bilinear form with matrix $u.v^t$.

2.1.3 Theorem

We now go back to the case where M is smooth. The map pr that associates with each point $p \in \mathbb{R}^3$ its closest point on M is called the *projection* on M . Note that the projection is defined everywhere except on the skeleton of $\mathbb{R}^3 \setminus M$. One can show that it is continuous wherever it is defined [50].

Definition 8 A triangulated surface S is said to be closely inscribed in M if :

- (i) its vertices lie on M
- (ii) the projection pr is defined on S and is a homeomorphism from S to M .

Theorem 13 Let $W \in \mathbb{R}^3$ be a volume whose boundary is a triangulated surface T closely inscribed in M . If B is the relative interior of a union of triangles of T , then :

$$\begin{aligned} |\phi_W^G(B) - \phi_V^G(pr(B))| &\leq C_M K \varepsilon \\ |\phi_W^H(B) - \phi_V^H(pr(B))| &\leq C_M K \varepsilon \\ ||\bar{H}_W(B) - \bar{H}_V(pr(B))|| &\leq C_M K \varepsilon \\ ||\bar{H}_W(B) - \bar{H}_V(pr(B))|| &\leq C_M K \varepsilon \end{aligned}$$

where C_M is a real number depending only on the maximum curvature of M and :

$$\begin{aligned} K &= \sum_{\{t \in T, t \subset \bar{B}\}} r(t)^2 + \sum_{\{t \in T, t \subset B, t \cap \partial B \neq \emptyset\}} r(t) \\ \varepsilon &= \max\{r(t) | t \in B\} \end{aligned}$$

$r(t)$ being the circumradius of triangle t .

In particular, this theorem holds when W is the Delaunay triangulation of an r -sample of M restricted to V , with $r < 0.06$. Indeed, such triangulations are closely inscribed, as shown by Nina Amenta et al.[6]. When the triangles of T have bounded aspect ratio, K boils down to $O(\text{area}(B) + \text{length}(\partial B))$. For restricted Delaunay triangulations, this is the case when the sampling is locally uniform in the sense of [58]. In these cases, theorem 13 shows that estimated curvature measures of W converge to ones of M linearly with respect to the sampling density. At the end of section 3.3.4, we show the convergence of our estimator in the case of restricted Delaunay triangulations, assuming a weaker local uniformity condition [7] on the sample points than the one used in [58].

2.2 Normal cycles and curvature measures

Introduced by Wintgen and Zähle [130, 126], the theory of normal cycles provides a unified way to define curvature for both smooth and polyhedral surfaces. Here is a very crude overview of their approach.

The first observation is that the curvature measures of a smooth surface actually are byproducts of an object associated to the surface, called the normal cycle of the surface. More precisely, the curvature measures of a surface can be easily recovered from the normal cycle of the surface. Second, the definition of the normal cycle of a surface has a unique natural extension to the polyhedral case. Finally, the curvature measures of a polyhedron are defined to be the measures recovered from its normal cycle. Before explaining what a normal cycle is, we shortly review an early approach to curvature measures and the required background.

2.2.1 A first approach

Historically, curvature measures were introduced by considering offsets of V , via the so-called tube formula. First assume that M is smooth.




Figure 2.3: *The tube formula*

Let ρ be the distance between M and the medial axis of the complement of V (ρ is sometimes called the *reach* of M) and

$$V_\varepsilon = \{p | p \notin V, d(p, V) < \varepsilon\} \subset \mathbb{R}^3$$

that is the ε -offset of V minus V . The tube formula then reads :

$$\text{Vol}(V_\varepsilon) = \text{area}(M)\varepsilon + \phi_V^H(M)\frac{\varepsilon^2}{2} + \phi_V^G(M)\frac{\varepsilon^3}{3}$$

for $\varepsilon < \rho$. Moreover, this formula can be localized : if one only considers the part $V_\varepsilon(B)$ of V_ε that projects on a subset B of M , then we have :

$$\text{Vol}(V_\varepsilon(B)) = \text{area}(B)\varepsilon + \phi_V^H(B)\frac{\varepsilon^2}{2} + \phi_V^G(B)\frac{\varepsilon^3}{3}$$

In the smooth case, the volume of $V_\varepsilon(B)$ is thus a polynomial in ε , and its coefficients are multiples of the curvature measures of B . H. Federer [49] actually showed that the volume of $V_\varepsilon(B)$ is always a polynomial in ε for $\varepsilon < \rho$, even if the boundary of V is not smooth. The coefficients of this polynomial thus provide a way to generalize the definition of curvature measures as soon as ρ is strictly positive. For instance, if V is a convex polyhedron, the obtained definitions agree with definition 5⁴.

Unfortunately, this approach breaks down as soon as ρ equals 0, which is already the case when V is a non convex polyhedron ; this is the reason why the theory of normal cycles was developed.

2.2.2 Background

The reader acquainted with exterior differential calculus might want to skip this section. [25] provides a good introduction to the subject.

2-differential forms

Definition Let $\mathcal{S}$ be a smooth manifold of dimension at least two embedded in some euclidean space $\mathbb{R}^k$. If f is a vector field on $\mathcal{S}$, we denote by $f_x \in T_x\mathcal{S}$ the vector associated with a point $x \in \mathcal{S}$. 2-differential forms are, in a certain sense, 2-dimensional analogs of vector fields :

⁴The case of a convex polyhedron is actually the first considered historically, by Jacob Steiner [116].

Definition 9 A 2-differential form ω on $\mathcal{S}$ associates with every point $x \in \mathcal{S}$ a skew-symmetric bilinear form on $T_x\mathcal{S}$, denoted by ω_x .

The following definition shows how a 2-differential form can be built from two vector fields :

Definition 10 The exterior product $f \wedge g$ of two vector fields f and g on $\mathcal{S}$, is the 2-differential form defined by :

$$(f \wedge g)_x(u, v) = (f_x \wedge g_x)(u, v) = \begin{vmatrix} f_x \cdot u & g_x \cdot u \\ f_x \cdot v & g_x \cdot v \end{vmatrix}$$

for all x in $\mathcal{S}$ and $(u, v) \in T_x\mathcal{S}$.

Exterior products are special cases of 2-differential forms. However, they provide a good intuition of the general case : any 2-differential form can actually be written as a linear combination of exterior products of vector fields. It can be seen from the definition of an exterior product that if A is a linear transformation of the plane P spanned by u and v , then $(f \wedge g)_x(Au, Av) = \det(A)(f \wedge g)_x(u, v)$. In particular, $(f \wedge g)_x(u, v) = (f \wedge g)_x(u', v')$ for any two direct orthonormal frames (u, v) and (u', v') of P . Note that this property extend to general 2-differential forms by linearity. Similarly, we have $f_x \wedge g_x(u, v) = f'_x \wedge g'_x(u, v)$ for any couple of orthonormal frames (f_x, g_x) and (f'_x, g'_x) spanning the same oriented plane. Important examples of exterior products are *area forms*. Area forms are a way to represent oriented surfaces as 2-differential forms. If $\mathcal{T} \subset \mathcal{S}$ is an oriented surface, then the area form of $\mathcal{T}$ is constructed as follows : for each point $x \in \mathcal{T}$, pick a direct orthonormal frame of the tangent plane $T_x\mathcal{T}$, say (u_x, v_x) . For $x \notin \mathcal{T}$, set $u_x = v_x = 0$. The area form of $\mathcal{T}$, denoted by $a_{\mathcal{T}}$, is the 2-differential form $u \wedge v$. Intuitively, area forms can be thought of as fields of surface elements : when applied to two vectors a and b in $T_x\mathcal{S}$, $a_{\mathcal{T}_x}$ yields the signed area of the parallelogram spanned by the projections of a and b on $T_x\mathcal{T}$.

Integration 2-differential forms can be integrated on oriented surfaces, in the same way vector fields can be integrated on oriented curves. To see how, let $\mathcal{T}$ be an oriented surface in $\mathcal{S}$ and, for each $x \in \mathcal{T}$, let (u_x, v_x) be a direct orthonormal frame of the tangent plane $T_x\mathcal{T}$. The integral of a 2-differential form ω on $\mathcal{T}$ is defined to be :

$$\int_{\mathcal{T}} \omega = \int_{\mathcal{T}} \omega_x(u_x, v_x) dx$$

For instance, one has $\int_{\mathcal{T}} a_{\mathcal{T}} = \text{area}(\mathcal{T})$, which is why area forms are called this way.

Change of variable A change of variable is merely a diffeomorphism $\phi : \mathcal{S}' \longrightarrow \mathcal{S}$ where $\mathcal{S}'$ is the manifold where the new variables live. Using such a map, a 2-differential form ω on $\mathcal{S}$ can be transformed into a 2-differential form on $\mathcal{S}'$, by a process called *pullback* :

Definition 11 The pullback of ω by ϕ , denoted by $\phi^* \omega$ is given by :

$$\phi^* \omega_x(u, v) = \omega_{\phi(x)}(D_x \phi(u), D_x \phi(v))$$

for all $x \in \mathcal{S}'$ and $u, v \in T_x \mathcal{S}'$.

In a certain sense, pulling a 2-differential form back amounts to expressing it in terms of the new variables. The change of variable formula relates the integral of a 2-differential form with the one of its pullback. The result turns out to be particularly simple :

$$\int_{\mathcal{S}'} \phi^* \omega = \int_{\phi(\mathcal{S}')} \omega \quad (2.3)$$

For example, if $\mathcal{S} = \mathcal{S}' = \mathbb{R}^2$ and h is an integrable function from $\mathcal{S}$ to $\mathbb{R}$, applying (2.3) to $\omega = h a_{\mathcal{S}}$ yields $\phi^* \omega = \text{Jac}(\phi) h \circ \phi a_{\mathcal{S}'} : (2.3)$ thus generalizes the classical change of variable formula. For this formula to hold, ϕ need actually not be a diffeomorphism from $\mathcal{S}'$ to $\mathcal{S}$; the only requirement is that ϕ should be a diffeomorphism from $\mathcal{S}'$ to $\phi(\mathcal{S}')$.

Integral 2-currents

Integral 2-currents generalize oriented surfaces [91]. They can be formally defined as linear combinations of oriented surfaces with integral coefficients. In particular, any oriented surface $\mathcal{T}$ can be considered as an integral 2-current, which we will abusively also denote $\mathcal{T}$. Integration of 2-differential forms is extended to integral 2-currents by linearity :

$$\int_{n\mathcal{T} + p\mathcal{T}'} \omega = n \int_{\mathcal{T}} \omega + p \int_{\mathcal{T}'} \omega$$

The surface that is setwise the same as $\mathcal{T}$ but with reverse orientation thus corresponds to the current $-\mathcal{T}$. Geometrically, integral 2-currents can be thought of as oriented surfaces with *multiplicities*. For instance, if $\mathcal{T}$ and $\mathcal{T}'$ are two oriented surfaces such that orientations of $\mathcal{T}$ and $\mathcal{T}'$ agree on $\mathcal{T} \cap \mathcal{T}'$, $\mathcal{T} + \mathcal{T}'$ can be represented as $\mathcal{T} \cup \mathcal{T}'$ endowed with the same orientation as $\mathcal{T}$ and $\mathcal{T}'$, points in $\mathcal{T} \cap \mathcal{T}'$ having a multiplicity equal to 2. If orientations of $\mathcal{T}$ and $\mathcal{T}'$ do not agree, then summing $\mathcal{T}$ and $\mathcal{T}'$ yields a cancellation on $\mathcal{T} \cap \mathcal{T}'$.




Figure 2.4: Sum of integral currents.

Invariant 2-forms

Now set $\mathcal{S} = \mathbb{R}^3 \times S^2$. $\mathcal{S}$ is obviously a subset of $\mathbb{R}^3 \times \mathbb{R}^3$. We will call the first factor of the latter product the *point space*, E_p , and the second one the *normal space*, E_n . The reason for this is that an element of $\mathcal{S}$ can be thought of as a point in space together with a unit normal vector. If u is a 3-vector, u^n will denote the vector $(0, u) \in E_p \times E_n$, and u^p the vector $(u, 0) \in E_p \times E_n$. Rigid motions of $\mathbb{R}^3$ can be naturally extended to $\mathcal{S}$: if g is such a motion, one can set $\hat{g}(p, n) = (g(p), \bar{g}(n))$, where $\bar{g}$ is the rotation associated with g . We now define two particular 2-differential forms on $\mathcal{S}$:

Definition 12 Let $(p, n) \in \mathcal{S}$ and $x, y \in \mathbb{R}^3$ such that (x, y, n) is a direct orthonormal frame of $\mathbb{R}^3$. We set :

$$\begin{aligned}\omega_{(p,n)}^H &= x^p \wedge y^n + x^n \wedge y^p \\ \omega_{(p,n)}^G &= x^n \wedge y^n\end{aligned}$$

One can actually check that these 2-forms do not depend on the choice of x and y . Moreover, they are *invariant under rigid motions*, that is satisfy $\hat{g}^* \omega = \omega$ for all rigid motion g . Geometric interpretations of these forms will be given in section 2.2.5. The dimension of the space of invariant forms is actually 4 [93].

2.2.3 Smooth case

The theory of normal cycles is inspired by the same ideas as the one presented in section 2.2.1, but transposed in a setting where they can be generalized : the theory of currents. Loosely speaking, normal cycles are a way to unfold offsets in a higher dimensional space :

Definition 13 The normal cycle $N(V)$ of V^5 is the current associated with the set :

$$ST^\perp V = \{(p, n(p)) \mid p \in M\} \subset E_p \times E_n$$

endowed with the orientation induced by the one of M .

M and $ST^\perp V$ are obviously diffeomorphic via the map :

$$\begin{aligned} i: M &\longmapsto ST^\perp V \\ p &\longmapsto (p, n(p)) \end{aligned}$$

The connection between normal cycles and curvature measures lies in the following lemma :

Lemma 14

$$\begin{aligned} \int_{N(V)} \omega_{|i(B \cap M)}^G &= \phi_V^G(B) \\ \int_{N(V)} \omega_{|i(B \cap M)}^H &= \phi_V^H(B) \end{aligned}$$

for all (Borel) set $B \in \mathbb{R}^3$.

Here $\omega_{|i(B \cap M)}$ denotes the restriction of ω to $i(B \cap M)$, that is the form that coincides with ω on $i(B \cap M)$ and vanishes elsewhere. In words, curvature measures of a surface can be recovered by integrating specific differential forms on its normal cycle.

Proof. By definition we have :

$$\int_{N(V)} \omega_{|i(B \cap M)}^G = \int_{i(B \cap M)} \omega^G$$

The change of variable formula now states that :

$$\int_{i(B \cap M)} \omega^G = \int_{(B \cap M)} i^* \omega^G$$

To prove the first claim, it is thus sufficient to show that :

$$i^* \omega^G = G a_M$$

Let (u, v) be a direct orthonormal frame of $T_x M$, where $x \in M$. By definition, we have :

$$\forall w \in T_x M \quad D_x i(w) = w^p + D_x n(w)^n$$

⁵we will sometimes abuse the terminology and write 'the normal cycle of the oriented surface M ' instead.

Expressing $\omega_{i(x)}^G$ in the frame (u^n, v^n, n_x^n) , we get

$$\begin{aligned} (i^* \omega^G)_x(u, v) &= \begin{vmatrix} u^n \cdot (u^p + D_x n(u)^n) & v^n \cdot (u^p + D_x n(u)^n) \\ u^n \cdot (v^p + D_x n(v)^n) & v^n \cdot (v^p + D_x n(v)^n) \end{vmatrix} \\ &= \begin{vmatrix} u \cdot D_x n(u) & v \cdot D_x n(u) \\ u \cdot D_x n(v) & v \cdot D_x n(v) \end{vmatrix} = G(x) \end{aligned}$$

The proof of the second equality is similar. We omit it here as we will prove a stronger result in section 2.3. $\square$

2.2.4 Convex case

When V is convex, a normal cycle can be defined even if M is not smooth. Indeed, in place of normal vectors, we can consider *normal cones* :

Definition 14 *The normal cone $NC_V(p)$ of a point $p \in V$ is the set of unit vectors v such that :*

$$\forall q \in V \quad \overrightarrow{pq} \cdot v \leq 0$$

Definition 15 *The normal cycle $N(V)$ of V is the current associated with the set*

$$\{(p, n) | p \in \partial V \quad n \in NC_V(p)\}$$

endowed with the orientation induced by the one of ∂V .

In particular, when V is convex and smooth, this definition agrees with the one given in the previous section. We now state a crucial property of the normal cycle, which we could have stated in the smooth case as well : additivity.

Proposition 15 *Let V_1 and V_2 be two convex sets in $\mathbb{R}^3$ such that $V_1 \cup V_2$ is convex. Then :*

$$N(V_1 \cap V_2) + N(V_1 \cup V_2) = N(V_1) + N(V_2)$$

Proof. It is sufficient to show that the multiplicities of any point (p, n) in $N(V_1 \cap V_2) + N(V_1 \cup V_2)$ and $N(V_1) + N(V_2)$ agree. If p does not belong to $\partial V_1 \cap \partial V_2$, this is obvious. If p lies in $\partial V_1 \cap \partial V_2$, one concludes easily by noticing that $NC_{V_1 \cap V_2}(p) = NC_{V_1}(p) \cup NC_{V_2}(p)$ and $NC_{V_1 \cup V_2}(p) = NC_{V_1}(p) \cap NC_{V_2}(p)$. $\square$

In figure 2.5 normal cycles are graphically represented by their image under the map sending $(p, n) \in E_p \times E_n$ to $p + n$.




Figure 2.5: *Additivity of the normal cycle*

2.2.5 Polyhedral case

Definition

Once we know what the normal cycle of a convex is, there is at most one way of defining the normal cycle of a polyhedron while keeping the additivity property. Indeed, if one is given a triangulation of the polyhedron V into tetrahedra $t_i, i = 1..n$, the normal cycle of V has to be :

$$N(V) = \sum_{n=1}^{\infty} (-1)^{n+1} \sum_{1 \leq i_1 < \dots < i_n \leq n} N(\cap_{j=1}^n t_{i_j})$$

by application of the inclusion-exclusion principle. We will give a geometric description of the obtained current that does not depend on the chosen triangulation, so that $N(V)$ is well-defined in the polyhedral case.

Simplices

Let us now describe the normal cycle of the polyhedron V . The way it is defined suggests to look first at the normal cycle of simplices. Remember that intuitively, these are unfolded versions of offsets of simplices. Just as their offsets, normal cycles of simplices can be decomposed into *spherical parts*, *cylindrical parts*, and *planar parts*. The difference is that these parts now live in $E_p \times E_n$. We will say that a subset A of $E_p \times E_n$ lies above a subset $B \subset \mathbb{R}^3$ if the projection of A on the point space is included in B . Let us now describe in turn each type of part for a simplex S of varying dimension :

- *spherical parts* lie above vertices of S . They are subsets of $\{(p, n) \mid \|n\| = 1\}$ where p is the considered vertex. If S is reduced to p , then the spherical part is a whole sphere. In case S is an edge, then each spherical part is a half sphere. When S is a triangle, they are spherical 2-gons, and if S is a tetrahedron, they are the spherical triangles spanned by the normals of neighboring facets. Edges of these spherical polygons are dual to edges of S , and the external angle between two incident spherical edges equals the angle between corresponding dual edges.
- *cylindrical parts* lie above edges of S . They are included in $\{(p, n) \mid p \in e, \|n\| = 1, n \cdot e = 0\}$, e being the considered edge. If S is reduced to e , the cylindrical part is a whole cylinder. If S is a triangle, it is a half-cylinder, and if S is tetrahedron, it is a portion of cylinder whose section is a circle arc joining the normals to incident facets.
- *planar parts* lie above facets. They have the form $\{(p, n) \mid p \in t\}$ where t is the considered triangle and n a unit normal vector to t . If S is reduced to t , both possible orientations for n have to be taken into account, whereas if S is a tetrahedron, one should only consider the outward normal.

General case

We can now go back to the case of a general polyhedron V . To begin with, for any point p lying in the interior of V , there is a triangulation of V such that p lies in the interior of a tetrahedron. Thus, there is nothing lying above the interior of V in $N(V)$. By a similar argument restricted to a face f of $M = \partial V$, the part of $N(V)$ lying above f is the planar part $f \times n$, where n is the outward normal to f . The two remaining cases are slightly more involved.

Above edges If e is a convex edge of V , V can be triangulated in such a way that e is included in only one tetrahedron. Above e , $N(V)$ thus coincides with the normal cycle of this tetrahedron : we get a cylindrical part delimited by the normals to faces incident on e , with multiplicity 1. If e is concave, one can find a triangulation of V such that e is an edge of exactly two tetrahedra t and t' . Above e , $N(V)$ is the sum of the cylindrical parts of $N(t)$ and $N(t')$ lying above e minus the cylindrical part of $N(t \cap t')$ lying above e .

The above picture shows a cross section along a plane perpendicular to the concave edge e . As can be seen, we get again a cylindrical part delimited by the normals to facets incident on e ,




Figure 2.6: *Normal cycle above a concave edge*

but with a multiplicity equal to -1 , that is with reverse orientation.

Above vertices Above a vertex p , the situation is more involved, as we obtain a linear combination of at least degree of p half-spheres, spherical 2-gons and spherical triangles. In magnitude, one can get arbitrarily large multiplicities if p is not supposed to be convex. A full description of the part of $N(V)$ lying above p can be given by computing the multiplicity $\mu_V(p, h)$ (or $\mu(p, h)$ for short) of (p, h) in $N(V)$ for each unit vector h :

Lemma 16

$$\mu_V(p, h) = \chi(St_V^+(p, h))$$

where χ is the Euler characteristic and $St_V^+(p, h)$ is the upper star of p , that is the union of relative interiors of cells⁶ of V incident on p and lying in the half plane $\{x \mid \vec{p}\vec{x} \cdot h \geq 0\}$

Proof. One checks easily that both sides coincide when V is a simplex. Indeed, $St_V^+(p, h)$ is the cone with apex p and with base (which we call b) the union of relative interiors of simplices of the triangle t opposite to p in V that lie in the half plane $\{x \mid \vec{p}\vec{x} \cdot h \geq 0\}$. If h is in the normal cone of V at p , then $St_V^+(p, h) = \{p\}$ so $\mu_V(p, h) = 1 = \chi(St_V^+(p, h))$. If not, then the plane $\{x \mid \vec{p}\vec{x} \cdot h = 0\}$ meets t . Thus, b is either a vertex of t or an edge of t . In any of these cases, $\chi(b) = 1$. As the Euler characteristic of a cone is 1 minus the one of the base, we have $\mu_V(p, h) = 0 = \chi(St_V^+(p, h))$.

The claimed result then follows as both sides have the additivity property with respect to V .

□

⁶The relative interior of a point is that point.

Note that this quantity, also called the *index* of p with respect to the direction h [8][46], is always smaller than 1 if p is regular in V .

Curvature measures for polyhedra

Curvature measures for a polyhedron V are defined as the outcome of the integration of corresponding invariant forms on $N(V)$, just like in the smooth case. Thanks to the structure of $N(V)$, it is sufficient to compute the integral of these forms on spherical, cylindrical and planar parts :

- a tangent plane to a planar part is spanned by two vectors of E^p . Applying ω^G or ω^H to a couple of two such vectors yields determinants with at least one zero column. Planar parts thus do not contribute to the curvature measures ϕ^H and ϕ^G , as could be expected.
- the tangent plane to a cylindrical part at a point (p, n) is spanned by u^p and v^n , where u is a vector parallel to the corresponding edge and v is orthogonal to u . For the same reason as above, ω^G vanishes when applied to (u^p, v^n) . u and v can be chosen so that (u, v, n) is a direct orthonormal frame. Expressing $\omega_{(p,n)}^H$ in the frame (u^p, v^n) , one obtains :

$$\omega_{(p,n)}^H(u^p, v^n) = \begin{vmatrix} u^p \cdot u^p & v^n \cdot u^p \\ u^p \cdot v^n & v^n \cdot v^n \end{vmatrix} + \begin{vmatrix} u^n \cdot u^p & v^p \cdot u^p \\ u^n \cdot v^n & v^p \cdot v^n \end{vmatrix} = 1$$

The integral of ω^H over a subset of a cylindrical part thus equals the area of this subset.

- a tangent plane to a spherical part is spanned by two vectors of E^n . Thus, integrating ω^H on a subset of such a part yields 0. Integrating ω^G yields the area of the subset, by a computation similar to the one given above.

The curvature measure $\phi_V^H(B)$ of a subset $B \in \mathbb{R}^3$ is the sum of the areas of cylindrical parts of $N(V)$ lying above B , weighted by their multiplicities. By the previously given description of $N(V)$ above edges, one obtains indeed the formula (2.2).

$\phi_V^G(B)$ is obtained by summing the areas of spherical parts lying above B weighted by their multiplicities. Let us do the computation for parts lying above a vertex $p \in B$. V can be triangulated such that all tetrahedra incident on p share an edge pq . These tetrahedra can be numbered in a circular order around pq , say $t_i, i = 1..n$. Let p_i be the common vertex of t_i, t_{i+1} and M , considering indices mod n , α_i be the angle $\widehat{p_{i-1}p p_i}$ and $\beta_i = \widehat{p_i p q}$. For each simplex S incident on p , the area of the spherical part $SP(S)$ of $N(S)$ lying above p is :

- $S = pq : 2\pi$
- $S = p_i pq : 2(\pi - \beta_i)$ as $SP(S)$ is a spherical 2-gon with angle $\pi - \beta_i$ at its vertices
- $S = t_i : 2\pi - \alpha_i - \beta_{i-1} - \beta_i$ by the formula giving the area of a spherical triangle as a function of its angles

Let us now apply the inclusion-exclusion principle to find the coefficient of each of the areas described above in the linear combination giving $\phi_V^G(p)$. Areas of $SP(t_i)$ appear once each. Intersecting two tetrahedra t_{i-1} and t_i yields the triangle $p_i pq$; as these are obtained exactly once $SP(p_i pq)$ has coefficient -1 . The remaining $(n^2 - 3n)/2$ pairwise intersections all equal pq . For $k \geq 3$, k -fold intersections also equal pq . Hence, the coefficient of $SP(pq)$ is :

$$-\frac{n^2 - 3n}{2} + \sum_{k=3}^{\infty} (-1)^{k+1} \binom{n}{k} = 1$$

Finally, we have :

$$\begin{aligned} \phi_V^G(p) &= \sum_{i=1}^n (2\pi - \alpha_i - \beta_{i-1} - \beta_i) - \sum_{i=1}^n 2(\pi - \beta_i) + 2\pi \\ &= 2\pi - \sum_{i=1}^n \alpha_i \end{aligned}$$

that is the classical definition of the angle defect at p . This computation thus agrees with the definition given by equation (2.1). Note that unlike the mean curvature measure, the gaussian curvature measure is independent on the orientation, in the sense that choosing the complement of V instead of V would yield the same measure.

2.3 The second fundamental form via the normal cycle

The concept of normal cycle was introduced to define mean and gaussian curvature measures for a general class of objects. In this section, we show that it can actually provide a complete description of the curvature of an object. Not surprisingly, integrating invariant forms on the normal cycle yields integrals of invariants of the Weingarten endomorphism, namely its trace and determinant. The basic idea here is to integrate non invariant forms in order to obtain integrals of each coefficients in the Weingarten endomorphism matrix.

2.3.1 Two more 2-differential forms

We now define, for each couple of 3-vectors X and Y , two 2-differential forms on $E^p \times E^n$ from which we will recover the second fundamental form.

Definition 16 Given a point $(p, n) \in E^p \times E^n$ we set :

$$\begin{aligned}\omega_{(p,n)}^{X,Y} &= (n^p \times_p X^p) \wedge Y^n \\ \tilde{\omega}_{(p,n)}^{X,Y} &= X^p \wedge (n^n \times_n Y^n)\end{aligned}$$

where $\times_n$ and $\times_p$ respectively denote cross products in E^n and in E^p .

Note that these two forms are bilinear in X and Y , but not symmetric. However, we will see that integrating them on normal cycles yields symmetric bilinear forms.

2.3.2 Smooth case

Lemma 17 If M is smooth, then :

$$\begin{aligned}\int_{N(V)} \tilde{\omega}_{|i(B \cap M)}^{X,Y} &= \bar{H}_V(B)(X, Y) \\ \int_{N(V)} \omega_{|i(B \cap M)}^{X,Y} &= \bar{H}_V(B)(X, Y)\end{aligned}$$

Proof. As in the the proof of lemma 14, we perform a change of variable in the left-hand side. To compute $i^* \tilde{\omega}^{X,Y}$ at a point $p \in M$, we consider the direct orthonormal frame (e_1, e_2, n) of $\mathbb{R}^3$ where e_1 and e_2 are principal directions and $n = n(p)$. If the principal curvatures associated with e_1 and e_2 are respectively λ_1 and λ_2 , we have :

$$\begin{aligned}i^* \tilde{\omega}_p^{X,Y}(e_1, e_2) &= \tilde{\omega}_{(p,n)}^{X,Y}(e_1^p + \lambda_1 e_1^n, e_2^p + \lambda_2 e_2^n) \\ &= \begin{vmatrix} X.e_1 & \det(n, Y, \lambda_1 e_1) \\ X.e_2 & \det(n, Y, \lambda_2 e_2) \end{vmatrix} \\ &= \begin{vmatrix} X_1 & -\lambda_1 Y_2 \\ X_2 & \lambda_2 Y_1 \end{vmatrix} \\ &= \lambda_2 X_1 Y_1 + \lambda_1 X_2 Y_2 \\ &= \bar{H}_V(p)(X, Y)\end{aligned}$$

where X_i and Y_i , $i = 1, 2, 3$ are the components of X and Y in (e_1, e_2, n) . The first claim thus follows. The second one can be proved in a similar way. $\square$

2.3.3 Polyhedral case

The fact that integrals of the second fundamental form can be recovered from the normal cycle of a smooth surface by integrating 2-differential forms enables us to define corresponding objects for polyhedral surfaces, as they also have a normal cycle. The next lemma justifies definition 7 :

Lemma 18 *If V is a polyhedron, then :*

$$\begin{aligned} \int_{N(V)} \tilde{\omega}_{|i(B \cap M)}^{X,Y} &= \bar{H}_V(B)(X, Y) \\ \int_{N(V)} \omega_{|i(B \cap M)}^{X,Y} &= \bar{H}_V(B)(X, Y) \end{aligned}$$

Proof. Clearly, $\tilde{\omega}^{X,Y}$ vanishes on planar and spherical parts. Let e be an edge of M , or a segment included in such an edge, and CP be the cylindrical part of $N(V)$ lying above e . The tangent plane to CP at $(p, n) \in CP$ has a direct orthonormal frame of the form (u^p, v^n) , where u is a unit vector parallel to e . We have :

$$\begin{aligned} \tilde{\omega}_{(p,n)}^{X,Y}(u^p, v^n) &= \begin{vmatrix} X.u & \det(n, Y, 0) \\ 0 & \det(n, Y, v) \end{vmatrix} \\ &= (X.u)(Y.u) \end{aligned}$$

As a function of X and Y , $\tilde{\omega}_{(p,n)}^{X,Y}(u^p, v^n)$ is thus a symmetric bilinear form with 1 as unique non-zero eigenvalue and u (or e) as associated eigenvector. It thus equals $\vec{e} \otimes \vec{e}$. Integration on CP yields :

$$\int_{CP} \tilde{\omega}^{X,Y} = \beta(e) \text{length}(e) \vec{e} \otimes \vec{e}$$

and the first result follows. The derivation of the second one follows the same lines and is left to the reader. $\square$

Both in the smooth and polyhedral case, anisotropic curvature measures generalize the mean curvature measure : indeed, the trace of $\bar{H}_V(B)(X, Y)$ or $\bar{H}_V(B)(X, Y)$ equals $\phi_V^H(B)$.

2.4 Sketch of proof of the theorem

The idea behind the proof of theorem 13 is roughly as follows. Let E denote the part of $N(W)$ lying above B and D be the part of $N(M)$ lying above $pr(B)$ (figure 2.7). Consider for simplicity that E and D are oriented surfaces, though it is not really accurate as they actually are currents. By lifting the projection pr to $E^p \times E^n$, one obtains a map f from E to D . Define C to be the union of all line segments joining points of E with their image under f . C is a volume whose boundary is the union of E , D , and a surface A which is the union of all line segments joining points of ∂E with their image under f . By applying Stokes theorem to C and a 2-differential form ω , one can express the difference between integrals of ω on E and D as the integral of ω on A plus an integral on C . In particular, when ω is a form associated with some curvature measure, this implies that the difference between the considered curvature measures of E and D is the sum of an integral on A plus an integral on C . In this particular case, the quantities to be integrated on E and D are bounded. Thus, to get a bound on the difference of curvature measures, it is sufficient to bound the volume of C and the area of A .

To do so we first bound the area of E and the length of ∂E , which respectively give rise to the two terms in the constant K involved in the theorem (see section 3.3.4).

Also, we show that the line segments from which C and A are built are short. More precisely, the distance between points of E and their image under f is $O(\varepsilon)$ with $\varepsilon = \max\{r(t) | t \in B\}$. Equivalently, the distance between a point $p \in B$ and its projection $pr(p)$ is $O(\varepsilon)$, as well as the angle between the normal vector at $pr(p)$ and any vector h such that (p, h) lies in the normal cycle of W (see section 3.3.4).

Loosely speaking, C (resp. A) is thus a distorted cylinder the height of which is $O(\varepsilon)$ and the base of which E (resp. ∂E) has area (resp. length) bounded by K . If C (resp. A) were a true cylinder, that is if all line segments were parallel, one could deduce that its volume (resp. area.) is $O(K\varepsilon)$. But this is not the case, so we have to take the distortion into account. It turns out that the effect of this distortion is to scale the volume at most by a number C_M depending only on the maximum curvature of M on $pr(B)$. One thus gets the desired bound $O(C_M K \varepsilon)$ claimed in theorem 13.




Figure 2.7: *Proof of theorem 13*

2.5 Results

In practice, in order to estimate the curvature tensor of a surface at a vertex of a mesh approximating that surface, we choose a small neighborhood B of triangles surrounding the vertex as averaging region, and return $\tilde{H}_T(B)$, which has a simpler expression than $\bar{H}_T(B)$. To get the principal directions, we compute the eigendirections of the obtained matrix. The one associated with the smallest eigenvalue in magnitude provides an estimation of the normal to the surface, and the remaining two others give approximations of the principal directions, provided associated eigenvalues are sufficiently different. These eigenvalues, divided by the area of B , give estimates of the principal curvatures. In the following, we take for B a fixed number of rings around the considered vertex. Note that theorem 13 does not ensure linear convergence with the sampling density in this case, as the averaging region shrinks as the sampling increases. Still, we will see that in practice linear convergence seems to hold, at least for a large enough number of rings.

Figure 2.8 shows the result of the estimation on a triangulated torus with 1000 vertices. This mesh was obtained by applying a Delaunay-based ([15]) reconstruction algorithm to a uniform random sampling of the surface of a torus. The estimated principal directions are close to the actual ones, that is lie along the meridians and parallels of the torus. However, a closer look shows that the estimation sometimes fails in the vicinity of what would be the topmost (or bottommost) part of the torus, if the torus would lie flat on a table. More precisely, the principal direction associated with the minimum curvature (in yellow) is correct, but the two other directions might not be (see the green rectangles in figure 2.8). The reason is that in the mentioned




Figure 2.8: *Frame of principal directions estimated on a meshed torus. For each vertex, the averaging domain used for computations is the 2-ring of that vertex.*

area, one of the principal curvature vanishes and thus the corresponding eigendirection can be mixed up with the normal direction, as they both lie in the same eigenspace of $\tilde{H}_M$. A possible way of solving this problem would be to combine this result with the one of the alternate tensor $\bar{H}_T(B)$, so as to recover a correct estimation of the principal direction associated with the maximum curvature.

Figure 2.9 shows a logarithmic plot of the mean estimation error as a function of the number of points of the torus. The solid curve corresponds to the angle (in degrees) between the true direction of minimum curvature and the estimated one. The dotted curve displays the estimation




Figure 2.9: Mean error on the minimal curvature and on the associated direction

error on the maximum curvature. From left to right, the three figures were respectively obtained using the 1, 2, and 3-ring of each vertex as averaging domain. Due to the high irregularity of the sampling, 1-rings are too small averaging regions to get a reliable curvature estimation. Using 3-rings, one sees that the logarithm of the error decreases with a slope roughly equal to -0.5 with respect to the logarithm of the number of points. This means that apparently, the estimated quantities converge linearly with the sampling density. On a torus with 10000 vertices, the observed average deviation on the direction of minimal curvature is 1.6° , and the mean error on the maximum curvature is less than 2%. For low resolution meshes, using 3-rings produces an unwanted blurring effect resulting in a less accurate estimation. For instance, on a torus with 200 vertices, 2-rings provide better results.

The estimator seems to be, to some extent, robust against noise : on a noisy cylinder, the estimated principal directions using 3-rings are globally correct (figure 2.10). Note however that the results are not as good if one uses only 1 or 2-rings as averaging domains. On this example, the noise added to the vertices of the cylinder is about one fourth of the typical spacing between samples.

Figure 2.13 shows the result on a mesh of Michelangelo's David. It is of course delicate to claim that the result is correct, as the actual curvatures are not available, and in fact not even defined -what is the curvature of a real world object? Still, the fact that displaying only the estimated minimal curvature directions is enough to recognize the masterpiece of Michelangelo proves in some sense the quality of the result.




Figure 2.10: *Estimated principal directions on a noisy cylinder using 3-rings*




Figure 2.11: *The principal directions estimated on an ellipse with 1442 vertices, are very similar to the actual ones, whose integral lines are shown in figure 2.12.*




Figure 2.12: *Lines of curvature of an ellipse*



Figure 2.13: *Directions of minimal curvature estimated on a mesh of Michelangelo's David. For each vertex, the averaging domain used for computations is the 2-ring of that vertex.*

Chapter 3

Approximation of normal cycles

Introduction

In his article *Euler characteristic and finitely additive Steiner measures* [89], John Milnor raises the following question : in what sense do two sets have to be close to each other, in order to guarantee that their curvature measures are close to each other? This open problem lies at the heart of our approach to curvature estimation, since we estimate the curvature measures of a surface by the ones of a piecewise linear approximation of that surface. Indeed, any answer to the problem leads to conditions on the approximating triangulation ensuring the accuracy of the estimation.

This line of research has already received some attention in the past. In the case of two convex sets, the situation is well understood. For instance, the difference between the mean curvature measures of two convex sets is bounded by a constant times their Hausdorff distance [67], as can be seen easily using the integral-geometric interpretation of curvature measures. Unfortunately, these considerations do not seem to apply successfully to the considerably more difficult case of non-convex sets.

If one drops the convexity assumption, it appears that the theory of normal cycles is a relevant tool to tackle the problem. As explained in the previous chapter, the normal cycle of a -sufficiently regular- set is a current that generalizes the notion of unit normal bundle to non-necessarily smooth objects[126], [128], [130],[131],[129]. The key fact about normal cycle theory is that the various curvature measures of an object can be recovered from its normal cycle by integration of well-chosen universal differential forms. Using this interpretation of

curvature measures, Joseph Fu [52] was able to prove, under certain assumptions, that if a sequence of polyhedra converge to a smooth submanifold of an Euclidean space, then their curvature measures converge to the ones of the submanifold. This theorem is established by showing that the normal cycles of the polyhedra converge, in the sense of currents, to the one of the submanifold. The proof of this convergence heavily relies on the compactness theorem for integral currents with bounded mass.

Let us mention for sake of completeness that Cheeger, Müller and Schröder [28] managed to obtain an approximation result for curvature measures, but in a different setting. They showed that the curvature measures of a Riemannian manifold can be well approximated by quantities defined by means of geodesic triangulations of the manifold. Note that this approach is intrinsic and therefore different from ours in spirit.

The goal of this chapter is to refine the result of Joseph Fu by giving a quantitative version of it. Indeed, his convergence result does not shed any light on the relationship between the curvature measures of a smooth submanifold and the ones of a given approximating polyhedron, which is the question we are interested in. We answer the question by giving an estimate of the flat norm of the difference of the normal cycle of a compact n -manifold V of $\mathbb{E}^n$ whose boundary is a smooth hypersurface, and the normal cycle of a -sufficiently regular- compact subset $\mathcal{C}$ in terms of the mass of the normal cycle of V , the Hausdorff distance between ∂V and $\partial \mathcal{C}$, the maximum angle between the normals to V and the “normals” to $\mathcal{C}^1$, and an *a priori* upperbound on the norm of the second fundamental form of ∂V . We thereby give an answer to the question raised by John Milnor in the special case where one of the two sets is smooth.

In the first section of this chapter, we give precise definitions relative to normal cycles, including the extension of anisotropic curvature measures to hypersurfaces of any dimension. In the second one, we prove the approximation result, and in the last one, we give practical corollaries of it such as the theorem 13 stated in the previous chapter. We refer to [32] for a treatment of the material of this chapter in the Riemannian framework.

3.1 Background on geometric measure theory

All the details can be found in [49].

¹this will be given a precise sense even when $\partial \mathcal{C}$ is not smooth.

General currents

Let $\mathcal{D}^m$ the $\mathbb{R}$ -vector space of C^∞ differential m -forms with compact support on $\mathbb{E}^n$. $\mathcal{D}^m$ can be endowed with a topology similar to the topology on the space of test functions used to define distributions, the so-called C^∞ topology. A sequence (ϕ_i) of elements of $\mathcal{D}^m$ converges to $\phi \in \mathcal{D}^m$ in the C^∞ topology if and only if the derivatives of any order of the ϕ_i converge to the corresponding derivatives of ϕ uniformly on every compact.

The topological dual of $\mathcal{D}^m$ is the $\mathbb{R}$ -vector space $\mathcal{D}_m$ of m -currents on $\mathbb{E}^n$. Equivalently, currents can be viewed as differential forms whose coefficients are distributions instead of smooth functions. The *support* $Spt(T)$ of a current T can be defined as the union of the support of its coefficients.

The subset of m -currents with compact support is denoted by $\mathcal{E}^m$. We endow $\mathcal{D}_m$ with the weak topology:

$$\forall \phi \in \mathcal{D}^m \quad \lim_{j \rightarrow \infty} T_j = T \iff \lim_{j \rightarrow \infty} T_j(\phi) = T(\phi)$$

Operations on currents

Basic notions relative to differential forms can be transposed to currents by dualization :

1. Boundary : to each m -current T one can associate a $m - 1$ -current ∂T called its *boundary*, defined by :

$$\forall \phi \in \mathcal{D}^{m-1} \quad \partial T(\phi) = T(d\phi)$$

where $d\phi$ denotes the exterior derivative of ϕ .

2. Push-forward : given an m -current T and a smooth map f defined on a neighborhood of the support of T , one can define the *push-forward* $f_\# T$ of T by f :

$$\forall \phi \in \mathcal{D}^m \quad f_\# T(\phi) = T(f^* \phi)$$

where $f^* \phi$ is the pull-back of ϕ by f . Note that this definition only makes sense when $f^* \phi$ is compactly supported. We thus have to assume that f is proper, that is $f^{-1}(K)$ is compact for every compact K . Actually, for the push-forward to be defined, it is sufficient that the restriction of f to the support of T is proper.

Since f^* commutes with exterior differentiation, $f_\#$ commutes with the boundary operator, so that the push-forward of a current without boundary is also without boundary.

Current representable by integration

We say that a current $T \in \mathcal{D}_m$ is *representable by integration* if there is a Borel regular measure $\|T\|$ on $\mathbb{E}^n$ finite on compact subsets and a unit m -vector fields $\vec{T}$ defined almost everywhere such that

$$\forall \phi \in \mathcal{D}^m \quad T(\phi) = \int \langle \vec{T}, \phi \rangle d\|T\|$$

Currents representable by integration are analogous to distributions of order 0. A current representable by integration T can be “restricted” to any $\|T\|$ -measurable set A (see [50] pp 356). The obtained current $T \llcorner A$ is defined by :

$$\forall \phi \in \mathcal{D}^m \quad T \llcorner A(\phi) = \int \langle \vec{T}, \phi \rangle \mathbf{1}_A d\|T\|$$

Rectifiable and integral currents

In particular, one can associate an m -current representable by integration to any oriented m -rectifiable subset S of dimension m of $\mathbb{E}^{n^2}$. It is a well-known fact that rectifiable sets of dimension m have a well-defined tangent space at $\mathcal{H}^m$ -almost every point. Let $\vec{S}$ be the unit m -vector field encoding these -oriented- tangent spaces. The current associated with S , still denoted by S , is defined by :

$$S(\phi) = \int_S \langle \vec{S}, \phi \rangle d\mathcal{H}^m$$

More general currents can be defined by incorporating integer multiplicities μ in the previous formula :

$$T(\phi) = \int_S \mu \langle \vec{S}, \phi \rangle d\mathcal{H}^m$$

If the support of S is compact, and $\int_S \mu d\mathcal{H}^m < \infty$, we say that T is *rectifiable*. The space of rectifiable currents is denoted by $\mathcal{R}_m$.

A current is said to be *integral* if it is rectifiable and if its boundary is rectifiable.

Mass and norms of currents

The norm of a m -differential form ϕ is the real number

²we refer to [50] for an exposition of the rather subtle notion of rectifiability.

$$||\phi|| = \sup_{p \in M^n} ||\phi_p||,$$

where, for each $p \in M^n$,

$$||\phi_p|| = \sup\{|\langle \phi_p, \zeta_p \rangle|, \zeta_p \in \Lambda^m T_p M^n, |\zeta_p| = 1\}.$$

There are different interesting seminorms on the space of currents $\mathcal{D}_m$. We mention the main ones:

- The mass of a current $T \in \mathcal{D}_m$ is the real number

$$\mathbf{M}(T) = \sup\{T(\phi), \text{ such that } \phi \in \mathcal{D}^m, ||\phi|| \leq 1.\}$$

For rectifiable currents of dimension m , the mass somehow generalizes the notion of m -volume : the mass of the current T defined by 3.1 is $\int_S \mu d\mathcal{H}^m$. Rectifiable currents thus have finite mass. Using general results on representation theory of geometric measure theory, it can be proved that if $\mathbf{M}(T) < \infty$, T is representable by integration.

- The flat norm of a current $T \in \mathcal{D}_m$ is the real number

$$\mathcal{F}(T) = \inf\{\mathbf{M}(A) + \mathbf{M}(B) \text{ such that } T = A + \partial B, A \in \mathcal{D}_m, B \in \mathcal{D}_{m+1}\}.$$

It can be shown that the flat norm can also be expressed in the following way :

$$\mathcal{F}(T) = \sup\{T(\phi), \text{ such that } \phi \in \mathcal{D}^m, ||\phi|| \leq 1, ||d\phi|| \leq 1\}.$$

The constancy theorem

We will use the following important result:

Theorem 19 *Let A be an oriented connected C^1 submanifold of $\mathbb{E}^n$ with boundary, and let T be an integral current whose support lies in A , and such that the support of ∂T lies in ∂A . Then, there exists an integer k such that $T = kA$.*

Note that this theorem actually holds for a larger class of currents, the so-called *real flat chains*.

3.2 Geometric measure theory and curvature measures

3.2.1 Normal cycle of a geometric subset of $\mathbb{E}^n$

Under certain regularity assumptions, one can associate with a compact set of $\mathbb{E}^n$ an integral $n - 1$ -current of (the total space of) $T\mathbb{E}^n$, called its normal cycle³, that generalizes the notion of unit normal bundle to non-necessarily smooth sets. In the previous chapter, we gave the definition of the normal cycle of such a compact set $\mathcal{C}$ in some particular cases : if $\partial\mathcal{C}$ is a C^2 m -dimensional manifold, the normal cycle $N(\mathcal{C})$ of $\mathcal{C}$ is just the integral $n - 1$ -current of $T\mathbb{E}^n$ associated with the outer unit normal bundle of $\mathcal{C}$. If $\mathcal{C}$ is convex, $N(\mathcal{C})$ can be defined in a similar way using normal cones instead of unit normals. Finally, if $\mathcal{C}$ is a polyhedron, one first start by decomposing $\mathcal{C}$ as a union of convex polyhedra $\mathcal{C}_i$. $N(\mathcal{C})$ is then defined by inclusion-exclusion :

$$N(\mathcal{C}) = \sum_{n=1}^{\infty} (-1)^{n+1} \sum_{1 \leq i_1 < \dots < i_n \leq n} N(\cap_{j=1}^n \mathcal{C}_{i_j})$$

By construction, the normal cycle is Euler-additive. These definitions provide a satisfactory way to generalize the notion of unit normal bundle to compact sets with C^2 boundary or compact polyhedra. Joseph Fu [57] proved that the concept of normal cycle can actually be further generalized to a very broad class of objects, which he calls *geometric sets*. Geometric sets are defined in a rather indirect way. Fu first exhibits some basic properties that any reasonable generalization of the normal cycle should satisfy. He then shows that for each compact subset of $\mathbb{E}^n$ there is at most one integral current satisfying the properties (theorem 20). If such a current S exists, the set is said geometric and its normal cycle is defined as being the current S . Before explaining what these properties are, let us give some notations.

We identify the total space of $T\mathbb{E}^n$ with $E \times F$ where E is the base space and F is the fiber⁴. Let $J : E \longrightarrow F$ be the canonical isomorphism between E and F . We endow $T\mathbb{E}^n$ with the dot product $\langle (e, f), (e', f') \rangle = \langle e, e' \rangle + \langle J^{-1}(f), J^{-1}(f') \rangle$. If u is a vector, u^* denotes the dot product by u . The *canonical 1-form* α on $T\mathbb{E}^n$ is defined by

$$\alpha(x, \xi) = \sum_{i=1}^n \xi_i e_i^*$$

³Here, the term “cycle” is used because normal cycles have no boundary.

⁴in chapter 2, we used to denote E by E^n and F by E^p . We change the notation to avoid confusion with the dimension.

where (e_i) is any orthonormal frame of E and ξ_i is the i -th coordinate of $J^{-1}(\xi)$ in (e_i) . The derivative ω of α is called the *symplectic form* on $T\mathbb{E}^n$. For any orthonormal frame (e_i) of E , we have :

$$\omega = e_1^* \wedge J(e_1)^* + \dots + e_n^* \wedge J(e_n)^*$$

Let S be an integral $(n-1)$ -current supported in $ST\mathbb{E}^n$. S is said to be *legendrian* if it cancels α and ω , that is if one has :

$$\forall \phi \in \mathcal{D}^{n-2} T(\phi \wedge \alpha) = 0$$

$$\forall \phi \in \mathcal{D}^{n-3} T(\phi \wedge \omega) = 0$$

The key theorem in Fu's definition of geometric sets is the following :

Theorem 20 (*Uniqueness theorem*) *Let $i : T\mathbb{E}^n \rightarrow \mathbb{R}$. There is at most one closed legendrian integral $(n-1)$ -current S whose support is a compact subset of $ST\mathbb{E}^n$ and such that :*

$$S(\phi(x, \xi)d\xi) = \int_{S^{n-1}} \sum_{x \in \mathbb{E}^n} \phi(x, \xi) i(x, \xi) d\mathcal{H}^{n-1} \xi$$

for any smooth function $\phi : ST\mathbb{E}^n \rightarrow \mathbb{R}$.

here $d\xi$ denotes the pull-back of the volume form of S^{n-1} by the projection $E \times F \rightarrow F$.

Definition 17 *A compact set $\mathcal{C} \subset \mathbb{E}^n$ is said geometric if the current S of theorem 20 exist when $i = i_{\mathcal{C}}$ is the function defined by :*

$$i_{\mathcal{C}}(x, \xi) = \lim_{r \downarrow 0} \lim_{s \downarrow 0} [\chi(\mathcal{C} \cap B(x, r) \cap \{p \mid (p-x) \cdot \xi \leq t\})]_{t=-s}^{t=s}$$

S is then called the normal cycle of $\mathcal{C}$ and denoted by $N(\mathcal{C})$.

When C is a stratified set, $i_{\mathcal{C}}(x, \xi)$ is just the index [23, 61] of x as critical point of the dot product by ξ . $i_{\mathcal{C}}(x, \xi)$ can be viewed as the multiplicity of (m, ξ) in $N(\mathcal{C})$ ([22]).

The definition of geometric sets does not give a practical way to decide whether a given set is geometric. Fortunately, important classes of sets have been shown to be geometric, besides smooth objects and polyhedra. The main examples are subanalytic sets [54], definable sets [14], Riemannian polyhedra [74], or sets with positive reach [57].

3.2.2 Curvature measures from normal cycles

As explained in the previous chapter, the crucial point about normal cycles is that the curvature measures of a smooth object arise as integrals of invariant differential forms on its unit normal bundle, that is on its normal cycle. This interpretation shows that curvature measures can be defined for any compact admitting a normal cycle. As we only gave the definitions of the invariant forms in the case $n = 3$, we now recall their construction in arbitrary dimension.

At any point (m, ξ) of $ST\mathbb{E}^n = \{(m, \xi) \in E \times F \mid \|\xi\| = 1\}$, consider an orthonormal frame $(e_1, \dots, e_{n-1})$ of $J^{-1}(\xi)^\perp$, and $(\varepsilon_1 = Je_1, \dots, \varepsilon_{n-1} = Je_{n-1})$ its image by J . On $ST\mathbb{E}^n$, we can build the $(n-1)$ -differential form

$$\Omega = (e_1^* + t\varepsilon_1^*) \wedge \dots \wedge (e_{n-1}^* + t\varepsilon_{n-1}^*)$$

This form clearly does not depend on the choice of the orthonormal frame (e_i) . The coefficient of t^i in this expression considered as a polynomial in the variable t is thus a well-defined $n-1$ form ω_i . Each ω_i is invariant under the action of the orthogonal group. The ω_i , together with the canonical 1-form and the symplectic form, actually span the algebra of all invariant forms on $ST\mathbb{E}^n$.

Definition 18 *Let $\mathcal{C}$ be a geometric compact subset of $\mathbb{E}^n$. The k -th curvature measure of $\mathcal{C}$, denoted by $\phi_{\mathcal{C}}^k$, associates with each Borel subset B of $\mathbb{E}^n$ the real number*

$$\phi_{\mathcal{C}}^k(B) = N(\mathcal{C}) \llcorner (B \times F)(\omega_k)$$

If $\mathcal{C}$ is the volume enclosed by a C^2 compact hypersurface, $\phi_{\mathcal{C}}^k(B)$ is just the integral over $B \cap \partial\mathcal{C}$ of the k -th symmetric function of the principal curvatures of $\partial\mathcal{C}$. We refer to [93] for a proof, which basically consists in a change of variable. With the notations of the previous chapter, $\phi^H = \phi^1$ and $\phi^G = \phi^2$.

3.2.3 Second fundamental form of geometric sets

In the previous chapter we showed how the normal cycle could be used to define a generalization of the notion of (integral of the) second fundamental form to polyhedra in 3-space : the anisotropic curvature measures. Among the two anisotropic curvature measures introduced in chapter 2, only one, $\bar{H}_{\mathcal{C}}$, seems to generalize naturally beyond the three dimensional case⁵ :

⁵Bernig [13] has recently studied a family of tensors-valued forms which includes this particular one.

Definition 19 Let $\mathcal{C}$ be a geometric compact subset of $\mathbb{E}^n$. The anisotropic curvature measure of $\mathcal{C}$, denoted by $\bar{H}_{\mathcal{C}}$, associates with each Borel subset B of $\mathbb{E}^n$ the following bilinear form on E :

$$\forall X, Y \in E^2 \quad \bar{H}_{\mathcal{C}}(B)(X, Y) = N(\mathcal{C})_{\perp}(B \times F)(\omega_{X,Y})$$

where :

$$\omega_{X,Y} = *_E(J^{-1}(\xi)^* \wedge X^*) \wedge J(Y)^*$$

here $*_E$ denotes the Hodge dual on E .

Just as in the three-dimensional case, when $\mathcal{C}$ is the volume enclosed by a C^2 compact hypersurface, $\bar{H}_{\mathcal{C}}(B)$ is the integral over $B \cap \partial\mathcal{C}$ of a symmetric bilinear form on E related to the second fundamental form of $\partial\mathcal{C}$. More precisely, this form $H_{\mathcal{C}}(m)$ coincides with the second fundamental form of $\partial\mathcal{C}$ on the tangent space of $\partial\mathcal{C}$, and vanishes on its orthogonal complement.

Proposition 21 If $\mathcal{C}$ is the volume enclosed by a C^2 compact hypersurface M , then for any Borel set $B \subset E$:

$$\bar{H}_{\mathcal{C}}(B) = \int_{B \cap M} H_{\mathcal{C}}(m) dm$$

The proof of this fact follows the same lines as in the case $n = 3$:

Proof. As usual, using the Gauss map G of M (see section 3.3.1), we perform a change of variable on the integral of $\omega_{X,Y}$ over $N(V)_{\perp}(B \times F)$ to transform it into an integral over B . We only have to show that $G^*\omega_{X,Y}$ is the bilinear form in X and Y described above times the volume form of M . Fix a point $p \in M$ and let $(e_1, \dots, e_{n-1})$ be an orthonormal frame of principal directions at p . Let i and j be two indices comprised between 1 and $n - 1$. By definition of $\omega_{X,Y}$ and because $DG(u) = (u, J(A_{\xi}(u)))$ for any vector u , we have :

$$\begin{aligned} G^*\omega_{e_i, e_j}(e_1, \dots, e_{n-1}) &= \sum_{k=1}^{n-1} (-1)^{k+1} \det(J^{-1}(\xi), e_i, e_1, \dots, e_{k-1}, e_{k+1}, \dots, e_{n-1}) \langle DG(e_k), J(e_j) \rangle \\ &= (-1)^{i+1} \det(J^{-1}(\xi), e_i, e_1, \dots, e_{i-1}, e_{i+1}, \dots, e_{n-1}) \langle DG(e_i), J(e_j) \rangle \\ &= \langle DG(e_i), J(e_j) \rangle = \langle A_{\xi}(e_i), e_j \rangle \end{aligned}$$

To conclude the proof we need to check that $G^*\omega_{e_i, J^{-1}(\xi)} = 0$ which is obvious from the preceding computation. $\square$

$\bar{H}_{\mathcal{C}}(B)$ is symmetric even if $\mathcal{C}$ is a non smooth geometric set, which is not obvious *a priori*, as $\omega_{X,Y} \neq \omega_{Y,X}$ if X and Y are not collinear. This a consequence of the fact that normal cycles, by definition, cancel the canonical symplectic form ω on $T\mathbb{E}^n$, that is satisfy :

$$\forall \phi \in \mathcal{D}^{n-3} T(\phi \wedge \omega) = 0$$

We recall that $\omega = e_1^* \wedge J(e_1)^* + \dots + e_n^* \wedge J(e_n)^*$ if (e_i) is an orthonormal frame of E . At a point (m, ξ) we will consider an orthonormal frame having the form $(e_1, \dots, e_n, J(e_1), \dots, J(e_n))$ with $e_n = J^{-1}(\xi)$. To prove that anisotropic curvature measures are symmetric, it is sufficient to show that for each i and j , $\omega_{e_i, e_j} - \omega_{e_j, e_i}$ is a wedge of ω with some $(n-3)$ -form ϕ . Actually, as normal cycles are supported in $ST\mathbb{E}^n$, we only need to prove that the restriction of $\omega_{e_i, e_j} - \omega_{e_j, e_i}$ to $ST\mathbb{E}^n$ can be written $\phi \wedge \omega$. First consider the case where neither i nor j equals n . Choose $\alpha = *_E(J^{-1}(\xi)^* \wedge e_i^* \wedge e_j^*)$. Assuming w.l.o.g. that $i < j$, α is $(-1)^{i+j+n}$ times the wedge of e_k^* for $k = 1, \dots, n-1$, i and j being omitted. Now set $\phi = \alpha_{\perp ST\mathbb{E}^n}$. We have :

$$\begin{aligned} \phi \wedge \omega &= \alpha \wedge \omega_{\perp ST\mathbb{E}^n} \\ &= \sum_{k=1}^{n-1} *_E(J^{-1}(\xi)^* \wedge e_i^* \wedge e_j^*) \wedge e_k^* \wedge J(e_k)^* \\ &= *_E(J^{-1}(\xi)^* \wedge e_i^* \wedge e_j^*) \wedge e_i^* \wedge J(e_i)^* + *_E(J^{-1}(\xi)^* \wedge e_i^* \wedge e_j^*) \wedge e_j^* \wedge J(e_j)^* \\ &= (-1)^{j+n+1} \wedge_{k=1, k \neq j}^{n-1} e_k^* \wedge J(e_i)^* + (-1)^{i+n} \wedge_{k=1, k \neq i}^{n-1} e_k^* \wedge J(e_j)^* \\ &= (n-2)!(\omega_{e_i, e_j} - \omega_{e_j, e_i})_{\perp ST\mathbb{E}^n} \end{aligned}$$

which is the desired result (The restriction of the rhs to $ST\mathbb{E}^n$ has been dropped in the intermediate lines of the above computation to simplify the notations). The case $j = n$ is simpler. Indeed, ω_{e_n, e_i} is obviously 0, and ω_{e_i, e_n} also, when restricted to $ST\mathbb{E}^n$.

3.3 An approximation result

In this paragraph, we shall evaluate an upper bound on the flat norm of the difference of the normal cycle of a compact n -manifold V of $\mathbb{E}^n$ the boundary of which is a smooth hypersurface M^{n-1} and the normal cycle of a geometric compact subset $\mathcal{C}$ the boundary of which $\partial\mathcal{C}$ is strongly close to M^{n-1} (see the definition below). We denote by $\delta(A, A')$ the Hausdorff distance between two subsets A and A' of $\mathbb{E}^n$.

3.3.1 Strongly close hypersurfaces

The Gauss map associated to M^{n-1}

We assume that M^{n-1} is the boundary of a compact n -manifold V . Let ξ be the outer unit normal vector field on the hypersurface M^{n-1} . We denote by G -in the last chapter it was denoted by i - the Gauss map associated to M^{n-1} :

$$G : M^{n-1} \hookrightarrow T\mathbb{E}^n$$

is defined by

$$G(m) = (m, \xi_m).$$

The derivative of G at m satisfies :

$$\begin{aligned} DG(m) : T_m M^{n-1} &\rightarrow T\mathbb{E}^n \simeq E \times F \\ X &\longmapsto (X, -A_\xi X) \end{aligned}$$

where A_ξ is the Weingarten endomorphism at m . Denoting by $\|h_m\|$ the norm of the second fundamental form of M^{n-1} at the point m , that is the largest magnitude of the principal curvatures, we thus have in particular :

$$\|DG(x)\| \leq \sup(1, \|h_m\|).$$

The projection on a smooth hypersurface

Since M^{n-1} is smooth, there exists a tubular neighborhood U (of varying radius) of M^{n-1} on which the orthogonal projection $\text{pr}|_U$ from U to M^{n-1} is well defined. By orthogonal projection we mean the function that maps a point to its closest point on M^{n-1} .

The following result is classical (see [50] for instance):

Proposition 22 *The map*

$$\text{pr}|_U : U \rightarrow M^{n-1}$$

is differentiable; moreover, at each point $p \in U$, its derivative is given by the following matrix, in a frame of principal vectors of M^{n-1} completed by the unit normal vector:

$$Dpr|_U(p) = \begin{pmatrix} \frac{1}{1+\delta\varepsilon\lambda_1} & \cdots & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ 0 & \cdots & \frac{1}{1+\delta\varepsilon\lambda_{n-1}} & 0 \end{pmatrix}$$

where $\lambda_1, \dots, \lambda_{n-1}$ are the principal curvatures of M^{n-1} , and $\varepsilon = \pm 1$. In particular,

$$\|Dpr|_U(p)\| \leq \frac{1}{1 - \delta(p, M^{n-1})\|h_{pr|_U(p)}\|}.$$

In what follows, we will say that a compact subset B lying in U is *close to M^{n-1}* . Moreover, if the orthogonal projection onto M^{n-1} induces a bijection between B and M^{n-1} we shall say that B is *strongly close to M^{n-1}* . If B is strongly close to M^{n-1} , then it is homeomorphic to it. In the following, we assume that $\mathcal{C}$ is a n -manifold with boundary which is geometric and whose boundary is strongly close to M^{n-1} .

Comparing the normals

To give a measure of the closeness between the tangent spaces to $\partial\mathcal{C}$ and M^{n-1} , one can compare, for every point m in M^{n-1} , the normal ξ_m to M^{n-1} with the set :

$$\{v \in F \mid (\text{pr}^{-1}(m), v) \in \text{spt } N(\mathcal{C})\}$$

that is the support of the normal cycle of $\mathcal{C}$ lying above $\text{pr}^{-1}(m)$. For every subset B of $\mathcal{C}$, we put :

$$\alpha_B = \sup_{m \in \text{pr}(B)} \sup\{\angle(\xi_m, v) \mid (\text{pr}^{-1}(m), v) \in \text{spt } N(\mathcal{C})\}$$

3.3.2 A homotopy between normal cycles

With the previous notations, consider the map f defined by the following diagram (p_1 denotes the projection on the first factor of $U \times F$) :

$$\begin{array}{ccc} U \times F & \xrightarrow{f} & \text{spt } N(V) \subset V \times F \\ p_1 \downarrow & & \uparrow G \\ U & \xrightarrow{\text{pr}} & M^{n-1} \end{array}$$

Let h be the affine homotopy between f and the identity, [49]:

$$h : (U \times F) \times [0, 1] \rightarrow \text{spt } N(V),$$

given by

$$h(x, X, t) = t f(x, X) + (1 - t)(x, X)$$

Let B' be a Borel subset and $B = B' \cap \partial \mathcal{C}$. To simplify the notations, we define the $(n-1)$ -currents D and D' by $D = N(\mathcal{C}) \llcorner (B \times F)$ and $D' = N(V) \llcorner (pr(B) \times F)$. We define also the n -current $C = h_{\#}(D \times [0, 1])$ and the $(n-1)$ -current A by $A = h_{\#}(\partial D \times [0, 1])$. Note that C and A are indeed defined as h is smooth and proper.




Figure 3.1: *Homotopy between normal cycles*

Proposition 23 *One has:*

$$\mathcal{F}(D - D') \leq (\mathbf{M}(D) + \mathbf{M}(\partial D)) \sup_{\text{spt } D} |f - Id| \sup_{\text{spt } D} (||Df||^{n-2}, ||Df||^{n-1}, 1)$$

Proof.

1. To begin with, we show that $f_{\#}(D) = D'$. First note that $B \times F = f^{-1}(pr(B) \times F)$, so that ([49] pp. 359) :

$$f_{\#}(D) \stackrel{\text{def}}{=} f_{\#}(N(\mathcal{C}) \llcorner (B \times F)) = f_{\#}(N(\mathcal{C})) \llcorner (pr(B) \times F)$$

$f_{\#}(N(\mathcal{C}))$ is an integral $(n-1)$ -current without boundary that is supported in the C^1 $(n-1)$ -manifold $\text{Spt}(N(V))$. Thus, by the constancy theorem (theorem 19), there is an integer c such that $f_{\#}(N(\mathcal{C})) = cN(V)$, which implies $f_{\#}(D) = cD'$. To prove that $c = 1$,

one can assume that $\partial\mathcal{C}$ contains an open set O on which it is smooth. Indeed, if it is not the case, one can modify locally $\partial\mathcal{C}$ leaving B unchanged in such a way that it holds. Then the restriction of p_1 to $O \times F$ is 1-1, thus $f|_{O \times F}$ is also 1-1. We conclude that $f_{\#}(N(\mathcal{C})_{\perp}(O \times F)) = N(V)_{\perp}(pr(O) \times F)$, that is $c = 1 : f_{\#}(D) = D'$.

2. In order to evaluate the flat norm of $D - D'$, we decompose $D - D' = D - f_{\#}(D)$ in a sum of a $(n - 1)$ -current and the boundary of a n -current, using the homotopy formula for currents (cf. [49]) :

$$D - D' \stackrel{def}{=} D - f_{\#}(D) = \partial C - A$$

By definition of the flat norm, we deduce immediately that the flat norm of $(D - D')$ satisfies

$$\mathcal{F}(D - D') \leq \mathbf{M}(C) + \mathbf{M}(A)$$

If D or ∂D has infinite mass, then the result is trivial. If not, they are representable by integration. In this case, we have ([49] 4.1.9.):

$$\mathbf{M}(C) = \mathbf{M}(h_{\#}(D \times [0, 1])) \leq \mathbf{M}(D) \sup_{\text{spt } D} |f - Id| \sup_{\text{spt } D} (||Df||^{n-1}, ||Id||^{n-1}),$$

and

$$\mathbf{M}(h_{\#}(\partial D \times [0, 1])) \leq \mathbf{M}(\partial D) \sup_{\text{spt } D} |f - Id| \sup_{\text{spt } D} (||Df||^{n-2}, ||Id||^{n-2}),$$

from which we deduce Proposition 23.

□

Proposition 24 *Let B' be a Borel subset of E^n and $B = B' \cap \partial\mathcal{C}$. Then*

1. $\sup_{\text{spt } D} |f - Id| \leq \max(\delta_B, \alpha_B)$, and
2. $\forall k \geq 1, \sup_{\text{spt } D} ||Df|| \leq \frac{\sup(1, ||h_B||)}{1 - \delta_B ||h_B||}$,

where $\delta_B = \delta(B, pr(B))$ is the Hausdorff distance between B and $pr(B)$ and $||h_B||$ is the maximum of the norm of the second fundamental form of M^{n-1} restricted to $pr(B)$.

Proof.

1. The first item is trivial;
2. For the second item, we remark that

$$Df = DG \circ Dpr \circ Dp_1.$$

The conclusion then follows since one has :

$$\sup_{pr(B)} \|DG\| \leq \sup_{pr(B)} (1, \|h_B\|), \|Dpr_B\| \leq \frac{1}{1 - \delta_B \|h_B\|}, \|Dp_1\| = 1.$$

□

To summarize, we have proved the following result :

Theorem 25 *Let M^{n-1} be a smooth hypersurface of $\mathbb{E}^n$ bounding a compact subset V and $\mathcal{C} \subset \mathbb{E}^n$ be a n -manifold with boundary that is geometric and the boundary of which is strongly close to M^{n-1} . Let B' be any Borel subset of $\mathbb{E}^n$ and $B = B' \cap \partial\mathcal{C}$. Then,*

$$\mathcal{F}(N(\mathcal{C})_{\perp}(B \times F) - N(K)_{\perp}(pr(B) \times F)) \leq \max(\delta_B, \alpha_B) \left(\frac{\sup_B (1, \|h_B\|)}{1 - \delta_B \|h_B\|} \right)^{n-1} (\mathbf{M}(N(\mathcal{C})_{\perp}(B \times F)) + \mathbf{M}(\partial N(\mathcal{C})_{\perp}(B \times F))),$$

where $\delta_B = \delta(B, pr(B))$, $\|h_B\|$ denotes the maximum of the norm of the second fundamental form h of M^{n-1} restricted to $pr(B)$.

3.3.3 Approximation of curvature measures

Once we have bounded the flat norm between the normal cycles of V and $\mathcal{C}$, we only need to bound the norms of the invariant forms and the ones of their exterior derivative to bound the differences between respective curvature measures, thanks to the alternate definition of the flat norm. Note however that as we only defined the forms ω_k on $ST\mathbb{E}^n$, we need first to extend them outside of $ST\mathbb{E}^n$, as the homotopy between the normal cycles described in section 3.3.2 is not supported in $ST\mathbb{E}^n$ but in $E \times B(0, 1)$ ($B(0, 1)$ denotes the unit ball of F). To do so, we consider a smooth function $\phi : \mathbb{R} \rightarrow \mathbb{R}$ supported in $[1/2, 3/2]$ and such that $\phi(1) = 1$. The extended forms ω'_k are then defined by

$$\omega'_k(m, \xi) = \phi(\|\xi\|) \omega_k(m, \xi / \|\xi\|)$$

The forms ω'_k are obviously bounded, as well as their derivatives, by some constants $C_1(n, k)$. We thus have :

Corollary 26 *Let M^{n-1} be a smooth hypersurface of $\mathbb{E}^n$ bounding a compact subset V and $\mathcal{C} \subset \mathbb{E}^n$ be a n -manifold with boundary that is geometric and the boundary of which is strongly close to M^{n-1} . Let B' be any Borel subset of $\mathbb{E}^n$ and $B = B' \cap \partial\mathcal{C}$. Then, for every $k, 0 \leq k \leq n-1$,*

$$|\phi_{\mathcal{C}}^k(B) - \phi_V^k(pr(B))| \leq C_1(n, k) \max(\delta_B, \alpha_B) \left(\left(\frac{\sup_B(1, \|h_B\|)}{1 - \delta_B \|h_B\|} \right)^{n-1} (\mathbf{M}(N(\mathcal{C})_{\perp}(B \times F)) + \mathbf{M}(\partial(N(\mathcal{C})_{\perp}(B \times F))) \right)$$

where $\delta_B = \delta(B, pr(B))$ is the Hausdorff distance between B and $pr(B)$ and $\|h_B\|$ is the maximum of the norm of the second fundamental form of M^{n-1} restricted to $pr(B)$.

We now proceed the same way with the anisotropic curvature measures, except that we do not need to extend the corresponding forms as their definition already makes sense at any point of $T\mathbb{E}^n$. Consider a frame (e_i) of $\mathbb{E}^n$. For any two indices i and j of $\mathbb{E}^n$, the coefficients of ω_{e_i, e_j} at (m, ξ) are linear functions of ξ . Thus they are bounded on $E \times B(0, 1)$, which is the domain of interest, since the homotopy between the normal cycles is supported in it. Moreover, $d\omega_{e_i, e_j}$ has constant coefficients. As a consequence, when X and Y are two unit vectors, the norms of $\omega_{X, Y}$ and its derivative on $E \times B(0, 1)$ are bounded by some constant $C_2(n)$. Hence :

Corollary 27 *Under the assumptions of corollary 26 and with the same notations :*

$$\|H_{\mathcal{C}}(B) - H_V(pr(B))\| \leq C_2(n) \max(\delta_B, \alpha_B) \left(\left(\frac{\sup_B(1, \|h_B\|)}{1 - \delta_B \|h_B\|} \right)^{n-1} (\mathbf{M}(N(\mathcal{C})_{\perp}(B \times F)) + \mathbf{M}(\partial(N(\mathcal{C})_{\perp}(B \times F))) \right),$$

here $\|\cdot\|$ denotes the operator norm.

When $\partial\mathcal{C} = N$ is C^2 and $B' = \mathbb{E}^n$, these results can be interpreted as follows. If N is closely inscribed in M and :

1. M and N are close (δ_B is small)
2. M and N have close normals (α_B is small)

3. The norm of the second fundamental form of M is not too big ($\|h_B\|$ is not too big)

4. The total curvature of N is not too big

then M and N have close curvature measures. Indeed, the mass of the normal cycle of $\mathcal{C}$ is in a certain sense a measure of the total curvature of N . For instance, if the dimension of N is 2 and k_1 and k_2 denote its principal curvatures, one has [56] :

$$\mathbf{M}(N(\mathcal{C})) = 2 \int_N (1 + k_1^2 + k_2^2 + k_1^2 k_2^2)^{1/2}$$

The requirement that the mass of the normal cycle of $\mathcal{C}$ is not too big cannot be removed, as the following example shows. Let M be a square in space, and N the surface obtained from M by adding conic spikes with height μ , slope θ , and centered on the vertices of a grid of size η . We assume that $2\mu < \theta\eta$, so that the spikes do not overlap. N is closely inscribed in M and when μ , θ , and η go to 0, δ_B and α_B also go to 0. However, one can tune the decay of these parameters in such a way that the total mean curvature H of N goes to infinity. H is simply the sum of the total mean curvatures of all the spikes of N . The total mean curvature of each spike is a function of μ and θ that is linear with respect to μ by homogeneity. Thus $H \simeq \mu/\eta^2 f(\theta)$ for some function f . Calculations show that $f(\theta) = \Omega(\theta)$ when θ goes to 0. Thus if one chooses $\theta = \eta^{1/3}$ and μ such that $2\mu \leq \theta\eta$ holds, e.g. $\mu = \eta^{4/3}/3$, then one has $H = \Omega(\eta^{-1/3})$. In this example, the total mean curvature does not converge because the mass of $N(\mathcal{C})$ is unbounded.

We have given conditions under which the curvature measures of a hypersurface are close to the ones of a given smooth hypersurface. Unfortunately, our approach breaks down as soon as both hypersurfaces are singular, in particular because the term $\|h_B\|$ then explodes. We leave to the reader the following question which, if the answer were positive, would settle the issue of approximation of curvature measure in the most general case :

Open problem 1 *Let M (resp. N) be a closed hypersurface of $\mathbb{E}^n$, and V (resp. W) the bounded component of the complement of M (resp. N). Assume V and W are geometric and let f be a homeomorphism between M and N . Let :*

$$1. \delta = \sup_{m \in M} d(m, f(m))$$

2. α be the maximum Hausdorff distance between the support of $N(V)_\perp(\{m\} \times F)$ and the one of $N(W)_\perp(\{f(m)\} \times F)$

Can one bound the difference between the curvature measures of M and N by a function ϕ of δ , α , and the masses of $N(V)$ and $N(W)$ that goes to 0 with δ and α ?

3.3.4 The case of triangulations in 3-space

We now apply the general curvature measures approximation results of the previous section to the particular case of a smooth surface $M = \partial V$ approximated by a triangulated surface $T = \partial W$. For simplicity, we will take as B the relative interior of a union of triangles of T , though for some applications other possibilities may be preferred, such as the intersection of T with a ball centered on it.

Triangulated mesh closely inscribed in a smooth surface

Following [52], we say that a triangulated mesh of $\mathbb{E}^3$ is *inscribed* in a smooth surface M if all its vertices belong to M . A triangulated mesh T is *closely inscribed* in a smooth surface M if it is both inscribed in it and strongly close to it. In what follows, we assume that T is closely inscribed in M . Here are some useful notations and definitions introduced in [92] to study the relationship between the geometry of a surface and the one of a mesh closely inscribed in it. Let t be a triangle of T .

- $\eta(t)$ denotes diameter of t , that is the length of its longest edge.
- $r(t)$ denotes the *circumradius* of a triangle t .
- The *straightness* of a triangle t is the real number

$$str(t) = \max_{p \text{ vertex of } t} |\sin(\theta_p)|,$$

where θ_p is the angle of t at p .

- The *relative height* of t with respect to M is the real number defined by:

$$\pi_M(t) = \sup_{m \in t} \eta(t) \|h_{pr(m)}\|$$

Let us now give explicit bounds on the various quantities involved in the bounds of section 3.3.3.

Lemma 28 *The Hausdorff distance δ_t between a triangle t and its projection $pr(t)$ on M is smaller than the diameter $\eta(t)$ of t .*

This is obvious since the vertices of the triangle lie on the surface.

Lemma 29 *For fixed M , we have :*

$$\alpha_t = O(r(t))$$

as $\eta(t)$ goes to 0.

In this case, α_t is nothing but the maximum angle between the normal to triangle t and the normals to M on $pr(t)$.

Proof. It is shown in [92] that if $\pi_M(t) \leq \frac{1}{2}$:

$$\sin \alpha_t \leq \left(\frac{4}{str(t)} + 2 \right) \pi_M(t)$$

As M is smooth and compact $||h||$ is bounded from below so $\pi_M(t)$ is a $O(\eta(t))$ so *a fortiori* a $O(r(t))$. Also as in any triangle we have :

$$r(t) = \frac{\text{length}(e)}{2|\sin(\theta_p)|}$$

p being any vertex of t and e being the edge opposite to p . Thus the vertex p such that $|\sin(\theta_p)|$ is the largest is also such that $\text{length}(e)$ is the largest, that is $\frac{\eta(t)}{str(t)} = 2r(t)$ and the conclusion follows. $\square$

This result is similar to the one obtained by Nina Amenta [6] in the case of restricted Delaunay triangulations.

The mass of the normal cycle of a triangulated mesh

We shall now evaluate the mass of the part $N(W)_\perp(B \times F)$ of $N(W)$ lying above B and the one of $\partial(N(W)_\perp(B \times F))$, the last two quantities that remain to bound to prove the theorem 13, which we recall for convenience.

Theorem 30 Let $W \in \mathbb{R}^3$ a volume whose boundary T is a triangulated surface closely inscribed in M . If B is the relative interior of a union of triangles of T , then :

$$\begin{aligned} |\phi_W^G(B) - \phi_V^G(pr(B))| &\leq C_M K \varepsilon \\ |\phi_W^H(B) - \phi_V^H(pr(B))| &\leq C_M K \varepsilon \\ ||\tilde{H}_W(B) - \tilde{H}_V(pr(B))|| &\leq C_M K \varepsilon \\ ||\bar{H}_W(B) - \bar{H}_V(pr(B))|| &\leq C_M K \varepsilon \end{aligned}$$

where C_M is a real number depending only on the maximum curvature of M and :

$$\begin{aligned} K &= \sum_{\{t \in T, t \subset \bar{B}\}} r(t)^2 + \sum_{\{t \in T, t \subset B, t \cap \partial B \neq \emptyset\}} r(t) \\ \varepsilon &= \max\{r(t) | t \in B\} \end{aligned}$$

In order to shorten notations we set :

$$\begin{aligned} s(B) &= \sum_{\{t \in T, t \subset \bar{B}\}} r(t)^2 \\ sd(B) &= \sum_{\{t \in T, t \subset B, t \cap \partial B \neq \emptyset\}} r(t) \end{aligned}$$

For the next two lemmas, we rely on the description of the normal cycle of a triangulated surface given in the previous chapter. We recall that the mass of an integral 2-current is the area of its support, locally weighted by the multiplicity.

Lemma 31 The mass of $N(W)_\perp(B \times F)$ is $O(s(B))$.

Proof. This mass can be decomposed in three terms : the mass lying above the interior of the triangles of T , M^t , the mass lying above the interior of the edges of T , M^e , and the mass lying above the vertices of T , M^v . M^t is merely the area of B , so it is $O(s(B))$. Let us now focus on M^e . We have :

$$M^e = \sum_{e \text{ edge of } B} |\beta(e)| \text{ length}(e)$$

Let e be an edge of T and t, t' be the triangles of T incident on e . The dihedral angle at e is $O(r(t) + r(t'))$ by 29, as well as the length of e . Thus we also have $M^e = O(s(B))$

The last quantity to consider is M^v . Let u be a vertex of T , and $u_i, i = 1..n$ its neighbors in circular order. If n_i is the unit normal to triangle uu_iu_{i+1} , then the mass lying above u is smaller than the sum of the areas of spherical triangles $n(u)n_in_{i+1}$. By lemma 29, the area of any such

triangle is $O((r(uu_iu_{i+1}) + r(uu_{i+1}u_{i+2}))^2)$. Summing on all $u \in B$, we get that $M^v = O(s(B))$. We thus proved the announced claim. $\square$

Lemma 32 *The mass of $\partial(N(W)_\perp(B \times F))$ is $O(sd(B))$.*

Proof. $\partial(N(W)_\perp(B \times F))$ decomposes into two parts : a union of line segments corresponding to the edges of ∂B , and a union of circle arcs corresponding to the edges of B meeting ∂B . The mass of the former is obviously $O(sd(B))$ and the one of the latter also by lemma 29. $\square$

From the preceding discussion together with the results of section 3.3.3, one deduces immediately theorem 13. Note that theorem 13 does not imply the convergence of curvature measures for any B , since the quantity $s(B) + sd(B)$ might go to infinity. An important case where this does not occur is when all triangles in B have all their angles bounded from below. As explained in chapter 2, the bound then boils down to the area of B plus the length of its boundary, times the maximum diameter of a triangle in B , times a constant. In particular, if B is the union of all triangles of T meeting a fixed compact set B' whose boundary is smooth and transverse to M , the area of B is bounded, as well as the length of its boundary. We thus get the convergence of respective curvature measures with a linear speed with respect to the maximum diameter of a triangle in B .

The case of restricted Delaunay triangulations

Among the triangulations of a -sufficiently dense- given point set lying on a surface, the Delaunay triangulation of these points restricted to the surface seems to be a particularly good one for the purpose of curvature estimation. First of all, the results of Nina Amenta [6] imply that these triangulations are closely inscribed provided that the point set is an ε -sample of the surface with $\varepsilon < 0.06$.

Second, one can expect that these triangulations tend to minimize the maximum circumradius of their triangles, which is involved in our bound. Indeed, this is true when the surface is a plane, since restricted Delaunay triangulations are in this case nothing but planar Delaunay triangulations, which are known to maximize [12] the granularity -the name sometimes given to the maximum circumradius of a triangulation. Of course, this does not hold for curved surfaces. Still, when the sampling density is high, the surface looks like a plane at the scale of a triangle. One can thus expect that in this case restricted Delaunay triangulations have nearly optimal granularity, though a precise argument still needs to be found. In any case, it follows

from the definitions that the granularity of the restricted Delaunay triangulation of an ε -sample is a $O(\varepsilon)$.

The quantity $s(B)$ is also optimized by Delaunay triangulations -in the plane, but it is more difficult to bound, even for restricted Delaunay triangulations. If no local uniformity assumption is made on the sampling, then one can build examples where $s(\mathbb{R}^3)$ is unbounded, such as what could be called the *Schwartz helicoid*⁶ (see figure 3.2).




Figure 3.2: A sequence of restricted Delaunay triangulations whose normal cycles have unbounded mass

The triangulation depicted in figure 3.2 is the restricted Delaunay triangulation of a particular sampling of a helicoid. This sampling is such that the density d_a of samples in the direction of the axis of the helicoid is much larger than the density d_f of samples along the fibers. If one chooses for instance $d_a = d_f^2$ and lets d_f go to infinity, then not only $s(\mathbb{R}^3)$ goes to infinity, but also do the masses of associated normal cycles. This shows that the problem in this example is not that $s(B)$ is a too weak bound.

A way to circumvent the problem is to require that the sampling is locally uniform in the sense of [58]. Indeed, in this case, it can be shown that the triangles of the restricted Delaunay triangulation have their smallest angle larger than a given constant, so that the conclusions of the previous section apply. However, the local uniformity as defined in [58] is rather restrictive. Typically, it does not allow samplings where some sample point is split into two very close samples, which can occur in practice. We thus consider a weaker notion of local uniformity introduced by D. Attali and J.D. Boissonnat in [7] :

⁶A famous similar example built with a cylinder instead of a helicoid is often called the *Schwartz lantern*

Definition 20 A set $\mathcal{P}$ is said to be a κ -light ε -sample of M if :

$$1 \leq \sharp(B(p, \varepsilon lfs(p)) \cap \mathcal{P}) \leq \kappa$$

for all point $p \in M$.

Loosely speaking, this definition allows to split each sample point into less than $\kappa + 1$ points. From now on, we assume that $\mathcal{P}$ is a κ -light ε -sample of M for some positive constants κ and ε . Let W be the Delaunay triangulation of $\mathcal{P}$ restricted to V and $T = \partial W$. We will also assume $\varepsilon < 0.1$ so that the closed ball property [43] is satisfied [5]. A triangle t in T then has a unique empty circumscribing ball centered on M , which we will denote by $B(t)$. The center and radius of $B(t)$ will be respectively referred to as $c(t)$ and $R(t)$.

Lemma 33 Let $\mathcal{P}$ be a κ -light ε -sample of $\mathcal{M}$, $p \in M$, and $l > 0$. If $2\varepsilon l < 1$ then :

$$\sharp(B(p, l\varepsilon lfs(p)) \cap \mathcal{P}) = O(\kappa l^2)$$

Proof. Because the lfs function is 1-lipschitz, its minimum on $B(p, l\varepsilon lfs(p))$ is larger than $lfs(p) - l\varepsilon lfs(p) > 0.5lfs(p)$. Now consider a maximal set of disjoint balls with radius $0.25\varepsilon lfs(p)$ and centered on $B(p, l\varepsilon lfs(p)) \cap M$. By a packing argument together with lemma 34, such a set has $O(l^2)$ elements. Doubling the radius of these balls yields a covering of $B(p, l\varepsilon lfs(p)) \cap M$. Each ball containing less than κ points by definition of the sampling condition, the assertion follows. $\square$

Lemma 34

$$area(M \cap B(t)) = \Theta(r(t)^2)$$

Proof. As the radius of curvature is bounded from below, we get that the maximum angle between the normal to M at $c(t)$ and the normal to M at any point in $M \cap B(t)$ is $O(r(t))$. Thus $M \cap B(t)$ is included in the complement of a cone with apex $c(t)$, axis along the normal to M at $c(t)$, and half-angle $\pi/2 - O(r(t))$. Consequently, the image of $M \cap B(t)$ under the orthogonal projection on the tangent plane to M at $c(t)$ contains a disk with center $c(t)$ and radius $\Omega(r(t))$. As orthogonal projections shrink areas, $area(M \cap B(t)) = \Omega(r(t)^2)$. Moreover, lemma 29 implies a lower bound on the jacobian of the projection, so that $area(M \cap B(t)) = O(r(t)^2)$. $\square$

Note that this implies that :

$$s(B) = O\left(\sum_{t \in T} \text{area}(M \cap B(t))\right)$$

since $r(t) \leq R(t)$

Lemma 35 *A vertex v of T has $O(\kappa)$ neighbors in $\mathcal{T}$.*

Proof. Let t be the triangle of T containing v such that $r(t)$ is maximal. Then the ball B with center $c(t)$ and radius $3r(t)$ contains all the neighbors of v . Applying lemma 33 yields the result. $\square$

We now show that if :

$$\bigcap_{j=1}^n M \cap B(t_j) \neq \emptyset$$

where the t_j are different triangles of T , then $n = O(\kappa^2)$. A point p in this intersection is in conflict with all the t_j . If p were inserted in T , its neighbors would thus be the vertices of the triangles t_j . As such a vertex can only belong to $O(\kappa)$ triangles, we get that the number of neighbors of p is $\Omega(n/\kappa)$. But $\mathcal{P} \cup \{p\}$ is a $(\kappa + 1)$ -light ε -sample of M , so $\kappa = \Omega(n/\kappa)$, and the conclusion follows. Since no more than $O(\kappa^2)$ sets of the form $M \cap B(t_j)$ can overlap, we have as a result :

Lemma 36

$$s(B) = O(\text{area}(M \cap \bigcup_{t \in B} B(t)))$$

Unfortunately, this quantity might be much larger than the area of B if some triangle of B sharing an edge with ∂B have a large angle at the vertex opposite to that edge (see figure 3.3).

However, if this is not the case, then it seems that $s(B)$ is a big-Oh of the area of B , and also that $sd(B)$ is a big-Oh of the length of ∂B . This would imply that the bound on the difference between the curvature measures of M and T have the same form as in the case of triangulations with bounded aspect ratio, that is $O(\varepsilon(\text{area}(B) + \text{length}(\partial B)))$.

Even if assuming that the sampling is κ -light does not imply that our bounds converge to zero for a general B , the bound on $s(B)$ obtained in lemma 36 shows that $s(\mathbb{R}^3)$ is bounded, from which one deduces :




Figure 3.3: The bound on $s(B)$ may be large

Theorem 37 Let $(\mathcal{P}_i), i \in \mathbb{N}$ be a sequence of κ -light ε_i -samples of M such that ε_i tends to zero as i tends to infinity, and W_i the Delaunay triangulation of $\mathcal{P}_i$ restricted to V . For all open set B such that $\text{area}(\partial B \cap M) = 0$, we have :

$$\begin{aligned} |\phi_{W_i}^G(B) - \phi_V^G(B)| &= o(1) \\ |\phi_{W_i}^H(B) - \phi_V^H(B)| &= o(1) \\ \|\bar{H}_{W_i}(B) - \bar{H}_M(B)\| &= o(1) \\ \|\bar{H}_{W_i}(B) - \bar{H}_M(B)\| &= o(1) \end{aligned}$$

Proof. Let us apply theorem 25 to M and W_i with $B = M$. As B has empty boundary, $sd(B) = 0$. By lemma 36, $s(B)$ is bounded by a multiple of the area of M . Moreover, δ_B and α_B tend to zero as for any restricted Delaunay triangulation. We deduce that the normal cycle of W_i converge to the one of V in the topology induced by the flat norm, and *a fortiori* in the weak topology. As mentioned in [52], this implies the weak convergence of curvature measures, which is exactly the claimed assertion. $\square$

Chapter 4

Application to anisotropic polygonal remeshing

Introduction

Despite a recent effort to make digital geometry tools robust to arbitrarily irregular meshes, most scanned surfaces need to undergo complete remeshing (alteration of the sampling and of the connectivity; see [122, 42, 73, 80, 79, 21, 3, 65]) before any further processing: results of finite element computations, compression, or editing rely heavily on an good description of the original geometry. Several techniques have been proposed over the last decade, with a wide variety of target applications. In [3], a thorough review shows that most existing methods combine mesh simplification and vertex optimization (see [72, 19] for example); others start with a complete resampling of the surface [122], mixed with connectivity optimization. However, even if this remeshing process has now been made both efficient and flexible, most techniques do not put any constraint on the local shape of the mesh elements: although vertex density is often required to depend on local curvatures, no condition is imposed on the resulting shape and orientation of the triangles or quads. Whenever we wish to align or stretch mesh elements with a certain direction field, we need *anisotropic remeshing*.

Such a specific remeshing is interesting for many reasons. While many elliptic partial differential equations ideally require meshes with quasi-equilateral triangles, elongated elements with large aspect ratio are often desired in the field of simulation, for fluid flow or anisotropic diffusion for instance. In these cases, a 2×2 matrix (referred to as a Riemannian metric tensor) traditionally indicates, for each point on the surface, the desired orientation and aspect ratio of

the mesh element locally desired [20].

Additionally, several researchers in approximation theory have proven that the same anisotropic requirement naturally arises when an optimal mesh is sought after: for a given number of elements, a mesh will “best” approximate a smooth surface (for the L^p norms with $p \geq 1$) if the anisotropy of the mesh follows (in non-hyperbolic regions) the eigenvalues and eigenvectors of the curvature tensor of the smooth surface regions [113, 34]. This can be intuitively noticed by considering a canonical example, such as an infinite cylinder: planar quads infinitely stretched along the lines of minimal curvature provide the best piecewise linear description. This similarity between applications in simulation and approximation is not surprising if we interpret both these results in terms of optimal error control. In this paper, we will explore the problem of anisotropic remeshing, and present a novel, efficient, and flexible stroke-based remeshing technique whose lines continuously follow intrinsic geometric properties across a model.

4.0.5 Previous Work

Because of the theoretical ubiquity of anisotropic meshes, algorithms for anisotropic remeshing have been proposed in several geometry-related fields.

Anisotropic Triangle Remeshing Bossen and Heckbert [20] proposed an anisotropic triangle meshing technique for flat, 2D regions on which a metric tensor is defined. They proceeded through successive vertex insertions, vertex removals, and iterative relaxations, that include edge flips to *align the edges* in accordance with the metric tensor. Shimada [112] used ellipse packing to introduce anisotropy in the remeshing; although this type of methods generates high quality anisotropic meshes whose elements conform precisely to the given tensor field, this accuracy is obtained at the price of rather slow computations, and results in very limited ways for a user to guide the design of the mesh.

Heckbert and Garland [70] made an interesting link between the quadric error metric used in their mesh simplification [59] and its asymptotic behavior on finely tessellated surfaces. In particular, they demonstrated that the triangles resulting from their mesh simplification technique will be more elongated along minimal curvature directions. Such remeshing-through-simplification methods provide fast results, but again, leave very little flexibility in the process. Moreover, the anisotropic behavior is only proven for fine meshes: the results show, however, a limited (and uncontrollable) amount of anisotropy on coarse meshes. Finally, notice that work on feature remeshing [21] has also pointed out the importance of using anisotropic triangles

in feature regions and of aligning their edges to the principal directions, although no complete anisotropic remeshing technique using these principles was proposed.

Anisotropic Quad Remeshing Several works have also focused on using *quadrangles* for remeshing, due to their appealing tensor-product nature. Borouchaki and Frey [18] described an anisotropic triangle mesh generation, and then transformed the resulting mesh into a quad-dominant mesh through a simple triangle-to-quad conversion. Shimada and Liao [111], on the other hand, proposed to directly use rectangle packing, where the rectangles are stretched according to a specified vector field on the surface. This computational intensive packing leads to a quad-dominant anisotropic mesh, aligned with the given vector field.

In Computer Graphics, there have also been recent attempts at finding anisotropic parameterizations [105, 66]. Gu *et al.* [65] showed how this could be used to provide a perfectly regular remeshing of surface meshes. However, no control over the alignment of the edges with specific directions is provided.

Lines of Curvatures and Curvature-based Strokes Even if anisotropy is a relatively recent research theme in mesh processing, this particularity of almost all shapes has long been noticed and used by artists. A caricaturist, for instance, only needs a few select strokes to convey strong geometric information. Similarly, a digital artist creates or edits a 3D model in a top-down fashion, using the main axes of symmetries and a few sparse strokes to efficiently design the mesh, contrasting drastically with the local point-sampling approach of most automatic remeshing techniques (Figure 4.1). In the scientific community, studies and previous non-photorealistic rendering techniques have also shown how much lines of curvatures are essential in describing the geometry [24, 71]: since local directions of minimum and maximum curvatures indicate respectively the slowest and steepest variation of the surface normal, these anisotropic, intrinsic quantities govern most lighting effects. In particular, many hatching techniques use strokes that are aligned along the principal curvatures: this results in a perceptually convincing display of complex surfaces [75, 76, 99, 60, 71].

4.0.6 Contributions

Although illustration and sketching techniques have been using principal curvature strokes to represent geometry, graphics techniques rarely even exploit anisotropy of a surface to drive the




Figure 4.1: *Artist-designed models (left) often conform to the anisotropy of a surface, contrasting with the conventional curvature-adapted point sampling used in most remeshing engines (right).*

remeshing process. Nevertheless, a straight edge on a coarse mesh naturally represents a zero-curvature line on the surface. It therefore seems appropriate (though non trivial!) to directly place edges parallel to the local principal directions in non-hyperbolic areas (see Figure 4.2, left), instead of first placing vertices to then slowly optimize their positions in order to align the induced edges.

In this paper, we propose a principal curvature stroke-based anisotropic remeshing method that is both efficient and flexible. Lines of minimum and maximum curvature are discretized into edges in regions with obvious anisotropy (Figure 4.2, left), while traditional point-sampling is used on isotropic regions and umbilic points where there is no favored direction (as typically done by artists; see Figure 4.2, right). This approach guarantees an efficient remeshing as it adapts to the natural anisotropy of a surface in order to reduce the number of necessary mesh elements. We also provide control over the mesh density, the adaptation to curvature, as well as over the amount of anisotropy desired in the final remeshed surface. Thus, our technique offers a unified framework to produce quad-dominant polygonal meshes ranging from isotropic to anisotropic, and from uniform to adapted sampling.

4.0.7 Overview

Figure 4.3 illustrates the main steps of our algorithm. We assume the original model to be a genus-0, non closed triangle mesh, possibly provided with tagged feature edges (non-zero




Figure 4.2: *Left: Skilled mesh designers tend to intuitively align edges with lines of minimum and maximum curvatures in anisotropic areas, as it provides a more compact representation of the local geometry. Right: Point sampling is, however, preferred in spherical areas where no particular direction is perceived.*




Figure 4.3: *From an input triangulated geometry, the curvature tensor field is estimated, then smoothed, and its umbilics are deduced (colored dots). Lines of curvatures (following the principal directions) are then traced on the surface, with a local density guided by the principal curvatures, while usual point-sampling is used near umbilic points (spherical regions). The final mesh is finally extracted by subsampling, and conforming-edge insertion. The result is an anisotropic mesh, with elongated quads aligned to the original principal directions, and triangles in isotropic regions. Such an anisotropy-based placement of the edges and cells makes for a very efficient and high-quality description of the geometry. A smooth surface can be obtained by quad/triangle subdivision of the newly generated model.*

genus input can be done on a per-chart basis). In a preliminary step, we build the *feature skeleton* [21, 3], representing all the tagged features (creases and corners) in a graph of adjacency.

The mesh is now ready to be remeshed:

- We first estimate the curvature tensor field of the surface at the vertices, and deduce the two **principal direction fields** stored as a 2D symmetric tensor field in a conformal parameter space. These fields are then smoothed, and the degenerate points (umbilics) are extracted (see Section 4.1).
- We then trace a *network of lines of curvature*, with a density guided by the local principal curvatures, in order to **sample the original geometry** appropriately along minimum and maximum curvatures, in agreement with asymptotic results from approximation theory. The isotropic regions (around the umbilic points, being either spherical or flat, are point-sampled since no obvious direction of symmetry is locally present (see Section 4.2).
- Finally, the vertices of the newly generated mesh are extracted from the *intersections* of lines of curvature on anisotropic areas, and a constrained Delaunay triangulation offers a convenient way to deduce the final edges from a subsampling of the lines of curvature (see Section 4.3). The output of our algorithm is a **quad-dominant anisotropic polygon mesh**, due to the natural orthogonality of the curvature lines.

We discuss the various computational geometry and numerical tools we used to significantly ease the implementation, as well as our results in Section 4.4.

4.1 Principal Direction Fields

Since we will base our remeshing method on lines of curvature, we first need to extract the principal curvatures. In this section, we describe how the curvature tensor field of the input surface is extracted, smoothed, and analyzed. Most of these steps are performed directly in parameter space, to speed up the computations.

4.1.1 Robust 3D Curvature Tensor Estimation

Due to the piecewise-linear nature of the input mesh, the very notion of curvature tensor, well known in Differential Geometry [63], becomes non trivial, and subject to various definitions [118, 88]. In order to have a *continuous* tensor field over the whole surface, we build a piecewise linear curvature tensor field by estimating the curvature tensor at each vertex and interpolating these values linearly across triangles. However, locally evaluating the surface curvature tensor at a vertex is not very natural. For every edge e of the mesh, on the other hand,




Figure 4.4: *Integration domain for curvature tensor estimation*

there is an obvious minimum (i.e., along the edge) and maximum (i.e., across the edge) curvature. A natural curvature tensor can therefore be defined at each point along an edge, as noticed recently in [33]. This line density of tensors can now be integrated (*averaged*, see Figure 4.4) over an arbitrary region B by summing the different contributions from B , leading to the simple expression:

$$\mathcal{T}(\mathbf{v}) = \frac{1}{|B|} \sum_{\text{edges } e} \beta(e) |e \cap B| \bar{e} \bar{e}^T \quad (4.1)$$

where $\mathbf{v}$ is an arbitrary vertex on the mesh, $|B|$ is the surface area around $\mathbf{v}$ over which the tensor is estimated, $\beta(e)$ is the signed angle between the normals to the two oriented triangles incident to edge e (positive if convex, negative if concave), $|e \cap B|$ is the length of $e \cap B$ (always between 0 and $|e|$), and $\bar{e}$ is a unit vector in the same direction as e . In our implementation, we evaluate the tensor at every vertex location $\mathbf{v}$, for a neighborhood B that approximates a geodesic disk around this vertex. This approximation is done by simply computing the disk around $\mathbf{v}$ that is within a sphere centered at $\mathbf{v}$. The sphere radius is specified by the user; a radius equal to $1/100^{\text{th}}$ of the bounding box diagonal is used by default. To remain consistent with our tensor field evaluation, the normal at each vertex can now be estimated by the eigenvector of $\mathcal{T}(\mathbf{v})$ associated with the eigenvalue of minimum magnitude. The two remaining eigenvalues $\kappa_{\min}$ and $\kappa_{\max}$ are estimates of the principal curvatures at $\mathbf{v}$. Notice that the associated directions are *switched*: the eigenvector associated with the minimum eigenvalue is the maximum curvature direction $\gamma_{\max}$, and vice versa for $\gamma_{\min}$ (see Figure 4.5). This curvature tensor evaluation procedure, in addition to being intuitive and simple to implement, has solid theoretical foundations, as well as convergence properties [33].




Figure 4.5: Principal directions γ_{min} and γ_{max} estimated at mesh vertices, scaled by their respective curvatures.

4.1.2 Flattening the Curvature Tensor Field

To allow for fast subsequent processing, we wish to 'flatten' the surface, along with its curvature tensor field. We use the discrete conformal parameterization recently presented in [84, 38] as the solution of choice for mapping the 3D surface to a 2D domain: based on a simple variational formulation, this parameterization automatically provides an angle-preserving mapping, without fixing any boundary positions, by simply solving a simple, sparse linear system. We also compute the induced area distortion as advocated in [3].

On this parameterization, we can now simply store the 2D curvature tensor (the normal component is no longer needed). For every vertex in this 2D parameterization, we thus compute the 2D curvature tensor $\mathbf{T}$ such as:

$$\mathbf{T} = \mathbf{P}^t \begin{pmatrix} \kappa_{min} & 0 \\ 0 & \kappa_{max} \end{pmatrix} \mathbf{P} \quad (4.2)$$

We do not need to compute the matrix $\mathbf{P}$ in practice. The tensor can be found simply by picking an edge from the 1-ring, projecting it onto the tangent plane, and computing the signed angle α between this projection and the eigenvector of the maximum eigenvalue: the quasi-conformality of our parameterization allows us to now find the projected eigenvector by starting

from the same edge in parameter space, and rotating it by α . The other eigenvector being orthogonal to the first one by definition, the symmetric matrix representing $\mathbf{T}$ can now be found explicitly.

Once we have $\mathbf{T}$ at each vertex, the 2D tensor field is then interpolated linearly, i.e., the matrix coefficients are linearly interpolated over each triangle (there are only three coefficients to interpolate, since the matrix is symmetric). Therefore, for any value (u, v) in the parameter space, we can return the value of the local tensor $\mathbf{T}(u, v)$.

4.1.3 Tensor Field Smoothing

Although the averaged nature of our tensor construction (Section 4.1.1) tends to remove local imperfections due to the piecewise-linear description of our input meshes, an additional pass of smoothing over the resulting 2D tensor field is often most needed. Indeed, if a coarse remeshing of the surface geometry is desired, we first have to smooth and simplify the tensor field in order to only capture the global geometry of the surface. However, if a very detailed remeshing is desired, no or little smoothing is needed.

A Gaussian filtering of the tensor (coefficient by coefficient) is performed directly in the parameter space. This is efficiently done by placing a small disk around each 2D vertex of our parameterization, with a radius inversely proportional to the local area distortion: the conformal nature of the parameterization will keep it a geodesic disk. We then convolve the field using this circular, isotropic support for the Gaussian function. Although this fast convolution is sufficient in most cases (see Figure 4.6), a more anisotropic smoothing of the three tensor coefficients can also be performed when higher geometric fidelity is required: the reader can refer to [71] or [88] for possible practical solutions. We finally get a smoothed, continuous curvature field that encodes the principal directions along with their associated curvatures as its eigenvectors and eigenvalues, respectively.

4.1.4 Tensor Field Umbilic Points

The topology of a tensor field is partially defined by its degenerate points, called *umbilic points*. Such degenerate points of a 2D symmetric tensor field are at locations (u_i, v_i) such as:

$$\mathbf{T}(u_i, v_i) = \begin{pmatrix} \lambda & 0 \\ 0 & \lambda \end{pmatrix}. \quad (4.3)$$




Figure 4.6: *Progressive smoothing of the principal direction fields. From left to the right: initial minimal curvature directions, the same region after 10 smoothing iterations, and another view of the smoothed field. Although the smoothing is computed in parameter space, the tensor field has been projected back onto the surface for illustration purposes. The color dots indicate umbilics.*

This corresponds to the regions of the mesh where the field is *isotropic*, *i.e.*, where the surface is locally spherical or flat. To find the umbilic points of our piecewise-linear tensor field, we follow Tricoche [120]: we define the deviator part $\mathbf{D}$ of our tensor field $\mathbf{T}$, obtained through:

$$\mathbf{D} = \mathbf{T} - \frac{1}{2} \text{tr}(\mathbf{T}) \mathbf{I}_2 = \begin{pmatrix} \alpha & \beta \\ \beta & -\alpha \end{pmatrix}, \quad (4.4)$$

where the special case $\alpha = \beta = 0$ corresponds to an umbilic point. Due to the linear interpolation within each triangle, only one umbilic point can exist per triangle, and it locally corresponds to either a *wedge* type, or a *trisector* type [120] as shown in Figure 4.7. All the umbilics can easily be found by going over each triangle and solving a 2×2 linear system. They are then classified using a third-order polynomial root-finding problem as described in [36]. We keep a list of all the types and 2D positions of these umbilics for further treatment. Notice finally that the smoothing of the tensor field described in the previous section drastically reduces the

number of umbilic points, as it also simplifies the topology of the extracted curvature tensor field.




Figure 4.7: *Trisector and wedge umbilic points are the only possible singularities of a piecewise-linear tensor field.*

4.1.5 Taking Care of Features

When tagged features are present on the input mesh, special care must be used during extraction, smoothing, and umbilic analysis. First, the averaged regions over which we integrate the curvature tensors must be clipped if they intersect a feature. Indeed, feature lines often represent a significant discontinuity in the geometry (as between two adjacent faces of a cube for instance), and a one-sided evaluation is therefore recommended. Second, the smoothing step must also perform the same clipping (in the 2D plane this time) during the Gaussian smoothing of a vertex $\mathbf{v}$ near a feature also to avoid “contamination” between separate regions; after the clipping is done, the contribution due to a feature vertex located within the support is set to be the average of the values of its neighbors on the same side of the feature as $\mathbf{v}$. These operations, simple to implement, are sufficient to deal correctly with features.

Once a smoothed tensor field is obtained, the next stage of our algorithm consists in resampling the original geometry stored as a 2D tensor field in parameter space, using both points and curvature-directed strokes.

4.2 Resampling

At this stage, we wish to anisotropically resample our geometry. Although a large majority of techniques perform resampling by spreading 0-elements (vertices, isotropic by nature) over the

surface, this way of proceeding does not qualify as anisotropic. However, 1-elements (edges) are, by nature, anisotropic as they represent a segment of zero curvature locally. Therefore, we propose to resample the geometry by what is known as *lines of curvatures* [63]: these lines are always along either the minimum, or the maximum curvatures. With a proper density in agreement with local curvatures, such a network of orthogonal curves will adequately discretize the object. The final edges will be found by subsampling these lines. Based on these observations, we show in this section how anisotropic areas are sampled with a set of curves aligned along principal directions, and how isotropic (*i.e.*, spherical) areas are simply discretized with points (see Figure 4.8).




Figure 4.8: *Point-based sampling vs. curve-based sampling: while most techniques spread vertices first before deducing edges and faces, we use lines of curvatures to find vertex positions, before simplifying these lines to straight edges, and then deducing faces.*

4.2.1 Curve-based Sampling for Anisotropic Areas

Our goal is to trace a network of orthogonal lines of curvature in anisotropic areas. We present the numerical approach we used to successfully tracing lines, before giving details on *where* the lines are traced on the surface.

Lines of Curvatures

By definition, a line of maximum (resp. minimum) curvature is a curve on a surface such as, at every point of the curve, the tangent vector of this curve is collinear with the principal direction of the surface that corresponds to the maximum (resp. minimum) curvature. Each line of

curvature either starts from an umbilic point and ends at another one, or has a closed orbit, or can enter and exit from the domain bounds. One can trace such a curve $\mathbf{C} : t \mapsto u(t), v(t)$ in the parameter space (u, v) of the surface (see Section 4.1.2) by integrating the following ordinary differential equation:

$$\begin{bmatrix} u'(t) \\ v'(t) \end{bmatrix} = \gamma(t), \quad (4.5)$$

where γ is an eigenvector of $\mathbf{T}(u(t), v(t))$. More precisely, γ is the eigenvector associated with the smallest (resp. largest) eigenvalue of $\mathbf{T}$ when computing a line of maximum (resp. minimum) curvature.

Numerical Integration of a Line

Equation (4.5) can be numerically solved with an embedded fourth-order Runge-Kutta integration with adaptive step [97] where the step length is weighted by the norm of the deviator (see Section 4.1.4), as recommended by Tricoche [120]. If a starting point $(u^{(0)}, v^{(0)})$ is chosen, the local tensor is directly evaluated on the parameterization and its associated eigenvector γ is computed on the fly: the integration routine provides the next point along the line of curvature. By iterating this process, we find a series of locations $(u^{(k)}, v^{(k)})$ that defines a piecewise-linear approximation of a line of curvature. Notice that once the line ends (at an umbilic point, at a feature line, at the boundary, or close to another line of curvature), we start again at (u_0, v_0) but in the opposite direction this time, to complete the line. We now turn to the problem of finding the local density required for these lines of curvature.

Local Density of Lines

Two pivotal questions at this point of the algorithm are: how many lines should be traced on the surface, and where should we trace them? A partial answer is to first compute the desired density of lines needed at any given point on the surface, or, inversely, the spacing distance between two lines. To achieve this, first consider two lines of curvature very close to each other. A cross section of the surface, normal to these two lines, will show an approximate arc of circle (the local osculating circle of the surface) with two points on it corresponding to the trace of these two lines. A linear approximation between these two points will be away from the actual osculating circle (i.e., the surface) by a small distance. If we want to *guarantee* that




Figure 4.9: *Notations*

this distance is less than ε in order to minimize the piecewise-linear reconstruction error, the distance d between the two points must be dependent on κ as follows (see Figure 4.9):

$$d(\kappa) = 2\sqrt{\varepsilon \left(\frac{2}{|\kappa|} - \varepsilon \right)}. \quad (4.6)$$

This means that for any point on a line of maximum (resp. minimum) curvature, an approximation of the optimal distance to the next line of same curvature is $d_{max} = d(\kappa_{min})$ (resp., $d_{min} = d(\kappa_{max})$). Notice that, in the limit (as element area goes to zero on a differentiable surface), Equation (4.6) leads to an aspect ratio of the rectangular elements equal to:

$$\frac{d_{max}}{d_{min}} \approx \sqrt{\frac{|\kappa_{max}|}{|\kappa_{min}|}}, \quad (4.7)$$

which coincides with the result obtained by [113] in approximation theory. The spacing between lines of curvature defined above thus provides, for fine meshes, optimal approximation of the underlying smooth surface. In our implementation, these theoretical distances are approximated quite well directly in parameter space: due to the conformal nature of the parameterization, multiplying such a distance by the local area stretching [3] will provide the distance in the parameter space.

Curve-based Sampling

Now that we know both how to trace lines of curvature and how spaced they should be, we can start the curve-based sampling per se. High-quality placement of *streamlines* have already been studied in other applications, for visualization of vector fields for instance. Different approaches, using image guidance [121], adapted seeding [77], and more recently flow-guided seeding [125], have been proposed, but always for regularly sampled fields. It is however a trivial matter to adapt them to our context: the technique we describe next is therefore a hybrid version of [77], and [125]. We will deal with the lines of minimum curvature and the lines of




Figure 4.10: *Streamline computation.*

maximum curvature *independently*. We first put all the umbilic points into a list of *potential seeds* for lines of curvatures. We then begin by tracing lines of maximum (resp. minimum) curvature originated from the umbilic point with maximum absolute curvature, as proposed in [125]. One line gets started if the umbilic point is a wedge, while three get started if it is trisector, to respect the local topology of the vector field (see Figure 4.7). If no umbilics were present, we start the line at the point with the largest $|\kappa_{min}|$ (resp. $|\kappa_{max}|$). After each integration step needed to trace the line of curvature, a pair of seeds, placed orthogonally to the current line at the ideal distance (computed locally as in Section 4.2.1), is added to the list of potential seeds [77] (see Figure 4.10).

The current line is traced until one of these cases happen:

- the line reaches another umbilic point;
- the line comes back close to its starting seed: in this case, a loop is created;
- the line crosses an edge of the feature graph or the domain boundary;
- or the line becomes too close to an existing line of maximum (resp. minimum) curvature.

The notion of closeness in the explanations above is relative to the local optimal distance d_{min} (resp., d_{max}) between lines. However, we artificially decrease the optimal distances near the umbilic points to allow for a higher-fidelity discretization. The set of potential seeds are put in a priority queue sorted by the difference between the local optimal distance at this seed and the actual distance to a streamline. The seed that best fits the local requirement is then used to start a new line, as described above. We perform this seed selection and the subsequent line tracing iteratively until a complete coverage is obtained. A final check is performed to make sure that no large areas are still uncovered. This is done by randomly sampling the parameterization space and evaluate desired distance vs. actual distances. Generally, only a handful of additional lines of curvatures get started this way.

Proximity Queries Since the algorithm described above makes heavy use of distance computations, we must handle all the proximity queries with care and efficiency. Due to the highly non-uniform distribution of samples used on the surface, a quad-tree data structure would not pay off. Instead, we opted for a conventional computational geometry tool, for which optimized implementations are readily available (such as in CGAL [48], the library we use): a *constrained Delaunay triangulation* (CDT). Indeed, a CDT allows for fast proximity queries to constraints; furthermore, exploiting the coherence of requests (as we advance along the line of curvature) through face caching results in near-linear complexity in the number of samples. We proceed as follows: we first enter each feature segment in a CDT. Then, while we trace one line of curvature, we cache each of its samples and perform the proximity queries in the current CDT, providing distances to existing lines and features. When we are done with this line, we incorporate all its constituting segments into the CDT as constraints, and start a new line.

Control Parameters The sampling process is made flexible by providing the user with three types of control. First, the parameter ε indicating the geometric accuracy of the remeshing (see Equation 4.6) is an easy way to guide the number of lines of curvature. Second, the user can also apply a transfer function F (as in [3]) to the curvatures, to tune the amount of curvature adaptation of the final mesh. Finally, the amount of isotropy vs. anisotropy is selected through a value $\rho \in [0; 1]$. We turn the optimal distance definitions from Equation 4.6 into: $d_{max} = d(\rho/2 | \kappa_{max}| + (1 - \rho/2) | \kappa_{min}|)$ and $d_{min} = d(\rho/2 | \kappa_{min}| + (1 - \rho/2) | \kappa_{max}|)$.

4.2.2 Point-based Sampling in Spherical Areas

In spherical and flat areas, the surface has no special direction of symmetry; placing edges in this case does not make sense. We therefore use a more traditional point sampling technique in these regions. Although efficient [3] or precise [4] point-sampling methods could be used, it must be noted that these regions are extremely rare: except for canonical shapes such as a plane or a sphere, the tensor smoothing we initially perform tends to reduce the spherical regions to single umbilic point, for which sampling is straightforward.

When a region has several umbilic points, we only pick a subset of them to sample the region according to desired spacing (computed using Equation (4.6) again). A score for each umbilic point is computed as a function of its desired distance and the actual distance to another selected sample or to a feature line¹. The best fit is selected, tagged as being an *isotropic*

¹This distance is computed through a proximity query to the CDT. Additionally, samples that are selected will

sample, and we iterate this process until we can no longer add samples. Notice that, occasionally, we use up all the umbilics without meeting the density requirement. This can only happen when large triangles in flat regions are present (since only one possible umbilic point was generated per triangle, a flat region may be undersampled). In these rare cases, we iteratively add more random samples in the triangles and proceed with the best-fit selection algorithm until saturation.

4.3 Meshing

The previous resampling stage has spread a series of lines of curvatures and isotropic samples over the surface. We now must deduce the final cells, edges and vertices of our remeshing process to complete our work. Principal curvatures being always orthogonal to one another, the network of lines of curvatures have created well-shaped quad regions all over the surface. We capitalize on this observation to extract a quad-dominant mesh as follows.

4.3.1 Vertex Creation

In anisotropic regions, we traced lines of curvature using polyline approximations while we used regular sample points for spherical and flat regions. The *vertices* will therefore be the intersections of curvature lines, and the isotropic samples that we spread. While the isotropic samples do not require any specific treatment, computing the line intersection has to be performed.

In order to perform these intersections quickly, as well as to prepare us for the next steps, we make use of a CDT again, in parameter space. We first enter all the features edges as constraints in a new CDT. We add all the little segments defining the lines of curvatures sequentially, as constraints as well. Finally, the isotropic samples are added as vertices in the CDT. The vertices, intersection of features or of the lines of curvatures, have *automatically* been added to the CDT since two intersecting edge constraints will generate a vertex insertion: the vertex creation phase is over.

Notice that the performance of this phase is, again, heavily affected by the order in which the constrained segments are added. We found, not surprisingly, that random insertion leads to slow performance. On the other hand, adding the segments sequentially along each line

be incorporated in the CDT in order to take them into account for future requests.

of curvature results in almost linear complexity, as the incremental CDT benefits from spatial coherence through caching. In our tests, the whole CDT process has been this way faster than any of the other algorithms dedicated to segment intersections we have tried without exploiting spatial coherence.

4.3.2 Edge Creation

The lines of curvatures must now be subsampled in order to extract the relevant edges. Although it could seem that simply joining the previously-extracted vertices would do, we must proceed with care to avoid folds on the mesh. We use a straightforward decimation process that safely removes all useless samples: going repeatedly over each vertex present in the CDT, we eliminate those which:

- are Runge-Kutta samples and have only one constraint segment attached (it will trim away all dangling curvature lines) (see Figure 4.11,A);
- have zero constrained segments attached and are not isotropic samples (vertices of this type appear during the decimation process, when a curvature line disappears totally for instance);
- have two constrained segments of same type attached (two minimum curvature line segments, two maximum curvature line segments, or two feature edges)—but only if removing these two segments and replacing them by a single constraint segment does *not* create any new intersections (see Figure 4.11,B). This last condition guarantees that our graph of region adjacencies stays planar: it will prevent *folding* in the final mesh.

This decimation is performed until we can no longer delete vertices. While this process has taken care of the anisotropic regions, we still do not have edges in isotropic regions. This is easily remedied by finally adding the CDT edges incident to the isotropic samples as constraints: it will provide a triangulation of each spherical or flat region (further edge-swaps can be performed later to reduce valence dispersion or approximation error; see [3]).

4.3.3 Polygon Creation

The last stage of our remeshing phase extracts a final polygonal mesh from the CDT by finding all regions entirely surrounded by constrained edges: these will be our polygons. This can be done efficiently by simply visiting each CDT triangle once and recursively visit its neighbors until constraint edges are reached (see Figure 4.11,C). These extracted polygons being possibly concave we perform a convex decomposition using an implementation of Greene’s dynamic




Figure 4.11: *Remeshing phase: a dome-like shape is sampled with lines of curvatures. All the curvature line segments (red/blue) and the feature edges (green) are added as constraints in a CDT in parameter space. The CDT creates a dense triangulation; a rapid vertex decimation (A,B) then suppresses most small edges, and leaves only few vertices, defining a coarse polygonal mesh. Adding constraint edges to the umbilic (center) point takes care of the near-spherical cap.*

programming algorithm [64] (also included in CGAL). We provide an additional option to bound the highest degree of the polygons to easily allow for quad/triangle mesh generation. This task is achieved through a recursive polygon partitioning algorithm that uses simple rules for conforming-edge insertion, as indicated in Figure 4.12.




Figure 4.12: *A hybrid quad/triangle mesh is generated by adding conforming edges to T-junctions in a systematic manner (this table is not exhaustive).*

4.4 Results and Discussion

Different remeshing examples for relatively simple shapes are illustrated in Figure 4.13. A dome-like shape (first row) exhibits a spherical area at the top, and anisotropic areas elsewhere. The lines of maximum curvature converge towards the umbilic point at the top, and the lines of minimum curvature are concentric, closed circles. The vertices on the boundary have been deduced from intersections between feature graph and lines of curvatures. Notice how the area nearby the umbilic point has been triangulated, while other areas have been tessellated with elongated four-sided elements. For illustration purposes, a quad/triangle subdivision algorithm [114, 81], designed to preserve the hybrid (quad/triangle) structure is applied to generate a smooth surface from the newly generated coarse mesh. Stretching the dome (second row) totally modifies the distribution of curvatures on the surface, generating rather elongated elements on highly anisotropic areas. Finally, a saddle-like shape exemplifies the various spacings happening as a function of curvatures.




Figure 4.13: *Top: A dome-like shape, its lines of curvatures, the output of our remeshing process, its limit surface after quad/triangle subdivision, with two close-ups of the cap; Bottom: A squeezed dome and a saddle shape exhibit high anisotropy.*

The model of a pig entirely remeshed with our technique is illustrated in Figure 4.14. The curvature-based sampling of our lines of curvatures produces elongated quads in anisotropic areas. The edges tend to follow the local directions of symmetry, as expected. Conforming edges have been added to the output polygonal model in order to obtain a hybrid quad/triangle model. The second row shows a close-up of the ear, along with a surface obtained by quad/triangle subdivision.

Finally, three other anisotropically remeshed models are shown in Figure 4.15. The *octa-flower* (A) is chosen to illustrate piecewise smooth anisotropic remeshing (G,H). The direction




Figure 4.14: *Remeshing a pig. Row 1, and right column: lines of minimum (blue) and maximum (red) curvature, and the anisotropic polygon mesh generated. Row 2: close-up on an ear showing the lines of curvatures, the resulting polygon mesh with conforming edges, the surface after quad/triangle subdivision (edges of the coarse model are superimposed), and the mesh after two iterations of subdivision.*

fields are estimated, then piecewise smoothed as described in Section 4.1.5 (B–F). The closeup (C) illustrates how the direction fields are not influenced by the features, or by each other across the sharp creases. Remeshing the *bunny head* with three resolutions is illustrated by Figure 4.15(I); notice the placement of the elements on the ears. The eye and the ear of the Michelangelo’s *David* model show the richness of the geometry: the lines of curvatures conform to all the details, creating a mesh adapted to the ‘anatomy’ of the original model. Note that we show the resulting polygonal mesh before insertion of conforming edges.

Timing Our current implementation allows us to process the hand model (Figure 4.3) in 0.4s for the tensor field computations, 60s for the sampling phase, and 1s for the final remeshing phase. These timings are typical of all other models, with the exception of the entire head of *Michelangelo’s David* that required 8 minutes to resample. Given that no post-optimization process is required, we regard these numbers as very reasonable.

Implementation As indicated through this paper, we have tried to systematically use numerical techniques and computational geometry tools optimized and readily available to decrease the difficulty of implementation. We strongly advise *against* an implementation “from scratch”




Figure 4.15: A-H: the octa-flower geometry illustrates the behavior of our remeshing technique for piecewise smooth surfaces. Principal direction fields are estimated and piecewise smoothed (C) (see Section 4.1.5). I: The bunny's head is remeshed with different mesh densities. J: Finally, Michelangelo's David is remeshed; close-ups on the eye and the ear show the complexity of the model, and how the lines of curvatures match the local structures. Below is another closeup, on the whole face this time, with lines of curvatures and final polygonal mesh.

of our technique: it would result in weeks of coding, with slow and brittle results. The use of numerical techniques polished over time, and of an optimized and robust computational geometry library guarantees a much easier implementation, as well as fast and robust results. For instance, the remeshing part of our technique requires only 200 lines of code when interfaced with CGAL with an appropriate filtered kernel [48], while earlier trials made for significant (ten times) larger code, and less robust and efficient results. For reference, the tensor field processing code requires 1000 lines, while the sampling process is 5000 lines. Notice also that being able to handle the David's head mesh is proof of numerical robustness: even very large area distortion due to flattening is accommodated for.

Limitations Due to the global parameterization used in this paper, the technique is limited to genus-0 patches. For closed or genus > 0 objects, this requires to go through chart construction

and surface cutting. Besides, the main bottleneck of our current approach is clearly the sampling stage. Although it is undeniably the most important stage, finding heuristics to improve it or to speed it up would be desirable. In addition, it would also be useful to develop a fast optimization phase, when higher quality bounds on the sampling density are needed. Finally, moving the remeshed vertices out of the original manifold could drastically improve the resulting error approximation, but this is not the focus of this work, and it will be explored at a later time.

4.5 Conclusions and Future Work

We have introduced a novel approach to remeshing, exploiting the natural anisotropy of most surfaces. Imitating artists' curvature strokes used in caricatures, we trace lines of curvatures onto the surface with a proper local curvature-dependent density before deducing a quad-dominant mesh, with elements naturally elongated along local minimum curvature directions. Resulting meshes are very efficient, in the sense that they capture the main geometric features with a very low number of elements. This method also offers control over the mesh quality and density. Obvious extensions include a user-guided selection of the lines of curvatures.

As future work we wish to find a way to sample and remesh directly on the manifold embedded in a three-dimensional space, without using a parameterization. Finally, exploring other resampling solutions is of interest. In particular, following the direction of minimum absolute curvature would be in complete agreement with approximation theory [34]. This approach leads to non-orthogonal edge intersections in hyperbolic regions, which is visually displeasing but optimal in terms of approximation error. We plan to investigate this alternate solution and evaluate its relevance to our community.

Part III

Meshing implicit surfaces with certified topology

Introduction

Implicit equations are a popular way to encode geometric objects [123]. Typical examples are CSG models, where objects are defined as results of boolean operations on simple geometric primitives. Given an implicit surface, associated geometric objects of interest, such as contour generators, are also defined by implicit equations. Another advantage of implicit representations is that they allow for efficient blending of surfaces, with obvious applications in CAD or metamorphosis. Finally, this type of representation is also relevant to other scientific fields, such as level sets methods or density estimation [40].

However, most graphical algorithms, and especially those implemented in hardware, cannot process implicit surfaces directly, and require that a piecewise linear approximation of the considered surface has been computed beforehand. As a consequence, polygonalization of implicit surfaces has been widely studied in the literature. Among the general classes of methods devoted to this problem, the most common one is the so-called extrinsic polygonalization method [123]. It consists in two steps : first build a tessellation of space, and then analyze the intersection of the considered surface with each cell of the tessellation to produce the approximation. The celebrated marching cube algorithm [83] belongs to this category. The goal of an implicit surface polygonizer is twofold : its output should be geometrically close to the original surface, and have the same topology. While the former is achieved by several polygonalization schemes [124], the latter has been barely addressed up to now.

Some algorithms achieve topological consistency, that is ensure that the result is indeed a manifold, by taking more or less arbitrary decisions when a topologically ambiguous configuration is encountered. This implies that their output might have a different topology from the original surface, except in very specific cases [82]. To the best of our knowledge, there is only one paper devoted to the more difficult problem of homeomorphic polygonalization [115]. The main theoretical tool used in this paper is Morse theory. The authors first find a level set of the considered function that can be easily polygonalized. This initial polygonalization is then progressively transformed into the desired one, by computing intermediate level sets. This requires in particular to perform topological changes when critical points are encountered. Unfortunately, this work is mostly heuristic, and the authors do not give any proof of the correctness of their algorithm.

In this chapter, we give the first certified algorithm for isotopic implicit surface polygonalization. Assuming the critical points of the function defining the surface are known, the whole algorithm can be implemented in the setting of interval analysis. We only assume that the considered isosurface is smooth, that is does not contain any critical point, which is generic by Sard's theorem [106]. Our polygonalization is the zero-set of the linear interpolation of the implicit function on a mesh of $\mathbb{R}^3$. We first exhibit a set of conditions on the mesh used for interpolation that ensure the topological correctness (section 1.6). Then, we describe an algorithm for building a mesh satisfying these conditions, thereby leading to a provably correct polygonalization algorithm (section 1.7).

1.6 A condition for isotopic meshing

Let f be a C^2 function from $\mathbb{R}^3$ to $\mathbb{R}$. We assume that $M = f^{-1}(0)$, the surface we want to polygonalise, is compact. In what follows, T denotes a triangulation of a domain $\Omega \subset \mathbb{R}^3$ containing M and $\hat{f}$ the function obtained by interpolating f linearly on T . A vertex v will be said *larger* (resp. *smaller*) than a vertex u if $f(v)$ is *larger* (resp. *smaller*) than $f(u)$; the sign of f at a vertex will be referred to as the sign of that vertex. We set $\hat{M} = \hat{f}^{-1}(0)$.

1.6.1 A glimpse at stratified Morse theory

Classical Morse theory

The topology of implicit surfaces is usually investigated through Morse theory [90]. Given a real function f defined on a manifold, Morse theory studies the topological changes in the sets $f^{-1}(]-\infty, a])$ (lower level-sets) when a varies. In our case, as f is defined on $\mathbb{R}^3$, this amounts to study how the topology of the part of the graph of f lying below a horizontal hyperplane changes as this hyperplane sweeps $\mathbb{R}^4$. Classical Morse theory assumes that f is of class C^2 . In this case, as is well known, these topological changes are related to the *critical points* of f , that is the points where the gradient ∇f of f vanishes. More precisely, the only topological changes occur when $f^{-1}(a)$ passes through a critical point p - a is then called a *critical value*. In the 2-dimensional case, the topology of $f^{-1}(]-\infty, a])$ can change in three possible ways, according to the type of critical point p (see figure 1.16).

In figure 1.16, the sets $f^{-1}(]-\infty, a])$ are displayed as striped regions. The leftmost column depicts the situation where p is a local maximum, that is when the Hessian of f at p is positive. In this case, $f^{-1}(]-\infty, a + \varepsilon])$ is obtained from $f^{-1}(]-\infty, a - \varepsilon])$ by gluing a topological disk along its boundary. In the case of a saddle point (i.e. the Hessian has signature $(1, 1)$), passing a critical value amounts to glue a thickened topological line segment (in gold) along its “thickened” boundary (in blue). Finally, passing through a local minimum (negative Hessian) just amounts to add a disk disconnected from $f^{-1}(]-\infty, a - \varepsilon])$. If p does not fall in any of these categories, that is if the Hessian at p is degenerate, then classical Morse theory cannot be applied. C^2 functions whose critical points all have non-degenerate Hessian are called *Morse functions*. From now on, we will assume that f is a Morse function. Also, we require that 0 is not a critical value of f , which implies that M is a manifold.




Figure 1.16: *Smooth Morse theory in 2D.*

Stratified Morse theory

As mentioned in the introduction, we chose to approximate the zero-set M of the smooth function f by the zero-set $\hat{M}$ of $\hat{f}$, which is piecewise linear. We thus need to be able to compare the topology of the level sets of $\hat{f}$ with the topology of those of f . Unfortunately, $\hat{f}$, being piecewise linear, falls out of the realm of classical Morse theory. Also, in the proof of lemma 54, we will need to apply Morse theory to a piecewise C^2 function. As a consequence, we have to resort to an extension of Morse theory developed by Goresky and MacPherson [61], called stratified Morse theory. This extension can handle a certain type of singular spaces, called *Whitney-stratified spaces*. Whitney-stratified spaces are unions of (open) smooth submanifolds of varying dimension, the strata, such that the boundary of each stratum is a union of lower dimensional strata². These spaces can be rather complicated. For our purpose, we can restrict ourselves to the case of a graph of a piecewise C^2 function g from $\mathbb{R}^3$ to $\mathbb{R}$. In this case, the 3-dimensional strata are the interior of the patches where the function is C^2 , and lower dimensional strata are lower dimensional faces of these patches. g should also satisfy some conditions³ for the theory to apply. In particular, the restriction of g to any stratum should be a Morse function. We will call such functions stratified Morse functions.

In stratified Morse theory, the critical points of a function are defined to be the critical points

²These spaces should also satisfy additional properties. For a precise definition, see [61].

³Basically, the height function restricted to the graph of g should be a Morse function in the sense of [61].

of the restriction of the function to a stratum. Note that points of 0-dimensional strata are by convention critical points. Just as in the classical case, the topology of the set $g^{-1}(]-\infty, a])$ changes only when a passes through a critical value, that is when $g^{-1}(a)$ passes through some critical point p . The difference is that the change in its topology can be much more involved than in the classical case. Still, like in the smooth case, it can be shown that the set $g^{-1}(]-\infty, a + \varepsilon])$ can always be obtained from $g^{-1}(]-\infty, a - \varepsilon])$ by gluing some set A along some subset $B \subset A$. The pair (A, B) is called the *local Morse data* of g at p . To put it more formally, if $B(p, \delta)$ denotes the ball centered on p and with radius δ , then one has :

$$A \simeq B(p, \delta) \cap g^{-1}([a - \varepsilon, a + \varepsilon])$$

and

$$B \simeq B(p, \delta) \cap g^{-1}(a - \varepsilon)$$

These definitions actually make sense, as one can show that the topology of each of the above spaces does not depend on ε and δ for $0 < \varepsilon \ll \delta \ll 1$. In the classical case, if critical point p has index λ , that is the Hessian of g at p has signature $(3 - \lambda, \lambda)$, then A is homeomorphic to the product of a λ -dimensional disk with a $(3 - \lambda)$ -dimensional one, and B is homeomorphic to the product of a $(\lambda - 1)$ -dimensional sphere with a $(n - \lambda)$ -dimensional disk (see figure 1.16).

Together with each critical point p of a Morse function g defined on a stratified space is associated an integer, called the index of g at p , and denoted by $ind(p, g)$ or simply by $ind(p)$ when no confusion is possible. The index is defined to be the increase in the Euler characteristic of $g^{-1}(]-\infty, a])$ when a goes from $g(p) - \varepsilon$ to $g(p) + \varepsilon$. If p is not a critical point, then its index is set to 0. Note that this index is different from the one classically used in the smooth setting, that is the number λ considered in the previous paragraph. When p is a critical point of a smooth function, one actually has $ind(p) = (-1)^\lambda$. From now on, by index we will mean the number $ind(p)$. Almost by definition, we get the following counterpart of Hopf's theorem in the stratified setting :

Theorem 38 *Let Y be a compact subset of $\mathbb{R}^3$ and $g : Y \rightarrow \mathbb{R}$ be a stratified Morse function. Then, χ denoting the Euler characteristic :*

$$\chi(Y) = \sum_{p \in Y} ind(p)$$

In the sequel, we will use the following consequence of this theorem :

Lemma 39 *Let f, g be two stratified Morse functions defined on $\mathbb{R}^3$ and Y be a compact subset of $\mathbb{R}^3$ such that $f|_Y$ and $g|_Y$ are stratified Morse functions. If f and g coincide in a neighborhood of ∂Y , then :*

$$\sum_{p \in Y} \text{ind}(p, f) = \sum_{p \in Y} \text{ind}(p, g)$$

Proof. We have $\sum_{p \in Y} \text{ind}(p, f|_Y) = \chi(Y) = \sum_{p \in Y} \text{ind}(p, g|_Y)$. Now the difference between $\sum_{p \in Y} \text{ind}(p, f)$ and $\sum_{p \in Y} \text{ind}(p, f|_Y)$ is the sum of $\text{ind}(p, f) - \text{ind}(p, f|_Y)$, where the sum runs over critical points of f lying on ∂Y , since both indices coincide for critical points lying in the interior of Y . As f and g coincide in a neighborhood of ∂Y , we have for each $p \in \partial Y$:

$$\text{ind}(p, f) - \text{ind}(p, f|_Y) = \text{ind}(p, g) - \text{ind}(p, g|_Y)$$

and the result follows. □

In the following, we will call the quantity $\sum_{p \in Y} \text{ind}(p, f)$ the index of f on Y . We recall that if $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ is a C^2 Morse function and $Y \subset \mathbb{R}^3$ is a 3-manifold with boundary, then ([62])

Lemma 40 *The index of f on Y is the degree of the map from ∂Y to the sphere S^2 that associates with each point $p \in \partial Y$ the normalized gradient of f at p .*

Obviously, there is no such result in the stratified setting, as the normalized gradient is not continuous any more, so its degree is not defined. However, there is a simple situation in which a result in the same spirit holds. Let $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ be a piecewise C^2 Morse function and p be a critical point of f .

Lemma 41 *Consider the set⁴ :*

$$C_\varepsilon = \text{convex hull}\{\nabla f(x) \mid x \in B(p, \varepsilon), \nabla f(x) \text{ is defined}\}$$

If for sufficiently small ε , $0 \notin C_\varepsilon$, then the lower-level set $f^{-1}(]-\infty, f(p) - \eta])$ is a strong deformation retract of $f^{-1}(]-\infty, f(p) + \eta])$ for sufficiently small η . In particular, the index of f at p is 0.

We recall that loosely speaking, a space B is a strong deformation retract⁵ of $A \supset B$ if A can be continuously collapsed to B without being torn. In particular, one has $\chi(A) = \chi(B)$. For a precise definition see any topology textbook, such as [68] or [41]. Lemma 41 is proved in [1] (proposition 1.2).

⁴The limit of the set C_ε as ε goes to 0 is known as the Clarke's subdifferential of f at p .

⁵In what follows, we write "deformation retract" for short.

PL case

We now apply stratified Morse theory to the simple case of the piecewise linear function $\hat{f}$. For piecewise linear functions, being a stratified Morse function means that no two neighboring vertices map to the same value by f , which we will assume from now on. We also assume that no vertex of T maps to 0 by f , which guarantees that $\hat{M}$ is a manifold. We refer to these two assumptions as *genericity assumptions*. Let us first recall some well-known definitions [46, 61] :

Definition 21 *The star of a vertex is the union of all simplices⁶ containing this vertex. The link of a vertex is the boundary of its star.*

Definition 22 *The lower star $St^-(v)$ of $\hat{f}$ at a vertex v is the union of all simplices incident on v all vertices of which but v are smaller than v . The lower link $Lk^-(v)$ of $\hat{f}$ at a vertex v is the union of all simplices of the link of v all vertices of which are smaller than v .*

Because $\hat{f}$ is linear on each simplex of T , its only critical points are the vertices of T . To guarantee that $\hat{M}$ is a manifold, we assume that no vertex of T maps to 0 by f (hyp. **b**). Again, this can be ensured by perturbing f slightly if necessary. We refer to hypothesis **a** and **b** as the *genericity assumptions*.

Proposition 42 *The local Morse data at a vertex v of T is homotopy equivalent to $(St^-(v), Lk^-(v))$.*

We recall that homotopy equivalence is a coarser relation than homeomorphy, allowing for instance for changes in the dimensions of the spaces involved. For precise definitions of homotopy equivalence of topological spaces and of pair of spaces, see [68] or [41].

Figure 1.17 shows the local Morse data in 2D in the case of a vertex with connected lower link (left), of a maximum (ie lower link equal to the link, middle), and of a “3-fold saddle” (lower link with 3 components, right). In the sense of stratified Morse theory, the vertex v in the left of figure 1.17 is a critical point, as any vertex. Still, no topological change in the lower level-sets occurs at such a point. This is what incited us to modify the definition of critical points in the PL case :

Definition 23 *A critical point of $\hat{f}$ is a vertex whose lower link is not contractible⁷. A vertex that is not a critical point of $\hat{f}$ will be called regular.*

⁶By simplex we mean a closed cell of T of any dimension.

⁷A topological space is contractible if it retracts by deformation to a point.




Figure 1.17: *Morse theory for PL functions in 2D. Plus and minus signs indicate whether neighbors of v are larger or smaller than v . Lower links are displayed in blue, sets $\hat{f}^{-1}(]-\infty, f(v) - \varepsilon])$ in red, and sets $\hat{f}^{-1}(]-\infty, f(v) + \varepsilon])$ in green.*

With this definition, any critical point induces a change in the homotopy type of lower level-sets. The index of a critical point v is 1 minus the Euler characteristic of $Lk^-(v)$ [8]. In figure 1.17 v respectively has index 0, 1, and -2 . In 2D the critical points are exactly the vertices with non-zero index. This is not true any more in 3D. For instance, vertices whose lower link has the topology of the disjoint union of an annulus and a disk are critical but have index 0. Still, regular points all have index 0. In 3D, a point is regular if and only if its lower link and its upper link (similarly defined) are connected, which yields an easy way detect critical points. Finally, remark that if a vertex meets the assumptions of proposition 41, then by proposition 42 its lower star retracts by deformation on its lower link, so that its lower star is contractible, i.e. the vertex is regular.

Before stating the theorem, we need one more topological notion :

1.6.2 Collapses

Loosely speaking, a collapse [100] is an operation which consists in removing cells from a simplicial complex without changing its connectivity. More precisely :

Definition 24 *If L is a simplicial complex and K a subcomplex of L , one says that there is an elementary collapse from L to K if there is a p -simplex s of L and a $(p - 1)$ -face t of s such that :*

- s is not a face of any simplex of L .
- t is not a face of any simplex of L other than s .
- $L = K \cup s$.
- $\partial s \setminus K$ is the relative interior of t .




Figure 1.18: *Elementary collapse.*

Definition 25 If L is a simplicial complex and K a subset of L , one says that L collapses to K if there is a subdivision L' of L such that K can be obtained from L' by a sequence of elementary collapses.

Definition 25 is illustrated in figure 1.19. In figure 1.19, the complexes in the middle and on the right do not collapse to the bold curve because they would need to be “torn” in order to do so. If a complex collapses to a subcomplex, then the subcomplex is a deformation retract of the complex, but the converse is not true in general. However, we have the following :

Lemma 43 A contractible subcomplex of a 2-sphere collapses to a point.




Figure 1.19: The grey complex on the left collapses to the bold curve (dashed edges represent the subdivision). This is not true for the two other complexes.

1.6.3 Main result

0. We assume that f does not vanish on any tetrahedron of T containing a critical point of f .

Theorem 44 Let W be a subcomplex of T .

If W satisfies the following conditions :

1. f does not vanish on ∂W .
2. W contains no tetrahedron of T containing a critical point of f .
- 2'. W contains no critical point of $\hat{f}$.
3. W collapses to $\hat{M}$.
4. f and $\hat{f}$ have the same index on each bounded component of $\Omega \setminus W$.

Then M and $\hat{M}$ are isotopic in W . Moreover, the Hausdorff distance between M and $\hat{M}$ is smaller than the “width” of W , that is the maximum over the components V of W of the Hausdorff distance between the subset of ∂V where f is positive and the one where f is negative.

In the conclusion of the theorem, isotopic in W means that M can be continuously deformed into $\hat{M}$ while remaining a manifold embedded in W , so that M could not be a knotted torus if $\hat{M}$ is an unknotted one, for instance. We first prove that under the conditions of the theorem, M and $\hat{M}$ are homeomorphic. Under the assumptions of the theorem, the fact that they actually are isotopic will be proved in the next section. Before proving the theorem, we first show by some examples that none of its assumptions can be removed. In the three following pictures, (local) minima of f are represented by min , (local) maxima by max , and saddle points by s . Critical points of $\hat{f}$ are represented similarly but with a caret. The sign preceding a critical point symbol indicates the sign of the considered function (f or $\hat{f}$) at the critical point.




Figure 1.20: Condition 0. is needed.

Figure 1.20 shows that condition 0. cannot be removed even in the 2D case. By allowing for critical points of f inside a triangle of T with positive vertices, one can build an example where M has an extra component w.r.t. $\hat{M}$ without violating conditions involving critical points and their indices. Indeed, in figure 1.20, f has index 0 on the triangle, since minima have index 1 and saddle points have index -1 .




Figure 1.21: *Critical points do not determine the topology of level-sets.*

Figure 1.21 is a 2D example of two zero-sets M (boundary of the gold region) and M' which are not homeomorphic, though their defining functions have the same critical points, with the same indices. The dashed curves represents a negative level-set (in green) of the function defining M' , and a positive one (in blue). Such an example can also be built such that $M' = \hat{M}$ for some mesh T . This shows the importance of the set W in the theorem. In particular, conditions **1.** and **3.** cannot be removed. Indeed, if one drops **1.**, taking for W any set satisfying **2.** and **3.** makes the theorem fail. On the other hand, if one drops **3.**, any W satisfying **2.** and **1.** also makes the theorem fail.




Figure 1.22: *Condition 2'. and 4. are needed.*

Figure 1.22 is a 3D example where M is a torus whereas $\hat{M}$ is a sphere. This is because $\hat{f}$ has an extra negative minimum inside $\hat{f}^{-1}(]-\infty, 0])$ whereas f has an index 1 saddle point outside

the bounding box Ω . Depending on whether this extra minimum lies in W or not (see the circle arc with arrows at both ends in figure 1.22), one obtains counterexamples to the theorem if assumptions 2'. or 4. are dropped. One can build similar examples showing that condition 2. is also needed.

We now return to the proof of theorem 44.

1.6.4 Proof of the homeomorphy

Lemma 45 *Let S and T be two subsets of a topological space X that meet (ie $S \cap T \neq \emptyset$).*

Assume the boundary of S is connected, as well as T and $X \setminus T$.

If the $X \setminus S$ and $X \setminus T$ meet but their boundaries do not, then S is contained in the interior of T or the other way around.

Proof. Let S and T be two such sets. ∂S is the disjoint union of $\partial S \cap \text{int}(T)$ and $\partial S \cap \text{int}(X \setminus T)$ since $\partial S \cap \partial T$ is empty. So we have a partition of ∂S in two relatively open sets. As it is connected, one has to be empty.

If $\partial S \cap \text{int}(T)$ is empty then $\partial S \subset \text{int}(X \setminus T)$ that is $T \cap \partial S$ is empty. As a consequence, T is included in $\text{int}(S)$ or in $\text{int}(X \setminus S)$ by connectedness. Since S and T meet, we have that $T \subset \text{int}(S)$. Now if $\partial S \cap \text{int}(X \setminus T)$ is empty then $X \setminus T$ is contained in $\text{int}(S)$ or in $\text{int}(X \setminus S)$ by connectedness again. Similarly as above it has to be contained in $\text{int}(X \setminus S)$, which means that $S \subset T$. Thus $\text{int}(S) \subset \text{int}(T)$ so $\partial S \supset S \setminus \text{int}(T) = S \cap \partial T$. If S would meet ∂T , then ∂S and ∂T would meet, which is impossible : S is included in the interior of T . $\square$

Lemma 46 *Let V be a connected component of W .*

$M \cap V$ is a connected smooth compact manifold without boundary.

Proof. Hypothesis 3 implies easily that V collapses to $\hat{M} \cap V$. Thus V contains a simplex having positive and negative vertices. As a consequence, f vanishes on V . Since f does not vanish on ∂W (1), M intersects V . Also, M does not meet the boundary of V (1), so $M \cap V$ is a smooth compact manifold without boundary.

Because V , which is connected, collapses to $\hat{M} \cap V$, $\hat{M} \cap V$ is a connected closed surface. Therefore, the complement of $\hat{M} \cap V$ has exactly two components, one of which is bounded. Because V collapses to $\hat{M}$, $\mathbb{R}^3 \setminus V$ also has exactly one bounded component which we denote by A and one unbounded component we denote by B . The complement of A , which is $B \cup V$,

is connected, because B and V are connected. For the same reason, $A \cup V$ is also connected. Moreover, since the complement of $A \cup V$ is B , it is also connected. In summary, A is connected as well as its complement, and the same is true for $A \cup V$.

Call now $M_i, i = 1..n$ the connected components of $M \cap V$. For each i , let N_i be the bounded component of $\mathbb{R}^3 \setminus M_i$. $M_i = \partial N_i$ does not meet $\partial(A \cup V) \subset \partial W$ (1), and $A \cup V$ is connected as is its complement. So N_i is included in $A \cup V$ thanks to lemma 45. Now N_i contains at least one critical point of f . But as $N_i \subset A \cup V$, such a point has to lie in A , by 2. So N_i meets A , but since $\partial N_i = M_i$ does not meet $\partial A \subset \bar{W}$, N_i contains A by lemma 45 again. Suppose $M \cap V$ is not connected. Then N_1 and N_2 both contain A so they intersect. Because M is smooth, their boundaries do not intersect. So one has w.l.o.g. $N_2 \subset N_1$. Now f vanishes on $\partial(N_1 \setminus N_2) = \partial N_1 \cup \partial N_2$, and therefore has an extremum in $N_1 \setminus N_2$, which is impossible because $N_1 \setminus N_2 \subset V$. $\square$




Figure 1.23: Proof of lemma 46.

So $M \cap V$ and $\hat{M} \cap V$ are connected compact surfaces without boundary. As seen in the preceding proof, A contains all critical points of f enclosed by $M \cap V$, with the same notations. Also, A contains all critical points of $\hat{f}$ enclosed by $\hat{M} \cap V$ by 2'. From condition 4., we deduce that the volumes enclosed by $M \cap V$ and by $\hat{M} \cap V$ have the same Euler characteristic, since the Euler characteristic of a lower level set is the index of the considered function on that lower level set (theorem 38). So $M \cap V$ and $\hat{M} \cap V$ have the same genus and are thus homeomorphic. To complete the proof that M and $\hat{M}$ are homeomorphic, it remains to check that :

Lemma 47 M is included in W .

Proof. Let D be some component of $\Omega \setminus W$. We claim that $M \cap D$ is empty. First $\hat{M} \cap D$ is empty by 3 so w.l.o.g vertices lying in the closure of D are all positive. If $M \cap D$ is not empty

then some component E of $f^{-1}(]-\infty, 0])$ meets D . Moreover, by condition **1**, ∂D does not meet E . Indeed, f is positive at vertices of ∂D and does not vanish on $\partial D \subset \partial W \cup \partial \Omega$. So E , being connected, is included in the interior of D . But then E is compact and thus f reaches its minimum on E : E contains a (negative) critical point of f . This is impossible since the tetrahedron containing this critical point would have negative vertices by condition **0**, though being included in D . $\square$

The proof of the bound on the Hausdorff distance between M and $\hat{M}$ is not difficult. Pick any point p in $\hat{M}$ and let V be the component of W containing it. Assume w.l.o.g. that $f(p) > 0$ and let p' be the closest point of p on the component of ∂V where f is negative. By the intermediate value theorem, the line segment pp' meets M at a point q . The distance between p and q is smaller than the distance between p and p' which is smaller than the Hausdorff distance between the two components of ∂V . This shows one half of the bound. The other half can be proved in a similar way.

1.6.5 Proof of the isotopy

Now that we know that M and $\hat{M}$ are homeomorphic, the fact that they are isotopic is a consequence of proposition 48, which is of independent interest. The background on 3-manifold topology required for the proof of this proposition can be found in [69]. In this section, all maps and manifolds considered are $\mathcal{C}^\infty$.

Proposition 48 *Let $\hat{S}$ be a orientable compact connected surface without boundary and let S be a surface such that*

- $\hat{S}$ is homeomorphic to S ,
- S is embedded in $V = \hat{S} \times [0, 1]$,
- $S \cap (\hat{S} \times \{0\}) = \emptyset$ and $S \cap (\hat{S} \times \{1\}) = \emptyset$,
- $V \setminus S$ has two connected components, one containing $\hat{S} \times \{0\}$ and the other one containing $\hat{S} \times \{1\}$.

Then S is isotopic to $\hat{S}$ in V .

To prove theorem 44, one applies proposition 48 to each component of M . More precisely, let M_i be any component of M , and let $\hat{M}_i$ (resp. W_i) be the corresponding component of $\hat{M}$ (resp. W). Now consider any regular neighborhood U of W_i (see [100]). Since W_i collapses to $\hat{M}_i$, U is also a regular neighborhood of $\hat{M}_i$ ([100] corollary 3.29). Now $\hat{M}_i$ has a regular neighborhood

N such that the pair $(\hat{M}_i, N)$ is homeomorphic to $(\hat{M}_i \times \{1/2\}, \hat{M}_i \times [0, 1])$. Indeed, such a neighborhood can be obtained as the union of a collar of $\hat{M}_i$ in $\hat{f}^{-1}(]-\infty, 0])$ and a collar of $\hat{M}_i$ in $\hat{f}^{-1}([0, +\infty[)$ (3.17 and 2.26 in [100]). Thus by the uniqueness result for regular neighborhoods (3.24 in [100]), the pair $(\hat{M}_i, U)$ is homeomorphic to $(\hat{M}_i \times \{1/2\}, \hat{M}_i \times [0, 1])$. Consider now a smooth surface $\hat{S}$ homeomorphic to $\hat{M}_i$ and let h be a homeomorphism between the pair $(\hat{M}_i, U)$ and $(\hat{S} \times \{1/2\}, \hat{S} \times [0, 1])$. Let also S be a smooth surface embedded in $\hat{S} \times [0, 1]$ isotopic to $h(M_i)$. S and $\hat{S}$ satisfy the assumptions of proposition 48, so they are isotopic, which implies that M_i and $\hat{M}_i$ are isotopic in U . As U can be chosen arbitrarily close to W_i , M_i and $\hat{M}_i$ actually are isotopic in W_i .

For technical reasons our proof of proposition 48 does not work when $\hat{S}$ is a sphere. Fortunately, isotopy always holds in this case, since there is no smooth knotted 2-sphere in $\mathbb{R}^3$ (this follows from Schoenflies theorem, see [98] p.34). From now on, we assume that $\hat{S}$ is not a sphere. The proof of proposition 48 is based upon the following theorem (see [69] p.16 for a proof).

Theorem 49 *Let $\tilde{V}$ be a connected compact irreducible Seifert-fibered manifold. Then any essential surface S in $\tilde{V}$ is isotopic to a surface which is either vertical, i.e. a union of regular fibers, or horizontal, i.e. transverse to all fibers.*

Let us explain the various terms involved in this theorem. A 3-manifold N is said *irreducible* if any 2-sphere embedded in N bounds a 3-ball embedded in N . A *Seifert-fibered manifold* is a 3-manifold that decomposes into a union of topological circles, the *fibers*, satisfying certain properties. In particular, the cartesian product of a surface S and a circle S^1 is a Seifert-fibered 3-manifold, with fibers the circles $\{x\} \times S^1$ for $x \in S$. We will not explain what a *regular* fiber is, but in the previous case, which will be ours, all fibers are regular. An orientable surface without boundary S embedded in a 3-manifold N is said *incompressible* if none of its components is homeomorphic to a sphere and if for any (topological) disk $D \subset N$ whose boundary is included in S , there is a disk $D' \subset S$ such that $\partial D = \partial D'$. Any disk D for which there is no D' is called a *compressing disk* for S (see figure 1.24 for an example of compressing disk). The notion of *essential surface* of a 3-manifold is similar to the one of incompressible surface, but more restrictive. However, when the 3-manifold has no boundary, both notions coincide.

In our setting, $\tilde{V}$ is the trivial Seifert-fibered manifold $\hat{S} \times S^1$, which we obtain by identifying the two boundary components of $V = \hat{S} \times [0, 1]$. We will still denote by S the surface corresponding to S in $\tilde{V}$. We first prove that $\tilde{V}$ and S fulfill the hypothesis of theorem 49 and

then deduce that S is isotopic to $\hat{S}$. Because we assume that $\hat{S}$ is not a 2-dimensional sphere, $\tilde{V} = \hat{S} \times S^1$ is irreducible ([69] prop 1.12 p.18). We now prove the following

Lemma 50 S is an essential surface in $\tilde{V}$.

Proof. Since $\tilde{V}$ has no boundary it is sufficient to prove that S is incompressible. Let $x \in S^1$ be the point corresponding to the endpoints of $[0, 1]$ and denote by $\hat{S}$ the section $\hat{S} \times \{x\}$ in $\tilde{V}$. Suppose S is compressible. So one can find a simple curve γ on S which does not bound a disk in S and which bounds an embedded disk D in $\tilde{V}$. Do the following surgery: cut S along γ and glue a disk homotopic to D along each of the two boundary components of $S \setminus \gamma$ (see figure 1.24). By doing so, one obtains a new surface with Euler characteristic greater than $\chi(S) = \chi(\hat{S})$. The previous surgery does not change the homology class: the new surface is homologous to S . Also, the surface S (with well chosen orientation) is homologous to $\hat{S}$, as $\hat{S}$ and S form the boundary of an open subset in $\tilde{V}$. On the other hand, it follows from Künneth formula ([G] p.198) that the homology class of $\hat{S}$ in $\tilde{V} = \hat{S} \times S^1$ is not zero. So one of the connected components S' of the new surface has a non zero homology class in $\tilde{V}$. Moreover, S' has a smaller genus than the one of S . Indeed, suppose it is not the case. As the new surface has a larger Euler characteristic than $\chi(S)$ and has at most two connected components, the only possibility is that this surface is the disjoint union of S' and a sphere. Considering the complement of the compressing disk in the sphere component shows that ∂D bounds a disk in S , which is a contradiction.




Figure 1.24: Surgery along a compressing disk.

Note that it is possible to choose D such that $D \cap \hat{S} = \emptyset$: among all the embedded disks with γ as boundary that meet $\hat{S}$ in a finite number n of simple closed curves, take as D the one such that n is minimum. Suppose that n is not zero. Among all these intersection curves, there is at least one curve α bounding a disk in $D \setminus (\hat{S} \cap D)$ (when the curves are nested, consider any innermost curve on D , see fig. 1.25 on the right). Also, the surface $\hat{S}$ is incompressible, since the injection of $\hat{S}$ in $\tilde{V}$ induces an injection between corresponding fundamental groups (see [69] p. 10). As a consequence, α bounds a disk in $\hat{S}$ and one can then make an isotopy to obtain a disk D' such that $D' \cap \hat{S} = (D \cap \hat{S}) \setminus \alpha$. This contradicts the minimality of n (see fig. 1.25).




Figure 1.25: *Decreasing the number of components of $D \cap \hat{S}$.*

The previous surgery cannot be iterated an infinite number of times, since the genus of S' decreases each time. Upon termination, one obtains a surface, called S' again, which is incompressible or the sphere S^2 , and which does not intersect the surface $\hat{S}$ because we chose compressing disks that do not meet $\hat{S}$. If S' is a 2-sphere, it does not bound a 3-ball because its homology class in $H_2(\tilde{V})$ is not zero. This implies that $\tilde{V}$ is not irreducible: contradiction. So S' is an incompressible surface. Applying theorem 49, one deduces that S' is isotopic to either a horizontal or a vertical surface.

Claim: S' is not isotopic to a vertical surface.

Proof: Suppose it is. Then there exists a surface S'' which is an union of fibers of $\tilde{V}$ and which is isotopic to S' . Choose one fiber ϕ included in S'' . Its intersection number with $\hat{S}$ is equal to

1 and has to remain constant during the isotopy. So S' contains a simple closed curve whose intersection number with $\hat{S}$ is equal to 1, namely the image of ϕ under the isotopy. But S' does not intersect $\hat{S}$: contradiction.

Hence S' is isotopic to a horizontal surface, which is a covering of $\hat{S}$ under the canonical projection of $\tilde{V}$. But this is not possible since $\text{genus}(S') < \text{genus}(S)$. So, S is incompressible, which concludes the proof of lemma 50. $\square$

Now, it follows from theorem 49 that S is isotopic to either a horizontal or a vertical surface. S does not intersect $\hat{S}$, so it cannot be isotopic to a vertical surface, by the same argument as above. So S is isotopic to a horizontal surface. This surface is a covering of $\hat{S}$ under the canonical projection of $\tilde{V}$. Because $\tilde{V} \setminus S$ is connected, it follows from [69] p.17-18 that the covering is trivial. Hence, S is isotopic to a horizontal surface which meets each fiber in one point. It is now a classical fact that this horizontal surface can be “pushed along the fibers” to construct an isotopy to $\hat{S}$ (see Fig. 1.6.5). Note that, using the same argument as the one used previously to prove that one can construct S' such that it does not intersect $\hat{S}$, the isotopy $f_t, t \in [0, 1]$ between $\hat{S}$ and S can be chosen so that $f_t(\hat{S}), t \in]0, 1]$ never intersects $\hat{S}$. So S is isotopic to $\hat{S}$ in V .




Figure 1.26: Pushing S to $\hat{S}$ along the fibers of $\tilde{V}$.

1.7 Algorithm

In the algorithm, we take as W a set that is related to the notion of watershed from topography. This set satisfies properties 2. and 3. by construction. In section 1.7.1, we give its definition,

basic properties, and construction algorithms. Section 1.7.2 describes the meshing algorithm itself, which ensures that V fulfills also conditions **0.**, **1.**, **2'.**, and **4.**, and proves its correctness.

1.7.1 PL watersheds

We first assume that the mesh T conforms to $\hat{M}$, i.e. $\hat{M}$ is contained in a union of triangles of T . We will see later how to remove this assumption, which is in contradiction with the genericity assumptions. Define W^+ as the result of the following procedure :

Positive Watershed Algorithm

```

set  $W^+ = \hat{M}$ .
mark all vertices of  $\hat{M}$ .
while there is a positive regular unmarked vertex  $v$  of  $T$  s.t. the vertices of  $Lk^-(v)$  are marked
do
    set  $W^+ = W^+ \cup St^-(v)$ .
    mark  $v$ .
end while
return  $W^+$ 

```

W^- is defined as the result of the same algorithm applied to $-f$. We set $W = W^+ \cup W^-$. Note that W contains no critical point of $\hat{f}$. Also, positive marked vertices are exactly the vertices of W^+ .




Figure 1.27: Construction of W^+ : lower stars of regular vertices (such as v_1) are added one by one. Lower stars of critical vertices (v_2) are discarded.

Lemma 51 *W collapses to $\hat{M}$.*

Proof. It is sufficient to show the result for W^+ . Let W_i^+ be the state of W^+ after i steps of the algorithm, and let v_i be the i -th marked vertex. As $W_0^+ = \hat{M}$, the only thing we have to show is that W_{i+1}^+ collapses to W_i^+ for all i . Let us first show that $Lk^-(v_i)$ is included in W_i^+ . If it is not the case, let u be the largest vertex of some simplex s of $Lk^-(v_i)$ that is not in W_i^+ . s is in $St^-(u)$ which is hence not included in W_i^+ . So u is either critical or not marked yet, which is a contradiction since v_i is marked. So $Lk^-(v_i) \subset W_i^+$. Now since v_i is regular, $Lk^-(v_i)$ collapses to a point (lemma 43). Consider a sequence of elementary collapses allowing to collapse $Lk^-(v_i)$ to p and let $s_j \subset Lk^-(v_i)$, $j = 1..n$ be the sequence of simplices defining these elementary collapses. The simplices $conv(s_j \cup v_i)$ and the edge pv_i define a valid sequence of elementary collapses allowing to collapse $W_{i+1}^+ = W_i^+ \cup St^-(v_i)$ to W_i^+ , which concludes the proof. $\square$

One may prefer a more intrinsic definition of W^+ . In the same spirit as in [47], one can define a partial order on the vertices of T by the closure of the acyclic relation $\prec$ defined by $u \prec v$ if $u \in St^-(v)$ or $u = v$. We will note this order $\prec$ again and say that v flows into u whenever $u \prec v$. The next lemma shows that the vertices of W^+ do not depend on the order the vertices are considered in the construction.

Lemma 52 *The vertices of W^+ are exactly the positive vertices that do not flow into any positive critical point.*

Proof. The vertices of W^+ have this property by construction. Let $p \notin W^+$ be a positive vertex and assume p does not flow into any positive critical point. In particular, p is regular. Hence, as $p \notin W^+$, the lower link of p , which is not empty, has to contain either a critical vertex or an unmarked one. It cannot contain a critical point because as T conforms to $\hat{M}$, vertices in $Lk^-(p)$ are all non-negative, and so p would flow into a positive critical point. There is thus an unmarked vertex p_1 in $Lk^-(p)$. If p_1 can be chosen positive, then p_1 satisfies the same assumptions as p so one can define p_2 in a similar way. By going on, one obtains a strictly decreasing sequence of positive vertices, that thus has to end. Let p_k its last term. $Lk^-(p_k)$ contains no positive unmarked vertices. But as T conforms to $\hat{M}$, vertices in $Lk^-(p_k)$ are all non-negative. As vertices of $\hat{M}$ are marked, we get a contradiction. $\square$

Note that W is the union of simplices with all their vertices in W . As a result, we get an intrinsic definition of W , and not only of its vertices. From an algorithmic point of view, it may

be efficient to examine the vertices in increasing order in the construction of W^+ . One can for instance maintain the ordered list of vertices neighboring W , always consider the first element of this list for marking, and discard it if it cannot be marked. Indeed, with this strategy, a vertex that cannot be marked at some point will never be marked.

Another consequence of lemma 52, which will be useful later, goes as follows. Call c the minimum of $|\hat{f}(v)| = |f(v)|$ over all critical points v of $\hat{f}$.

Lemma 53 *W contains all vertices whose image by $|f|$ is smaller than c .*

Proof. Let p be such that $|f(p)| < c$. Without loss of generality, assume that p is positive. Any critical point v in which p flows satisfies $f(v) < f(p)$. So it cannot be positive by definition of c : by lemma 52, p lies in W^+ . $\square$

Non conforming case. We now drop the assumption that T conforms to $\hat{M}$ and assume genericity again. From T and $\hat{M}$ one can build a mesh S that is finer than T , conforms to $\hat{M}$, and has all its extra vertices on $\hat{M}$. Indeed, it suffices to triangulate the overlay of $\hat{M}$ and T without adding extra vertices except those of $\hat{M} \cap T$. This can be done as the cells of the overlay are convex. The construction of W described above can then be applied to S . A positive vertex of T has its lower link in S containing only vertices of $\hat{M}$ if and only if its lower link in T contains only negative vertices. Thus, in order to find the say positive vertices of $W \cap T$, one can apply the positive watershed algorithm described above to T , if at the initialization step one marks all negative vertices having a positive neighbor instead of those of $\hat{M}$. Still, note that if a negative critical point has a positive neighbor, then this neighbor will not be marked by this modified algorithm, whereas it could have been marked by the standard algorithm applied to S . However, if we assume that vertices having a neighbor of opposite sign are regular (condition **a**), then this does not happen and the result W' of the modified algorithm is equal to W . In our meshing algorithm, we will not build the mesh S , but rather make sure condition **a** holds, and apply the modified algorithm.

Updating W' . The intrinsic definition of W —or W' — given above yields an efficient way of updating W when T undergoes local transformations. It is sufficient to describe the algorithm for updating the vertices of W^+ . Let T_1 be a mesh obtained from T by removing some set of tetrahedra E and remeshing the void left by E . Call A the set of positive critical points of the linear interpolation of f on T_1 that lie in E . Then the vertex set of the positive watershed W_1^+ associated with T_1 can be computed from the vertex set of W^+ by performing the following two operations. To begin with, the set of vertices of T_1 that flow into A must be

removed from W^+ (lemma 52), which amounts to a graph traversal. Remaining vertices all belong to W_1^+ . Then, mark these vertices and apply the positive watershed algorithm loop to get the vertex set of W_1^+ .

Remark. The presented definition of a watershed seems quite well-behaved and leads to an easy construction algorithm, but it is not fully satisfactory. In particular, the watershed we compute is in general strictly included in the 'true watershed'. The 'true watershed' seems hard to compute, though, and can intersect a triangle in a very complicated way. There might be interesting intermediate definitions between ours and the true one, for instance based on the PL analog of the Morse complex introduced in [46].

1.7.2 Main algorithm

Assume the critical points of f are given. Theorem 44 enables us to build a mesh isotopic to M using only one simple predicate, *vanish*. *vanish* takes a triangle or a box and returns true if f vanishes on that triangle or that box. We actually not even need a predicate, but rather a filter. More precisely, *vanish* may return true even if f does not vanish on the considered element, but not the other way around. Still, we require that *vanish* returns the correct answer if the input triangle or box is sufficiently small. Such filters can be designed using interval analysis.

Our algorithm also requires to build a refinable triangulation of space such that $\hat{f}$ (resp. $\nabla \hat{f}$) converges to f (resp. ∇f) when the size of elements tends to 0. As noticed by Shewchuk [110], this is guaranteed provided all tetrahedra have dihedral and planar angles bounded away from π . In [11], Bern, Eppstein and Gilbert described an octree-based algorithm yielding meshes whose angles are bounded away from 0. In our case, which is much easier, the desired triangulation can simply be obtained by adding a vertex at the center of each square and each cube of the octree, triangulating the squares radially from their center, and doing the same with the cubes. Indeed, resulting planar and dihedral angles are all bounded away from 180° . One can expect that this scheme does not produce too many elements upon refinement, because the size of elements is allowed to change rapidly as we do not require that these have a bounded aspect ratio (see figure 1.28). The main algorithm uses an octree O , the associated triangulation T , the watershed W' . We will say that two (closed) boxes of O are neighbors if they intersect. O is initialized to a bounding box Ω of M . Such a bounding box can be found by computing the critical points of the coordinate functions restricted to M , if possible, or using interval analysis. Besides, we maintain four sets of boxes ordered by decreasing size. *Critical* contains all boxes containing a critical point of $\hat{f}$ that is not in a box containing a critical point of f .




Figure 1.28: Octree and triangulation used in the algorithm. In this 2D example, only the edges of the triangulation of the box on the right are shown (dashed).

Index contains all boxes neighboring a box b containing a critical point of f and such that f and $\hat{f}$ have different indices on b . *Boundary1* contains all boxes containing two neighbors -in T - of opposite signs one of which is critical for $\hat{f}$ (condition **a**, see paragraph **Non conforming case**). Finally, *Boundary2* contains all boxes that contain a triangle t of $\partial W'$ such that $vanish(t)$ is true and that are not included in W' .

Main Algorithm

Initialization Refine O until $vanish(b)$ is false for all boxes containing at least one critical point of f .

compute T and W' , and the four sets *Critical*, *Boundary1*, *Boundary2*, and *Index*.

while (true) do

 update T , W' , and the four sets.

if $Critical \neq \emptyset$ **then**

 split its first element.

else if $Boundary1 \neq \emptyset$ **then**

 split its first element.

else if $Boundary2 \neq \emptyset$ **then**

 split its first element.

else if f and $\hat{f}$ have different indices on some bounded component of $\Omega \setminus W'$ **then**

 split the first element of *Index*.

else

return $\hat{M}$

end if

end while

Thanks to theorem 44 applied to W' , the correctness of this algorithm almost amounts to its termination. The only problem is that W' might contain some critical point of f , thereby violating condition 2.. It thus seems that the definition of W' needs to be slightly modified. The modification consists in taking as W'^+ vertices -and the same for W'^- - the positive vertices that do not flow into positive critical points of $\hat{f}$ nor into vertices lying in a box containing a positive critical point of f . With this modification, lemma 51 still holds and lemma 53 holds if one replaces c by the minimum c' of c and the minimum of $|f|$ on the boxes containing a critical point of f . Also, c' is positive as f does not vanish on these boxes.

We now show that the main algorithm terminates. First note that after the initialization step, no box containing a critical point of f is split. The magnitude of ∇f is thus larger than a certain constant g_{min} on the complement C of the union of these boxes. Let us show that the size of the boxes of *Critical* that are split at some point is bounded from below. As $\nabla \hat{f}$ converges to ∇f , there is a number s_1 such that for each tetrahedron with diameter smaller than s_1 , $\|\nabla f - \nabla \hat{f}\|$ is smaller than $g_{min}/2$ on the interior of that tetrahedron. If the tetrahedron is included in C , this implies that $\nabla \hat{f}$ and ∇f make an angle smaller than $\pi/6$.

Lemma 54 *Let $A \subset \mathbb{R}^3$ be such that ∂A is a manifold included in C and containing no vertex of T . Suppose that all boxes meeting ∂A are smaller than s_1 . Then f and $\hat{f}$ have the same index on A .*

Proof. Let $p \in \partial A$ and $d(p)$ denote the local feature size of p with respect to the 2-skeleton of T , as defined -in 2D- by Ruppert [101]. Simplices of T that meet the open ball centered at p of radius $d(p)$ all share a vertex $v(p)$ -by definition, $d(p)$ is the largest number such that this holds. We call d_{min} the minimum of d , which is known to be positive, and set k equal to the minimum of d_{min} and e , the half of the distance from ∂A to the closest box that does not meet ∂A .

Let us now consider a smooth nonnegative function $\phi : \mathbb{R}^3 \rightarrow \mathbb{R}$ with support included in the open ball centered at 0 of radius k . The convolution of $\hat{f}$ and ϕ is a smooth function $\tilde{f}$. Let p be a point at distance less than e from ∂A . The gradient of $\tilde{f}$ at p is a weighted average of the gradients of $\hat{f}$ at points lying in the open ball centered at p and with radius k . All gradients involved in this average are gradients of $\hat{f}$ on tetrahedra incident on $v(p)$. Moreover, the size of these tetrahedra is smaller than s_1 because $k \leq e$. As a consequence, considered gradients all




Figure 1.29: *Proof of lemma 54.*

make an angle smaller than $\pi/6$ with the gradient of f at $v(p)$. As the weights in the average are nonnegative, we have that the angle between $\nabla \tilde{f}(p)$ and $\nabla f(v(p))$ is smaller than $\pi/6$. Also, the angle between $\nabla f(v(p))$ and $\nabla f(p)$ is less than $\pi/3$ since both vectors make an angle smaller than $\pi/6$ with the gradient of $\hat{f}$ on some tetrahedron containing p and $v(p)$. Finally, we get that $\nabla \tilde{f}(p)$ and $\nabla f(p)$ make a positive dot product.

Let now U_1 be a neighborhood of ∂A whose closure does not contain any vertex of T and U_2 be an open set such that $U_1 \cup U_2 = \mathbb{R}^3$. We also require that the Hausdorff distance between U_1 and ∂A is smaller than e and that $U_2 \cap \partial A = \emptyset$. Denote by $\{u_1, u_2\}$ a partition of unity subordinate to the covering $\{U_1, U_2\}$. This means that for $i = 1..2$, u_i is a non negative smooth function defined on $\mathbb{R}^3$, with support in U_i , and such that $u_1 + u_2$ is identically 1. In particular, u_2 equals 1 on the complement of U_1 , and vice versa. So the function $g = u_2 \hat{f} + u_1 \tilde{f}$ coincide with $\hat{f}$ on $\mathbb{R}^3 \setminus U_1$ and with $\tilde{f}$ on $\mathbb{R}^3 \setminus U_2 \supset \partial A$. Now recall that $\nabla \tilde{f}$ and ∇f make a positive dot product on ∂A . Hence the linear homotopy between both vector fields does not vanish on ∂A : by normalization, one gets a homotopy between $\nabla \tilde{f}/\|\nabla \tilde{f}\|$ and $\nabla f/\|\nabla f\|$, considered as maps from ∂A to the unit sphere. Because the degree is invariant under homotopy [68], we deduce that these maps have the same degree, which shows that f and $\tilde{f}$ have the same index on A . Now as g and $\tilde{f}$ coincide in a neighborhood of ∂A , f and g have the same index on A . To complete the proof, it thus suffices to show that g and $\hat{f}$ also have the same index on A . Now the critical points of $\hat{f}$ are critical for g , with the same index, as U_1 contains no such point. Potential other critical points of g can only lie in U_1 . But the gradient of g at any point p of U_1 where it is defined is a convex combination of $\nabla \tilde{f}(p)$ and $\nabla \hat{f}(p)$: it thus makes a positive dot product with $\nabla f(p)$. As a consequence, 0 is not in the convex hull of the image of a small neighborhood of p by ∇g , which implies that g has index 0 at p (lemma 41). We thus proved the announced claim. $\square$

Suppose that some box b of *Critical* of size smaller than s_1 is split. Let v be a critical point of $\hat{f}$ included in b . All the boxes containing v are in *Critical* and their size is smaller than s_1 since we consider boxes in decreasing order. Now the gradients of $\hat{f}$ on tetrahedra incident on v all make a positive dot product with ∇f which is a contradiction with lemma 41 which implies that v is not critical. So the conclusion is that *Critical* becomes -at least temporarily- empty after a finite number of consecutive splittings of boxes in *Critical*.

Now if the algorithm splits a box b in *Boundary1*, then b contains a say positive critical point of $\hat{f}$, which belongs to a box containing a critical point of f as *Critical* is empty. So the maximum of $|f|$ on b is larger than the minimum of $|f|$ on the boxes containing a critical point of f (i.e. c'). On the other hand, f vanishes on b since b contains a negative vertex. This cannot happen if the size of b is below a certain value, so that boxes in *Boundary1* cannot be split eternally.

Suppose that the algorithm splits arbitrarily small boxes in *Boundary2*. If a small enough box b is split, then b contains a triangle t of W' on which f vanishes. So, if the size of b is small enough, the maximum of $|f|$ on b will be smaller than c' . By lemma 53, all vertices of b belong to W' so $b \subset W'$ which is a contradiction. Thus the size of split boxes in *Boundary2* is also bounded from below.

To complete the proof of termination, we need to prove that *Index* does not contain too small boxes. This is true by applying lemma 54 to small offsets of the boxes containing critical points of f . Finally :

Theorem 55 *The main algorithm returns an isotopic piecewise linear approximation of M .*

Furthermore, if one wishes to guarantee that the Hausdorff distance between M and its approximation is less than say ε , it suffices to modify the positive watershed algorithm so as to control that the width of W is smaller than ε , thanks to theorem 44.

1.8 Conclusion

We have given an algorithm that approximates regular level sets of a given function with piecewise linear manifolds having the same topology. Moreover, our algorithm can be modified so as to ensure geometric closeness in the Hausdorff sense. Though no implementation has been carried out yet, we believe that it should be rather efficient due to the simplicity of the involved predicates and the relative coarseness of the required space decomposition.

The main drawback of our algorithm is that it requires, as is, the knowledge of the critical points of the considered function. A closer look shows that we almost only need to find a set of boxes containing all the critical points, and on which the function does not vanish. This task, corresponding to the initialization step in the main algorithm, can be done in a certified way using interval analysis. Once this is done, the only remaining problem is to compute the index of the function on these boxes in a robust way. In a forthcoming version of this work, we will show how this can be done in the framework of interval analysis, thereby giving a complete solution to the problem. Also, we plan to adapt the algorithm to the case of surfaces with boundaries, which is useful for instance when one wants to study the considered level set inside a user-specified bounding box.

Conclusion

Surface discretization is a broad area of research, raising a variety of problems of seemingly different natures. A global treatment of this topic, analogous to sampling theory for signals, still seems out of reach. In this thesis, we have studied three particular problems related to surface discretization : the construction of conforming Delaunay triangulations, the estimation of the curvature tensor of a smooth surface from an approximating mesh, and the polygonalization of implicit surfaces with guaranteed topology. Our work on these topics leaves many gaps wide open.

We have designed a certified algorithm building conforming Delaunay triangulations having a reasonable number of vertices, which is an improvement on previous methods. Still, obviously, the meshes produced by our algorithm are not competitive for the main application of 3D meshing, namely finite elements simulations. In particular, we do not control the grading of elements, nor their shape and orientation, which crucially determine the accuracy of simulations. A challenging problem could be to adapt the ideas introduced in part II chapter 4 for anisotropic surface remeshing to the problem of 3D mesh generation.

In part II, we have shown how the notion of curvature tensor could be generalized to non-necessarily smooth surfaces, such as polyhedra. This generalization consists in a tensor-valued measure, which we call anisotropic curvature measure. As most important result, we have proved that these anisotropic curvature measures behave nicely under approximations. In particular, a high resolution mesh of a smooth surface will have anisotropic curvature measures close to the ones of the smooth surface. In part III, we have given an algorithm yielding topologically correct polygonal approximations of implicitly defined surfaces. This algorithm outputs the zero-set of the linear interpolation of the function considered on a well-chosen mesh. Its correctness relies on a result ensuring the existence of an isotopy between the zero-set of a

smooth function and the one of its linear interpolation on a mesh, under certain conditions on the mesh.

The Parts II and III, though focusing on different problems, thus belong to the same line of research : given an “object” defined both in the smooth and the piecewise linear setting, such as anisotropic curvature measures for surfaces or level-sets for real functions, what discretization conditions should one require in order to guarantee that the smooth “object” is “close” to its PL analog? From a practical point of view, this line of research is motivated by the increasing need to process discretized objects. On a more fundamental side, we believe it can also provide a deeper understanding of each of the geometric problems considered. As a possible perspective, many other geometric objects could be studied along these lines : conformal parameterizations, geodesics, Morse complexes, or Reeb graphs for example.

The discretization condition used in part II to ensure the accuracy of curvature estimation requires not only that the smooth surface and its PL approximation are close in the Hausdorff sense, but also that the normals to both objects are close. Similarly, the proof of correctness of our implicit surface polygonalization algorithm requires that the considered implicit function and its PL approximation have close -or rather not too different- gradients. As a consequence, our results are limited to the approximation of smooth objects. In particular, they do not handle noisy data. Still, for instance, there should exist a notion of curvature behaving continuously under corruption by noise, at least up to a certain noise intensity. Indeed, everybody agrees that the curvature of a lens can be defined in a very accurate way, whereas at the atomic scale, the surface of a lens is highly irregular. This example suggests that a definition of curvature robust against noise might be found in a multi-resolution framework. A satisfactory multi-resolution framework for discrete geometric objects would be of great interest in many respects. Some attempts were made to adapt wavelet theory to surfaces, but many classical properties of wavelets are then lost. Another possible track would be to adapt scale-space theory to triangulated surfaces. Finally, a very attractive track, and also more intrinsically geometric, could be to use ideas of Edelsbrunner et al. [45] on persistence of critical points. It has already been suggested that persistence could be used for level-set denoising. We believe this concept could prove helpful in other situations as well.

Bibliography

- [1] A.A. Agrachev, D. Pallaschke, and S. Scholtes, *On Morse Theory for Piecewise Smooth Functions*, Journal of Dynamical and Control Systems 3, pp. 449-469, 1997.
- [2] P. Alliez, D. Cohen-Steiner, M. Desbrun, O. Devillers and B. Lévy, *Anisotropic Polygonal Remeshing*, to appear in SIGGRAPH 2003.
- [3] ALLIEZ, P., MEYER, M., AND DESBRUN, M. 2002. Interactive Geometry Remeshing. *ACM Transactions on Graphics* 21(3), 347–354. ACM SIGGRAPH conference proceedings.
- [4] ALLIEZ, P., COLIN DE VERDIÈRE, É., DEVILLERS, O., AND ISENBURG, M. 2002. Isotropic Surface Remeshing. In *Shape Modeling International Conference Proceedings*. To appear.
- [5] Nina Amenta and Marshall Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.*, 22(4):481-504, 1999.
- [6] Nina Amenta, Sunghee Choi, Tamal Dey and Naveen Leekha. A simple algorithm for homeomorphic surface reconstruction. *Intl. Journal on Computational Geometry & Applications*, vol. 12, (2002), 125–141.
- [7] D. Attali, J.D. Boissonnat and A. Lieutier. *Complexity of the Delaunay Triangulation of Points on Surfaces: The Smooth Case*, In Proc. 19th Annu. ACM Sympos. Comput. Geom., pp 237-246, 2003.
- [8] Th. Banchoff, *Critical points and curvature for embedded polyhedra*, J. Diff. Geom 1 (1967) 245 – 256.
- [9] M. Berger, B. Gostiaux, *Géométrie différentielle : variétés, courbes et surfaces*, Presses Universitaires de France.

- [10] J. Berkmann and T. Caelli *Computation of surface geometry and segmentation using covariance techniques* IEEE Trans. Pattern Analysis and Machine Intelligence 16. 11, 1114-1116 (1994).
- [11] M. Bern, D. Eppstein, and J. Gilbert, *Provably Good Mesh Generation*, Journal of Computer and System Sciences 48 (1994) 384-409.
- [12] M. Bern, D. Eppstein, *Mesh Generation And Optimal Triangulation* Computing in Euclidean Geometry, Edited by Ding-Zhu Du and Frank Hwang, World Scientific, Lecture Notes Series on Computing – Vol. 1 (1992)
- [13] A. Bernig, *Distributional invariants and curvature tensors of singular spaces*, in preparation.
- [14] A. Bernig and L. Bröcker, *Courbures intrinsèques dans les catégories analytico-géométriques*, Annales de l’Institut Fourier 53 (2003), 1897-1924.
- [15] J.D. Boissonnat and F. Cazals. *Smooth Surface Reconstruction via Natural Neighbour Interpolation of Distance Functions*, Comp. Geometry Theory and Applications, pp 185–203, 2002.
- [16] J.D. Boissonnat and S. Oudot, *Provably good surface sampling and approximation*, Proceedings of SGP 2003.
- [17] J.D. Boissonnat and M. Yvinec, *Algorithmic Geometry*, Cambridge University Press, 1998.
- [18] BOROUCHAKI, H., AND FREY, P. 1998. Adaptive Triangular-Quadrilateral Mesh Generation. *Intl. J. Numer. Methods Eng.* 41, 915–934.
- [19] BOROUCHAKI, H. 1998. Geometric Surface Mesh. In *Int. Conf. on Integrated and Manufacturing in Mechanical Engineering*, 343–350.
- [20] BOSSEN, F., AND HECKBERT, P. 1996. A Pliant Method for Anisotropic Mesh Generation. In *5th Intl. Meshing Roundtable*, 63–76.
- [21] BOTSCH, M., AND KOBBELT, L. 2001. Resampling Feature and Blend Regions in Polygonal Meshes for Surface Anti-Aliasing. In *Eurographics proceedings*, 402–410.

- [22] L. Bröcker and A. Bernig, *Lipschitz-Killing invariants*, Preprint, 2002.
- [23] L. Bröcker and M. Kuppe, *Integral geometry of tame sets*, Preprint, 2000.
- [24] BRADY, M., PONCE, J., YUILLE, A., AND ASADA, H. 1985. Describing Surfaces. *Journal of Computer Vision, Graphics, and Image Processing* 32, 1–28.
- [25] H. Cartan, *Cours de calcul différentiel*, Hermann.
- [26] F. Cazals and M. Pouget *Estimating Differential Quantities using Polynomial fitting of Osculating Jets* in Proc. Symposium on Geometry Processing, Aachen (2003).
- [27] J. Cheeger, W. Müller, R. Schrader, *Kinematic and tube formulas for piecewise linear spaces*, Indiana Univ. Math. J. 35 (1986) 737-754.
- [28] J. Cheeger, W. Müller, R. Schrader, *On the curvature of piecewise flat spaces*, Comm. Math. Phys. 92 (1984) 405 – 454.
- [29] X. Chen and F. Schmitt *Intrinsic surface properties from surface triangulation* in Proc. of European Conference on Computer Vision. 739-743 (1992).
- [30] B.Y. Chen, *Geometry of submanifolds*, Dekker, 1973, New-York.
- [31] L.P. Chew, *Guaranteed-quality triangular meshes*, Technical Report TR-89-983, Department of Computer Science, Cornell University, Ithaca, NY, 1989.
- [32] D. Cohen-Steiner and J.M. Morvan, *Approximation of the second fundamental form of a hypersurface of a Riemannian manifold*, INRIA Research report 4868.
- [33] COHEN-STEINER, D., AND MORVAN, J.-M., 2003. Restricted Delaunay Triangulations and Normal Cycle. In *ACM Symp. on Computational Geometry*.
- [34] D’AZEVEDO, E. F. 2000. Are Bilinear Quadrilaterals Better Than Linear Triangles? *SIAM Journal on Scientific Computing* 22(1), 198–217.
- [35] L.H. de Figueiredo, J. de Miranda Gomez, D. Terzopoulos, and L. Velho, *Physically-base methods for polygonalization of implicit surfaces*. In Graphics Interface ’92, pp. 250-257, May 1992.

- [36] DELMARCELLE, T., AND HESSELINK, L. 1994. The Topology of Symmetric, Second-Order Tensor Fields. In *IEEE Visualization Proceedings*, 140–145.
- [37] M. Desbrun, M. Meyer, P. Schröder and A. Barr *Discrete differential-geometry operators in nD* , preprint, the Caltech Multi-Res Modeling Group.
- [38] DESBRUN, M., MEYER, M., AND ALLIEZ, P. 2002. Intrinsic Parameterizations of Surface Meshes. In *Proceedings of Eurographics*, 209–218.
- [39] T.K. Dey, *Curve and surface reconstruction*, Chapter in Handbook of Discrete and Computational Geometry, Goodman and O’ Rourke eds., CRC press, 2nd edition.
- [40] D.P. Dobkin, Silvio V. F. Levy, William P. Thurston and Allan R. Wilks, *Contour Tracing by Piecewise Linear Approximations*, ACM Transactions on Graphics, Vol.9, No.4, October 1990, Pages 389-423.
- [41] A. Dold, *Lectures on algebraic topology*, Springer 1972.
- [42] ECK, M., DEROSE, T., DUCHAMP, T., HOPPE, H., LOUNSBERY, M., AND STUETZLE, W. 1995. Multiresolution Analysis of Arbitrary Meshes. In *ACM SIGGRAPH Conference Proceedings*, 173–182.
- [43] H. Edelsbrunner and N. R. Shah. Triangulating topological spaces. *Int. J. on Comp. Geom.*, 7:365-378, 1997.
- [44] H. Edelsbrunner and T. S. Tan, *An upper bound for conforming Delaunay triangulations*, Discrete Comput. Geom., vol 10, 2, pp 197-213, 1993.
- [45] H. Edelsbrunner, D. Letscher and Afra Zomorodian, *Topological Persistence and Simplification*, in Proc. 41st Annu. IEEE Sympos. Found. Comput. Sci. 2000, pp. 454-463.
- [46] Herbert Edelsbrunner, John Harer, and Afra Zomorodian. Hierarchical Morse complexes for piecewise linear 2-manifolds. In *Proc. 17th Annu. ACM Sympos. Comput. Geom.*, pages 70-79, 2001.
- [47] H. Edelsbrunner, *Surface reconstruction by wrapping finite point sets in space*, Ricky Pollack and Eli Goodman Festschrift, ed. B. Aronov, S. Basu, J. Pach and M. Sharir, Springer-Verlag, to appear.

- [48] FABRI, A., GIEZEMAN, G.-J., KETTNER, L., SCHIRRA, S., AND SCHÖNHERR, S. 2000. On the Design of CGAL, a Computational Geometry Algorithms Library. *Softw. – Pract. Exp.* 30, 11, 1167–1202. www.cgal.org.
- [49] H. Federer, *Curvature measure theory*, Trans. Amer. Math. Soc **93** (1959) 418 – 491.
- [50] H. Federer, *Geometric Measure Theory*, Springer-Verlag, New York, 1983.
- [51] R. Forman, *A User’s Guide to Discrete Morse Theory*, Seminaire Lotharingien de Combinatoire, 48 (2002).
- [52] J. Fu, *Convergence of curvatures in secant approximations*, J.Differential Geometry 37 (1993) 177 – 190.
- [53] J. Fu, *Curvature measures and generalized Morse theory*, J. Diff. Geom. 30 (1989) 619-642.
- [54] J. Fu, *Curvature measures of subanalytic sets*, Amer. J. Math, 116, (819 – 880).
- [55] J. Fu, *Curvature of Singular Spaces via the Normal Cycle*, Amer. Math. Soc. **116** (1994) 819 – 880.
- [56] J. Fu, *Intrinsic diameter and curvature integrals of compact surfaces immersed in $\mathbb{R}^3$* , University of Georgia Mathematics preprint series, preprint No.11 Volume 6(1998).
- [57] J. Fu, *Monge-Ampère functions 1*, Indiana Univ. Math. J. 38 (1989), 745-771.
- [58] S. Funke and E.A.Ramos *Smooth-Surface Reconstruction in Near-Linear Time*, to appear in SODA 2002.
- [59] GARLAND, M., AND HECKBERT, P. 1998. Simplifying Surfaces with Color and Texture using Quadric Error Metrics. In *IEEE Visualization Conference Proceedings*, 263–269.
- [60] GIRSHICK, A., INTERRANTE, V., HAKER, S., AND LEMOINE, T. 2000. Line Direction Matters: an Argument for the use of Principal Directions in 3D Line Drawings. In *International Symposium on Non Photorealistic Animation and Rendering*.
- [61] M. Goresky and R. MacPherson, *Stratified Morse Theory*, Springer Verlag, 1988.

- [62] D. Gottlieb and G. Samaranayake, *The index of discontinuous vector fields*, New-York Journal of Mathematics, 1, pp 130-148, 1995.
- [63] GRAY, A., Ed. 1998. *Modern Differential Geometry of Curves and Surfaces*. Second edition. CRC Press.
- [G] M.J. Greenberg, *Lectures on algebraic topology*, W.A. Benjamin Inc New York, 1967.
- [64] GREENE, D. H. 1983. The Decomposition of Polygons into Convex Parts. In *Computational Geometry*, F. P. Preparata, Ed., vol. 1 of *Adv. Comput. Res.* JAI Press, Greenwich, Conn., 235–259.
- [65] GU, X., GORTLER, S., AND HOPPE, H. 2002. Geometry Images. In *ACM SIGGRAPH Conference Proceedings*, 355–361.
- [66] GUSKOV, I. 2002. An Anisotropic Mesh Parameterization Scheme. In *Proceedings of 11th International Meshing Roundtable*, 325–332.
- [67] H. Hadwiger, *Vorlesungen über Inhalt, Oberfläche und Isoperimetrie*, Springer, Berlin, 1957.
- [68] A. Hatcher, *Algebraic topology*, <http://www.math.cornell.edu/hatcher>.
- [69] A. Hatcher, *Notes on Basic 3-Manifold Topology*, <http://www.math.cornell.edu/hatcher>.
- [70] HECKBERT, P., AND GARLAND, M. 1999. Optimal Triangulation and Quadric-Based Surface Simplification. *Journal of Computational Geometry: Theory and Applications* 14(1-3) (nov), 49–65.
- [71] HERTZMANN, A., AND ZORIN, D. 2000. Illustrating Smooth Surfaces. In *ACM SIGGRAPH Conference Proceedings*, 517–526.
- [72] HOPPE, H., DEROSE, T., DUCHAMP, T., McDONALD, J., AND STUETZLE, W. 1993. Mesh Optimization. In *ACM SIGGRAPH Conference Proceedings*, 19–26.
- [73] HOPPE, H. 1996. Progressive Meshes. In *ACM SIGGRAPH Conference Proceedings*, 99–108.

- [74] D. Hug and Rolf Schneider, *Kinematic and Crofton formulae of integral geometry : recent variants and extensions* in Homenatge al professor Santalo i Sors (C. Barcelo i, ed.), Universitat de Girona, 2002, pp. 51-80.
- [75] INTERRANTE, V., FUCHS, H., AND PIZER, S. 1996. Illustrating Transparent Surfaces with Curvature-directed Strokes. In *IEEE Visualization*.
- [76] INTERRANTE, V. 1997. Illustrating Surface Shape in Volume Data via Principal Direction-Driven 3D Line Integral Convolution. In *ACM SIGGRAPH Conference Proceedings*, 109–116.
- [77] JOBARD, B., AND LEFER, W. 1997. Creating Evenly-Spaced Streamlines of Arbitrary Density. In *Proceedings of the Eurographics Workshop on Visualization in Scientific Computing*, 45–55.
- [78] S. Karbacher and G. Hausler *A new approach for shape modeling and smoothing of scattered data* in Three-Dimensional Image Capture and Applications, R. Ellson and J. Nurre, Eds. SPIE Proceedings, vol.3313. The international Society for Optical Engineering, 168-177 (1998).
- [79] KOBBELT, L., VORSATZ, J., LABSIK, U., AND SEIDEL, H.-P. 1999. A Shrink Wrapping Approach to Remeshing Polygonal Surfaces. *Computer Graphics Forum, Eurographics '99 issue 18*, 119–130.
- [80] LEE, A. W. F., SWELDENS, W., SCHRÖDER, P., COWSAR, L., AND DOBKIN, D. 1998. MAPS: Multiresolution Adaptive Parameterization of Surfaces. In *ACM SIGGRAPH Conference Proceedings*, 95–104.
- [81] LEVIN, A. 2003. Polynomial generation of quasi-interpolation bin stationary non-uniform subdivision. *Computer Aided Geometric Design* 10(1), 41–60.
- [82] Adriano Lopez and Ken Brodlie, *Improving the Robustness and Accuracy of the Marching Cubes Algorithm for Isosurfacing*, IEEE Transactions on Visualization and Computer Graphics, Vol.9, No.1, January-March 2003.
- [83] W.E. Lorensen and H.E. Cline, *Marching Cubes : A high resolution 3D surface construction algorithm*, Computer Graphics, 21(4):163-169, 1987.

- [84] LÉVY, B., PETITJEAN, S., RAY, N., AND MAILLOT, J. 2002. Least Squares Conformal Maps for Automatic Texture Atlas Generation. In *ACM SIGGRAPH Conference Proceedings*, 362–371.
- [85] R. Martin *Estimation of principal curvatures from range data* Int. J. of Shape Modeling 4, 3,4, 99-109 (1998).
- [86] A. McIvor and R. Valkenburg *A comparison of local surface geometry estimation methods*, Machine Vision and Applications 10. 1, 17-26 (1997).
- [87] D. Meek and D. Walton, *On surface normal and Gaussian curvature approximation given data sampled from a smooth surface*, Computer-Aided Geometric Design 17, 521-543.
- [88] MEYER, M., DESBRUN, M., SCHRÖDER, P., AND BARR, A. H., 2002. Discrete Differential-Geometry Operators for Triangulated 2-Manifolds. *Proceedings of VisMath*.
- [89] J. Milnor *Euler characteristic and finitely additive Steiner measures* Collected Papers, Vol 1, Ed Publish or Perish (1994)
- [90] J. Milnor, *Morse theory*, Ann. of Math. studies 51, Princeton University Press, Princeton N.J. 1963.
- [91] F. Morgan, *Geometric measure theory*, Acad. Press, INC. 1987.
- [92] J.M. Morvan, B. Thibert, *On the approximation of the normal vector field of a smooth surface*, RR 4476 INRIA Sophia Antipolis.
- [93] J.M. Morvan, *On generalized curvatures*, in preparation.
- [94] M. Murphy, D. M. Mount and C. W. Gable, *A Point-Placement Strategy for Conform-
ing Delaunay Tetrahedralization*, International Journal of Computational Geometry and Applications, vol 11, 6, pp 669-682, 2001.
- [95] Y. Ohtake, A. Belyaev, M. Alexa, G. Turk and H.P. Seidel, *Multi-level Partition of Unity
Implicits*, Proceedings of SIGGRAPH 2003.
- [96] S. Petitjean, *A survey of methods for recovering quadrics in triangle meshes*, ACM Computing Surveys, number 34, volume 2, pages 211-262, 2002.

- [97] PRESS, W., FLANNERY, B., TEUKOLSKY, S., AND VETTERLING, W. 1994. *Numerical recipes in C – The art of scientific programming*, 2nd ed. Cambridge University Press, UK.
- [98] D. Rolfsen, *Knots and Links*, Math. Lecture Series 7, Publish or Perish, Inc. 1990.
- [99] ROSSL, C., AND KOBBELT, L. 2000. Line-art Rendering of 3D Models. In *Proceedings of Pacific Graphics*.
- [100] ROURKE AND SANDERSON Introduction to Piecewise-Linear Topology. Springer-Verlag, 1982.
- [101] J. Ruppert, *A Delaunay refinement algorithm for quality 2-dimensional mesh generation*, Journal of algorithms, vol 18, pp 548-585, 1995.
- [102] J. Ruppert, *A Delaunay refinement algorithm for quality 2-dimensional mesh generation*, Journal of Algorithms, 18:548–585, 1995.
- [103] A. Saalfeld, *Delaunay edge refinements*, Proceedings of the 3rd Canadian Conference on Computational Geometry, pp 33-36, 1991.
- [104] P. Sander and S. Zucker *Inferring surface trace and differential structure from 3-D images* IEEE Trans. Pattern Analysis and Machine Intelligence 12. 9, 833-854 (1990).
- [105] SANDER, P., GORTLER, S., SNYDER, J., AND HOPPE, H. 2002. Signal-specialized parametrization. In *Eurographics Workshop on Rendering 2002*.
- [106] A. Sard, *The measure of the critical values of differentiable maps*, Bull. Amer. Math. Soc. 48 (1942), 883-890.
- [107] J.R. Shewchuk, *Mesh Generation for Domains with Small Angles* , Proceedings of the Sixteenth Annual Symposium on Computational Geometry, pp 1-10, 2000.
- [108] J.R. Shewchuk, *Sweep Algorithms for Constructing Higher-Dimensional Constrained Delaunay Triangulations*, Proceedings of the Sixteenth Annual Symposium on Computational Geometry, pp 350-359, 2000.
- [109] J.R. Shewchuk, *Tetrahedral Mesh Generation by Delaunay Refinement*, Proceedings of the 14th ACM Symposium on Computational Geometry, pp 86-95, 1998.

- [110] J.R. Shewchuk, *What Is a Good Linear Finite Element? Interpolation, Conditioning, Anisotropy, and Quality Measures*, in preparation.
- [111] SHIMADA, K., AND LIAO, J. 1998. Quadrilateral Meshing with Directionality Control through the Packing of Square Cells. In *7th Intl. Meshing Roundtable*, 61–76.
- [112] SHIMADA, K. 1996. Anisotropic Triangular Meshing of Parametric Surfaces via Close Packing of Ellipsoidal Bubbles. In *6th Intl. Meshing Roundtable*, 63–74.
- [113] SIMPSON, R. B. 1994. Anisotropic Mesh Transformations and Optimal Error Control. *Appl. Num. Math.* 14(1-3), 183–198.
- [114] STAM, J., AND LOOP, C., 2002. Quad/triangle subdivision. Preprint.
- [115] B.T. Stander and John C. Hart, *Guaranteeing the Topology of an Implicit Surface Polygonizer for Interactive Modeling*, Proceedings of SIGGRAPH 97, pages 279-286.
- [116] J. Steiner, *Jber Preuss. Akad. Wiss.* 114-118,(1840). In *Gesammelte Werke*, vol 2, New York, Chelsea 1971.
- [117] T.-S. Tan, *An Optimal Bound for High-Quality Conforming Triangulations* , Discrete Comput. Geom., vol 15, 2, pp 169-193, 1996.
- [118] TAUBIN, G. 1995. Estimating the Tensor of Curvature of a Surface from a Polyhedral Approximation. In *Proceedings of Fifth International Conference on Computer Vision*, 902–907.
- [119] G. Taubin, *Estimating the Tensor of Curvature of a Surface from a Polyhedral Approximation*, Fifth International Conference on Computer Vision (ICCV’95).
- [120] TRICOCHÉ, X. 2002. *Vector and Tensor Field Topology Simplification, Tracking, and Visualization*. PhD thesis, Universität Kaiserslautern.
- [121] TURK, G., AND BANKS, D. 1996. Image-Guided Streamline Placement. In *ACM SIGGRAPH Conference Proceedings*, 453–460.
- [122] TURK, G. 1992. Re-Tiling Polygonal Surfaces. In *ACM SIGGRAPH Conference Proceedings*, 55–64.

- [123] Luiz Velho, Jonas Gomes, Luiz Henrique de Figueiredo, *Implicit Objects in Computer Graphics*, Springer-Verlag, 2002.
- [124] L. Velho. Simple and efficient polygonization of implicit surfaces. *Journal of Graphics Tools*, 1(2):5-24, 1996. ISSN 1086-7651.
- [125] VERMA, V., KAO, D. T., AND PANG, A. 2000. A Flow-guided Streamline Seeding Strategy. In *IEEE Visualization*, 163–170.
- [126] P. Wintgen, *Normal cycle and integral curvature for polyhedra in Riemannian manifolds*, *Differential Geometry* (Gy. Soos and J. Szenthe, eds.), North-Holland, Amsterdam, 1982.
- [127] T. Yoshimi and F. Tomita *Robust estimation for range image segmentation and reconstruction* in Proc. of IAPR Workshop on Machine Vision Applications, Kawasaki, Japan. 506-509 (1994).
- [128] M. Zähle, *Curvature measures and Random sets*, 1,2, *Math. Nachr.* 119, (1984), 327-339. 557 – 567.
- [129] M. Zähle, *Curvatures and currents for union of sets with positive reach*, *Geometriae Dedicata* 23 (1987) 155-171.
- [130] M. Zähle, *Integral and current representations of Federer’s curvature measures*, *Arch. Math. (Basel)* 46, (1986), 557-567.
- [131] M. Zähle, *Polyhedron theorems for non smooth cell complexes*, *Math.nachr.* 131 (1987), 299-310.

QUELQUES PROBLÈMES LIÉS À LA DISCRÉTISATION DES SURFACES

Un nombre croissant d'applications nécessite d'opérer des traitements algorithmiques sur des objets tridimensionnels. Le plus souvent, ceux-ci sont représentés par des surfaces triangulées. Cette thèse aborde trois problèmes posés par la manipulation de ces surfaces. On donne d'abord un algorithme qui, étant donnée une surface triangulée, construit une triangulation de Delaunay volumique la contenant comme sous-complexe. De telles triangulations sont utiles par exemple pour le calcul scientifique. Puis, on définit une généralisation de la courbure s'appliquant à des surfaces non nécessairement lisses, donc en particulier aux surfaces triangulées, et on étudie sa stabilité. Celle-ci est ensuite utilisée dans un algorithme de remaillage de surfaces triangulées visant à optimiser le rapport complexité/distorsion. Enfin, on donne un algorithme de maillage de surfaces implicites garantissant que l'approximation produite a la même topologie que la surface initiale.

Mots-clefs : surfaces triangulées, géométrie différentielle, topologie différentielle, approximation.

TOPICS IN SURFACE DISCRETIZATION

A rapidly growing number of applications requires to deal with three-dimensional objects on a computer. These objects are usually represented by triangulated surfaces. This thesis addresses three problems one encounters when dealing with such surfaces. We first give an algorithm which builds a volumic Delaunay triangulation containing a given triangulated surface as a sub-complex. Such triangulations are useful for numerical simulations for instance. Then, we introduce a generalization of curvature which applies to non-necessarily smooth objects, thus in particular to triangulated surfaces, and we study its stability. This generalization is then used to design an algorithm for remeshing triangulated surfaces while aiming to reach an optimal complexity/distortion ratio. Finally, we give an algorithm for meshing implicit surfaces which guarantees that the output has the same topology as the input surface.

Keywords : Triangulated Surfaces, Differential Geometry, Differential Topology, Approximation.

Thèse préparée à l'INRIA dans le projet GÉOMÉTRICA