A coordinate system on a surface: definition, properties and applications

Jean-Daniel Boissonnat, Julia Flötotto

HAL Id: inria-00072058

https://inria.hal.science/inria-00072058

Submitted on 23 May 2006

INRIA

A coordinate system on a surface: definition, properties and applications

Jean-Daniel Boissonnat — Julia Flötotto

N° 4530

Août 2002

THÈME 2

Rapport de recherche

A coordinate system on a surface: definition, properties and applications

Jean-Daniel Boissonnat , Julia Flötotto [*]

Thème 2 — Génie logiciel
et calcul symbolique
Projet Prisme

Abstract: Coordinate systems associated to a finite set of sample points have been extensively studied, especially in the context of interpolation of multivariate scattered data. Notably, Sibson proposed the so-called natural neighbor coordinates that are defined from the Voronoi diagram of the sample points. A drawback of those coordinate systems is that their definition domain is restricted to the convex hull of the sample points. This makes them difficult to use when the sample points belong to a surface. To overcome this difficulty, we propose a new system of coordinates. Given a closed surface S, i.e. a $(d-1)$-manifold of $\mathbb{R}^d$, the coordinate system is defined everywhere on the surface, is continuous, and is local even if the sampling density is finite. Moreover, it is inherently $(d-1)$-dimensional while the previous systems are d-dimensional. No assumption is made about the ordering, the connectivity or topology of the sample points nor of the surface. We illustrate our results with an application to interpolation over a surface.

Key-words: computational geometry, sampled surfaces, scattered data interpolation

[*] Email: FirstName.LastName@sophia.inria.fr

Système de coordonnées sur une surface: définitions, propriétés et applications.

Résumé : Nous définissons un système de coordonnées pour une surface, c'est-à-dire une variété $(d-1)$-dimensionnelle de $\mathbb{R}^d$, associé à un échantillon de points de la surface. Ce problème est bien étudié si le domaine de définition du système de coordonnées est restreint à l'enveloppe convexe des points de l'échantillon. Notamment, Sibson a proposé les coordonnées naturelles. Boissonnat et Cazals ont montré que les coordonnées naturelles ont les bonnes propriétés si la densité de l'échantillonnage de la surface tend vers l'infini. Dans ce travail, on propose un système de coordonnées qui est défini partout sur la surface, il est continu et il est local même si la densité de l'échantillonnage est finie. En plus, il est $(d-1)$-dimensionnel et pas d-dimensionnel comme les systèmes habituels. Nous appliquons ce système de coordonnées à l'interpolation d'une fonction définie sur une surface.

Mots-clés : gomtrie.algorithmique

1 INTRODUCTION

Surfaces represented by a set of unordered sample points are encountered in many application areas such as computer graphics, computer aided design (CAD) and reverse engineering, image processing, and scientific computation. Many algorithms that are applied to sampled surfaces rely on the definition of a local neighborhood on the surface. Reconstructing the surface from the sample points is one way to respond to this demand. However, it might be unnecessary and also too costly to establish a global approximation of the surface. Differently, our method defines, for any point on a sampled surface, a set of coordinates associated with some neighboring sample points. If the sampling is sufficiently dense, this coordinate system is provably local on the surface and has good continuity properties. Moreover, it can be computed efficiently because locality allows efficient filtering methods. We do not impose any restriction on the genus of the surface, the number of connected components, nor any other global features of the surface. Uniform sampling is neither required, and we allow the sampling density to be related to the local curvature of the surface.

1.1 Related Work

In this section, we describe how our work is related to previous work. It is divided in two parts: first, we outline the work on natural neighbor coordinate systems and, second, we give an introduction to scattered data interpolation on a surface, which is the application we develop at the end of the paper.

1.1.1 Natural neighbor coordinate systems

Natural neighbor interpolation has been introduced by Sibson [24] to interpolate multivariate scattered data. Given a set of points $\mathcal{A} = \{A_1, \ldots, A_n\}$, the natural neighbor coordinate system associated to $\mathcal{A}$ is defined from the Voronoi diagram of $\mathcal{A}$. Various papers ([24], [15], [22], [10],[19]) show that it satisfies the following definition by Brown [10].

Definition 1.1 ([10]) *A system of coordinates over $\mathcal{U} \subseteq \mathbb{R}^d$ associated with $\mathcal{A}$ is a set of continuous functions $s_i : \mathcal{U} \to \mathbb{R}, i = 1..n$, such that for all $X \in \mathcal{U}$,*

(i) $X = \sum_{i=1}^{n} s_i(X) A_i$ *(local coordinate property).*

(ii) For any $i \leq n, s_i(A_j) = \delta_{ij}$, where δ_{ij} is the Kronecker symbol.

(iii) $\sum_{i=1}^{n} s_i(X) = 1.$

The major drawback to applying the natural neighbor coordinate system to points issued from a surface comes from the fact that its definition is limited to the convex hull of the sample points. To avoid this problem, a common solution consists of adding a box enclosing the object. Obviously, this solution causes problems, e.g. the choice of the size of the bounding box, the number of sample points taken from it, artifacts arising from the bounding box points, and the augmented computation cost. In [10], Brown has enlarged the coordinate definition outside the convex hull to cover the union of the Delaunay balls, which is still too much restrictive in many applications.

A second drawback is that a point is likely to have neighbors that are far away from the point. Boissonnat and Cazals have shown that the sum of the coordinates associated to those neighbors that are far away tends to zero when the sampling density increases [6]. However, even though the influence of the far neighbors is small, the fact that Sibson's coordinates are not local affects not only the beauty and rigor of the result, but the time required to compute the coordinates and the exactness of an interpolation scheme.

For points issued from a sphere, Brown proposes a solution in [9]. While the above discussed coordinates are defined with respect to the d-dimensional Voronoi diagram of $\mathcal{A}$, Brown defines natural neighbor coordinates with respect to the geodesic Voronoi diagram on the sphere. This definition generalizes Sibson's coordinates in a straightforward manner. Therefore, the basic properties of definition 1.1 are fulfilled, except the local coordinate property *(i)* which cannot be fulfilled since points on the sphere do not belong to the convex hull of their neighbors. The obvious difficulty in enlarging Brown's approach to general surfaces is that geodesic Voronoi diagrams are much more complicated than Euclidean diagrams and difficult to compute [20]. Moreover, in many applications the surface is not known and neither is the geodesic Voronoi diagram.

In this paper, we suggest another system of coordinates for points on a surface. It is closely related to natural neighbor coordinates, yet instead of considering the geodesic Voronoi diagram on the surface, as Brown, or the Euclidean d-dimensional Voronoi diagram of the sample set, as Sibson, it is defined in the intersection of the tangent plane of each surface point with the Euclidean Voronoi diagram of the sample set. If the tangent planes are not given as part of the input, they can be easily estimated from the sample points. The resulting coordinate system is local and it is inherently $(d-1)$-dimensional.

1.1.2   Scattered data interpolation on a surface

In the last part of the paper, we apply the new coordinate system to interpolate a function defined on a surface. More exactly, we want to approximate $\Phi : \mathcal{S} \to \mathbb{R}$ where $\mathcal{S} \subset \mathbb{R}^d$ is a smooth surface, knowing a sample set $\{(A_i, z_i) : A_i \in \mathcal{S}, z_i = \Phi(A_i)\}$ and a query point $X \in \mathcal{S}$ at which we want to interpolate Φ.

This problem, which is also called 'scattered data fitting' or 'surface on surface' problem, arises in a variety of settings. For example in geodesy, geophysics, and meteorology, $\mathcal{S}$ is some model of the earth, and the function to interpolate from a number of discrete measurements represents temperature, rainfall, pressure, etc. In other contexts, $\mathcal{S}$ might be a complicated surface, e.g., the surface of some mechanical piece in CAD, a molecular surface or the wing of an airplane [5]. Several methods exist to solve this problem. One of the most popular is to enlarge the definition of splines to treat the case of a non-planar parameter domain. This was first done for the spherical case in [3]. With this achievement, it suffices to partition a general surface $\mathcal{S}$ into a collection of non-overlapping surface patches, i.e. geodesic triangles, and to define a globally smooth interpolation function as piecewise polynomials on the patches that are carefully accorded at the boundaries. See [21] for an introduction to splines on surfaces.

Other methods are radial basis functions, variational methods or multi-resolution methods. See [16] for a survey of the principal methods for scattered data fitting on the sphere. Foley et al propose

in [18] to map $\mathcal{S}$ to the sphere, to apply a known interpolation method on the sphere, e.g. [17], and to apply the inverse mapping back from the sphere to the surface to get the solution.

In the context of interpolation, our method has several advantages in addition to those mentioned at the beginning of this section. It applies directly to the point samples without need of a prior subdivision or triangulation of the surface. The quality of the result can be expressed in terms of the curvature of the surface and of the sampling density. If the surface is locally planar, the interpolant has linear precision. If, additionally, the gradient $\nabla(\Phi(A_i))$ of Φ at the sample points $A_i, i = 1..n,$ is known, we define an interpolant that reproduces exactly a quadratic function – again if the surface is locally planar. Our method generalizes easily to interpolation of vector-valued functions, i.e. functions $\Phi : \mathcal{S} \to \mathbb{R}^k, k = 2, 3....$ We simply interpolate each coordinate of the result vector independantly as if it were a scalar function. The interpolants are tested in different settings. Examples are shown in section 5.

1.2 Paper outline

After the introduction, we proceed in section 2 with the definition of some basic concepts, and we recall some known results that are needed in the sequel. In section 3, we define the T-neighbors of a surface point X with respect to a sample of the surface. We show that all T-neighbors of X lie in a small neighborhood around X if the surface is well sampled. In section 4, we define the T-coordinate system on the surface with respect to the sample. We show the main properties of the T-coordinates; in particular, we show that the coordinate functions have compact supports and are continuously differentiable almost everywhere on the surface. Although a point on a surface cannot, in general, be expressed as a convex combination of other points on the surface, we show in subsection 4.4, that the local coordinate property is approximately satisfied, with an error that depends on the local curvature of the surface and on the sampling density. In section 5, we describe the applications of T-coordinates to scattered data interpolation on a surface, and we show some experimental results. Perspectives and a conclusion are given in the last section.

2 BASIC NOTATIONS AND RESULTS

In this section, we give the definition of the main ingredients of our framework and recall some basic results. Notably, we give a short introduction to natural neighbor coordinates for the general case of points in $\mathbb{R}^d$, we define power diagrams and show their relationship to sections of Voronoi diagrams, and we recall results on sampled surfaces.

2.1 Voronoi diagrams and natural neighbors

Let $\mathcal{A} = \{A_1, ..., A_n\}$ be a set of points in $\mathbb{R}^d$. Without real loss of generality, we can assume that no $d + 2$ points lie on the same sphere. The Voronoi cell of A_i is $V(A_i) = \{X \in \mathbb{R}^d : ||X - A_i|| \leq ||X - A_j|| \quad \forall j = 1, ..., n\}$ where $||X - Y||$ denotes the Euclidean distance between points $X, Y \in \mathbb{R}^d$. The collection of Voronoi cells is called the *Voronoi diagram* of $\mathcal{A}$ or $Vor(\mathcal{A})$.

Let $\mathcal{A}'$ be a subset of points of $\mathcal{A}$ whose Voronoi cells have a non-empty intersection. The convex hull $conv(\mathcal{A}')$ is called a Delaunay face, and the collection of Delaunay faces is called the Delaunay triangulation of $\mathcal{A}$, denoted $Del(\mathcal{A})$. Under the general position assumption, the d-dimensional faces of $Del(\mathcal{A})$ are simplexes, hence the name triangulation. See figure 1.
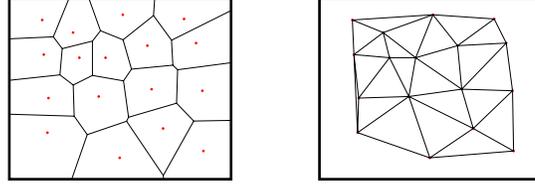

Figure 1: (a) a Voronoi diagram, (b) the dual Delaunay triangulation.

Given a point $X \notin \mathcal{A}$, we define $Vor^+ = Vor(\mathcal{A} \cup \{X\})$, $Del^+ = Del(\mathcal{A} \cup \{X\})$ and $V^+(X)$ and $V^+(A_i)$ to be the Voronoi cells of X and A_i in Vor^+.
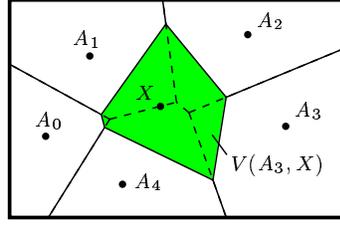

Figure 2: X has five natural neighbors $A_0, ..., A_4$.

Let $V(X, A_i) = V^+(X) \cap V(A_i)$, the part of the Voronoi cell $V^+(X)$ that has been 'stolen' from $V(A_i)$ at the insertion of X. If $V(X, A_i) \neq \emptyset$, A_i is a natural neighbor of X. Let $\nu_i(X)$ be the volume of $V(X, A_i)$ and $\nu(X)$ be the volume of $V^+(X)$. Figure 2 shows an example. Observe that if X belongs to $conv(\mathcal{A})$, the convex hull of $\mathcal{A}$, $\nu(X)$ is bounded.

Definition 2.1 *The natural neighbor coordinates associated to $\mathcal{A}$ of a point $X \in conv(\mathcal{A})$ are the functions $\sigma_i(X) = \frac{\nu_i(X)}{\nu(X)}, i = 1..n$.*

As already mentioned, the natural neighbor coordinates satisfy the three conditions of Definition 1.1. Moreover, the support Δ_i of the natural neighbor coordinate σ_i, i.e. the set $\{X \in \mathcal{U} | \sigma_i(X) \neq 0\}$, is the interior of the union of the spheres circumscribing the Delaunay simplexes adjacent to A_i.

Definition 2.2 *The natural neighbors of a point X with respect to $\mathcal{A}$ are the points $A_i \in \mathcal{A}$ with $\sigma_i(X) \neq 0, i = 1..n$*

Notice that, the natural neighbors of X are exactly the vertices other than X of the simplices of Del^+ incident to X.

2.2 Natural neighbor coordinates in power diagrams and sections of Voronoi diagrams

In this section, we recall the definition of a power diagram, and we point out the relationship between k-dimensional sections of Voronoi diagrams in $\mathbb{R}^d$ and k-dimensional power diagrams. We state some observations concerning orthogonal spheres, and, at last, we define natural neighbor coordinates in power diagrams and prove their continuity properties.

2.2.1 Power diagrams and sections of Voronoi diagrams

Let $\mathcal{B} = \{B_1, B_2, ..., B_n\}$ be a set of weighted points (or sites) in $\mathbb{R}^d$, $B_i = (p_i, w_i) \in \mathbb{R}^d \times \mathbb{R}$. A weighted point B_i can also be considered as a sphere with center p_i and radius $\sqrt{w_i}$. Notice, that w_i might be negative and the associated sphere imaginary. The power distance $\Gamma_i(p)$ of a point $p \in \mathbb{R}^d$ with respect to B_i is defined as $\Gamma_i(p) := ||p_i - p||^2 - w_i$. A weighted point (p_1, w_1) (or the corresponding sphere) is orthogonal to another weighted point (p_2, w_2) (or sphere) if $\Gamma_1(p_2) = w_2$. Without real loss of generality, we can assume that the weighted points are in general position. This means that no $d + 2$ weighted points are orthogonal to the same sphere. The power cell of B_i is $P(B_i) = \{p \in \mathbb{R}^d : \Gamma_i(p) \leq \Gamma_j(p) \quad \forall j = 1, ..., n\}$. The collection of power cells is called the *power diagram* of $\mathcal{B}$ or $Pow(\mathcal{B})$. The dual of the power diagram is called the *regular triangulation* or $Reg(\mathcal{B})$. The vertices of $Reg(\mathcal{B})$ belong to $\{p_1, ..., p_n\}$ and, under the general position assumption, the d-dimensional faces of $Reg(\mathcal{B})$ are simplexes. See e.g. [8] for further properties of power diagrams and regular triangulations.

Let H be a hyperplane in $\mathbb{R}^d$, and let the Voronoi diagram $Vor(\mathcal{A})$ be defined as in section 2.1. We call $Vor'(\mathcal{A})$ the intersection of the Voronoi diagram with H, $Vor'(\mathcal{A}) := Vor(\mathcal{A}) \cap H$, $V'(A_i) := V(A_i) \cap H$.

Observation 2.3 $Vor'(\mathcal{A})$ *is the power diagram of the points A'_i that are the projection of the sample points $A_i \in \mathcal{A}$ onto H weighted with $w_i = -||A_i - A'_i||^2$.*

Proof: Because A'_i is the orthogonal projection of A_i on H, we know that $\forall X \in H : ||X - A_i||^2 = ||X - A'_i||^2 + ||A_i - A'_i||^2$. Therefore, $V'(A_i) = \{X \in H : ||X - A_i||^2 \leq ||X - A_j||^2 \, \forall j = 1, ..., n\} = \{X \in H : ||X - A'_i||^2 + ||A_i - A'_i||^2 \leq ||X - A'_j||^2 + ||A_i - A'_j||^2 \, \forall j = 1, ..., n\} =: P(A'_i, w_i)$. $\square$

2.2.2 Properties of orthogonal spheres

We call *orthosphere* of a d-simplex of $Reg(\mathcal{B})$ the sphere that is orthogonal to the $d + 1$ weighted points associated to the vertices of the face. Its center is a vertex of $Pow(\mathcal{B})$. Let the bisector hyperplane of two weighted points B_i and B_j be the hyperplane $\mathcal{H} = \{p \in \mathbb{R}^d : \Gamma_i(p) = \Gamma_j(p)\}$. A $(d - 1)$-dimensional face of $Pow(\mathcal{B})$ is contained in the bisector hyperplane of its defining sites.

In the context of this paper, we use power diagrams that correspond to sections of Voronoi diagrams. Consequently, we only consider sites with negative or zero weight. They have the following two properties:

Proposition 2.4 *Any sphere orthogonal to a site B_i with $w_i \leq 0$ is real. Furthermore, it contains p_i in its interior if $w_i < 0$ or on its boundary if $w_i = 0$.*

Proof: Let $O = (p_o, w_o)$ be a sphere orthogonal to B_i with radius r_o. We have $||p_o - p_i||^2 - w_o = w_i \leq 0$, thus, $w_o \geq ||p_o - p_i||^2 \geq 0$, and $||p_o - p_i|| \leq r_o = \sqrt{w_o}$. If $w_i = 0$: $||p_o - p_i|| = r_o$.$\square$

Proposition 2.5 *All spheres orthogonal to two weighted points B_1 and B_2 with $w_1 \leq 0$ and $w_2 \leq 0$ meet the line $(p_1 p_2)$ passing through p_1 and p_2 in the same two points y_1 and y_2.*

Proof: We denote by $\mathcal{C}$ the (infinite) set of spheres orthogonal to B_1 and B_2. By proposition 2.4, the spheres of $\mathcal{C}$ are real. Furthermore, notice that they are centered on the bisector hyperplane $\mathcal{H}$ of B_1 and B_2 which is orthogonal to $(p_1 p_2)$. See also figure 3.
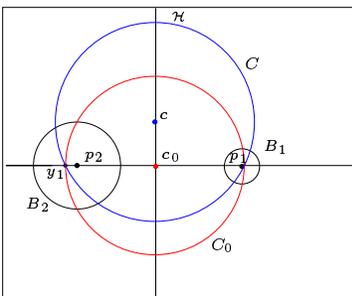

Figure 3: Orthogonal spheres to B_1 and B_2.

Let $C_0 \in \mathcal{C}$ have its center c_0 on $(p_1 p_2)$ and y_1 be the intersection point of C_0 with $(p_1 p_2)$ that is further from p_1. Since C_0 is real, its weight is $||c_0 - y_1||^2$. Let C be any other sphere of $\mathcal{C}$ with center c and weight w_c. Since C_0 and C are both orthogonal to B_1, we get:

$$||c - p_1||^2 - w_c = ||c_0 - p_1||^2 - ||c_0 - y_1||^2 = w_1. \tag{1}$$

Applying Pythagore's theorem to the triangle (c, c_0, p_1) yields $||c - p_1||^2 - ||c_0 - p_1||^2 = ||c - c_0||^2$. Thus, equation (1) becomes $w_c = ||c_0 - y_1||^2 + ||c - c_0||^2$. By Pythagore's theorem in triangle (c, c_0, y_1), we deduce $w_c = ||c - y_1||^2$ which implies that y_1 lies on C. Similar arguments apply for y_2.$\square$

2.2.3   Natural coordinates in power diagrams

Natural neighbor coordinates in power diagrams are defined very similarly to natural neighbor coordinates in ordinary Voronoi diagrams. The major difference between the coordinates in Voronoi

diagrams and in power diagrams is due to the fact that a site of a power diagram might have an empty cell. This implies that the cell of some sites may disappear when inserting a new site, let's say $X = (p_X, w_X)$. Nevertheless, the coordinate can be defined as the proportion of volume of the power cell $P(X)$ which is stolen from some site B_i at the insertion of X into $Pow(\mathcal{B})$.

We define $Pow^+(\mathcal{B}) := Pow(\mathcal{B} \cup \{X\})$ and $Reg^+(\mathcal{B}) := Reg(\mathcal{B} \cup \{X\})$. $P^+(X)$ is the power cell of X in $Pow^+(\mathcal{B})$ and $P(B_i, X) := P^+(X) \cap P(B_i)$. The volume of $P(B_i, X)$ is denoted by $\pi_i(X)$ and the volume of $P(X)$ is denoted by $\pi(X)$. $\mathcal{C}$ is the convex hull of the $p_i, i = 1..n$.

Definition 2.6 *The natural neighbor coordinates associated to $\mathcal{B}$ of a point $X \in \mathcal{C}$ with $\pi(X) \neq \infty$ and $\pi(X) \neq 0$ are the functions $\lambda_i(X) = \frac{\pi_i(X)}{\pi(X)}, i = 1..n$.*

Definition 2.7 *The natural neighbors associated to $\mathcal{B}$ of a point $X \in \mathcal{C}$ with $\pi(X) \neq \infty$ and $\pi(X) \neq 0$ are the points $B_i \in \mathcal{B}$ with $\lambda_i(X) \neq 0, i = 1..n$.*

Notice, that in this case, the natural neighbors of X include not only the vertices other than X of the simplices of $Reg^+(\mathcal{B})$ incident to X but also the vertices that disappear from $Reg(\mathcal{B})$ when inserting X.

Sibson proved that the natural neighbor coordinates in power diagrams satisfy the three conditions of Definition 1.1 [23]. The following lemma generalizes the work of Piper [22] about the continuity properties of the coordinate function in Voronoi diagrams to power diagrams. It is restricted to point sets with non-positive weights as it makes use of proposition 2.4. For the following lemmas, let $X = (p_X, w_X)$ with $w_X \leq 0$, $\pi(X) \neq \infty$, and $\pi(X) \neq 0$, and $\mathcal{B}$ be a set of weighted points with non-positive weights.

Lemma 2.8 *The natural neighbor coordinate $\lambda_i(X)$ of X associated to $\mathcal{B}$ is C^0 continuous over $\mathcal{C}$ and C^1 continuous over $\mathcal{C}$ except at a finite set of points.*

Proof: As Farin noticed in [15], the natural neighbor coordinate is a rational function of X within a cell of the arrangement of the Delaunay spheres or orthospheres in our case (which are real by proposition 2.4). This is due to the fact that, in each cell, the natural neighbors are fixed, and the vertices of $P(B_i, X)$ are rational functions of X. Therefore, $\pi_i(X)$ which is the volume of $P(X, B_i)$ is a rational function of X. Since by assumption, the volume function $\pi(X) = \sum \pi_i(X)$ does not vanish, the differentiability of the normalized coordinate λ_i follows from the differentiability of the π_i.

When X crosses the boundary of an orthosphere, the rational function changes. To prove that π_i is C^1 continuous on the orthosphere, it suffices to prove that the restriction of X to a line that intersects the orthosphere transversally is C^1 continuous. Following Piper [22], we restrict π_i to a straight line l through p_i. For power diagrams in general this is not sufficient because such a line might be tangent to an orthosphere orthogonal to B_i. However, this is different if the weights are non-positive because we know from proposition 2.4 that p_i lies inside or on the boundary of any sphere orthogonal to B_i.

Piper's proof can be generalized as follows: We consider the volume $v_i(X)$ of the intersection of the bisector hyperplane of B_i and X with $P(B_i)$. This volume varies continuously when p_X moves on l except when the bisector contains a $(d-1)$-dimensional face of $P(B_i)$. This happens

on one point ϵ_{ij} on each ray from p_i to p_j where p_j is the center of a neighbor B_j of B_i in $Reg(\mathcal{B})$. Precisely, ϵ_{ij} is the intersection point of the spheres orthogonal to B_i and B_j with the line $(p_i p_j)$ further to p_i as described in proposition 2.5. It is a vertex of the arrangement of orthospheres. If $w_X = w_j$, $\epsilon_{ij} = p_j$. Since $\pi_i(X)$ is obtained by integrating $v_i(X)$ over l, the C^0 continuity of the restriction of v_i to l implies the C^1 continuity of the restriction of π_i to l. It follows that π_i and consequently λ_i are C^1 continuous on $\mathcal{C} \setminus \mathcal{E}_i$ with $\mathcal{E}_i := \{\epsilon_{ij} | \ B_j$ neighbor of B_i in $Reg(\mathcal{B})\}$.

In order to prove the C^0 continuity at the points $\mathcal{E}_i$, the same argument as for natural coordinates in Voronoi diagrams can be used [15] : $\pi_i(X)$ is continuous everywhere except if $p_i = p_X$ and $w_i = w_X$. Then, $\pi_i(X)$ is equal to the volume of $P(B_i)$. When X approaches B_i but $B_i \neq X$, the bisector of B_i and X intersects $P(B_i)$, and $\pi_i(X)$ is just a fraction of the volume of $P(B_i)$. However, the normalized coordinate function λ_i is continuous because $\lim_{X \to B_i} \lambda_j(X) = \lim_{X \to B_i} \pi_j(X) = 0$, as follows from the definition of π_j, for all $j \neq i$. Since $\sum_{i=1}^{n} \lambda_i(X) = 1$, we deduce $\lim_{X \to B_i} \lambda_i(X) = 1$, which shows that λ_i is continuous everywhere in $\mathcal{C}$. $\square$

2.3   Sampled surfaces: definitions and results

In this section, we consider the case where the sample points are taken from a smooth surface $\mathcal{S}$, i.e. a twice-differentiable surface. We assume the first and the second derivative of $\mathcal{S}$ at X to be continuous and the third derivatives to exist for all $X \in \mathcal{S}$. We give a definition of what we call the sampling density of $\mathcal{S}$ and we recall several results about the local behavior of the surface samples, notably from [4] and [6].

2.3.1   Voronoi diagram on a surface

We first define the Voronoi diagram of a set of points restricted to a surface, following previous work by Chew [12] and Edelsbrunner and Shah [14].

Definition 2.9 (Restricted Voronoi diagram, Chew)
The Voronoi diagram of $\mathcal{A}$ restricted to $\mathcal{S}$ is the (curved) cell complex obtained by intersecting each face of $Vor(\mathcal{A})$ with $\mathcal{S}$. We denote it by $Vor_{\mathcal{S}}(\mathcal{A})$.

We denote by $V_{\mathcal{S}}(A_i)$ the Voronoi cell of $Vor_{\mathcal{S}}(\mathcal{A})$ consisting of the points of $\mathcal{S}$ that are closer to A_i (for the Euclidean distance) than to any A_j, $j \neq i$, . A vertex of $V_{\mathcal{S}}(A_i)$ is the intersection of an edge of $V(A_i)$ with $\mathcal{S}$. Hence it is the center of a ball passing through d points of $\mathcal{A}$ and not enclosing other points of $\mathcal{A}$.

Definition 2.10 (Restricted Delaunay triangulation, Chew)
The Delaunay triangulation of $\mathcal{A}$ restricted to $\mathcal{S}$ is the subcomplex of $Del(\mathcal{A})$ consisting of the faces of $Del(\mathcal{A})$ whose dual Voronoi edges intersect $\mathcal{S}$. We denote it by $Del_{\mathcal{S}}(\mathcal{A})$.

2.3.2 Medial axis and local feature size

Let $\mathcal{F} \subset \mathbb{R}^d$ be a compact object, $\mathcal{S}$ its boundary. We call $\overrightarrow{n_X}$ the outward unit normal to $\mathcal{S}$ at X. The following definitions allow to characterize a sampling of a surface. See [25] and [4] for basic results on properties of the medial axis.

Definition 2.11 (Amenta & Bern)

1. **(Medial axis)** *The medial axis of a surface $\mathcal{S}$ in $\mathbb{R}^d$ is the closure of the set of points with more than one closest point on $\mathcal{S}$.*

2. **(Medial ball)** *A ball that is centered on the medial axis, tangent to the surface and whose interior does not intersect the surface is called a medial ball.*

3. **(Local feature size)** *The local feature size $\mathrm{lfs}(X)$ of a point $X \in \mathcal{S}$ is the least distance from X to the medial axis of $\mathcal{S}$.*

4. **(ϵ-sample)** *A set of sample points $\mathcal{A}$ of $\mathcal{S}$ is said to be an ϵ-sample of $\mathcal{S}$, $\epsilon < 1$, if every point $X \in \mathcal{S}$ has a sample point at distance at most $\epsilon\,\mathrm{lfs}(X)$.*

The following proposition from Boissonnat and Cazals will be useful in section 3.2. [6]

Lemma 2.12 [6, proposition 14]
Let $B(X, r)$ be a ball centered at X with radius r. For any $X \in \mathcal{S}$ and any $r < \mathrm{lfs}(X)$, $\mathcal{S} \cap B(X, r)$ is a topological $(d-1)$-ball.

2.3.3 Properties of well sampled surfaces

In this section, we mainly recall some results of Amenta and Bern [4].

The first lemma states that the local feature size is Lipschitz.

Lemma 2.13 [4, lemma 1]
For any two points $X, Y \in \mathcal{S}$, $\mathrm{lfs}(X) \leq \mathrm{lfs}(Y) + ||X - Y||$.

It follows that, if $\mathcal{A}$ is a ϵ-sample, the maximum distance between $X \notin \mathcal{A}$ and the closest sample point $A_i \in \mathcal{A}$ is $\frac{\epsilon}{1-\epsilon}\mathrm{lfs}(A_i)$.

Considering two points on the surface that are close, Amenta and Bern show that the angle between the line segment connecting the two points and the normal at the points is large, whereas the angle between two normals is small.

Lemma 2.14 [4, lemma 2]
For any two points X and Y on $\mathcal{S}$ with $||X - Y|| \leq \rho\,\mathrm{lfs}(X)$, the smaller angle between the line segment $[XY]$ and the normal to $\mathcal{S}$ at X is at least $\frac{\pi}{2} - \arcsin(\frac{\rho}{2})$.

Lemma 2.15 [4, lemma 3]
For any two points X and Y on $\mathcal{S}$ with $||X - Y|| \leq \rho\,\mathrm{lfs}(X)$, $\rho < \frac{1}{3}$, the angle between the normals to $\mathcal{S}$ at X and at Y is at most $\frac{\rho}{1-3\rho}$.

Since the Voronoi diagram of an ϵ-sample consists of long and skinny cells, the normal direction at a point $X \in \mathcal{S}$ can be estimated from the Voronoi diagram of $Vor(\mathcal{A} \cup \{X\})$. To be more specific, we recall the definition of a pole.

Definition 2.16 *The* pole p_X *of* X *is the Voronoi vertex of* $V^+(X)$ *which is furthest from* X.

As stated in the next lemma, the line passing through X and its pole provides a good approximation of the (non-oriented) direction of $\overrightarrow{n_X}$, the normal to $\mathcal{S}$ at X. Consequently, the plane $\widetilde{T}_X$ that contains X and is orthogonal to $\overrightarrow{v_X}$ approximates well the tangent plane to $\mathcal{S}$ at X.

Lemma 2.17 [4, lemma 5] *The smaller angle between* $\overrightarrow{n_X}$ *and the line passing through* X *and its pole is at most* $2\arcsin(\frac{\epsilon}{1-\epsilon})$.

In the sequel, we will need the following lemma that states that the tangent plane of a point $X \in \mathcal{S}$ cannot be parallel to the bisector of two sample points that are at distance at most $\rho\mathrm{lfs}(X)$ from X, for $\rho < \frac{1}{3}$.

Lemma 2.18 *The angle between the tangent plane* T_X *to* $\mathcal{S}$ *at* $X \in \mathcal{S}$ *and the bisector of two sample points* A_i *and* A_j *that are at distance at most* $\rho\mathrm{lfs}(X)$ *from* X, $\rho \leq \frac{1}{3}$, *is at least* $\pi/2 - \arcsin(\frac{\rho}{1-\rho}) - \frac{\rho}{1-3\rho}$.
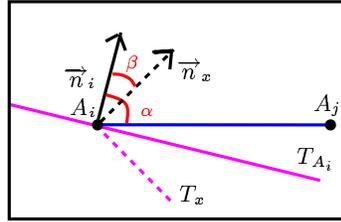

Figure 4: Bounding the angle between the normal n_x and the segment $[A_iA_j]$.

Proof: See figure 4 for notations. We derive a lower bound for the angle α between the vector $\overrightarrow{A_iA_j}$ and the surface normal $\overrightarrow{n_i}$ at A_i and an upper bound for the angle β between the normals at A_i and X. The angle between $\overrightarrow{n_X}$ and $\overrightarrow{A_iA_j}$ is at least $\alpha - \beta$.

In order to bound α, we apply lemma 2.14. The distance between the sample points is bounded by $\|A_i - A_j\| \leq \|A_i - X\| + \|A_j - X\| \leq 2\rho\,\mathrm{lfs}(X)$. Using lemma 2.13, we get $\mathrm{lfs}(X) \leq \mathrm{lfs}(A_i) + \rho\mathrm{lfs}(X)$ and $\mathrm{lfs}(X) \leq \frac{1}{1-\rho}\mathrm{lfs}(A_i)$. Hence, $\|A_i - A_j\| \leq \frac{2\rho}{1-\rho}\mathrm{lfs}(A_i)$. It follows from Lemma 2.15 that α is at least $\pi/2 - \arcsin(\frac{\rho}{1-\rho})$. The angle β between the normals at A_i and X is at most $\frac{\rho}{1-3\rho}$ if $\rho < \frac{1}{3}$ (lemma 2.15). Concluding, we get $\alpha - \beta \geq \pi/2 - \arcsin(\frac{\rho}{1-\rho}) - \frac{\rho}{1-3\rho}$, provided that $\rho < \frac{1}{3}$. $\square$

Lemma 2.19 *For any $\rho < \frac{\pi}{2+3\pi}$, the projection onto T_X of the intersection of the ball $B(X, \rho \mathrm{lfs}(X))$ with S is 1-1.*

Proof: By Lemma 2.15, for any point Y at distance at most $\rho \mathrm{lfs}(X)$ from X, the angle between the normals to S at X and at Y is at most $\frac{\rho}{1-3\rho}$. The projection is 1-1 if $\frac{\rho}{1-3\rho} < \frac{\pi}{2} \iff \rho < \frac{\pi}{2+3\pi} \approx 0.27$.

3 SURFACE NEIGHBORS

With this section begins the core part of the paper in which we define a local neighborhood of a point X of a smooth surface S without boundary with respect to an ϵ-sample $\mathcal{A}$ of S. In order not to compute geodesic Voronoi diagrams on S, we approximate S locally by the tangent plane T_X of X to S. We determine the natural neighbors of X in the Voronoi diagram $Vor(\mathcal{A})$ restricted to T_X, and we call them T-neighbors of X. In the remainder of this section, we formally define the T-neighbors of a point $X \in S$ and prove that they are close to X.

3.1 Definition of T-neighbors

In a first time, we assume that for each point $X \in S$, the normal $\overrightarrow{n_X}$ to S at X is known, and therefore, the tangent plane T_X at X.

Let $Vor'(\mathcal{A})$ be the intersection of $Vor(\mathcal{A})$ with the tangent plane T_X. The Delaunay triangulation restricted to T_X that consists of the faces of $Del(\mathcal{A})$ whose dual Voronoi edges intersect T_X is called $Del'(\mathcal{A})$. Alternatively, $Vor'(\mathcal{A})$ is the $(d-1)$-dimensional power diagram of the points A_i' that are the projection of the sample points $A_i \in \mathcal{A}$ onto T_X weighted with $w_i = -\|A_i - A_i'\|^2$. Let $Reg'(\mathcal{A})$ be the regular triangulation dual to $Vor'(\mathcal{A})$. Since two cells of $Vor'(\mathcal{A})$ are adjacent iff their corresponding cells in $Vor(\mathcal{A})$ are adjacent and intersect T_X, $Reg'(\mathcal{A})$ is the projection of $Del'(\mathcal{A})$ onto T_X.

Definition 3.1 (T-neighbor of X) *The T-neighbors associated to $\mathcal{A}$ of a point $X \in S$ are the sample points $A_i \in \mathcal{A}$ such that their projection A_i' is a natural neighbor of X in $Vor'(\mathcal{A})$.*

To see that the concept of T-neighbors is well-defined, we make two observations: First, the definitions of section 2.2 assume general position of the point sites. The case that two sites have the same position and the same weight is excluded. In our context, this occurs if the bisector of two T-neighbors A_k and A_j of X coincides with T_X: A_k and A_j are projected at the same position, and they have the same weight because they are at the same distance to the tangent plane but on opposite sides. However, without real loss of generality, we can assume that no bisector of two sample points is tangent to S.

Notice, still, that we can easily show that the angle between T_X and the bisector of A_k and A_j is strictly positive, if $\epsilon \leq \frac{1}{9}$: in the remainder of this section, we show that A_k and A_j are at distance at most $\frac{2\epsilon}{\sqrt{1-2\epsilon}} \mathrm{lfs}(X)$ to X (lemma 3.3), thus, lemma 2.18 applies with $\rho = \frac{2\epsilon}{\sqrt{1-2\epsilon}}$.

Second, we need the following lemma to show that X lies in the convex hull of the projection of its T-neighbors. Equivalently, X lies in the convex hull of its natural neighbors in $Vor'(\mathcal{A})$, which in turn is equivalent to the fact that $V'(X)$ is bounded.

Lemma 3.2 *X belongs to the convex hull of the projection of its T-neighbors on T_X.*

Proof: For a contradiction, assume that $V'(X)$ is unbounded. Then the interior of $V'(X)$ contains a point at infinity p_∞. Since, among the points of $\mathcal{A} \cup \{X\}$, X is the closest to p_∞, the halfspace H^+ that contains p_∞ and is limited by the hyperplane H passing through X and normal to Xp_∞ does not contain any point of $\mathcal{A}$. See figure 5.
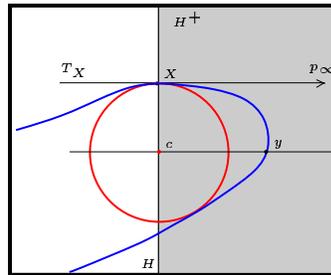
Figure 5: H^+ must contain sample points.

The medial ball passing through X and lying in the region limited by $\mathcal{S}$ is centered at a point $c \in H$. c is a point of the medial axis of $\mathcal{S}$. Let y be a point on the intersection of $\mathcal{S}$ with the ray issued from c, orthogonal to H and contained in H^+. Such a point exists if $\mathcal{S}$ has no boundary. We have $||y - c|| \leq \mathrm{lfs}(y)$ and for any $A_i \in \mathcal{A}, ||y - A_i|| > ||y - c||$. This contradicts the fact that $\mathcal{A}$ is an ϵ-sample with $\epsilon < 1$. $\square$
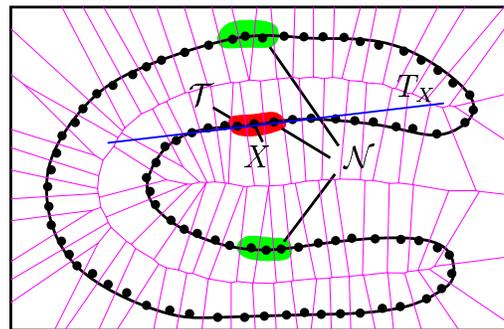

Figure 6: Natural neighbors $\mathcal{N}$ and T-neighbors $\mathcal{T}$ of X

In figure 6, a schematic drawing illustrates that some of the natural neighbors of X which are highlighted in the picture can be far away from X on "the other side" of $\mathcal{S}$ – separated from X by the medial axis of $\mathcal{S}$. The T-neighbors are the subset of the natural neighbors that are close to X on $\mathcal{S}$. This is what we prove in the next section.

If the normal at X is not known (this may be the case when the surface is only known at a finite set of points), we can approximate the tangent plane T_X by the plane $\widetilde{T}_X$ that passes through X and is orthogonal to the vector joining X to its pole. We can define $\tilde{T}$-neighbors in very much the same way as T-neighbors : the only difference is that T_X is replaced by $\widetilde{T}_X$.

3.2 Locality of the T-neighbors

We now derive a bound on the distance between a point $X \in \mathcal{S}$ and its T-neighbors with respect to the local feature size $\mathrm{lfs}(X)$. The proof of the lemma is closely inspired from [4, lemma 5].

Lemma 3.3 *Let $\mathcal{A}$ be an ϵ-sample of $\mathcal{S}$ and $X \in \mathcal{S}$.*
(a) The T-neighbors of X are all contained in a ball of radius $\frac{2\epsilon}{\sqrt{1-2\epsilon}}\mathrm{lfs}(X)$ centered at X.
(b) The $\tilde{T}$-neighbors of X are contained in a ball of radius

$$\frac{2\epsilon}{\cos(3\arcsin(\frac{\epsilon}{1-\epsilon}))} \, \mathrm{lfs}(X)$$

centered at X.
(c) In both cases, the T-neighbors of X are contained in a ball of radius $2\epsilon(1 + O(\epsilon))\,\mathrm{lfs}(X)$.

Proof: Let v be a vertex of $V'(X) = V^+(X) \cap T_X$. We derive an upper bound on the distance between X and v. Because $V'(X)$ is bounded by the (intersection with T_X of the) bisectors of X and its T-neighbors, the distance between X and its T-neighbors is at most twice the distance between X and v.

Let B_1 and B_2 be the two balls of radius $\mathrm{lfs}(X)$ that are tangent to $\mathcal{S}$ at X. Assume without loss of generality that B_1 and v are on opposite sides of $\mathcal{S}$. Let α be the angle $\angle v m_1 X$ where m_1 is the center of B_1. We find the same angle $\alpha = \angle v X q'$ where q' is the orthogonal projection of X onto the line segment $v m_1$. See figure 7.
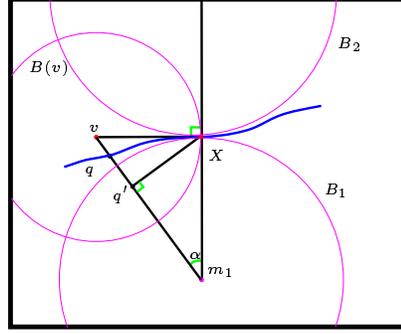
Figure 7: For the proof of lemma 3.3(a)

Since v and m_1 lie on different sides of $\mathcal{S}$, the line segment vm_1 must intersect $\mathcal{S}$. Let q be such an intersection point. We call $B(v)$ the ball with center v and radius $||X - v||$. Since $v \in V^+(X)$, $B(v)$ is empty of sample points. Because B_1 is also empty of sample points, X is the point of $\mathcal{A} \cup \{X\}$ that is closest to q. It follows that $||X - q|| \leq \frac{\epsilon}{1-\epsilon} \mathrm{lfs}(X)$. On the other hand, $||X - q|| \geq ||X - q'|| = \sin\alpha \, \mathrm{lfs}(X)$. Hence, $\alpha \leq \arcsin(\frac{\epsilon}{1-\epsilon})$. From the triangle vXq', we know that $||X - v|| = \frac{||X - q'||}{\cos\alpha} \leq \frac{||X-q||}{\cos\alpha}$. With $\cos\alpha \geq \cos(\arcsin(\frac{\epsilon}{1-\epsilon})) = \frac{\sqrt{1-2\epsilon}}{1-\epsilon}$, we get $||X - v|| \leq \frac{\epsilon}{\sqrt{1-2\epsilon}} \mathrm{lfs}(X) = \epsilon(1 + O(\epsilon))\mathrm{lfs}(X)$.
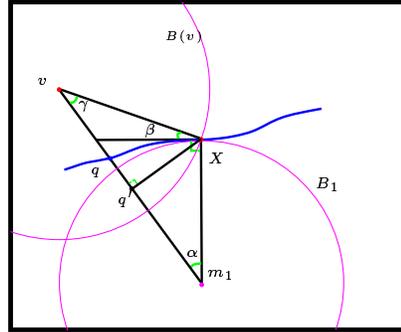
Figure 8: For the proof of lemma 3.3(b)

(b) We consider a vertex v of $\widetilde{V}'(X) = V^+(X) \cap \widetilde{T}(X)$. Let β be the angle between $T(X)$ and $\widetilde{T}_X$. From lemma 2.17, we know that $\beta \leq 2\arcsin(\frac{\epsilon}{1-\epsilon})$. See figure 8 for notations. We define q as before, and we obtain the same bounds $||X - q|| \leq \frac{\epsilon}{1-\epsilon}\mathrm{lfs}(X)$ and $\alpha = \angle vm_1X \leq \arcsin(\frac{\epsilon}{1-\epsilon})$.

Let $\gamma = \angle m_1vX = \frac{\pi}{2} - \alpha - \beta \geq \frac{\pi}{2} - 3\arcsin(\frac{\epsilon}{1-\epsilon})$ and $\sin\gamma \geq \cos(3\arcsin(\frac{\epsilon}{1-\epsilon}))$.

From triangle vXq', we know that $||X - v|| = \frac{||q' - X||}{\sin \gamma} \leq \frac{||X - q||}{\sin \gamma} \leq \frac{\epsilon}{\cos(3 \arcsin(\frac{\epsilon}{1 - \epsilon}))} \operatorname{lfs}(X) = \epsilon(1 + O(\epsilon)) \operatorname{lfs}(X)$. $\square$

The following lemma states that X and all its T-neighbors are contained in a small topological ball, if $\epsilon < 0.3$. It is a direct consequence of lemma 3.3.

Lemma 3.4 *Let $B(X, r)$ be a ball centered at X with radius $r = \frac{2\epsilon}{\sqrt{1 - 2\epsilon}} \operatorname{lfs}(X)$. If $\epsilon < \frac{\sqrt{5} - 1}{4}$, $B(X, r) \cap \mathcal{S}$ is a topological $(d - 1)$-ball that contains all the T-neighbors of X.*

Proof: By lemma 3.3, all the T-neighbors of X are contained in a ball $B(X, r)$ centered at X of radius $r = \frac{2\epsilon}{\sqrt{1 - 2\epsilon}} \operatorname{lfs}(X)$. With lemma 2.12, $\mathcal{S} \cap B(X, R)$ is a topological $(d-1)$-ball if $r < \operatorname{lfs}(X)$, which is true for $\epsilon < \frac{\sqrt{5} - 1}{4} \approx 0.3$. $\square$

4 SURFACE COORDINATE SYSTEM

We can now define a surface coordinate system associated to a sample set $\mathcal{A}$. Similarly to the T-neighbors, it is defined in the tangent plane of each surface point and consequently called T-coordinate system. In the rest of the section, we show some properties of T-coordinate systems : we show that the support of the T-coordinates is local in subsection 4.2, study their continuity properties in subsection 4.3, and prove in subsection 4.4 that the local coordinate property is approximately satisfied, with an error that depends on the local curvature of the surface and on the sampling density. The results for $\widetilde{T}_X$-coordinates, the generalization of T-coordinates to estimated tangent planes, are asymptotically identical. We do not give the proofs because they can be easily deduced by applying lemma 3.3(b) and (c).

4.1 Definition and basic properties

The same definitions as in section 3 apply.

Definition 4.1 (T**-coordinate associated to** $\mathcal{A}$) *The T-coordinate $\tau_i(X)$ of a point X of $\mathcal{S}$ is the natural neighbor coordinate $\lambda_i(X)$ of X associated to A_i' in the power diagram $Vor'(A)$, $i = 1, ..., n$.*

By construction, the T-coordinates τ_i fulfill properties (ii) and (iii) of a *system of coordinates over $\mathcal{S}$ associated to $\mathcal{A}$* as they are listed in definition 1.1. The local coordinate property *(i)* is satisfied for the projected sample points A_i'. With respect to $\mathcal{A}$, the local coordinate property is only true if the surface is locally planar so that all T-neighbors of X lie in the tangent plane T_X.

4.2 Locally bounded support

Let Δ_i denote the support of τ_i, i.e. the subset of the points $X \in \mathcal{S}$ such that $\tau_i(X) \neq 0$. In order to show the locality of Δ_i on $\mathcal{S}$, we apply the bound on the distance between a point $X \in \mathcal{S}$ and its T-neighbors.

Lemma 4.2 *The support Δ_i of τ_i is contained in a ball of radius $\frac{2\epsilon}{1-3\epsilon}\text{lfs}(A_i)$ centered at A_i.*

Proof: Applying lemma 3.3 and lemma 2.13, we obtain $||X - A_i|| \leq \delta\,\text{lfs}(X) \leq \frac{\delta}{1-\delta}\text{lfs}(A_i)$ with $\delta = \frac{2\epsilon}{\sqrt{1-2\epsilon}}$, which is at most $\frac{2\epsilon}{1-3\epsilon}\text{lfs}(A_i)$.□

4.3   Continuity of the coordinate function

In this section, we study the continuity of the function τ_i when X moves on $\mathcal{S}$. Let us first state the lemma:

Lemma 4.3 *The T-coordinate $\tau_i, i = 1, ..., n$, is continuous everywhere on $\mathcal{S}$.*

Proof: Assume that $\mathcal{S}$ is parameterized by (u, v). The coordinate function $\tau_i(X)$ is, by definition, equal to the natural neighbor coordinate $\lambda_i(X(u, v))$ in the power diagram $Vor'(\mathcal{A} \cup X)$. When X moves on $\mathcal{S}$, the projected sample points change their position and their weight continuously. Indeed, since $\mathcal{S}$ is smooth, the projection onto the tangent plane is a smooth mapping. The weight of A_i is given by :

$$w_i(u, v) = -||A_i - A'_i(u, v)||^2 = -(\overrightarrow{A_i X(u,v)} * \overrightarrow{n_{X(u,v)}})^2.[1]$$

By assumption, $X(u, v)$ is twice differentiable. The normal $\overrightarrow{n_X(u, v)}$ is C^1 continuous.[2] The same is true for the position of $A'_i(u, v)$ since $A'_i(u, v) = A_i - ||A_i - A'_i(u, v)||\overrightarrow{n_X(u, v)}$. Recall also from lemma 3.2, that X is always in the convex hull of the projected sample points. Consequently, the T-coordinates are continuous on all of $\mathcal{S}$ because the natural neighbor coordinates are continuous as described in section 2.2.□

Lemma 4.4 *The T-coordinate $\tau_i, i = 1, ..., n$, is continuously differentiable everywhere on $\mathcal{S}$ except at the sample points and at the points $X \in \mathcal{S}$ such that T_X contains a $(d - 2)$-dimensional Voronoi face of the Voronoi cell $V^+(X)$ in $Vor(\mathcal{A} \cup \{X\})$.*

Proof: Consider the power diagram $Vor'(\mathcal{A} \cup \{X\})$. By lemma 2.8, τ_i is C^1 continuous except at a finite number of points. Recall from the proof of lemma 2.8 that there is a C^1 discontinuity at the sample points. Still, we want to characterize all points of C^1 discontinuity with respect to $Vor^+(\mathcal{A})$.

For convenience, we adopt the terminology associated to $\mathbb{R}^3$ and call a $(d - 2)$-dimensional face an edge. The natural neighbor coordinate τ_i is not continuously differentiable at a point X such that the bisector of X and A'_i contains an edge of $V'(\mathcal{A})$ say $V'(A'_i) \cap V'(A'_j)$ (see also the proof of lemma 2.8). This means that a point v of this edge has equal power with respect to X, A'_i, and A'_j. But, by definition of $Vor'(\mathcal{A} \cup \{X\})$, this means also that v is at equal distance from A_i, A_j, and X. Consequently, the Voronoi edge of $V'(X)$ is an edge of $Vor(\mathcal{A} \cup \{X\})$. This edge is therefore contained in T_X.□

---

[1] We denote the scalar product of two vectors by $\overrightarrow{p} * \overrightarrow{q}$.

[2] $\overrightarrow{n_{X(u,v)}} = \frac{X_u \wedge X_v}{|X_u \wedge X_v|}$ where $X_u = \frac{\partial X}{\partial u}$ ($X_v = \frac{\partial X}{\partial v}$) is the first derivative of $X(u, v)$ with respect to u (with respect to v).

4.4 The local coordinate property

In this section, we bound the error committed with respect to the local coordinate property, that is $e(X) = ||X - \sum_i \tau_i(X) A_i||$. This corresponds to bounding the distance of a T-neighbor of a point $X \in \mathcal{S}$ to the tangent plane T_X, since $X = \sum_i \tau_i(X) A_i' = \sum_i \tau_i(X)(A_i + ||A_i - A_i'||\overrightarrow{n_x})$ or $X = \sum_i \tau_i(X)(A_i - ||A_i - A_i'||\overrightarrow{n_x})$ depending on the sign of $\overrightarrow{A_i' A_i} \overrightarrow{n_x}$.

Proposition 4.5

$$X = \sum_i \tau_i(X) A_i + O(\epsilon^2)\text{lfs}(X) \overrightarrow{n_x}.$$

Proof: It is sufficient to prove that the distance between a T-neighbor $A_i \in \mathcal{A}$ of $X \in \mathcal{S}$ and its projection A_i' on the tangent plane T_X, is bounded by

$$||A_i - A_i'|| \leq \frac{2\epsilon^2}{1 - 2\epsilon}\text{lfs}(X) = O(\epsilon^2)\text{lfs}(X).$$

It is easy to see that $||A_i - A_i'|| = ||A_i - X|| \sin \theta$ where $\theta = \angle A_i X A_i'$. Since A_i does not belong to the two balls of radius $\text{lfs}(X)$ centered at X, $\sin \theta \leq \frac{||A_i - X||}{2\text{lfs}(X)}$. With lemma 3.3, we have $||A_i - X|| \leq \frac{2\epsilon}{\sqrt{1-2\epsilon}}\text{lfs}(X)$ which implies the inequality above and the proposition. $\square$

5 APPLICATION: INTERPOLATION

As mentioned in the introduction, we apply the T-coordinates to interpolate a function which is defined on a surface but only known at some points. In section 5.1, we define three types of interpolants. They are tested in different experimental settings as described in section 5.2. In section 5.2.4, we visualize the coordinate function τ_i. Some notes on the implementation are given in section 5.3.

5.1 Interpolating a function on S

In this section, the previously defined T-coordinate system is used to define an interpolant for functions defined on smooth surfaces. More exactly, we want to approximate $\Phi : \mathcal{S} \to \mathbb{R}$ where $\mathcal{S} \subset \mathbb{R}^d$ is defined as before. We assume that an ϵ-sample $\mathcal{A}$ of $\mathcal{S}$ with function values $\{(A_i, z_i) : A_i \in \mathcal{A}, z_i = \Phi(A_i)\}$ is given. The interpolation is carried out for a point $X \in \mathcal{S}$.

5.1.1 Linear precision

For a point $X \in \mathcal{S}$, we compute its T-coordinates within its tangent plane T_X. The interpolation of $\Phi(X)$ is given as the linear combination of the function values of the T-neighbors weighted by the coordinates:

$$I^0(X) = \sum_i \tau_i(X) z_i.$$

Assume that $\mathcal{S}$ is locally flat and identical to a plane H. At any point X of S where the T-neighbors of X are all in H, the interpolant has linear precision. Indeed, if $z_i = a + bA_i$ for all T-neighbors of X, we have

$$I^0(X) = \sum_i \tau_i(X)(a + bA_i) = a + bX$$

by the local coordinate property. If S is not flat around X, we saw in section 4.4 how the error can be bounded with respect to $\mathrm{lfs}(X)$.

Let $\Phi(A_i) = a + bA_i$ be a linear function. We have $I^0(X) = \sum_i \tau_i(X)(a + bA_i) = a + bX + b\,e(X)\overrightarrow{n_X}$ where $e(X) = \sum_i \tau_i(X)(\overrightarrow{A_i'A_i} * \overrightarrow{n_x}) = O(\epsilon^2)\mathrm{lfs}(X)$. Hence,

$$I^0(X) = \Phi(X) + O(\epsilon^2).$$

5.1.2 Quadratic precision

We can also define an interpolant that reproduces a quadratic function $\Phi(X) = a + b^t X + X^t \mathcal{Q} X$. As additional input, we assume the gradients $g_i = \nabla(\Phi(A_i))$ of Φ at the sample points $A_i, i = 1..n$, to be known. The interpolant is defined as follows:

$$I^1(X) = \sum_i \tau_i(X)(z_i + \frac{1}{2} g_i^t (X - A_i))$$

Assume again that $\mathcal{S}$ is locally flat and identical to a plane H. At any point X of S where the T-neighbors of X are all in H, the interpolant has quadratic precision. Indeed, if $z_i = \Phi(A_i) = a + b^t A_i + A_i^t \mathcal{Q} A_i$, $g_i = b + 2\mathcal{Q} A_i$, and applying the local coordinate property, we obtain:

$$I^1(X) = a + b^t X + X^t \mathcal{Q} X.$$

If S is not flat around X, we have :

$$I^1(X) = a + b^t X + (X + e(X)\overrightarrow{n_X})^t \mathcal{Q} X.$$

where $e(X)$ is defined as in the previous case. Thus, the error of the interpolant is $\|\Phi(X) - I^1(X)\| = e(X)\overrightarrow{n_X}^t \mathcal{Q} X$, which is $O(\epsilon^2)$ inside any bounded domain.

In this context, we also implemented the so-called Z^1 interpolant proposed by Sibson [24] which we adapted to our setting by replacing the natural neighbor coordinates with the T-coordinates. This interpolant does not recapture general quadratic functions but only spherical quadrics of the form $\Phi(X) = a + b^t X + X^t X$. Sibson showed that it is C^1 continuous with gradient g_i at A_i. His proof relies on the local coordinate property of the natural neighbor coordinates. Consequently, in our case, an error is introduced depending on the local feature size. Using our terminology, Sibson's interpolant writes:

$$Z^1(X) = \frac{\frac{\sum_i \tau_i(X)\|A_i - X\|}{\sum_i \frac{\tau_i(X)}{\|A_i - X\|}} I^0(X) + \frac{\sum_i \tau_i(X)\|A_i - X\|^2}{\sum_i \frac{\tau_i(X)}{\|A_i - X\|}} \sum_i \frac{\tau_i(X)(z_i + g_i(X - A_i))}{\|A_i - X\|}}{\frac{\sum_i \tau_i(X)\|A_i - X\|}{\sum_i \frac{\tau_i(X)}{\|A_i - X\|}} + \sum_i \tau_i(X)\|A_i - X\|^2}.$$

5.2   Experimental results

5.2.1   Functions in $3D$

We assume that we are given a real-valued function $f(X)$ that can be evaluated at all points $X \in \mathcal{S}$. For example, we interpolated the following three functions on the sphere:

- $f_1(x, y, z) = 4x^2 + y^4 + 6z^4,$

- $f_2(x, y, z) = e^{-2((x-1)^2 + y^2 + z^2)} + 0.5\, e^{-4(x^2 + (y-0.7)^2 + (z-0.7)^2)}$
  $- 0.25\, e^{-4(x^2 + (y+0.7)^2 + (z+0.7)^2)},$

- $f_3(x, y, z) = 1 + x^8 + e^{2y^3} + e^{2z^2} + 10xyz.$

The functions f_2 and f_3 are test functions from [18] and [2].

   To visualize the result, we deformed the sphere at each point along its normal by the amount of the corresponding interpolation result. In figures 9(a), 10(a), and 11(a), the error statistics are depicted together with the exact result of the function application where the function is evaluated at the 6000 points of the sphere model and the sphere is deformed correspondingly. For each interpolant, mean and maximum errors are given with respect to the absolute difference between the actual and the interpolated function value on the 6000 evaluation points. In the plots, the three lower curves correspond to the mean error while the three upper curves correspond to the maximum error. In all cases, the linear interpolant I^0 is the least accurate concerning mean as well as maximum error. The quadratic interpolants I^1 and Z^1 achieve comparable results, yet, Z^1 yields to be slightly better than the simpler quadratic interpolant I^1. Figures 9, 10, and 11 (b) and (c) show the result of the interpolation on the 6000 sphere points given function values at 50 and 250 of the 6000 points.
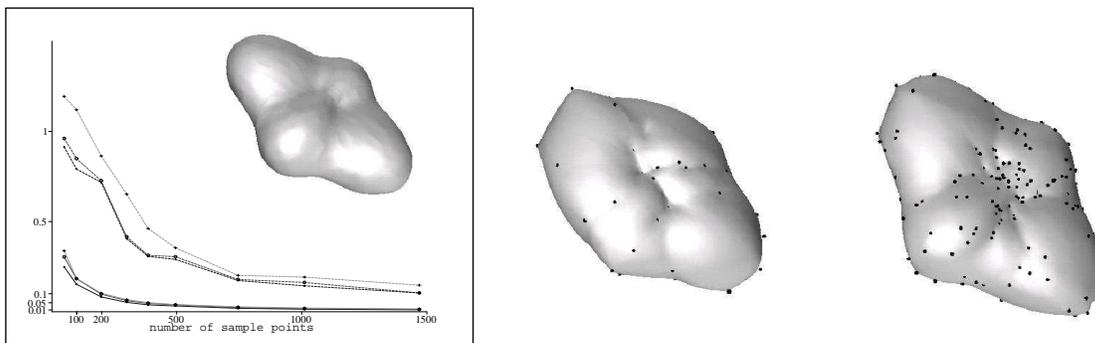
Figure 9: f_1 (a) exact model and error statistic, interpolation with (b) 50, (c) 250 sample points.
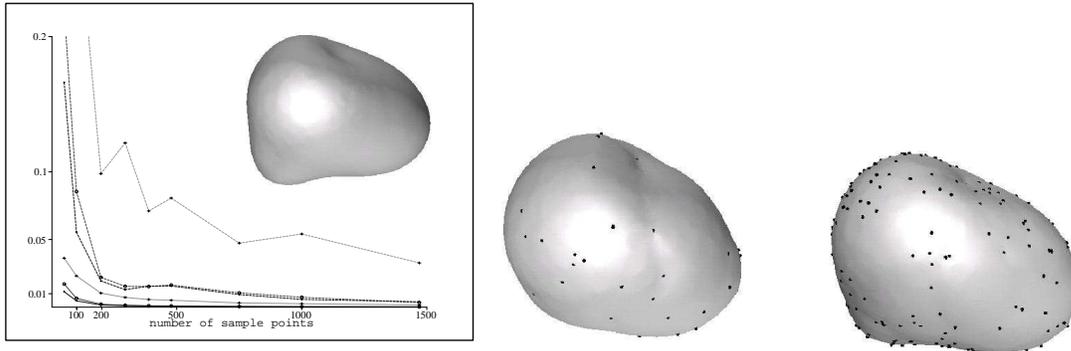
Figure 10: f_2 (a) exact model and error statistic, interpolation with (b) 50, (c) 250 sample points.
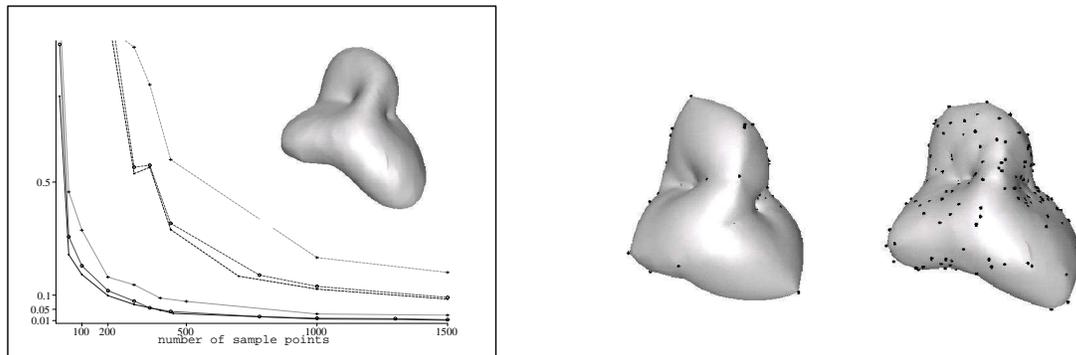
Figure 11: f_3 (a) exact model and error statistic, interpolation with (b) 50, (c) 250 sample points.

5.2.2   Functions on parameterized surfaces

In this section, we suppose that we are given a parameterized surface $\mathcal{S}$ and a function $h(u, v)$: $\mathcal{S} \rightarrow \mathbb{R}$ defined on $\mathcal{S}$. We evaluate $\mathcal{S}$ on a regular grid of 40000 points (which is for numerical reasons slightly perturbed). The function is evaluated on a random subset of the grid points (in parametric space). The gradient needs to be expressed with respect to the Cartesian coordinates $(g_i = (\frac{\partial u}{\partial x} \frac{\partial h}{\partial u}(u_i, v_i) + \frac{\partial v}{\partial x} \frac{\partial h}{\partial v}(u_i, v_i), \frac{\partial u}{\partial x} \frac{\partial h}{\partial u}(u_i, v_i) + \frac{\partial v}{\partial y} \frac{\partial h}{\partial v}(u_i, v_i), \frac{\partial u}{\partial z} \frac{\partial h}{\partial u}(u_i, v_i) + \frac{\partial v}{\partial x} \frac{\partial h}{\partial v}(u_i, v_i))$. We use interpolants I^0, I^1 and Z^1 to estimate the function value on the remaining grid points. Some

of the results as well as an error statistic are shown in figure 12. The surfaces are deformed by the function value in direction of the surface normal. The error statistic has the same interpretation as those in section 5.2.1.
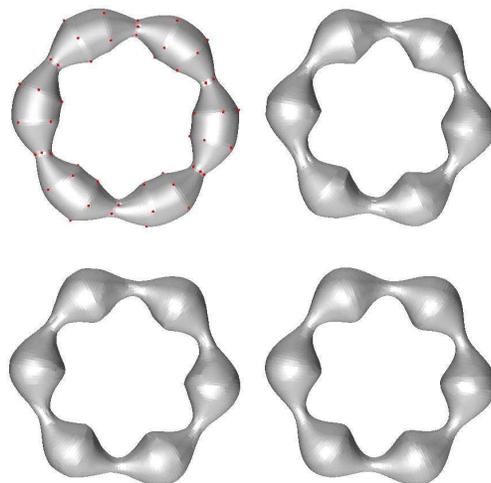


Figure 12: Interpolation of $h(u, v) = 0.6\cos(6v)$ on the torus with (a) 100 regularly spaced sample points (shown in red), (b) 500, (c) 1000 and (d) 4000 random sample points.
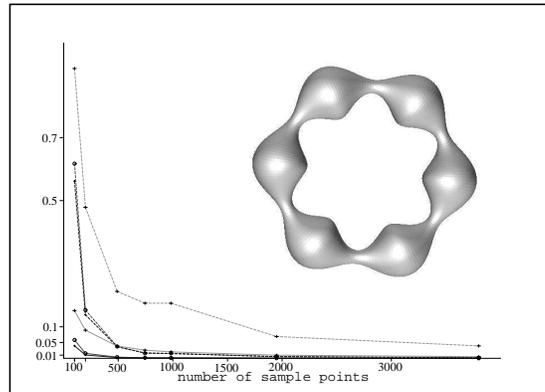

Figure 13: Exact model and error statistic of figure 12.

5.2.3   Vector fields on parameterized surfaces

The interpolation of vector fields, i.e. functions that are defined from the surface to a higher dimensional space, can be treated in the same way as scalar functions by interpolating each coordinate of the result separately with one of the interpolants presented in section 5.1. For example, let $v : \mathcal{S} \to \mathbb{R}^3$ with $v(u, v) = (v_x(u, v), v_y(u, v), v_z(u, v))$, then, v_x, v_y, and v_z are interpolated independently. The error of the interpolation is measured by the squared distance between the vector obtained by applying the function on a point X and the interpolation result at X. Figure 14 shows the interpolation of $v(u, v)$ with $v_x(u, v) = -cos(u)\ cos(v)$, $v_y(u, v) = -cos(u)\ sin(v)$, and $v_z(u, v) = -sin(u)$ on the cylinder. The setting is the same as in the previous section: the function as well as the function gradient is known on a subset of the 40000 grid points, and the interpolants are evaluated on the rest of the grid points. For surfaces with boundary, all boundary points are part of the sample points. To visualize, we translate each grid point by the vector resulting from the interpolation or the function application. Concerning the error statistics, figure 15, this time the mean error of the I^0 interpolant is worse than the maximum error of I^1 and Z^1. Other than that, the interpretation is the same as in the previous sections.
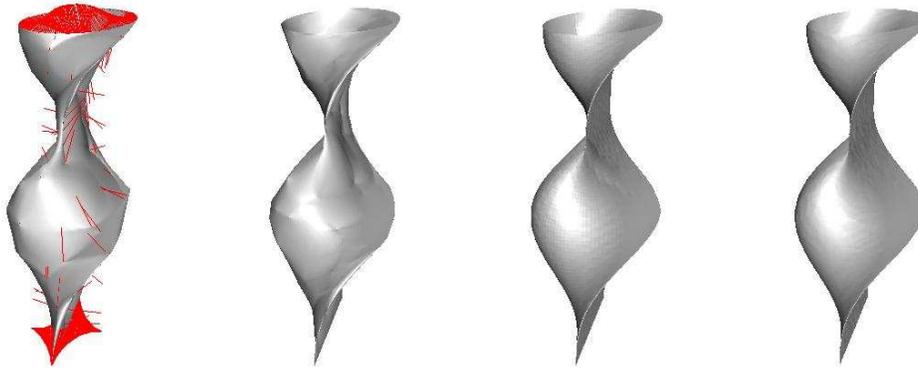
Figure 14: Interpolation of $h(u, v) = 0.6 \cos(6v)$ on the torus with (a) 100 regularly spaced sample points (shown in red), (b) 500, (c) 1000 and (d) 4000 random sample points.
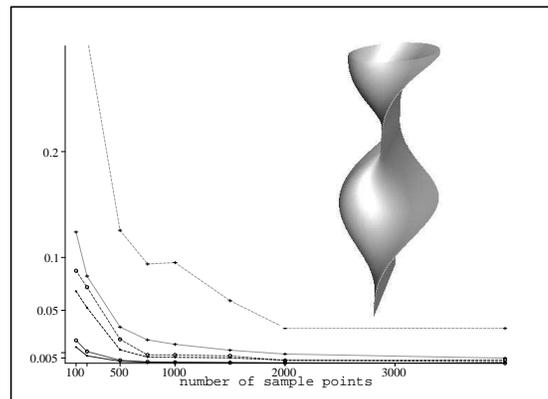
Figure 15: Exact model and error statistic for figure 12.

5.2.4 Visualization of the coordinate function

In order to demonstrate the locality and the smoothness of the T-coordinate, we visualize the coordinate function τ_i with respect to a sample point A_i on a parameterized surface $\mathcal{S}$. Figure 16 shows an example of the cylinder from two different viewpoints and of the torus. τ_i is computed on a (perturbed) regular grid of 40000 points with respect to a subset of 100 regularly placed sample points. Each grid point is translated by the value of the coordinate τ_i in direction of the surface normal at that point. Of course, A_i itself has the highest value for τ_i ($\tau_i(A_i) = 1$). Notice the locality of the coordinate function, as well as the C^1 discontinuity on the sample point A_i itself.
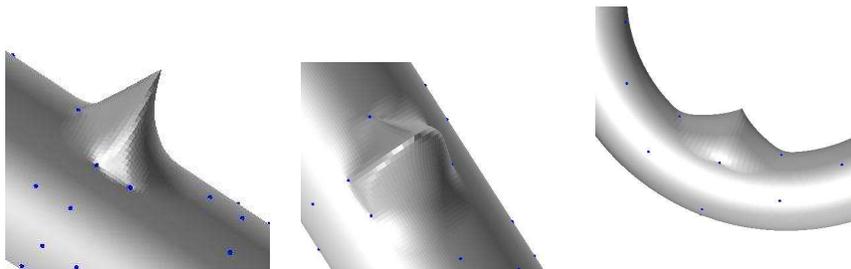
Figure 16: The coordinate function τ_i on the cylinder and the torus.

5.3   Implementation

The implementation is based on the Computational Geometry Algorithms Library $CGAL$ [11]. It makes use of the two-dimensional regular triangulation, the three-dimensional Delaunay triangulation as well as the polyhedral surface class provided by $CGAL$.

Table 1 indicates the running time for the examples depicted in section 5.2.2. The models contain 40000 grid points. The number of sample points varies from 100 to 4000. For comparison, the running time of the Delaunay triangulation of the 40000 grid points is shown in the second column. The remaining columns show the running time for the interpolation on the 40000 grid points with respect to n sample points, n= 100, 500, 1000, 2000, and 4000. The experiments were run on a PC computer with Pentium III bi-processors at 730 MHZ, 640 MB main memory. Further details on the implementation can be found in a companion paper [7] that will appear soon.

Model	DT	100	500	1000	2000	4000
Cylinder	468	227	362	439	551	766
Sphere	309	566	571	496	445	392
Torus	277	185	285	333	394	470

Table 1: Running time in seconds.

6   CONCLUSION AND FUTURE WORK

The T-neighborhood and the T-coordinate system are associated to a set of sample points from a surface equipped with normal vectors. We show that the neighbors as well as the coordinate system are local on the surface when the surface is well sampled. The coordinate system is continuous everywhere on the surface. It is continuously differentiable except at a finite number of surface points. Furthermore, we show the application of the coordinate system to different problems related to interpolation. In the near future, we plan several extensions of our method: In order to simplify the presentation in this framework, we restricted ourselves to $(d-1)$-manifolds, yet, the same definitions apply to all smooth k-manifolds, $k < d$, of $\mathbb{R}^d$. For the same reason, we did not generalize the definitions to manifolds with boundary. Also, in [19] coordinate systems with higher continuity were defined which generalize Sibson's coordinates. The same technique can be applied to T-coordinates. In addition, we will concentrate on the question of efficiency: due to the proved locality of the T-neighbors, efficient filtering methods will considerably speed up computation time. See [13] for relevant work on this question. Last but not least, we plan to develop further applications e.g. around texture mapping, surface reconstruction, and point set surfaces as they are recently emerging in computer graphics, see for example [1].

Acknowledgements

References

[1] Marc Alexa, Johannes Behr, Daniel Cohen-Or, David Levin, Shachar Fleishman, and Claudio T. Silva. Point set surfaces. In *Proceedings of Visualization 2001*, to appear 2001.

[2] P. Alfeld, M. Neamtu, and L. Schumaker. Fitting scattered data on sphere-like surfaces using spherical splines, 1996.

[3] Peter Alfeld, Marian Neamtu, and Larry L. Schumaker. Bernstein-bézier polynomials on spheres and sphere-like surfaces, 1996.

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

[5] R. Barnhill, B. Piper, and K. Rescorla. Interpolation to arbitrary data on a surface. In G. Farin, editor, *Geometric Modeling: Algorithms and New Trends*, pages 281–289. SIAM, Philadelphia, 1987.

[6] J-D. Boissonnat and F. Cazals. Natural neighbor coordinates of points on a surface. *Computational Geometry: Theory and Applications*, 19(2-3):155–173, 2001.

[7] Jean-Daniel Boissonnat and Julia Flötotto. Robust and generic implementation of a surface coordinate system and applications. in preparation.

[8] Jean-Daniel Boissonnat and Mariette Yvinec. *Algorithmic Geometry*. Cambridge University Press, UK, 1998. Translated by Hervé Brönnimann.

[9] J. L. Brown. Natural neighbor interpolation on the sphere. In P.-J. Laurent, A. Le Méhauté, and L. L. Schumaker, editors, *Wavelets, Images, and Surface Fitting*, pages 67–74. A. K. Peters, Wellesley MA, 1994.

[10] J. L. Brown. Systems of coordinates associated with points scattered in the plane. *Comput. Aided Design*, 14:547–559, 1997.

[11] The CGAL reference manual, August 2001. Release 2.3. http://www.cgal.org.

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

[13] Tamal K. Dey, Stefan Funke, and Edgar A. Ramos. Surface reconstruction in almost linear time under locally uniform sampling. In *Abstracts 17th European Workshop Comput. Geom.*, pages 129–132. Freie Universität Berlin, 2001.

[14] H. Edelsbrunner and N. R. Shah. Triangulating topological spaces. In *Proc. 10th Annu. ACM Sympos. Comput. Geom.*, pages 285–292, 1994.

[15] G. Farin. Surfaces over Dirichlet tesselations. *Comput. Aided Geom. Design*, 7:281–292, 1990.

[16] G. Fasshauer and L. Schumaker. Scattered data fitting on the sphere. In T. Lyche M. Daehlen and L. L. Schumaker, editors, *Methods for Curves and Surfaces II*, pages 117–166. Vanderbilt University Press, 1998.

[17] T. Foley. Interpolation to scattered data on a spherical domain. In J. C. Mason and M. G. Cox, editors, *Algorithms for Approximation II*, pages 303–310. Chapman and Hall (London), 1990.

[18] Thomas A. Foley, David A. Lane, Gregory M. Nielson, Richard Franke, and Hans Hagen. Interpolation of scattered data on closed surfaces. *Computer Aided Geometric Design*, 7(1-4):303–312, 1990.

[19] Hisamoto Hiyoshi and Kokichi Sugihara. Voronoi-based interpolation with higher continuity. In *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, pages 242–250, 2000.

[20] Greg Leibon and David Letscher. Delaunay triangulations and Voronoi diagrams for Riemannian manifolds. In *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, pages 341–349, 2000.

[21] Marian Neamtu. Splines on surfaces. http://citeseer.nj.nec.com/neamtu01splines.html, 2001.

[22] B. Piper. Properties of local coordinates based on dirichlet tesselations. *Computing Suppl.*, 8:227–239, 1993.

[23] R. Sibson. A vector identity for the Dirichlet tesselation. *Math. Proc. Camb. Phil. Soc.*, 87:151–155, 1980.

[24] R. Sibson. A brief description of natural neighbour interpolation. In Vic Barnet, editor, *Interpreting Multivariate Data*, pages 21–36. John Wiley & Sons, Chichester, 1981.

[25] F. E. Wolter. Cut locus and medial axis in global shape interrogation and representation. Technical Report 92-2, MIT, Dept. Ocean Engg., Design Lab, Cambridge, MA 02139, USA, January 1992.

Contents

6 CONCLUSION AND FUTURE WORK   26